

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, revised October 2024

We are grateful to Vittorio Bassi, Susanna Berkouwer, Chris Blattman, Josh Dean, Kyle Emerick, Xavier Gine, Rachel Glennerster, Michael Greenstone, Faraz Hayat, Koichiro Ito, Kelsey Jack, Florence Kondylis, Michael Kremer, Jeremy Magruder, Craig McIntosh, Mushfiq Mobarak, Andrew Robertson, Mark Rosenzweig, 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, Y-RISE, and NBER Development 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, Rathana 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Fiona Burlig, Amir Jina, Erin M. Kelley, Gregory V. Lane, and Harshil Sahai. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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, revised October 2024

JEL No. D81,O12,O13,Q12,Q54

ABSTRACT

Climate change increases weather variability, preventing farmers from tailoring investments to the upcoming monsoon. In theory, accurate, seasonal forecasts overcome this challenge. We experimentally evaluate monsoon onset forecasts in India, randomizing 250 villages into control, forecast, and benchmark insurance groups. Forecast farmers update their beliefs and behavior: receiving “good news” relative to a farmer’s prior increases cultivation, farm inputs, farm profits (for those unaffected by flooding) and reduces business; receiving “bad news” reduces cultivated land and farm profits but increases business. Overall, forecasts raise a welfare index by 0.06 SD. Unlike insurance, forecasts reduce climate risk by enabling tailoring.

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

Gregory V. Lane
Harris School of Public Policy
University of Chicago
1307 E. 60th Street
Chicago, IL 60637
and NBER
gvplane@gmail.com

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

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

Erin M. Kelley
World Bank
erinmkelley@worldbank.org

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)). Farming 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)). Absent full insurance, weather risk causes farmers to make fewer profitable investments (Rosenzweig and Binswanger (1993); Donovan (2021)). Moreover, when 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.

In this paper, we use a cluster-randomized experiment to estimate the causal effects of a novel approach to helping farmers cope with 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, enabling them to decide how much land to cultivate, adjusting the crop mix, ordering inputs in advance, and changing off-farm activities to maximize overall welfare. Our empirical results focus on a forecast that provides information about when the Indian Summer Monsoon — whose variability is predicted to increase under climate change (Katzenberger et al. (2021); Prabhu and Chitale (2024)) — will arrive. Using historical data, we show that monsoon onset timing is an important determinant of agricultural outcomes in India, with earlier monsoons (i.e., longer growing seasons) leading to higher yields both on average and for cash crops in particular, which can impact both on- and off-farm decisions. However, the timing of monsoon onset is highly variable, making these features difficult for farmers to predict and optimize around in the status quo (Kumar et al. (2013)).

We overcome this challenge with a novel forecast of the onset of the monsoon. While accurate forecasts could enable substantial behavior change (FAO (2019)), their use is not widespread in practice, in part because existing forecasts of this type have limited accuracy (Mase and Prokopy (2014); Rosenzweig and Udry (2019)), and because forecast dissemination in low-income countries is traditionally poor (Yegbemey and Egah (2021)).¹ In contrast, our forecast, first described in Stolbova et al. (2016) and published by the Potsdam Institute for Climate Impact Research (PIK), is extremely accurate, locally-resolved, and can be provided to farmers well in advance of the monsoon’s arrival.² Released approximately 40 days before onset, the PIK forecast enables farmers to make early decisions about key farm inputs such as land use, crops, labor supply, and fertilizer

¹Seasonal forecasts are likely to be particularly valuable for the nearly two-thirds of the global population living in the monsoonal climate systems that broadly characterize the tropics (Wang et al. (2021)). Moreover, 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 growing-season level, such as what crops to plant and how much land to cultivate.

²This forecast relies on recent improvements in weather modeling (e.g., Rajeevan et al. (2007)), and statistically identifies “tipping points” that are relevant for monsoon rainfall onset in a particular location.

purchases (Gine et al. (2015)), as well as how much to engage in non-agricultural business (as many farmers rely on other sources of income). This forecast has particular accuracy over Telangana, the site of our experiment: in this region, the forecasted onset date has been accurate to within one week in each of the past 10 years prior to our experiment.

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. The forecast reduces risk by providing farmers with information about the upcoming growing season, allowing them to tailor their inputs accordingly. In contrast, insurance – the canonical risk-coping instrument– enables farmers to shift consumption across states but provides no information, making it a useful comparison. We sample 5-10 farmers per village for inclusion in the experiment. To avoid bias from spillovers, all main sample farmers in a given village receive the same treatment. To ensure that farmers view the forecast as credible, we partner with the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT), a respected international organization based in Hyderabad.

We ask four main research questions. First, how does the forecast change farmers’ beliefs about monsoon onset? We anticipate that receipt of the forecast should shift farmers’ beliefs closer to the forecasted monsoon date.

Second, how does the forecast impact farmers’ agricultural inputs and farm outcomes? To generate predictions, we develop a theoretical model which shows that the effects of a forecast should depend on farmers’ prior beliefs. In particular, if the forecasted onset aligns with a farmer’s prior (such that the farmer receives “neutral news” from the forecast), they should not substantially adjust their input decisions. If the forecasted onset is earlier (later) than the farmer’s prior—henceforth “good news” (“bad news”)—the profit-maximizing farmer should respond by increasing (reducing) their investments in risky agricultural production. To generate a prediction about what should happen in this setting we use historical data on crop yields and rainfall to show that an earlier monsoon onset (i) leads to higher yields of two key crops, staple rice and cash crop cotton, and (ii) differentially improves cotton yields. We therefore expect good-news farmers to increase overall agricultural investment and to shift into cash cropping, and bad-news farmers to invest less in agriculture.

We predict that agricultural outputs should generally align with on-farm investments, such that higher investments should lead to increased production, yields, and farm profits on average. However, as Rosenzweig and Udry (2020) point out, agriculture is an inherently stochastic process, meaning that in any given year, there may be an imperfect correspondence between inputs and outputs. Growing season shocks that are unrelated to the forecast (e.g., extreme rainfall or drought) could therefore cause a mismatch between farm inputs and outputs.

Third, how do forecasts impact non-agricultural business? Forecasts likely enable farmers to optimize their off-farm activity to the coming rainfall realization, as Rosenzweig and Udry (2019) point out the broad economy is sensitive to the monsoon. However, the direction of our theoretical prediction is ambiguous, and depends on whether off-farm profits are positively or negatively correlated with agriculture.

Fourth, how do forecasts impact farmer well-being? Theory predicts that forecasts should (weakly) improve welfare for all farmers, regardless of prior beliefs. By enabling farmers to tailor their activities and optimize their investments, accurate information should lead to improvements in standard welfare metrics such as consumption, asset ownership (including livestock), and savings net of debt, noting that in rural villages, saving excess cash rather than and/or avoiding high-interest-rate loans is often welfare-improving.

Our empirical results broadly follow these predictions. Figure 4 presents our results on beliefs, and Figure 5 summarizes our main findings on agricultural inputs and outputs, non-agricultural business, and aggregate welfare. First, farmers update their beliefs in response to our forecast.³ After receiving the information, farmers in forecast villages have beliefs about the monsoon onset date that are 26% closer to the forecasted onset date than farmers in control villages. In addition, farmers demand the forecast. Using a Becker et al. (1964) mechanism (henceforth “BDM”) to elicit willingness-to-pay (WTP), we find that average WTP for the forecast is comparable to the average WTP for our index insurance product.⁴

Second, farmers alter their agricultural investments—land use, crop choice, and input expenditures—in response to the forecast. Farmers who receive bad news meaningfully reduced land under cultivation (−22% of the control mean), with a point estimate suggesting an approximately 10% decline in expenditure. Farmers who receive neutral news did not alter their investments. Farmers who receive good news increased land under cultivation (21%) and total expenditure considerably (31%), and were 33% percent more likely to plant cash crops. These results are consistent with farmers optimizing in response to the forecast, with bad-news farmers avoiding investing too much in agriculture during a worse-than-expected season, and good-news farmers taking advantage of a better-than-expected season by investing more than they would have otherwise. Summarizing these outcomes in an index, we find that farmers who receive bad news reduce investment by 0.08 SD relative to farmers *with similar priors* in the control group; we see no impact on neutral-news farmers’ investments; and good-news farmers increase agricultural investment by a standardized effect of 0.31 SD (while the former is imprecise, the difference between bad- and good-news responses and the good-news effect itself are highly significant).

These meaningful and agronomically-appropriate investment changes in turn impact farm production. Bad-news farmers reduced their agricultural output (and the value of this output) by 25% (22%), in line with these farmers’ reduction in land under cultivation, and neutral-news farmers experience no change in agricultural output, consistent with their (lack of) input response. While imprecisely estimated, good-news farmers increased their agricultural output by 22%, also in line with their investment effects. We fail to reject zero yield effects for all groups.

However, we find an incomplete mapping between agricultural input decisions and profits. Bad-news farmers have even lower farm profits (−\$400, or −40%, s.e. \$236) than their input changes would

³There is substantial within-village heterogeneity in priors, with 46% of the total variation in mean prior remaining after removing village fixed effects.

⁴We 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.7).

imply. Similarly, we see no statistically significant impacts on neutral- and good-news farmers, but the good-news point estimate in particular is quite close to zero (-\$65, or -6%, s.e. \$345). Thus, while we do find that reductions (increases) in investments led to changes in output, these did not perfectly translate into changes in agricultural profit.

The gap between farmers’ agricultural investments and their farm profits suggests that they may have been negatively affected by a shock that was unrelated to the forecast. Though the forecast correctly predicted the timing of the monsoon’s arrival, heavy flooding – outside of the scope of the forecast – hit Telangana in early July (Business Line (2022); The New Indian Express (2022)). The likelihood of flood exposure was balanced between our forecast group and our treatment group. Nevertheless, we estimate that good-news farmers report crop losses of \$200 more than their control counterparts, consistent with this group having planted more valuable crops. Though estimates are imprecise, if we add the value of this lost output to our observed profits, we estimate that profits including losses declined by only 17% in the bad-news group, increased by 6% in the neutral-news group, and increased by 9% in the good-news group – much more closely aligned with our observed input changes. In a final analysis among only those farmers who were not flood-affected, we estimate a pattern of agricultural profit treatment effects (-\$341 for bad-news farmers, -\$96 for neutral-news farmers, and +\$498 for good-news farmers) that is consistent with the agricultural outcomes we would expect the forecast to generate in the absence of a shock.

Next, we measure treatment effects on off-farm business activity. All of our estimates are noisy, but we find patterns consistent with forecast farmers treating non-agricultural businesses as a substitute for agriculture. Bad-news farmers increase non-farm business operation by 43%, while good-news farmers decrease business operation by 36%. While neither point estimate is statistically different from zero, we reject equality between the two estimates. This pattern is mirrored in non-agricultural investment and business profits. Bad-news farmers increase investment by 17%, translating into an \$80 increase in profits. In contrast, good-news farmers reduce business investments by 78% on average, with a corresponding reduction in profits of \$44.

Fourth, we measure impacts on farmer welfare. In doing so, we pool across news types, since our theory predicts improvements for all groups regardless of prior. Forecasts increase per-capita food consumption by 7%. While imprecise, our estimates imply that non-food consumption somewhat, asset value and net savings rise, and we find no change in livestock. We summarize these results in an index, estimating an overall welfare improvement of 0.06 SD ($p < 0.05$). If anything, effects are largest for bad-news farmers (0.14 SD, s.e. 0.06), null for neutral-news farmers (0.00 SD, s.e. 0.04), and weakly positive for good-news farmers (0.05 SD, s.e. 0.05), consistent with bad- and good-news farmers having made investment changes in response to the forecast.⁵ Taken together, these results suggest that – even in the presence of a meaningful orthogonal shock – forecasts increased overall farmer welfare.

As a final exercise, we compare the impacts of the forecast against index insurance. First, we

⁵For context, though the components that make up each index varies by paper, our estimate falls between the value of emergency loans (Lane (2024), 0.02 SD) and that of irrigation access (Jones et al. (2022), 0.11 SD).

reproduce the standard result that index insurance causes farmers to increase investment (0.13 SD increase in our investment index). Second, and per our model, we demonstrate that these effects are concentrated among “optimistic” farmers with early priors – those for whom the forecast would have been bad news (0.18 SD increase with insurance vs. 0.08 SD decrease under the forecast). Third, unlike good-news forecast farmers, insurance farmers are only moderately more likely to grow cash crops – and we can reject equality between these two groups, demonstrating that insurance is unlikely to have a tailoring effect. Fourth, we only find positive effects on agricultural profits for insurance farmers who were unaffected by floods, despite impacts on farm investments. This corroborates our agricultural profit finding for forecast farmers, further highlighting the importance of stochasticity in determining agricultural outcomes (Rosenzweig and Udry (2020)). Fifth, insurance farmers increase non-agricultural business activity, consistent with insurance increasing investment but not enabling tailoring. Finally, we find a null result on our well-being index (-0.06 SD, $p > 0.63$). These results show that while both forecasts and insurance influence farmer investments and resulting outcomes, they do so in fundamentally different ways. Forecasts reduce climate risk by enabling farmers to tailor their inputs to the coming growing season, whereas insurance leads to increased investment, but does not induce tailoring.

Taken together, our results demonstrate that long-range monsoon forecasts can help 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.⁶ 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.⁷ Our treatment causes farmers to update their beliefs in the direction of the forecast, resulting in meaningful changes in both investments and outcomes. Our results shed light on the mechanism through which forecasts work: enabling farmers to tailor their behavior to the coming monsoon 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 seminal work by Rosenzweig and Udry (2019), who use a farmer fixed-effect design to study the Indian

⁶See 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 agriculture, 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)) and longer-range forecasts of the El Niño Southern Oscillation (Shrader et al. (2023); Downey et al. (2023)) and the global climate (Schlenker and Taylor (2021)), all of which highlight the value of forecasting under climate change.

⁷In doing so, we build on Patel (2024), who estimates a model of farmer learning in the context of salinity forecasts in Bangladesh to shed light on how beliefs are formed *in situ*.

Meteorological Department’s (IMD) monsoon forecast, and argue that while the IMD’s forecast has remarkably low accuracy, an accurate long-range forecast of the Indian summer monsoon has the potential to be worth tens of billions of rupees.⁸ Our results highlight the promise of such a forecast, which will only become more valuable as the climate continues to change.

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 how much land to cultivate, what to plant, how to adjust inputs, and whether to engage in non-farm business. 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 (Lybbert and Sumner (2012); Karlan et al. (2014); Cole and Xiong (2017); Carter et al. (2017a)). We show that this canonical approach allows farmers to smooth risk across states of the world, but does not enable tailored investment.⁹ 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.¹⁰ Additionally, as forecasts can be widely distributed via SMS, and because we find farmers exhibit demand for this information, forecasts have the potential to be adopted at scale.

Finally, by empirically demonstrating the effectiveness of a specific climate adaptation technology – the forecast – we advance the 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 smaller body of newer work which highlights the importance of adaptation (e.g., Hultgren et al. (2022); Carleton et al. (2022)) but does not examine the role of specific adaptation strategies.¹¹ In contrast, we experimentally evaluate a broadly 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, raising welfare.

⁸In 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.

⁹A 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)).

¹⁰Prior work has focused on the low demand for insurance (Mobarak and Rosenzweig (2014); Jensen and Barrett (2017); Carter et al. (2017b)), highlighting the large subsidies needed to increase uptake. 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.

¹¹A 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 infrastructure and can be disseminated at low cost.

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 7 compares forecasts to insurance. Section 8 concludes.

2 Research context

Two thirds of the global population, including most of India, lives in a monsoonal climate system (Wang et al. (2021)), characterized by an intense wet season during which most of the year’s rain falls and a dry season with limited precipitation. Agriculture in these regions is thus heavily linked with the rainy season. However, the timing of the onset of the rainy season is highly variable, making it difficult for farmers to optimize.

2.1 Historical impacts of monsoon onset timing on agricultural yields

We begin by documenting the importance of monsoon onset timing for Indian agriculture. Prior work has shown that 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)). We build on this work, using historical data on rainy-season (kharif) agriculture across India to document both that monsoon onset delays negatively impact crop production, and also that these damaging impacts are substantially worse for cotton – a key cash crop in our setting – than for rice – a key staple crop.

Specifically, we use district-level yield data across the country from the Indian Ministry of Agriculture and Farmers’ Welfare and daily gridded precipitation data from the European Centre for Medium Range Weather Forecasting Reanalysis data (ERA-5) spanning 2001 to 2018 to estimate the historical effect of monsoon onset delay on crop yields.¹² We estimate a simple panel fixed effects regression, with a preferred specification of:

$$\log(\text{Yield})_{dy} = \beta \text{Onset}_{dy} + \alpha_d + \delta_t + \varepsilon_{dt} \quad (1)$$

where the outcome variable is log yield of cotton or rice in district d in year y , Onset_{dy} is standardized onset, α_d are district fixed effects, δ_t are year fixed effects, and ε_{dt} is an error term, clustered at the state level. Table 1 presents the results, including robustness to alternative specifications, with our preferred specification in Column (2). The identifying assumption – similar to that in a large literature on the impacts of weather on agricultural outcomes (e.g., Deschênes and Greenstone (2007), Schlenker and Roberts (2009), Hultgren et al. (2022)) – is that, conditional on location- and time-specific fixed effects, monsoon onset timing is plausibly exogenous.

¹²Appendix B provides more detail about the data and estimation. We define monsoon onset following Moron and Robertson (2014), and restrict the sample to districts in the monsoonal region of India (excluding the northern, southern, and eastern tips of the country).

Table 1: Effect of monsoon onset timing on rice and cotton yield

	(1)	(2)	(3)
	Log(Yield)	Log(Yield)	Log(Yield)
Panel A: Rice			
Onset (std. dev.)	-0.024** (0.011)	-0.016** (0.008)	-0.039*** (0.012)
Panel B: Cotton			
Onset (std. dev.)	-0.047*** (0.024)	-0.038** (0.025)	-0.092*** (0.046)
N (rice)	2321	2321	2321
N (cotton)	1098	1098	1098
State FEs	Yes		
District FEs		Yes	Yes
Year FEs	Yes	Yes	
State $\times$ Year trend			Yes

Notes: This table presents the effect of monsoon onset timing on yields of rice (panel A) and cotton (panel B), estimated using Equation (1). The outcome in each column is crop yield in logs, and the independent variable is monsoon onset in standard deviations, both observed at the district-by-year level. Higher onset values indicate later monsoon arrival. We define monsoon onset per Moron and Robertson (2014), and restrict the sample to monsoonal regions of India (see Appendix B for more details). Standard errors are clustered by state. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Later monsoon onset leads to lower district-level yields for both rice (Panel A, 1.6 percent decline per SD of onset delay in our preferred specification) and cotton (Panel B, 3.8 percent decline per SD of onset delay). Importantly, the yield decline is 2 to 3 times larger for cotton than rice. With an average onset standard deviation of approximately 1.5 weeks, a monsoon that arrives three weeks late would cause rice yields to fall by 3.2% and cotton yields to fall by 7.6%. These results are robust across three different specifications; if anything, changing fixed effects increases our point estimates. These results have two main implications, which we take to the data from our experiment below. First, we predict that farmers who expect to face an earlier monsoon should increase their investments in agriculture. Second, these results suggest that farmers who expect to face an earlier monsoon ought to increase their investments in cash crops in particular. While this analysis uses 20 years of historical data to show how farmers ought respond to an earlier monsoon *on average*, whether this strategy would lead to higher yields and profits in any given year will depend on unpredictable weather conditions (Rosenzweig and Udry (2020), Suri and Udry (2022)).

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 improve farmer welfare, because they enable farmers to materially

alter investment decisions such as land use, crop mix, input expenditure, and non-farm activities in advance of the monsoon’s arrival.

We rely on a new long-range forecast of the monsoon’s onset produced by PIK, and described in Stolbova et al. (2016).¹³ 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.¹⁴ 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 (Gine et al. (2015)), in addition to changing off-farm investments. 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.¹⁵ Moreover, the IMD forecast only arrives 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 national 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 input changes.¹⁶

¹³See Appendix H for more details on monsoon forecasting.

¹⁴At 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.

¹⁵Unlike 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.

¹⁶Seasonal 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.

2.3 Agriculture in Telangana

We conduct our experiment in Telangana. 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, compared to the national average of 46%, the sector provides only 15% of the Gross State Value Added, equal to the national average (Government of Telangana (2020); Ministry of Agriculture and Farmers Welfare (2023)). 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. Appendix Figure A.1 demonstrates that there are substantial year-over-year fluctuations in both the amount of land under cultivation and in the share of land planted to rice, cotton, and other crops over time, making the state an ideal place to evaluate the effect of forecasts on crop land and crop choice.

Telangana, like much of central India, is dependent on the monsoon for agriculture, with approximately 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. Consequently, weather risk is a substantial concern for agriculture in the state, as it rests in one of the most variable areas of India’s monsoonal region. Farmers receive some government assistance, though both formal and informal insurance are far from full. 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 ceased to operate in the state.¹⁷ Private insurance exists, but is severely underutilised. At baseline, only 0.75% of farmers in our sample had heard of rainfall insurance.

In the status quo, farmers’ information about the weather is also limited. While 65% of farmers in our sample report having received some information about the upcoming Kharif season at baseline (conducted prior to planting in early May; see Figure 2), the reliability of their sources is unclear. Appendix Figure A.2 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%).

¹⁷The program initially required all agricultural loan-holders to purchase insurance, but when the government subsequently made this condition voluntary, demand collapsed.

3 Model

In this section we present a simple two-period model of farmers' decision-making under risk, which we use to illustrate the effects of the monsoon forecasts and insurance product.¹⁸ 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 $i \in \{1, \dots, S\}$, 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 $x > 0$.¹⁹ 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.²⁰

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

$$\begin{aligned} \max_{s, x} \quad & u(c_1) + \beta \sum_{i=1}^S u(c_2^i | \epsilon_i) g(\epsilon_i) \\ \text{s.t.} \quad & c_1 = y - s - p \cdot x \text{ \& } c_2^i = f(x, \epsilon_i) + s \end{aligned} \tag{2}$$

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 C.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


¹⁸We provide extended model details in Appendix C.

¹⁹For 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, as shown by Amale et al. (2023) and confirmed by our own data in Section 2.

²⁰The 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.

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.

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



Notes: This figure plots the simulated relationship in our model between the treatment effect of forecasts and insurance on optimal investment and the farmer's prior. 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 when farmers believe the monsoon will arrive. The figure shows differential investment responses between the forecast, as shown here for an average monsoon (purple), and the insurance product (black) for farmers with different priors. See Appendix C.3 for simulation details.

Figure 1 illustrates the key results of the model, plotting the treatment effect of a forecast (purple) on agricultural investment. This figure depicts responses to a forecast of an average monsoon.²¹ In this case, 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. Farmers with late (and therefore overly-pessimistic) priors receive good news from the forecast, and increase their investment. This figure therefore illustrates how forecasts help farmers to tailor their behavior to the coming growing season.

We also plot responses to an insurance product (black), which delivers a payout in sufficiently bad states. Regardless of farmers' priors on the upcoming growing season conditions, insurance – which conveys no information about the state – causes all farmers to weakly increase their

²¹Appendix Figure C.1 plots farmer responses to forecasts of early or late monsoons, under which the shape of the treatment effects is broadly preserved but the level shifts.

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. This highlights the different mechanisms behind forecasts and insurance.

Finally, while our discussion of the model has centered around predictions about agricultural input, we also predict that welfare should (weakly) rise for all forecast farmers, regardless of prior beliefs. Early-prior farmer welfare is expected to rise because their returns to agriculture should go up relative to the control, while late-prior farmer welfare is expected to rise because they can reallocate money that would have gone to agricultural investment into other forms of profitable investments, such as non-agricultural business, assets (including livestock) or into savings, both of which can be consumed in the second period.

4 Experimental design and data

4.1 Experimental design

Informed by our theoretical framework, we designed a randomized controlled trial to estimate the impacts 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).

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 and insurance being offered in our experiment.

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 E.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 the forecast, the enumerator would provide the farmer with the following information (which predicts an average monsoon):

"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 each farmer a follow-up SMS with the same text.

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 H.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)).²² These payout thresholds were set only using historical data, and were fully independent of the forecast.

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)).²³ Farmers in the insurance treatment arm received an information sheet covering these details (Figure E.2). As with the forecast product, we offered farmers this insurance product through a BDM mechanism in order to elicit willingness-to-pay, which we describe in more detail

²²See 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.

²³For this calculation, as for all others in the paper, we use an exchange rate of \$1 = INR 82.

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).²⁴ We present takeup of the forecast and insurance product in Appendix Figure A.3 and Appendix Table A.4. Takeup is over 85 percent for both treatment groups.²⁵ 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. 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. In addition, this figure shows evidence of heavy rainfall in some areas during July, consistent with popular press reports of flooding during this time period (Business Line (2022); The New Indian Express (2022)).

4.2 Data


Outcome data We collected detailed data on four main categories of outcomes: beliefs, agricultural investment and output, non-farm business, and welfare.

Our first outcome of interest is farmer beliefs about the arrival of the coming monsoon. 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

²⁴For more details on our BDM, which was modeled on Berkouwer and Dean (2022), see Appendix D.

²⁵Appendix Figure A.3 and Appendix Table A.3 document that the later a farmer thinks the monsoon is likely to be, the more likely they are to purchase each product when offered.

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.


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 both 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.5 demonstrates that prior beliefs matter—in the direction predicted by theory—for farmers’ agricultural investment decisions. Among farmers in the control group only, we find that later priors are associated with less land under cultivation, a lower probability of planting cash crops, lower total input expenditure, and reductions in a standardized investment index.²⁶

The second main category of outcomes are farmers’ agricultural investment decisions and resulting farm outputs. We consider a number of choices that may be affected by our treatments, including the amount of land cultivated, crop mix, and expenditures on inputs such as labor, seeds,

²⁶Though we do not have sharp predictions on the relationship between prior beliefs and farmer characteristics, we present these correlations in Appendix Table A.6. We find that households with heads residing in their birth village have later priors, while farmers with more landholdings have earlier priors.

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 kertes (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.

fertilizer, and irrigation. 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. In terms of farm production, the primary outcomes of interest are agricultural output, value of production, yield, and farm profits.

Third, we measure farmers’ engagement in non-agricultural business. We collect data on whether a farmer operates a non-farm enterprise, how much money they invest in this enterprise, and business profits.²⁷

Finally, we collect several measures of economic well-being. First, we measure per-capita spending on food and non-food consumption.²⁸ Next, we capture the value of a farmers’ assets and the number of livestock in their possession. Third, we measure savings and debts, both indicators of farmers’ financial position.²⁹

²⁷We collect off-farm business data over the last 30 days to increase recall, and scale these estimates to the length of the growing season for comparability with the agricultural outcomes.

²⁸We exclude medical spending from our main welfare analysis, as it is unclear whether increased medical spending implies higher (e.g., households can afford to spend more money on treatment) or lower (e.g., households have negative health shocks that require expenditure) welfare.

²⁹In addition to these standard welfare metrics, we consider impacts of our treatments on mental health, using the PHQ-8 screening tool, a standard and locally-validated depression metric (Bhat et al. (2022)). We also measure migration by capturing how many individuals from the household migrated elsewhere over the cropping season and the value of remittances they sent home.

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).³⁰ 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*. Our pre-registered analysis includes splitting our forecast treatment effects by prior beliefs. 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 F.

5 Forecast effects on beliefs

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 posterior (measured during baseline round II) vs. prior (measured during baseline round I) beliefs in the forecast treatment group against the control and insurance groups. Since the insurance group did not receive the forecast, it serves as a

³⁰Of 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 2: Effect of the forecast and insurance on beliefs

	(1) posterior – forecast	(2) posterior – prior	(3) K-S Stat
Forecast	-0.180** (0.083)	-0.239** (0.094)	-0.050* (0.027)
Insurance	-0.024 (0.096)	-0.095 (0.111)	-0.020 (0.032)
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 (3). 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, enumerator 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 G.19.

placebo. Specifically, we estimate:


$$Y_{iv} = \beta_0 + \beta_1 \text{Forecast offer}_v + \beta_2 \text{Insurance offer}_v + \gamma \mathbf{X}_{iv} + \eta_{iv} \quad (3)$$

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, enumerator fixed effects, and a set of controls chosen by double-selection LASSO, and η_{iv} is an error term, clustered at the village level.³¹

Table 2 presents the results. We find that the absolute difference between the forecast and the prior is 26% 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 (27% lower than control, Column 2) and in the Komolgorov-Smirnov test (11% 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 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. An important caveat to these posterior belief estimates is that these measurements were conducted between baseline round I and the expected monsoon arrival. Timing was dictated by survey logistics, rather than by the timeline of farmer decision-making. Thus, we

³¹Because 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 G.6, 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 kertes (a local unit of time that is approximately two weeks long), and elicited via the beans tasked described in Section 4. We then plot the mean of each farmer’s prior and posterior distributions. 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 (3), estimated using the posterior mean as the outcome variable (coefficient of -0.197 and p -value 0.031 without controlling for prior beliefs, coefficient -0.195 and p -value 0.033 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: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

take these results simply as evidence that farmers did in fact update their beliefs in response to the monsoon forecast, and do not use them for further analysis.

Willingness-to-pay We find that farmers’ willingness to pay (WTP) for forecasts is similar to their WTP for insurance, which provides \$190 in case the monsoon is delayed by 30 days or more, suggesting that farmers find forecasts valuable, though WTP for both products is relatively low (Appendix Figure A.5). Lending credence to our WTP measure, we find suggestive evidence that the strength of farmers’ priors is negatively correlated with WTP (Appendix Table G.5).³² Nevertheless, we interpret these results with some caution. As forecast information is a public good which can be readily disseminated within the village, farmers may offer a lower price in the BDM game compared to their true valuation.

³²In 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.

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 nobody did). Appendix Table A.7 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.).

6 Forecast effects on agriculture, off-farm business, and welfare

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

$$Y_{iv} = \beta_0 + \sum_{b=1,2,3} \beta_1^b \text{Forecast offer}_v \times [\text{Prior bin} = b]_i + \beta_2 \text{Insurance offer}_v + \rho_b [\text{Prior bin} = b]_i + \gamma \mathbf{X}_{iv} + \eta_{iv} \quad (4)$$

where $[\text{Prior bin} = b]_i$ are indicators which divide farmers into terciles on the basis of their priors.³³ 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 good news from the forecast). All other variables are as defined in Equation (3) above.³⁴ In this section, we discuss the treatment effects of the forecast on agricultural inputs and farm outputs, non-agricultural business activity, and farmer welfare. In Section 7, we compare these impacts to those of insurance.

The realized forecast was for an average monsoon, close to the mean of farmer beliefs (Figure 3). According to our theoretical model and historical yield analysis, we therefore expect (bad-news) farmers in the first tercile to *reduce* investment in agriculture, (neutral-news) farmers in the second tercile not to change their behavior, and (good-news) farmers in the third tercile to *increase* agricultural activity, including investing in cash crops. Given the stochasticity inherent to agriculture, in any given year, these expected changes in investment may or may not lead to

³³In our pre-analysis plan, we specified that we would split the sample into bad-news and good-news farmers. Given that the monsoon was average, with a large mass of farmers with priors around the forecast, we divide the sample into terciles to better reflect this heterogeneity. We present continuous treatment effects on our summary investment index in Figure 6. In Appendix G, we also present results from a specification where we pool all forecast farmers. As expected, the results tend to aggregate to zero across a variety of outcome variables, as they average negative and positive treatment effects.

³⁴Because 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 agricultural investment choices, agricultural productivity measures, off-farm business, and welfare measures.


substantial changes in output (Rosenzweig and Udry (2020), McCullough et al. (2020), and Suri and Udry (2022)). In the absence of orthogonal shocks, however, we predict that agricultural outputs should generally align with up-front investments, such that higher investments should lead to increased production, yields, and farm profits.

The effect of the forecast on off-farm enterprise is theoretically ambiguous. If farmers' off-farm business are highly linked to agriculture, or are otherwise dependent on the monsoon, then we would expect to see a similar pattern of response for these outcomes. On the other hand, if farmers treat off-farm business as a substitute for agriculture, we would expect to see activity rise for bad-news farmers and fall for good-news farmers. Which of these effects dominates is not clear *ex ante*.

We conclude our main analysis with estimates of treatment effects on farmer well-being. Because an accurate forecast in theory allows all farmers, regardless of prior, to tailor their on- and off-farm decisions to the coming monsoon season, theory predicts that overall welfare should (weakly) improve as a result of the forecast.³⁵ Our results, which we summarize in Figure 5 (standardizing all effects for comparability across outcomes), are broadly in line with these predictions.

³⁵To estimate effects on welfare outcomes, we therefore use the pooled specification described in Equation (3).

Figure 5: Summary of main results



Notes: This figure summarizes our main results. All effects are in standard deviations. The top row plots agricultural inputs; the middle row plots farm outputs and non-farm business; and the bottom row plots welfare outcomes. In the top two rows, we divide the distribution of farmer priors into terciles. Tercile 1 (dark blue) farmers were the most optimistic about the monsoon onset, and received bad news from the forecast. Tercile 2 (light blue) farmers' beliefs were essentially in line with the forecast, and therefore received neutral news. Tercile 3 (green) farmers were the most pessimistic about the monsoon onset, and received good news from the forecast. In the bottom row, we present pooled effects of the forecast (purple triangle) and insurance (gray circle). Coefficients and 95% confidence intervals are plotted for the forecast and insurance treatments, estimated using Equation (4) in the top two rows, where we interact the forecast treatment with the prior belief terciles. In the bottom row, we estimate point estimates and standard errors 95% confidence intervals using Equation (3). Overbraces indicate the significance level of tests for equality between coefficients. Green overbraces show tests between bad news and good news; gray overbraces show tests between good news and insurance; and purple overbraces show tests between pooled forecasts and insurance. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6.1 Effects on agriculture

Land and crop choice We first investigate the impact of our treatments on land use and crop choice (Table 3), and find evidence in support of our theory. Farmers who received bad news *reduce* land under cultivation by 22% of the control mean, and were 31% less likely to add a crop type from last year to this year.³⁶ While they were also less likely to change crops, this effect is not statistically significant. Farmers who received neutral news do not change their land under cultivation (point estimate of -3.3%), or their crop choices.

In contrast, farmers who received good news *increase* land under cultivation by 21%. They were also 17 percentage points more likely to grow a cash crop, 13 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, 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, suggesting that the changes we see were driven by new crops being added to the mix, rather than substitution.

We find statistically significant differences between farmers who received good news and bad news on land cultivation (p -value 0.001), cash cropping (p -value 0.032), changing crops from last year (p -value 0.023), and adding a crop between last year and this year (p -value 0.004). These results are consistent with the forecast enabling tailored investments: farmers in this treatment group adjusted their crop mix to match their updated expectations about the upcoming growing season.

³⁶Throughout 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 3: 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	-0.475*** (0.161)	0.015 (0.049)	-0.053 (0.053)	-0.113* (0.059)	0.012 (0.045)
Forecast × Ind Bin 2	-0.070 (0.147)	0.045 (0.037)	0.043 (0.051)	0.011 (0.047)	0.014 (0.038)
Forecast × Ind Bin 3	0.435* (0.233)	0.171*** (0.061)	0.130** (0.064)	0.135* (0.072)	0.027 (0.054)
Insurance	0.167 (0.133)	0.061 (0.037)	0.043 (0.046)	0.042 (0.048)	-0.004 (0.037)
q-val Tercile 1	0.034	1.000	1.000	0.316	1.000
q-val Tercile 2	1.000	1.000	1.000	1.000	1.000
q-val Tercile 3	0.078	0.051	0.078	0.078	0.229
q-val Insurance	0.335	0.302	0.483	0.483	0.962
Test Tercile 1=3	0.001	0.032	0.023	0.004	0.830
Test Insur. = Ter. 3	0.278	0.093	0.207	0.210	0.604
Control Mean	2.12	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 (4). 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 Kharif season of the experiment than in the prior season. Added Crop is an indicator for planting at least one additional crop in Kharif season of the experiment compared to the prior season. Sub Crop is an indicator for planting at least one fewer crop in the Kharif season of the experiment than in the prior season. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in Tables 3 and 4 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Appendix Table G.20 presents an IV analogue.

Farm inputs We next investigate the impact of our treatments on agricultural input expenditures (Table 4). While imprecise, points estimates for farmers who received bad news show that these farmers reduced their input expenditures, though somewhat less than they reduced land under cultivation.³⁷ Bad-news farmers reduced their total expenditures by 9%. This decline is driven by roughly proportional decreases in spending on fertilizer (8%), seeds (10%), irrigation and labor (7% and 6%).³⁸ We find no effect on farmers who received neutral news. However, good-news farmers increase their investments substantially, with total expenditures increasing by 31% of the control mean, driven by significant changes in labor expenditure and fertilizer (35% and 26% of the control mean, respectively), and positive but noisy impacts for seed and irrigation. We reject equality between good-news and bad-news farmers on total spending (p -value 0.019).³⁹

We also create an overall farm investment index from outcomes in Table 3 (land cultivation and

³⁷Appendix Table G.10 contains treatment effects on per-acre input use. Broadly, we find an increase in total per-acre inputs for bad news farmers. This is consistent with bad-news farmers decreasing land area cultivated by 22% but total inputs by 10%. We do not find changes in per-acre input use for neutral- or good-news farmers.

³⁸Appendix Table G.9 splits labor into “early” and “late” relative to the growing season.

³⁹We also reject equality between good-news and bad-news farmers for fertilizer expenditure (p -value 0.058), and labor expenditure (p -value 0.026).

cash cropping) and Table 4 (total input expenditure). The point estimate for bad-news farmers shows that these farmers reduced investment by 0.08 standard deviations (s.e. 0.07). We find no impacts on farmers who received neutral news, with a standardized treatment effect on the investment index of 0.03 SD (s.e. 0.06). However, good-news farmers increased investments by 0.31 SD effect (s.e. 0.09). We reject equality between good-news and bad-news farmers at the 1% level. All but one of our treatment effects on agricultural investment (adding a crop for bad-news farmers) remain statistically significant at conventional levels after correcting for multiple hypothesis testing.

These treatment effects suggest that the impact of forecasts differs significantly across farmers with different prior beliefs. Farmers who receive bad news from the forecast are more likely to reduce agricultural investment, farmers who receive neutral news do not change behavior, and farmers who receive good news increase agricultural investments and change crops. This pattern of results is consistent with our theoretical model. Even in cases where we cannot reject zero for the bad- or good-news group on their own, we typically reject equality between them.

Table 4: 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	-30.93 (42.17)	-0.68 (2.61)	-1.98 (7.78)	-42.36 (85.23)	-130.18 (160.48)	-0.08 (0.07)
Forecast × Ind Bin 2	-28.94 (39.37)	-2.01 (1.60)	-1.21 (4.91)	-44.45 (67.58)	-57.46 (123.92)	0.03 (0.06)
Forecast × Ind Bin 3	96.40* (55.58)	2.20 (3.35)	9.90 (7.99)	263.23** (105.15)	441.92** (185.33)	0.31*** (0.09)
Insurance	95.98** (42.79)	-0.93 (1.34)	-0.14 (5.67)	109.85* (63.60)	256.27** (128.51)	0.13** (0.05)
q-val Tercile 1	1.000	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	0.091	0.206	0.121	0.055	0.055	
q-val Insurance	0.302	0.575	0.962	0.302	0.302	
Test Tercile 1=3	0.058	0.493	0.304	0.026	0.019	0.000
Test Insur. = Ter. 3	0.994	0.368	0.273	0.203	0.373	0.065
Control Mean	372.80	7.22	26.81	761.96	1443.49	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 (4). 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. We exclud it 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in Tables 3 and 4 (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 G.21.

Agricultural output We examine three measures of agricultural production: total crop output in kilograms, the value of this production, and crop yield per hectare (Table 5).⁴⁰ Our forecast treatment effects follow the broad pattern we documented for agricultural inputs. Farmers who received bad news reduced agricultural output, including a 25% decline in production and a 22% decline in crop value — consistent with having reduced land area and total expenditures. We find close to zero effects for farmers receiving neutral news, in line with their having not changed agricultural inputs. Though both are imprecisely estimated, we find that good-news farmers increased agricultural output by 22% and the value of production by 16% — following from their increases in cultivated land, input expenditure, and the probability of cash cropping. We reject equality between bad- and good-news farmers on both production ($p = 0.017$) and value of production ($p = 0.056$). Finally, we see no meaningful changes in yield for any group. Thus, changes in output are likely due to expansions or contractions of in production rather than intensification.

Table 5: Effect of the forecast and insurance on agricultural output

	(1) Prod (Kg)	(2) Value Prod (\$)	(3) Yield
Forecast × Ind Bin 1	-16.90** (8.17)	-534.76* (284.68)	-6.59 (4.35)
Forecast × Ind Bin 2	-10.75 (7.50)	-188.75 (235.47)	-0.73 (3.61)
Forecast × Ind Bin 3	14.52 (10.99)	390.05 (415.29)	-0.49 (4.12)
Insurance	2.33 (6.77)	125.42 (222.88)	-1.66 (2.59)
q-val Tercile 1	0.218	0.218	0.218
q-val Tercile 2	0.837	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000
q-val Insurance	1.000	1.000	1.000
Test Tercile 1=3	0.017	0.056	0.261
Test Insur. = Ter. 3	0.283	0.519	0.776
Control Mean	66.91	2419.93	35.37
Observations	1201	1201	1170

Notes: This table presents estimates of the treatment effects of forecasts and insurance on agricultural output, estimated using Equation (4). Prod (Kg) is total agricultural production in kilograms. Value Prod (\$) is the value of all crops produced in USD, whether they were sold or not, using median district prices for each crop. Yield is kilograms of production per hectare. 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, enumerator 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 coefficients. 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 G.22.

⁴⁰We use district median prices to value production, to avoid selection in which farmers had actually sold their crop by the time of the survey from biasing our results. District median prices for key crops (e.g., cotton and paddy) are in line with local administrative data on market, or mandi, prices for Telangana (Allen and Atkin (2022); Kochhar and Song (2024)) during the year of our study.

Agricultural profits Agricultural profits should generally align with on-farm investments, such that higher (lower) investments should lead to higher (lower) profits on average. Table 6 presents impacts of the forecast on agricultural profits. We find that bad-news farmers have meaningfully lower farm profits ($-\$401$, s.e. $\$236$) than their counterparts in the control group. This is a more substantial decrease than their input changes would imply. We do not find statistically significant impacts on profits for either neutral- or good-news farmers. Though both point estimates are negative, the effect for good-news farmers is quite close to zero ($-\$64$, s.e. $\$344$), inconsistent with increases in inputs and production.

The fact that changes in investments did not perfectly translate into change in agricultural profits suggests that some farmers may have been negatively affected by a shock that was unrelated to monsoon onset timing (and thus to the forecast). Indeed, Telangana was hit by heavy flooding in early July (Business Line (2022); The New Indian Express (2022)). While the likelihood of flood exposure is balanced between treatment and control (Appendix Table A.10), because good-news farmers spent more on inputs and farmed more valuable crops, the same flood may have led to greater losses among this group than in the control. We conduct a (non-pre-specified) analysis of these shocks in the remainder of Table 6. We begin with the value of crop losses (i.e., production that was lost to shocks, valued at district-median prices). Bad-news farmers saw a loss of $\$54$ (not different from zero), while neutral- and good-news farmers experienced meaningful crop losses of approximately $\$217$ (p -value 0.077) and $\$208$ (p -value 0.165), respectively. We next present a counterfactual profit calculation, adding the value of these losses as revenue. Though all estimates are noisy, we find a pattern that is more in line with farmers’ input adjustments: bad-news farmers see a profit reduction of 18%, neutral-news farmers see approximately no change, and good-news farmers see an increase of 9%. Finally, we show agricultural profits, excluding farmers who reported losses from flooding or cyclones from the sample.⁴¹ Again, all estimates are noisy – as the sample is considerably smaller – but we find evidence consistent with agricultural profits responding as expected. The point estimates imply that agricultural profits for bad-news farmers fell by 35%, neutral-news farmers fell by 10%, and good-news farmers increased by 50%.

These results are consistent with the external validity point made in Rosenzweig and Udry (2020): agriculture is an inherently stochastic process. While the forecast appears to have led treatment farmers to make choices that would have been agronomically-appropriate on average, the occurrence of an orthogonal flood shock reduced profits for many of these farmers during our experiment. Nevertheless, Appendix Figure A.6 shows that farmers’ self-reported trust in the forecast increased substantially over the course of the growing season, demonstrating that farmers understand the distinction between monsoon onset and other growing season realizations.⁴²

⁴¹Because there were no documented cyclones in Telangana in 2022, we interpret “cyclone” as heavy rain or flooding.

⁴²If anything, the average *ex post* trust is *higher* for farmers who experienced the flood shock (7.1 on a 1-10 scale) than for those who did not (6.8).

Table 6: Effect of the forecast and insurance on agricultural profit

	(1) Ag Profit	(2) Loss	(3) Profit + Loss	(4) Ag Profit Non-Flood
Forecast × Ind Bin 1	-401.08* (235.82)	54.39 (135.26)	-296.60 (322.67)	-341.47 (427.31)
Forecast × Ind Bin 2	-118.98 (193.99)	217.10* (122.55)	99.46 (221.29)	-96.46 (253.28)
Forecast × Ind Bin 3	-64.98 (344.24)	207.89 (149.54)	154.50 (373.06)	498.33 (518.06)
Insurance	-150.43 (182.75)	195.48** (90.98)	-1.21 (207.64)	407.19 (298.57)
q-val Tercile 1	0.218	0.298	0.243	
q-val Tercile 2	1.000	0.837	1.000	
q-val Tercile 3	1.000	1.000	1.000	
q-val Insurance	1.000	0.237	1.000	
Test Tercile 1=3	0.400	0.411	0.338	0.222
Test Insur. = Ter. 3	0.805	0.935	0.670	0.871
Control Mean	970.62	661.07	1654.24	1052.59
Observations	1201	1201	1201	554

Notes: This table presents estimates of the treatment effects of forecasts and insurance on agricultural profit and loss, estimated using Equation (4). Ag Profit is the value of production (evaluated at district-median prices) less total expenditure. Loss is the value of reported crop losses (evaluated at district-median prices). Profit + Loss is the value of production plus the value of crop losses, less total expenditure in USD. Ag Profit Non-Flood is agricultural profits for the sample of households that did not report crop losses due to flooding or cyclones. All outcomes are 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, enumerator 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 coefficients. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. We exclude Column (4) from our MHT because it is a sub-sample analysis of Column (1). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. We present an IV analogue in Appendix Table G.23.

6.2 Effects on non-agricultural business

Table 7 presents the effects of the forecast on non-agricultural business. We find 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 42% more likely than control to own a non-agricultural business, increased non-agricultural investment by 17%, and increased business profits by \$80. 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 76% (significant at the 10% level), and saw a \$44 decline in business profits. These results, which are in the opposite direction to our agricultural input findings, are consistent with farmers treating business as a substitute for agriculture.⁴³

⁴³In addition to estimates for non-agricultural business, in Appendix A, we present results for other income sources (Appendix Table A.8) and migration (Appendix Table A.9). Forecasts have no significant impacts on other income-generating activities, though we see evidence that bad-news farmers saw reduced labor income – consistent with choosing not to work on others’ farms in the face of a poor growing season. If anything, livestock income declines, consistent with (especially good-news) farmers selling fewer assets. Finally, we see that the forecast 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 migrant labor depends on a good monsoon, bad news may therefore depress migration (Rosenzweig and Udry, 2020).

Table 7: Effect of the forecast and insurance on business activity

	(1) Non-Ag Bus.	(2) Non-Ag Invest	(3) Bus Profit
Forecast × Ind Bin 1	0.06 (0.04)	26.46 (71.79)	79.97 (75.99)
Forecast × Ind Bin 2	0.01 (0.03)	6.48 (58.57)	17.15 (48.14)
Forecast × Ind Bin 3	-0.05 (0.04)	-122.57* (68.31)	-43.70 (90.28)
Insurance	0.09*** (0.03)	101.44 (63.13)	104.95* (55.09)
q-val Tercile 1	0.784	0.784	0.784
q-val Tercile 2	1.000	1.000	1.000
q-val Tercile 3	0.287	0.281	0.501
q-val Insurance	0.028	0.078	0.061
Test Tercile 1=3	0.060	0.130	0.267
Test Insur. = Ter. 3	0.007	0.002	0.151
Control Mean	0.14	157.98	165.51
Observations	1197	1199	1197

Notes: This table presents estimates of the treatment effects of forecasts and insurance on non-farm business activity estimated using Equation (4). 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. 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, enumerator 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 coefficients. 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 G.25.

6.3 Effects on welfare

Finally, we measure impacts on farmer well-being in Table 8. Because theory predicts that forecasts should increase overall welfare for all farmers, we estimate pooled treatment effects of the forecast using Equation (3). Forecasts increase per-capita food consumption by 7% of the control mean (statistically significant at the 5 percent level), and we find no impacts on other consumption (point estimate of $-\$0.52$, s.e. $\$0.82$).⁴⁴ Forecasts raise asset value by 8%. We see no effect on livestock count. Forecasts increase savings net of debt by $\$184$, compared to a control group mean of $-\$1,083$, largely driven by a decrease in debt, though these estimates are all imprecise.⁴⁵ We aggregate these economic well-being measures into an inverse-covariance weighted index, and estimate that forecasts raise overall well-being by 0.06 SD ($p = 0.048$). Treatment effects on the index are largest for bad-news farmers (0.14 SD, s.e. 0.06), zero for neutral-news farmers (0.00 SD, s.e. 0.04), and weakly positive for good-news farmers (0.05 SD, s.e. 0.05), consistent with responses to our forecast treatment having been concentrated in the bad- and good-news groups (Appendix

⁴⁴Appendix Table A.12 shows that this effect rises to 9.7% in non-flooded villages. We present effects on further-disaggregated consumption categories in Appendix Table A.14.

⁴⁵We provide additional results on household finances in Appendix Table A.15. We estimate impacts on mental health in Appendix Table A.16. We find no impact on overall mental health, though forecast farmers do report increases in poor appetite and/or overeating. We also find evidence of somewhat worse mental health for bad-news farmers, consistent with stress from learning bad news.

Table A.13. Together, these results demonstrate that the forecast improved welfare, even in the presence of negative shocks to agricultural production.

Table 8: Effect of the forecast and insurance on economic well-being

	(1) Food cons	(2) Other cons	(3) Asset value	(4) Livestock	(5) Net savings	(6) Welfare index
Forecast	0.87** (0.42)	-0.52 (0.82)	113.68 (134.04)	0.01 (0.02)	184.33 (145.61)	0.06** (0.03)
Insurance	0.46 (0.47)	1.92* (1.01)	-123.01 (160.45)	0.01 (0.03)	-405.18* (219.70)	-0.02 (0.05)
q-val Forecast	0.248	0.944	0.944	0.944	0.700	
q-val Insurance	0.499	0.196	0.499	0.868	0.196	
Control Mean	13.22	9.93	1503.10	0.22	-1031.41	0.00
Observations	1201	1201	1201	1201	1129	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on economic well-being. The estimation follows Equation (3). Food cons is food consumption per household member in USD over the past 30 days. Other cons is non-food consumption (other than medical expenses). Asset value is the value of assets in USD. Livestock is the count of livestock in Tropical Livestock Units. Net savings is savings less outstanding debt in USD. Welfare index is an inverse covariance weighted index of the other five outcomes in the table. We exclude the index from the MHT correction as it is a composite of outcomes already included. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7 Forecasts vs. insurance

As a final exercise, we compare forecasts to insurance, the canonical risk-coping tool. In principle, forecasts and insurance should work in fundamentally different ways. Accurate forecasts provide farmers with information, allowing them to tailor their investments to a given state of the world, thereby reducing risk. In contrast, insurance allows farmers to shift consumption between states of the world but provides no information, increasing welfare without reducing risk. We therefore predict that insurance farmers should increase investments in general, but should not match these changes to the coming growing season. Figure 5 summarizes the effects of insurance, and compares these impacts to those of the good-news forecast group, on twelve key outcomes. Our results, which we discuss in more detail below, broadly align with these predictions.

Land, crop choice, and inputs Tables 3 and 4 documents that insurance led farmers to engage in more agricultural activity. Though imprecisely-estimated, these farmers increased their land under cultivation by approximately 9 percent. This effect is similar in magnitude to the impacts of index insurance in other settings (Karlan et al. (2014)), and we cannot reject equality between insurance and good-news-forecast farmers (p -value 0.279). We also find that insurance caused farmers to increase total expenditures by 18%, driven by a 15% increase in labor expenditures and a 25% increase in fertilizer spending, and again cannot reject equality between good-news farmers in the forecast group and insurance on any expenditure category.


However, we do find meaningful differences in farm inputs between the forecast and insurance. Specifically, we find that good-news farmers changed crop choice, while insurance farmers are much less likely to have done so. Though the insurance point estimates on cash crop (6.1 percentage points), changing the crop mix from last year to this year (4.3 percentage points), and adding a crop (4.2 percentage points) are positive, they are meaningfully smaller (just one-third the size) than the good-news effects. We reject equality between the good-news and insurance coefficients on growing cash crops (p -value 0.093). Our overall investment index increases by 0.13 SD for the insurance group, but this is also meaningfully smaller than the good-news group (p -value on equality 0.065). Taken together, these results demonstrate that while insurance increases agricultural investments, it lacks the forecast’s ability to target investments to the coming growing season.

Heterogeneity by prior beliefs Our forecast results emphasize heterogeneity by prior beliefs. While not typically modeled in the insurance literature, these same priors may also change how farmers respond to insurance. To illustrate this, we add an insurance product to the model presented in Section 3.⁴⁶ The model predicts (Figure 1, black line) that while insurance should weakly increase investments for all farmers, these impacts will be largest among “optimistic” farmers with early priors – those who would have received bad news from the forecast – and smallest among “pessimistic” farmers with late priors – those who would have received bad news from the forecast. Intuitively, insurance enables optimistic (but risk-averse) farmers to substantially increase agricultural activity in anticipation of a favorable season by protecting against downside risk, but does not increase the appeal of agriculture for a pessimistic farmer who believes that agricultural investments are likely to go to waste in an unfavorable season.

Figure 6 presents an empirical test of these predictions, using the agricultural investment index as our key outcome of interest. As shown above, early-prior (bad news) respond to the forecast by investing less, while late-prior (good news) farmers respond to the forecast by investing more, both compared to control farmers with similar priors. In contrast, early-prior (optimistic) insurance farmers invest more than the control, while late-prior (pessimistic) farmers do not change their investments. Appendix Table A.11 reports effects of the insurance treatment by prior tercile on our core farm inputs. For early-prior insurance farmers, we find large effects on land under cultivation (23% of the control mean), input expenditures (29%), and the investment index (0.17 SD), while for late-prior insurance farmers, we find null results on all outcomes, summarized in an investment index effect of 0.02 SD, though we cannot reject equality. We do not find economically meaningful or statistically significant impacts on crop change for early- or late-prior farmers, consistent with insurance providing no information. Taken together, these results demonstrate that insurance is more effective at encouraging investment among optimistic farmers than pessimistic farmers. These findings contrast sharply with the forecast, which helps to correct farmer beliefs, reducing investment among overly-optimistic farmers and encouraging investment among overly-pessimistic farmers.

⁴⁶To do so, we model farmers as gaining a fixed amount of income if the realized state falls below some pre-determined threshold. See Appendix C for more details.

Figure 6: 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.

Agricultural output As with forecast farmers, Table 5 reveals a mismatch between agricultural inputs and farm output for the insurance group. Despite increasing farm investments, insurance led to virtually no change in production (3%), the value of production (5%), or yields (−5%). Table 6 shows that insurance reduced agricultural profits by −\$150 (though this is not different from zero), revealing that insurance farmers faced meaningful crop losses (\$195). In contrast to the forecast farmers, we do see that the insurance group was 12 percentage points more likely to report negative shocks, driven by cyclones (Appendix Table A.10).⁴⁷ Among farmers reporting neither flood nor cyclone shocks, insurance increased farm profits by \$407 (39%), though this is imprecisely estimated. These results thus demonstrate that the discrepancy between farm input and agricultural output effects is not unique to forecasts, but rather likely driven by exogenous weather shocks beyond what these two risk-coping instruments were designed to address.

Non-farm business While we do not model this explicitly, theory would suggest that insurance should lead to increased investment both on and off the farm. In keeping with this prediction, Table 7 shows that insurance meaningfully increased non-agricultural business ownership (64%

⁴⁷ As above, we could find no corroborating evidence of cyclones in the media or scientific reporting, so we interpret this as farmers reporting heavy rainfall.

relative to the control), investment in non-agricultural business (64%, not different from zero), and business profits (\$105). These effects are similar to those in the *bad-news* group (and the former two are statistically distinguishable from the good-news group), suggesting that insurance farmers increased investment across the board, rather than treating farming and business as substitutes — again consistent with insurance lacking the tailoring properties of the forecast.

Welfare Finally, we broadly estimate that insurance did not appear to have large impacts on farmer well-being (Table 8). We estimate point estimates of 3% increases in food consumption, a 20% increase in other consumption, 8% declines in asset value, no change in livestock, but only other consumption is significant at the 10% level. We do estimate that insurance farmers have lower net savings (a reduction of \$405 on a mean of $-\$1,031$), which aligns with insurance farmers having made more investments and been hit with an (uninsured) flood shock, leading to losses. We estimate a null effect of insurance on our summary welfare index (point estimate -0.02 SD, s.e. 0.05). These results are consistent with prior work demonstrating that while insurance can be expected to improve welfare on average (Mobarak and Rosenzweig (2014); Carter et al. (2017a); Cole and Xiong (2017)), in any given year, farmers may see negative impacts (Karlan et al. (2014)) — particularly if, as we discuss above, insurance induces investments from overly-optimistic farmers.

Discussion We document that both forecasts and index insurance meaningfully alter farmers’ planting behavior, but that these instruments operate along different margins. As a result, forecasts and insurance may be complements rather than substitutes. Access to a forecast potentially makes insurance more valuable, enabling farmers to invest heavily knowing their downside risk is covered. Access to insurance protects farmers from the possibility of an incorrect forecast, 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 or vice versa, and that both products together could substantially increase farmer welfare.⁴⁸ Measuring this interaction is therefore an important topic for future research.

8 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 of the onset timing of the Indian Summer Monsoon well in advance of its arrival. Our simple theoretical model predicts that such a forecast will enable farmers to tailor their decisions to the coming monsoon season, and that these responses should depend on farmers’ prior beliefs. Our estimated treatment effects confirm these predictions.

The forecast caused bad-news farmers to reduce agricultural inputs, had no impact on neutral-news farmers, and caused good-news farmers to increase farm inputs. In keeping with these input

⁴⁸Because 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.

changes, we broadly see negative, close-to-zero, and positive point estimates on agricultural outcomes for the three groups of forecast farmers. Turning to farm profits, while we find negative impacts for bad-news farmers, we estimate close-to-zero effects for good-news farmers. We document that this is likely due to severe flooding, a growing-season shock that was orthogonal to the forecast; both agricultural profits net of crop losses and profits among flood-unaffected farmers are positive in the good-news group. Our results on non-agricultural business are the mirror-image of the impacts on agriculture, with positive point estimates for bad-news farmers, and negative point estimates for good-news farmers. Finally, we find that the forecast increased welfare, with increases in food consumption, asset value, net savings, and an overall welfare index.

Our findings demonstrate that forecasts are a useful tool for coping with a variable climate, as they reduce agricultural risk, which will become increasingly important as the climate changes further. Our evidence also reveals important areas for future work. Additional research is needed to develop accurate, long-range forecasts of growing season conditions throughout the tropics. Our findings demonstrate that these forecasts have the potential to improve farmer decision-making in the face of variable weather. In addition, while we find strong evidence that forecasts enable farmers to tailor their investments to the upcoming monsoon season, information about onset timing does not shield farmers from the full set of shocks that make agriculture inherently risky. Future research should therefore seek to evaluate the extent to which complimentary policies, such as crop insurance or flood forecasts, can increase the welfare impacts of forecasts 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 writ large. There already exist similar forecasts elsewhere in India, and advances in climate science are enabling their wider development. Broadly representing the global meteorological, 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 REZAEI (2023): “Forecasts: Consumption, production, and behavioral responses,” Working paper.
- ALLEN, TREB AND DAVID ATKIN (2022): “Volatility and the gains from trade,” *Econometrica*, 90 (5), 2053–2092.
- 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.
- BUSINESS LINE (2022): “Floods in Telangana severely impact infrastructure, says post-event analysis,” .
- 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 (2017a): “Index insurance for developing country agriculture: A reassessment,” *Annual Review of Resource Economics*, 9, 421–438.
- (2017b): “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.
- 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.
- DOWNEY, MITCH, NELSON LIND, AND JEFFREY G. SHRADER (2023): “Adjusting to rain before it falls,” *Management Science*, 69 (12), 7151–7882.
- 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 KONDYLLIS, 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.
- KATZENBERGER, ANJA, JACOB SCHEWE, JULIA PONGRATZ, AND ANDERS LEVERMANN (2021): “Robust increase of Indian monsoon rainfall and its variability under future warming in CMIP6 models,” *Earth System Dynamics*, 12 (2), 367–386.
- 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.
- KOCHHAR, RAJAT AND RUOZI SONG (2024): “Does market power in agricultural markets hinder farmer climate change adaptation?” Working paper.
- 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.

- MCCULLOUGH, ELLEN B., JULIANNE D. QUINN, ANDREW M. SIMONS, LESLIE J. VERTERAMO, AND JOSHUA WOODARD (2020): “Modeling climate smart soil health investments in sub-Saharan Africa,” Working paper.
- 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 AGRICULTURE AND FARMERS WELFARE (2023): “Lok Sabha Starred Question No. 228,” .
- 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.
- PATEL, DEV (2024): “Learning about a warming world: Attention and adaptation in agriculture,” Working paper.
- PINDYCK, ROBERT S. (2013): “Climate change policy: What do the models tell us?” *Journal of Economic Literature*, 51 (3), 860–872.
- PRABHU, SHRAVAN AND VISHWAS CHITALE (2024): “Decoding India’s changing monsoon patterns: A tehsil-level assessment,” .

- 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.
- ROSENZWEIG, MARK R AND CHRISTOPHER UDRY (2020): “External Validity in a Stochastic World: Evidence from Low-Income Countries,” *The Review of Economic Studies*, 87 (1), 343–381.
- SCHLENKER, WOLFRAM AND MICHAEL J. ROBERTS (2009): “Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change,” *Proceedings of the National Academy of Sciences*, 106 (37), 15594–15598.
- SCHLENKER, WOLFRAM AND CHARLES A. TAYLOR (2021): “Market expectations of a warming climate,” *Journal of Financial Economics*, 142 (2), 627–640.
- SHRADER, JEFFREY G. (2023): “Improving climate damage estimates by accounting for adaptation,” Working paper.
- SHRADER, JEFFREY G., LAURA BAKKENSEN, AND DEREK LEMOINE (2023): “Fatal errors: The mortality value of accurate weather forecasts,” NBER working paper No. 31361.
- 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.
- SURI, TAVNEET AND CHRISTOPHER UDRY (2022): “Agricultural Technology in Africa,” *Journal of Economic Perspectives*, 36 (1), 33–56.
- THE NEW INDIAN EXPRESS (2022): “Telangana floods: Godavari in spate, houses collapse, streets turn into streams,” .
- 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.
- YEBEMEY, 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.
- YEBEMEY, ROSAINE N. AND JANVIER EGAH (2021): “Reaching out to smallholder farmers in developing countries with climate services: A literature review of current information delivery channels,” *Climate Services*, 23, 100253.

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

Online appendix

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

Contents


A Appendix tables and figures	A3
A.1 Context	A3
A.2 Attrition and balance	A5
A.3 Takeup	A8
A.4 Forecast accuracy	A10
A.5 Beliefs	A11
A.6 Belief spillovers	A13
A.7 Willingness-to-pay	A14
A.8 Income-generating activities and migration	A15
A.9 Trust	A17
A.10 Shocks	A18
A.11 Insurance and prior beliefs	A19
A.12 Additional welfare results	A20
B Panel analysis: Effect of onset timing on crop yields	A25
B.1 Data	A25
B.2 Defining monsoonal regions	A26
C Model details	A28
C.1 Setup	A28
C.2 Optimal farmer investment and saving decisions	A28
C.3 Parametrization for simulations	A31
C.4 Model predictions for alternative forecast realizations	A32
D Becker et al. (1964) appendix	A35
D.1 Methodological overview	A35
D.2 Distribution of BDM prices	A37

E	Information sheets	A38
F	Deviations from our pre-analysis plan	A40
G	Additional pre-specified results	A43
G.1	Pooled forecast treatment results	A43
G.2	Correlates of willingness-to-pay	A46
G.3	Belief heterogeneity	A48
G.4	Additional farm input results	A49
G.5	Heterogeneity	A51
G.6	Local average treatment effects	A59
H	Seasonal climate forecasts	A66

A Appendix tables and figures


A.1 Context

Figure A.1: Variability of cultivation in Telangana



Notes: This figure presents statistics on the use of agricultural land in Telangana over time, using data from the Ministry of Agriculture and Farmers' Welfare. The left panel shows total land under cultivation. The right panel shows the percent of agricultural land area cropped to the top five crops: cotton (the main cash crop), rice (the main staple crop), maize, arhar, and soyabean.

Figure A.2: Sources of information about the 2022 monsoon at baseline



Notes: This figure presents farmers' reported sources of information on monsoon onset timing for the kharif season studied in the experiment. Data were collected at baseline. Farmers were able to report the use of multiple sources. In village is farmers in the respondent's village; Elsewhere is farmers in other villages; IMD is the government forecast; Extension is other extension services; Ecological is ecological signals (such as animal behavior); TV/radio, Newspaper, Other, Don't know, and Refuse are self-explanatory.

A.2 Attrition and balance

Table A.1: Differential attrition by treatment group

	(1)
Forecast	-0.010 (0.016)
Insurance	-0.038*** (0.014)
Control mean	0.04
Observations	1240

Notes: This table presents attrition (defined as being present in the first baseline round but not present in *either* baseline round II or endline) by treatment status. The regression includes strata fixed effects. Errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.2: Correlates of attrition (control only)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2022 onset prior	0.023 (0.024)							
2022 onset SD		0.154** (0.071)						
# of households			-0.000 (0.000)					
# of farmers				-0.000 (0.000)				
% area rain-fed					0.000 (0.000)			
% area irrigated						-0.000 (0.001)		
Cultivated area (ha)							-0.000 (0.000)	
District = Medak								-0.056** (0.024)
Ctrl. mean indep. var.	4.91	1.00	411.89	449.61	55.61	30.69	364.30	0.41
Observations	495	495	495	495	495	495	495	495

Notes: This table presents correlates of attrition (defined as being present in the first baseline round but not present in *either* baseline round II or endline). We restrict the sample to control group households only. 2022 onset prior (SD) is the mean (SD) of a household's prior belief distribution (elicited using the beans task described in Section 4 and measured in kartes), and are measured at the individual level. All other covariates are measured at the village level. Errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Table A.3: Balance across treatment arms

	(1)	(2)	(3)	Difference		
	Control	Forecast	Insurance	(2)-(1)	(3)-(1)	(2)-(3)
<i>Panel A: Village characteristics</i>						
# of households	413.82 [367.61]	470.45 [647.08]	378.68 [249.78]	56.63 (74.51)	-35.14 (51.08)	91.77 (73.76)
# of farmers	453.16 [526.19]	480.57 [461.82]	549.70 [615.04]	27.41 (70.21)	96.54 (101.59)	-69.13 (98.26)
Cultivated area (ha)	365.67 [375.22]	362.94 [356.27]	420.78 [451.81]	-2.73 (51.88)	55.10 (74.04)	-57.84 (73.00)
% area rain-fed	55.63 [23.15]	56.47 [23.67]	59.65 [21.39]	0.84 (3.32)	4.02 (3.81)	-3.19 (3.84)
% area irrigated	30.77 [19.84]	29.73 [20.16]	32.17 [19.37]	-1.05 (2.84)	1.39 (3.38)	-2.44 (3.40)
Observations	100	100	50			
<i>Panel B: Household-level characteristics</i>						
HH size	5.39 [2.52]	5.30 [2.35]	5.25 [2.07]	-0.08 (0.18)	-0.14 (0.20)	0.06 (0.20)
HH head age	47.99 [12.31]	47.48 [11.67]	46.43 [11.78]	-0.51 (0.93)	-1.57 (1.24)	1.06 (1.20)
HH head educ	6.05 [5.12]	6.03 [5.05]	6.45 [5.04]	-0.03 (0.38)	0.39 (0.50)	-0.42 (0.51)
# of plots	2.01 [1.20]	1.98 [1.09]	2.07 [1.12]	-0.03 (0.10)	0.06 (0.12)	-0.10 (0.11)
Total land (ha)	2.71 [4.75]	2.32 [2.38]	2.54 [2.24]	-0.38 (0.28)	-0.16 (0.31)	-0.22 (0.26)
Observations	472	481	247			
<i>Panel C: Beliefs about the monsoon</i>						
2022 onset mean	4.89 [0.63]	4.84 [0.50]	4.86 [0.51]	-0.05 (0.07)	-0.03 (0.09)	-0.03 (0.08)
2022 onset SD	0.98 [0.32]	0.89 [0.27]	0.90 [0.29]	-0.09** (0.04)	-0.08* (0.04)	-0.01 (0.04)
Historical onset mean	4.84 [0.56]	4.82 [0.49]	4.96 [0.46]	-0.01 (0.07)	0.12 (0.08)	-0.14* (0.07)
Historical onset SD	0.82 [0.19]	0.77 [0.19]	0.79 [0.19]	-0.05** (0.02)	-0.03 (0.03)	-0.01 (0.03)
Observations	472	481	247			

Notes: This table presents tests for balance across the three treatment arms. Panel A presents balance at the village level. Panels B and C present balance at the household level. All outcomes in Panel C are measured in kartes using the beans task described in Section 4. Columns (1) – (3) show means and [standard deviations]. The remaining columns present the pair-wise differences and (standard errors). Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.3 Takeup

Figure A.3: Takeup of forecasts and insurance



Notes: This figure presents takeup of the forecast (purple) and insurance (gray) products. The top panel shows takeup as a share of households in each treatment arm, while the bottom plots takeup against the mean of the prior distribution, measured in kartes. The dashed line presents the prior distribution. Priors are winsorized at the 3rd and 97th percentile.


Table A.4: Effect of forecast and insurance offers on takeup

	(1) Forecast takeup	(2) Insurance takeup	(3) Forecast Bin 1	(4) Forecast Bin 2	(5) Forecast Bin 3	(6) Insurance takeup
Forecast	0.878*** (0.021)	0.004 (0.007)				
Insurance	0.024 (0.016)	0.865*** (0.031)	0.018 (0.012)	0.002 (0.005)	0.002 (0.002)	0.866*** (0.031)
Forecast × Ind Bin 1			0.820*** (0.043)	-0.004 (0.007)	0.003 (0.003)	0.022 (0.017)
Forecast × Ind Bin 2			-0.000 (0.010)	0.891*** (0.026)	0.004 (0.003)	0.003 (0.013)
Forecast × Ind Bin 3			0.011 (0.014)	0.002 (0.006)	0.926*** (0.024)	-0.025* (0.013)
Control Mean	0.00	0.00	0.00	0.00	0.00	0.00
Observations	1201	1201	1201	1201	1201	1201

Notes: This table presents the effect of offering the forecast and insurance treatments on takeup of those treatments. We produce Columns (1) and (2) by estimating Equation (3) with forecast takeup and insurance takeup as the outcome variables, respectively. Columns (3) through (6) present results estimated using Equation (4), with the interaction between forecast takeup and prior bins 1–3 (Columns 3–5), and insurance takeup (Column 6) as the outcome variable. 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, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.4 Forecast accuracy

Figure A.4: Rainfall realizations and forecast accuracy in our sample



Notes: This figure shows rainfall over our sample, measured at each of our 25 rain gauges (light gray lines). Following standard practice in climate science, each line plots rainfall amounts calculated in moving cumulative 5-day sums (or pentads). The solid black line plots the mean over all 25 gauges. The purple shaded area shows the monsoon onset window predicted by the forecast, during which time *all* 25 gauges reported non-zero rainfall. The gray shaded area shows the subsequent dry spell predicted by the forecast. Finally, the dashed horizontal line shows the rainfall threshold used to determine insurance payouts. We use a very generous insurance payout rule. Insurance payments were triggered if rainfall had not reached 30mm of precipitation over a 5-day period before the trigger date (and if there was a dry spell within 30 days of the first rains lasting 10 days with less than 5mm of cumulative rainfall). This ensured that as many people as possible would be paid. Using this threshold, 13 out of 25 gauges triggered insurance payouts, even though all of these rain gauges saw rainfall during the forecasted onset period.

A.5 Beliefs

Table A.5: Association between control-group priors and agricultural investment

	(1) Land Ha.	(2) Cash Crop	(3) Total Exp.	(4) Invest Index
Tercile 2	-0.200 (0.152)	-0.153*** (0.053)	-77.770 (159.138)	-0.175** (0.075)
Tercile 3	-0.347* (0.209)	-0.168*** (0.062)	-303.730* (178.404)	-0.238** (0.096)
Test Tercile 2=3	0.483	0.796	0.177	0.463
Tercile 1 mean	2.29	0.59	1517.10	0.12
Observations	473	473	473	473

Notes: This table reports the relationship between control group prior beliefs and farm inputs. Land Ha. is land under cultivation in hectares, Cash Crop is an indicator for the farmer planting at least one cash crop, Total Exp. is the total amount spent on all inputs (in USD), and Invest Index is an inverse covariance weighted index of the previous four variables. We regress each outcome on indicators being in the 2nd and 3rd tercile of mean prior beliefs, with tercile 1 as the omitted category. Priors are elicited using the bean task described in Section 4. “Test Tercile 2 = 3” is the p -value on the test of equality between the two coefficients. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. The sample includes control-group farmers only. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.6: Correlation between beliefs and farmer characteristics

	Mean of prior belief distribution (kartes)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HH size	0.008 (0.008)							
HH head age		0.001 (0.002)						
HH head education			0.003 (0.004)					
HH head home village (1/0)				0.225** (0.091)				
# of plots					-0.018 (0.019)			
Total land (ha)						-0.011** (0.005)		
Cash crops 2021 (1/0)							-0.035 (0.054)	
Risk aversion								-0.002 (0.007)
Ctrl. mean indep. var.	5.39	47.99	6.05	0.92	2.01	2.71	0.52	4.64
Observations	1202	1202	1202	267	1202	1202	1202	1202

Notes: This table presents the correlation between farmers' prior beliefs (measured in kartes, using the beans task described in Section 4) and baseline characteristics. HH size is the number of household members (including the head), HH head age is the age of the household head in years, HH head education is the household head's years of schooling. HH head home village is an indicator for whether the household head was born in the village in which they currently reside. # of plots is the number of plots farmed by the household. Total land (ha) is acres of land farmed by the household. Cash crops 2021 (1/0) is an indicator for having farmed cash crops in the kharif season prior to the experiment. Risk aversion measures the farmer's choice in an incentivized risk game, where higher values indicate that the farmer is more risk averse. Ctrl. mean indep. var. is the mean of the independent variable in the control group. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.6 Belief spillovers


Table A.7: Effect of forecast on untreated farmer beliefs (spillover sample)

	(1) Arrival Date	(2) Arrive Ontime
Forecast Village	0.066 (2.139)	-0.007 (0.007)
Control Mean	1.27	0.00
Observations	303	304

Notes: This table presents the effect of information spillovers on beliefs. Forecast Village is an indicator for being an untreated farmer (i.e., not in the main sample) in a forecast offer village. Arrival Date is the date that the farmer expected the monsoon to arrive in kartes. Arrive On time is an indicator for whether the farmer believed the monsoon would arrive on time, using their own criteria. The sample includes only untreated farmers in control villages and in forecast villages. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.7 Willingness-to-pay

Figure A.5: Willingness-to-pay for forecasts and insurance



Notes: This figure presents willingness-to-pay curves for the forecast (purple) and insurance product (gray), elicited using the BDM mechanism described in Section 4 and Appendix D. Mean WTP for the forecast (insurance) is \$1.08 (\$1.29). The area under the demand curve for forecasts (insurance) is \$1.42 (\$1.56).

A.8 Income-generating activities and migration

Table A.8: Effect of the forecast and insurance on other income-generating activities

	Panel A: Forecast vs. Insurance		
	(1) Labor Inc.	(2) Livestock Inc.	(3) Remittance Inc.
Forecast	-44.45 (33.59)	-81.30 (70.69)	-2.77 (2.37)
Insurance	-29.11 (40.93)	-41.70 (113.22)	-5.72** (2.53)
q-val Forecast	0.338	0.338	0.338
q-val Insurance	0.907	0.907	0.079

	Panel B: Forecast Terciles		
	(1) Labor Inc.	(2) Livestock Inc.	(3) Remittance Inc.
Forecast × Ind Bin 1	-114.22** (54.72)	-0.18 (100.40)	-1.37 (4.72)
Forecast × Ind Bin 2	7.33 (44.15)	-143.86 (101.26)	-2.10 (2.26)
Forecast × Ind Bin 3	-20.22 (67.89)	-177.44 (129.70)	-6.79** (3.41)
q-val Tercile 1	0.126	1.000	1.000
q-val Tercile 2	1.000	0.900	0.900
q-val Tercile 3	0.353	0.211	0.163
Test Tercile 1=3	0.264	0.290	0.323
Insur. = Ter. 3	0.895	0.405	0.691
Control Mean	324.53	496.70	7.46
Observations	1199	125	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on other income-generating activities, estimated using Equations (3, panel A) and (4, panel B). Labor Inc. is labor income in the last 12 months, Livestock Inc. is income from selling livestock in the last 12 months, and Remittance Inc. is income from remittances received in past 30 days, all 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.9: Effect of the forecast and insurance on migration


Panel A: Forecast vs. Insurance				
	(1) Any Migrant	(2) Num Temp Mig.	(3) N. Female	(4) N. Male
Forecast	-0.03 (0.02)	-0.09** (0.04)	-0.03** (0.01)	-0.05* (0.03)
Insurance	-0.00 (0.03)	-0.08** (0.04)	-0.04** (0.02)	-0.04 (0.03)
q-val Forecast	0.084	0.057	0.057	0.062
q-val Insurance	0.288	0.111	0.111	0.111

Panel B: Forecast Terciles				
Forecast × Ind Bin 1	-0.05 (0.04)	-0.15** (0.06)	-0.06** (0.02)	-0.08* (0.04)
Forecast × Ind Bin 2	-0.03 (0.03)	-0.04 (0.05)	-0.02 (0.02)	-0.02 (0.04)
Forecast × Ind Bin 3	0.00 (0.04)	-0.07 (0.06)	-0.02 (0.03)	-0.04 (0.05)
q-val Tercile 1	0.074	0.037	0.037	0.037
q-val Tercile 2	1.000	1.000	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.329	0.401	0.245	0.534
Insur. = Ter. 3	0.863	0.851	0.546	0.963
Control Mean	0.15	0.22	0.06	0.15
Observations	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on migration, estimated using Equations (3, panel A) and (4, panel B). Any Migrant is an indicator for any migrant having left the household in the past 12 months. Num Temp Mig. is the number of temporary migrants who left the household in the last 12 months. N.Female and N. Male are the number of temporary female and male migrants, respectively. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.9 Trust

Figure A.6: Farmer trust in the forecast



Notes: This figure presents farmers' stated trust in the forecast, measured on a 1–10 scale. The solid histogram presents trust in the forecast when farmers received the information, while the hollow histogram presents trust after the monsoon had arrived. The sample includes only farmers in the forecast treatment group.

A.10 Shocks

Table A.10: Shock realizations across treatments

	Panel A: Forecast vs. Insurance				
	(1) Flood	(2) Drought	(3) Animal	(4) Cyclone	(5) Any
Forecast	-0.04 (0.03)	0.01 (0.02)	0.03 (0.03)	0.02 (0.03)	0.03 (0.04)
Insurance	0.02 (0.04)	0.04* (0.02)	-0.03 (0.03)	0.17*** (0.04)	0.12*** (0.04)
q-val Forecast	1.000	1.000	1.000	1.000	1.000
q-val Insurance	0.274	0.093	0.197	0.001	0.008
	Panel B: Forecast Terciles				
	(1) Flood	(2) Drought	(3) Animal	(4) Cyclone	(5) Any
Forecast × Ind Bin 1	-0.07 (0.04)	0.01 (0.03)	0.01 (0.04)	0.04 (0.04)	0.01 (0.05)
Forecast × Ind Bin 2	-0.05 (0.04)	0.00 (0.02)	0.05 (0.03)	0.04 (0.05)	0.05 (0.06)
Forecast × Ind Bin 3	0.03 (0.06)	0.03 (0.03)	0.01 (0.05)	-0.05 (0.07)	0.04 (0.06)
q-val Tercile 1	1.000	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	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.184	0.504	0.929	0.208	0.674
Test Insur. = Ter. 3	0.915	0.845	0.383	0.001	0.222
Control Mean	0.24	0.07	0.12	0.31	0.67
Observations	1201	1201	1201	1201	1201

Notes: This table presents estimates of the difference in shock realizations across treatment groups estimated using Equations (3, panel A) and (4, panel B). All outcomes are indicators for self-reported crop damage resulting from a particular shock type. Flood is an indicator for flood damage, Drought for damage from too little rain, Animal for damage from animals eating or trampling crops, Cyclone for damage from wind or excessive rain, and Any is an indicator for suffering losses from any of these four shocks. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.11 Insurance and prior beliefs

Table A.11: Effect of insurance on inputs by prior terciles

	(1) Land Ha.	(2) Cash Crop	(3) Total Inputs	(4) Invest Index
Insurance × Ind Bin 1	0.486** (0.226)	0.023 (0.065)	415.311* (217.148)	0.172* (0.095)
Insurance × Ind Bin 2	0.110 (0.172)	0.092* (0.051)	274.439 (172.141)	0.157** (0.076)
Insurance × Ind Bin 3	-0.073 (0.303)	0.048 (0.077)	59.705 (243.760)	0.017 (0.127)
q-val Insure Ter. 1	0.092	0.321	0.092	
q-val Insure Ter. 2	0.210	0.201	0.201	
q-val Insure Ter. 3	1.000	1.000	1.000	
Test Tercile 1=3	0.140	0.802	0.254	0.323
Control Mean	2.12	0.51	1443.49	0.00
Observations	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of insurance on farm inputs by prior belief terciles. Land Ha. is cultivated land in hectares. Cash Crop is an indicator for growing at least one cash crop. Total Inputs is the total amount spent on all inputs, including all previous outcomes and any other costs reported by farmers, in USD. Invest Index is an inverse covariance weighted index of land cultivated, cash crop cultivation, and total input expenditure. We exclude it 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 tercile 1 were the most optimistic. Prior bin 2 had average (correct) beliefs. Prior bin 3 were the most pessimistic. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Additionally, each regression also controls for the forecast treatment, with separate controls by prior belief, as in Equation (4). Sharpened q -values are adjusted across all outcomes shown (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.12 Additional welfare results

Table A.12: Effect of the forecast and insurance on economic well-being, no flood shock

	(1) Food cons	(2) Other cons	(3) Asset value	(4) Livestock	(5) Net savings	(6) Welfare index
Forecast	1.32** (0.65)	-0.59 (1.06)	42.70 (160.89)	0.01 (0.03)	25.87 (205.65)	0.04 (0.05)
Insurance	0.80 (0.75)	2.35 (1.88)	-55.39 (229.95)	-0.02 (0.05)	-779.87*** (275.25)	-0.06 (0.06)
q-val Forecast	0.274	1.000	1.000	1.000	1.000	
q-val Insurance	0.618	0.618	0.913	0.913	0.025	
Control Mean	13.57	9.83	1546.97	0.19	-839.15	0.02
Observations	554	554	554	554	516	554

Notes: This table presents estimates of the treatment effects of forecasts and insurance on well-being, restricting the sample to those who did not report losses due to flooding or cyclones. The estimation follows Equation (3). Food cons is food consumption per household member in USD over the past 30 days. Other cons is non-food consumption (other than medical expenses). Asset value is the value of assets in USD. Livestock is the count of livestock in Tropical Livestock Units. Net savings is savings less outstanding debt in USD. Welfare index is an inverse covariance weighted index of the other five outcomes in the table. We exclude the index from the MHT correction, as it is a composite of outcomes already included. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted across all outcomes shown (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.13: Effect of forecast and insurance on economic well-being by prior tercile

	(1) Food cons	(2) Other cons	(3) Asset value	(4) Livestock	(5) Net savings	(6) Welfare index
Forecast ×	0.32	-0.48	349.12	0.07*	372.18	0.14**
Ind Bin 1	(0.65)	(1.45)	(264.78)	(0.04)	(252.85)	(0.06)
Forecast ×	1.19**	-0.80	29.35	-0.04	59.12	-0.00
Ind Bin 2	(0.53)	(0.97)	(153.37)	(0.03)	(228.37)	(0.04)
Forecast ×	1.14	-0.28	-135.55	0.02	121.49	0.05
Ind Bin 3	(0.84)	(1.45)	(145.77)	(0.05)	(253.75)	(0.05)
Insurance	0.45	1.91*	-115.78	0.01	-406.54*	-0.02
	(0.47)	(1.00)	(158.40)	(0.03)	(217.14)	(0.05)
q-val Tercile 1	0.455	0.455	0.455	0.455	0.455	
q-val Tercile 2	0.143	1.000	1.000	0.928	1.000	
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000	
q-val Insurance	0.492	0.182	0.536	0.782	0.182	
Test Tercile 1=3	0.412	0.917	0.094	0.466	0.515	0.288
Test Insur. = Ter. 3	0.449	0.147	0.906	0.840	0.064	0.285
Control Mean	13.22	9.93	1503.10	0.22	-1031.41	0.00
Observations	1201	1201	1201	1201	1129	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on well-being. The estimation follows Equation (4). Food cons is food consumption per household member in USD over the past 30 days. Other cons is non-food consumption (other than medical expenses). Asset value is the value of assets in USD. Livestock is the count of livestock in Tropical Livestock Units. Net savings is savings less outstanding debt in USD. Welfare index is an inverse covariance weighted index of the other five outcomes in the table. We exclude the index from the MHT correction, as it is a composite of outcomes already included. Bins 1–3 indicate the prior tercile for a respondent. Prior tercile 1 were the most optimistic. Prior bin 2 had average (correct) beliefs. Prior bin 3 were the most pessimistic. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted across all outcomes shown (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.14: Effect of the forecast and insurance on per-capita consumption (disaggregated)

Panel A: Forecast vs. Insurance									
	(1) Cereals	(2) Milk	(3) Tab / Alc	(4) Meat	(5) Mobile	(6) Clothing	(7) Medicine	(8) Celebration	(9) Total
Forecast	0.64** (0.29)	0.19* (0.11)	-0.68** (0.27)	-0.07 (0.17)	0.08 (0.06)	0.14 (0.35)	-0.27 (0.64)	-0.23 (0.36)	-1.02 (1.72)
Insurance	0.11 (0.32)	0.11 (0.13)	-0.36 (0.38)	0.20 (0.20)	0.11 (0.07)	0.40 (0.42)	-0.55 (0.72)	0.60 (0.48)	0.27 (1.92)
q-val Forecast	0.133	0.212	0.105	0.626	0.307	0.626	0.626	0.626	0.626
q-val Insurance	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Panel B: Forecast Terciles									
Forecast × Ind Bin 1	0.16 (0.41)	0.05 (0.17)	-1.00*** (0.38)	0.11 (0.27)	0.03 (0.09)	1.36** (0.58)	-0.28 (1.03)	-0.17 (0.70)	-1.76 (2.93)
Forecast × Ind Bin 2	0.99*** (0.38)	0.15 (0.14)	-0.61* (0.36)	-0.15 (0.22)	0.09 (0.08)	-0.76 (0.48)	0.01 (0.84)	-0.46 (0.43)	-0.76 (2.24)
Forecast × Ind Bin 3	0.84 (0.54)	0.48** (0.22)	-0.30 (0.54)	-0.25 (0.34)	0.11 (0.10)	-0.28 (0.60)	-0.85 (1.13)	0.10 (0.55)	-0.81 (2.82)
q-val Tercile 1	1.000	1.000	0.084	1.000	1.000	0.084	1.000	1.000	1.000
q-val Tercile 2	0.088	0.514	0.432	0.740	0.514	0.432	0.792	0.514	0.792
q-val Tercile 3	0.982	0.350	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.297	0.104	0.272	0.383	0.513	0.041	0.693	0.761	0.807
Test Insur. = Ter. 3	0.208	0.112	0.909	0.206	0.923	0.276	0.784	0.464	0.716
Control Mean	7.28	1.96	3.23	3.80	1.55	2.64	6.34	1.71	32.46
Observations	1200	1201	1201	1201	1201	1200	1200	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on disaggregated per-capita consumption expenditure categories, estimated using Equations (3, panel A) and (4, panel B). Cereals is spending on rice, millet, suji, ragi, or any other grain. Milk is spending on dairy products. Tab / Alc is spending on tobacco or alcohol. Meat is spending on chicken, beef, goat, fish, or eggs. Mobile is spending on phone credit. Clothing is spending on any clothing for household members. Medicine is spending on medical expenses. Celebrations is spending on celebrations or festivals. All outcomes are in USD spent per household member during the past 30 days, and are winsorized at the 5th and 95th percentile. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.15: Effect of the forecast and insurance on household finances

Panel A: Forecast vs. Insurance					
	(1) Savings	(2) Took Loan	(3) Debt Out	(4) Missed Payment	(5) Farm Loan
Forecast	-14.20 (23.91)	-0.06 (0.04)	-193.13 (137.96)	-0.11* (0.06)	-0.09** (0.04)
Insurance	-47.90** (21.52)	0.18*** (0.04)	374.72* (206.98)	-0.01 (0.06)	0.18*** (0.04)
q-val Forecast	0.254	0.175	0.194	0.175	0.085
q-val Insurance	0.027	0.001	0.046	0.211	0.001

Panel B: Forecast Terciles					
Forecast × Ind Bin 1	-48.58 (33.78)	-0.07 (0.05)	-449.03* (238.27)	-0.15* (0.09)	-0.09* (0.05)
Forecast × Ind Bin 2	-0.84 (31.61)	-0.06 (0.05)	-47.90 (206.07)	-0.17* (0.09)	-0.11** (0.05)
Forecast × Ind Bin 3	19.67 (41.85)	-0.02 (0.07)	-30.84 (256.21)	0.02 (0.14)	-0.04 (0.07)
q-val Tercile 1	0.171	0.171	0.171	0.171	0.171
q-val Tercile 2	0.644	0.265	0.644	0.176	0.176
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.163	0.582	0.258	0.254	0.566
Test Insur. = Ter. 3	0.127	0.008	0.163	0.879	0.006
Control Mean	149.23	0.50	1173.75	0.43	0.47
Observations	1129	1201	1201	269	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on household finances, estimated using Equations (3, panel A) and (4, panel B). Savings is total savings in USD, Took Loan is an indicator for whether the household took a loan in the last 12 months, Debt Out is the amount of outstanding debt in USD, Missed Payment is an indicator for having missed a loan payment in the last 12 months, and Farm Loan is an indicator for having taken a farm loan in the last 12 months. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.16: Effect of the forecast and insurance on mental health

Panel A: Forecast vs. Insurance									
	(1) Q1	(2) Q2	(3) Q3	(4) Q4	(5) Q5	(6) Q6	(7) Q7	(8) Q8	(9) Norm
Forecast	0.04 (0.09)	-0.01 (0.08)	0.10 (0.07)	0.00 (0.09)	0.17** (0.07)	0.10 (0.09)	-0.07 (0.08)	0.11 (0.07)	0.07 (0.05)
Insurance	0.14 (0.14)	-0.17** (0.09)	-0.11 (0.08)	-0.06 (0.09)	-0.03 (0.07)	0.03 (0.09)	0.14 (0.09)	0.03 (0.08)	-0.01 (0.05)
q-val Forecast	0.868	0.868	0.654	0.868	0.207	0.704	0.800	0.654	0.654
q-val Insurance	0.956	0.627	0.698	1.000	1.000	1.000	0.698	1.000	1.000

Panel B: Forecast Terciles									
Forecast × Ind Bin 1	0.06 (0.21)	0.00 (0.12)	0.19 (0.12)	0.11 (0.14)	0.26** (0.12)	0.06 (0.13)	-0.03 (0.16)	0.23** (0.11)	0.14* (0.08)
Forecast × Ind Bin 2	-0.00 (0.09)	-0.05 (0.11)	0.10 (0.10)	-0.02 (0.10)	0.16** (0.08)	0.11 (0.11)	-0.12 (0.09)	0.10 (0.09)	0.04 (0.07)
Forecast × Ind Bin 3	0.05 (0.18)	0.09 (0.16)	-0.06 (0.16)	-0.11 (0.17)	0.03 (0.14)	0.14 (0.20)	-0.00 (0.16)	-0.04 (0.12)	0.02 (0.13)
q-val Tercile 1	1.000	1.000	0.264	0.725	0.210	1.000	1.000	0.210	0.210
q-val Tercile 2	1.000	1.000	1.000	1.000	0.733	1.000	1.000	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.981	0.670	0.218	0.310	0.175	0.739	0.916	0.106	0.420
Test Insur. = Ter. 3	0.606	0.120	0.779	0.785	0.678	0.605	0.362	0.582	0.887
Control Mean	1.56	0.97	0.73	1.29	0.47	0.67	0.66	0.50	-0.02
Observations	1201	1201	1201	1201	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on mental health, estimated using Equations (3, panel A) and (4, panel B). We measure outcomes using the PhQ-8 screening tool, a standard and locally-validated depression metric (Bhat et al. (2022)). Outcomes for questions 1-8 are measured in the number of days in the past seven that the respondent agreed with the question prompt. Norm (column 9) is standardized PhQ-8 score. Q1 asked about having little pleasure in doing things, Q2 feeling depressed or hopeless, Q3 having trouble sleeping or sleeping too much, Q4 having little energy, Q5 having poor appetite or overeating, Q6 feeling bad about yourself or think you have little worth, Q7 trouble concentrating, and Q8 moving or speaking slowly so others do not notice. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B Panel analysis: Effect of onset timing on crop yields

This Appendix provides additional details about our data and sample construction for our historical analysis of the effect of monsoon onset timing on crop yields across India (described in Section 2).

B.1 Data

Monsoon onset data Our precipitation data come from the European Centre for Medium-Range Weather Forecasting Reanalysis dataset (ERA5). To convert this precipitation information into data on monsoon onset at the grid-cell level, we follow Moron and Robertson (2014)’s definition: “the first wet day (≥ 1 mm) of the first 5-day wet sequence from April 1st that receives at least the 5-day wet spell interannual mean in April – October for that pixel.” To avoid false positives — cases when a wet spell is followed by drought — an onset date cannot be followed “by a 10-day dry spell (receiving less than 5 mm) in the following 30 d from the onset.” We adjust the first-day rainfall threshold to from 1mm to 4mm, to better match the Indian context. Specifically, our algorithm for defining monsoon onset is:

0: p_{idy} , our unit of observation, is precipitation in mm/day in pixel i on day d in year y .

I: For each observation between April and October, check if they can potentially be “the first wet day” observations: check if $p_{idy} \geq 4$. Denote these possible first wet days as $p_{i^*d^*y^*}^{*}$.⁴⁹

II: For each potential “first wet day” check if the subsequent 5-day wet sequence “receives at least the climatological 5-day wet spell amount”: check if

$$\sum_{k=0}^4 p_{i^*,d^*+k,y^*}^{*} \geq \frac{1}{N} \sum_y \sum_{d=Apr,1}^{D=Oct,27} \sum_{k=0}^4 p_{i^*,d+k,y}$$

where N is the number of all day-year pairs. In other words, for each pixel, we check if 5-day wet spell is greater or equal than the average 5-day April-October spell in all years.

III: For each potential “first wet day” check if the subsequent wet spell is not interrupted by a 10-day dry spell (receiving less than 5 mm) in the following 30 days from the onset: check if

$$\min_{j \in \{0, \dots, 20\}} \sum_{k=0}^9 p_{i^*,d^*+k+j,y^*}^{*} \geq 5$$

⁴⁹The asterisk denotes a constant, not a parameter.

IV: Choose a subset of observations p_{i*d*y}^* for which inequalities from *Step II* and *Step III* hold.

Now, for each year-pixel pair we can take its own subset and we denote monsoon onset as the smallest d among this subset. If the subset is empty, we say that monsoon is not defined.

Finally, we construct district-level onset data by taking the area-weighted-average onset date of all pixels that lie (fully or partially) within each district. We standardize each district’s onset dates according to the district-specific mean and variance of the onset.


Agriculture data We obtain district-level kharif season crop yield data from the Indian Ministry of Agriculture and Farmers’ Welfare. Our data spans 2001–2018. Districts are defined according to the 2001 Census of India. We focus our analysis on two major crops: rice, a key staple, and cotton, a key cash crop.

B.2 Defining monsoonal regions

We are interested in the causal effect of monsoon onset timing on crop yields. While most of India is characterized by monsoonal rainfall patterns, with little-to-no rain outside the main monsoon season, some regions have other rainfall patterns. We therefore use a k -means clustering algorithm to classify pixels as monsoonal or non-monsoonal based on the ERA5 data, and restrict our analysis to the monsoonal regions.

We first apply a low-pass recursive first-order filter with a cutoff of $\frac{1}{30}$ to the ERA5 gridded daily precipitation data. We use a Butterworth filter, though our monsoonal area definition is robust to alternatives. Second, we standardize the filtered data to have a mean of zero and variance of one. Third, we find a leading empirical orthogonal function for standardized filtered data for each pixel. Finally, we run a standard k -means clustering algorithm on these data. We use 5 clusters following Moron et al. (2017) but monsoonal area is mostly robust if we decrease the number of clusters. Figure B.1 shows the results of this k -means clustering on a grid level (left panel) and the low-passed averaged rainfall of all points inside the cluster (in mm/d) with ± 1 standard deviation (right panel). Based on these patterns, we decide to treat either grey areas or grey and red areas as monsoonal. We restrict the sample in our regression analysis to these monsoonal regions.

Figure B.1: Defining India's monsoonal regions: k -means clustering



Notes: This figure plots our definition of India's monsoonal regions, defined using our k -means clustering approach. The left panel shows the results of the k -means clustering of Indian rainfall regimes into 5 different clusters. The right panels show the average rainfall pattern during the Kharif season for these clusters, with a strong seasonal pattern in rainfall — characteristic of the monsoon — seen only for the red and grey clusters.

C Model details

C.1 Setup

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 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 $i \in \{1, \dots, S\}$, 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 $x > 0$. 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.⁵⁰

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:

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

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), and p is the price of the input x , and β is the discount factor.

We next turn to optimal farmer behavior, and then study how forecasts and insurance would affect these decisions.

C.2 Optimal farmer investment and saving decisions

We present first-order conditions to illustrate how beliefs affect farmers' decisions.

⁵⁰The investment level x can also be interpreted as a continuum of crop choices, with varying levels of productivity. These productivities depend on the state and are correlated with how expensive each crop is to plant. In that sense, for any given state, there is an optimal crop choice x that would maximize production subject to budget constraints.

Savings The first-order condition for savings s implies the following Euler equation:⁵¹

$$\beta = \frac{u'(c_1)}{\mathbf{E}[u'(c_2)]} \quad (\text{C.2})$$

where $\mathbf{E}[u'(c_2)]$ is the expected consumption in the second period:

$$\mathbf{E}[u'(c_2)] = \sum_i u'(c_2^i, \epsilon_i) g(\epsilon_i) \quad (\text{C.3})$$

Thus, conditional on investment level x , farmers choose savings such that the ratio of marginal utilities between the first and second period equals the patience parameter (discount factor) β .

Investment The first-order condition for investment x implies that investment prices should equal a weighted marginal product:

$$p = \mathbf{E}[w f'(x)] \quad (\text{C.4})$$

where $\mathbf{E}[w f'(x)]$ is the (weighted) expected marginal product of investment level x :

$$\mathbf{E}[w f'(x)] = \sum_i w(c_1, c_2^i, \epsilon_i) f'(x, \epsilon_i) g(\epsilon_i) \quad (\text{C.5})$$

with weights:

$$w(c_1, c_2^i, \epsilon_i) = \beta \frac{u'(c_2^i, \epsilon_i)}{u'(c_1)} = \frac{u'(c_2^i, \epsilon_i)}{\mathbf{E}[u'(c_2)]} = w(c_2^i, \epsilon_i), \quad (\text{C.6})$$

where the second equality comes from plugging in the FOC for savings in (C.2).

The farmer thus sets investment levels to at expected marginal products over all states, weighting states by their relative marginal utility of consumption. While the investment decision deals with smoothing consumption across states in the second period, the savings decision smooths consumption across periods.

Forecasts Consider first a forecast that shifts beliefs from late G_l to early G'_e . In other words, G'_e puts higher probability ϵ_i for higher i . Suppose the farmer was previously solving the problem with G_l , setting optimal investment levels at x^l :

$$\mathbf{E}_{G_l}[w f'(x^l)] = p \quad (\text{C.7})$$

⁵¹The results are qualitatively unchanged with additional constraints that limit borrowing and savings:

$$\underline{s} \leq s \leq \bar{s}$$

Conditional on weights w , the previous investment level x^l has larger marginal product under the new beliefs G'_e :

$$\mathbf{E}_{G'_e}[wf'(x^l)] > \mathbf{E}_{G_l}[wf'(x^l)] = p \quad (\text{C.8})$$

This is because the new beliefs are weighted toward higher states, which have higher marginal product at any x (f' rises with ϵ). To meet the optimal marginal product of p , the farmer must thus lower the marginal product by raising x (f' is concave). Thus, the optimal investment level increases:

$$x^e > x^l \quad (\text{C.9})$$

By symmetry, a forecast that shifts beliefs from early G_e to late G'_l would *decrease* investment levels.

The argument above is conditional on weights w , that capture the relative marginal utility of consumption across states. To the degree farmers are risk averse, they will reduce investment levels x so as to smooth consumption across states. Suppose now that farmers shift beliefs from G_l to G'_e . For any given investment level of x , the farmer's beliefs shift the expected w toward higher states, which have lower marginal utility. While the marginal product is higher in higher states, the weights are higher in lower states. This mechanism would thus *lower* the weighted marginal product $\mathbf{E}_{G'_e}[wf'(x^l)]$ in contrast to the mechanism above. Thus, changes in investment from forecasts are dampened by the degree of risk aversion (concavity of u).

Insurance To incorporate insurance, we now include an additional payout b that occurs in the second period, depending on the state:

$$c_2^i = f(x, \epsilon_i) + s + b \cdot 1\{\epsilon_i \in S_I\},$$

where E is the set of (low) states for which the insurance payout applies. Note that because this additional term is not a function of either investment or savings, the first-order conditions are unchanged.

Under insurance, the following changes occur *ceteris paribus*: for low states, c_2^i increases from the payouts, causing $u'(c_2^i)$ to fall by concavity the weights; for high states, c_2^i is unchanged; on net, $\mathbf{E}[u'(c_2^i)]$ falls. Thus, the weights $w(c_2^i, \epsilon_i)$ in (C.6) will fall for low states (because $u'(c_2^i)$ falls) and rise for high states (because $\mathbf{E}[u'(c_2^i)]$ falls). Intuitively, for the investment decision, farmers now place relatively higher weight on higher states, as insurance allows them to smooth relatively more. Because higher states are more productive, this raises the optimal level of investment.

Note that these effects are heterogeneous. If farmers have *early* priors, they place higher prob-

ability weight on *low* states, dampening the above channel. Thus, insurance would cause these farmers to increase investment relatively *less*. In contrast, if farmers have later priors, they will increase investment relatively more in response to insurance.

C.3 Parametrization for simulations

To quantitatively simulate farmer behavior under various counterfactuals, we impose functional form assumptions.

Utility Farmers' preferences have constant relative risk aversion (CRRA):

$$u(c) = \frac{c^{1-r} - 1}{1-r} \quad (\text{C.10})$$

Production The technology is Cobb-Douglas in investment:

$$f(x, \epsilon) = \bar{z} \cdot z(\epsilon) \cdot x^\alpha \quad (\text{C.11})$$

where $z(\epsilon) \in (0, 1)$ is a (logistic) productivity shock that increases with the state ϵ :

$$z(\epsilon) = \frac{1}{4k} \exp\left(-\frac{\epsilon}{k}\right) \left[1 + \exp\left(-\frac{\epsilon}{k}\right)\right]^{-2} \quad (\text{C.12})$$

The scale parameter k governs how states map into productivity, with lower values driving larger productivity differences across states.

Beliefs and updating The set of possible states S is discrete with 40 possible values $\epsilon_1, \dots, \epsilon_{40}$. This is distributed according a (rescaled) normal distribution with mean μ and standard deviation parameter σ that is unknown to the farmer:

$$\bar{g}(\epsilon) = \frac{\phi(\epsilon, \mu, \sigma)}{\sum_i \phi(\epsilon_i, \mu, \sigma)} \quad (\text{C.13})$$

where $\phi(\cdot, \mu, \sigma)$ is the PDF of a normal distribution. Farmers have (potentially incorrect) prior beliefs with mean μ_p and SD σ_p :

$$g(\epsilon) = \frac{\phi(\epsilon, \mu_p, \sigma_p)}{\sum_i \phi(\epsilon_i, \mu_p, \sigma_p)} \quad (\text{C.14})$$

The forecast distribution is centered around the actual mean μ with SD σ_f that reflects forecast accuracy:

$$h(\epsilon) = \frac{\phi(\epsilon, \mu, \sigma_f)}{\sum_i \phi(\epsilon_i, \mu, \sigma_f)} \quad (\text{C.15})$$

Upon receiving forecast h , the farmer updates from prior g to posterior g' in a Bayesian fashion:

$$g'(\epsilon) = \frac{\phi(\epsilon, \mu', \sigma')}{\sum_i \phi(\epsilon_i, \mu', \sigma')} \quad (\text{C.16})$$

where the posterior mean μ' is a variance-weighted average of the prior and forecast means:

$$\mu' = \frac{\sigma_f^2 \mu_p + \sigma_p^2 \mu}{\sigma_p^2 + \sigma_f^2} \quad (\text{C.17})$$

and the posterior SD σ' scales down the prior in proportion to the (relative) forecast SD:

$$\sigma' = \frac{\sigma_p \sigma_f}{\sqrt{\sigma_p^2 + \sigma_f^2}} \quad (\text{C.18})$$

The parameters are set according to Table C.1 below. Note that we choose parameters such that even the most optimistic farmers believe they face some agriculture risk. This is necessary for the strictly decreasing relationship between insurance treatment effects and priors.

C.4 Model predictions for alternative forecast realizations


Appendix Figure C.1 plots treatment effects of a forecast in our model under a forecast of an average monsoon (as depicted in Figure 1 in the main text), a forecast of an early monsoon, or a forecast of a late monsoon. The central curve replicates the effects of a forecast of an average monsoon. 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.

Table C.1: Parameters for model simulation

Parameter	Description	Value
<i>Panel A: Utility Parameters</i>		
r	Relative risk aversion	0.5
β	Discount factor	0.95
y	Starting wealth	5
p	Input price	1
<i>Panel B: Production Parameters</i>		
α	Production function curvature	0.6
$\bar{z}$	Max productivity	3
k	Scale parameter of productivity	2
<i>Panel C: State Parameters</i>		
S	Possible states	$-10, -9.5, -9, \dots, 9.5, 10$
μ	Mean of actual & forecast distribution	0
σ_f	SD of forecast (accuracy)	2
σ_p	SD of farmer beliefs	5
<i>Panel D: Insurance Parameters</i>		
S_I	States for insurance payout	$-10, -9.5, \dots, -3.5$
b	Insurance payout	3

Notes: This table presents the parameters used in our model simulation, as well as their assumed values (or range of values).

Figure C.1: Investment choice with a forecast, alternative realizations (model)



Notes: This figure plots the simulated relationship in our model between the treatment effect of forecasts on optimal investment and farmers' priors prior. 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 when farmers believe the monsoon will arrive. This plot 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 upper 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.

D Becker et al. (1964) appendix

To elicit WTP for the given product, we use a Becker et al. (1964) (BDM) mechanism. We explain a two-step procedure to the household. In the first step, the household states their WTP. Then, the enumerator reveals an INR value written on the tablet. If the value listed on the tablet is above the household’s stated WTP, the household does not get to purchase the product and their cash is returned. If the value is below the household’s WTP, the household purchases the product and the cash goes to the enumerator. Because it is vital that this procedure is thoroughly understood by households before they begin, the enumerator plays a “practice” round with a common household product (e.g., a bar of soap). Therefore, any misunderstanding about the process will be resolved before the BDM procedure for the product of interest (i.e., the forecast or insurance) is started.

D.1 Methodological overview

The BDM mechanism is an incentive compatible process through which a rational participant should reveal their true maximum WTP. We implement the BDM procedure using the following steps, modeled closely after Berkouwer and Dean (2022):

1. Prior to the baseline visit, we assign each participant a random BDM price drawn from either the forecast or insurance distribution of BDM prices (described below).
2. Each enumerator is then given a sealed envelope that contains that BDM price (in INR) for the participants they are visiting that day. The enumerators are not aware of the assigned prices.
3. When the BDM procedure begins, the enumerator places the sealed envelope so that participant can see it.
4. Beginning with a starting price of INR 500 for both the forecast and insurance, the enumerator asks if the participant would commit to purchasing the respective product at that price. If the participant agrees, the enumerator subsequently increases the price by INR 500 and asks again if the participant would be willing to purchase the product at this new price. If the participant again agrees to purchase the product, the price is again raised by INR 500. If the participant declines this new price, the enumerator reduces the prices by INR 250.

Instead, if the participant declines to buy the product at the initial price, the enumerator lowers the price by half (to 250) and asks again if the participant would be willing to purchase at this new, lower price. This process is repeated 11 times with the relevant intervals shrinking

each iteration (or until the relevant interval drops below 1 rupee), so that by the end of the process we approach the participant's true WTP.

For concreteness, we illustrate the beginning iterations of this process:

- (a) If the envelope said the price was INR 500, would you choose to purchase the forecast / insurance?
 - i. If yes: If the envelope said the price was INR 1,000, would you choose to purchase the forecast / insurance?
 - A. If yes: If the envelope said the price was INR 1,500, would you choose to purchase the forecast / insurance?
 - Etc.
 - B. If no: If the envelope said the price was INR 1,250 would you choose to purchase the forecast / insurance?
 - Etc.
 - ii. If no: If the envelope said the price was INR 250, would you choose to purchase the forecast / insurance?
 - A. If yes: If the envelope said the price was INR 375, would you choose to purchase the forecast / insurance?
 - Etc.
 - B. If no: If the envelope said the price was INR 125, would you choose to purchase the forecast / insurance?
 - Etc.

At the end of this process, the enumerator confirms that the participant fully understands their decision and the consequences of once the envelope is opened. They then ask that the participant retrieves the agreed upon amount in cash and place the bank notes next to the envelope containing the price. Finally, they will allow the participant a final chance to change their mind before the envelope is opened.

- 5. Once the participant has confirmed the price and has placed the cash, the participant and the enumerator together open the envelope and reveal the price.
- 6. If the participant's maximum WTP is lower than the BDM price in the envelope, the participant will not be able to purchase the forecast / insurance and will instead take back their cash.

7. If the participant's maximum WTP is at least as high as the BDM price in the envelope, the participant purchases the forecast / insurance, paying the price that was written inside the envelope out of their cash.


D.2 Distribution of BDM prices

We set the distribution of BDM price draws to low values so that nearly all farmers with positive willingness to pay will ultimately purchase the forecast or insurance product. In this way, we will increase power by maximizing adoption of each product without compromising the incentive compatibility of the BDM procedure. To this end, neither the participants nor the enumerators will be informed about the underlying price distribution. We choose the following distributions for each product:

- For the forecast product, 95% of participants will receive a price of zero while the remaining 5% of prices will be drawn from a uniform distribution ranging from 1 to 100 INR.
- For the insurance product, 95% of participants will receive a price of zero while the remaining 5% of prices will be drawn from a uniform distribution ranging from 1 to 100 INR.


E Information sheets

Figure E.1: Forecast information sheet



Notes: We provided farmers with this information sheet about the forecast when offering them the product through the BDM mechanism described in Section 4 and Appendix D.

Figure E.2: Insurance information sheet



Notes: We provided farmers with this information sheet about the insurance product when offering them the product through the BDM mechanism described in Section 4 and Appendix D.

F Deviations from our pre-analysis plan

This experiment was pre-registered with the AEA as Trial No. AEARCTR-0008846 and accepted by the *Journal of Development Economics* via pre-results review. We have endeavored to follow the PAP as closely as possible, but have nevertheless had some deviations, which we list here. Changes to regression specifications are noted with footnotes in the main text.

- **Data.** Due to time constraints, we left out several variables from our baseline survey: information on time preferences and intra-household bargaining, both of which we had planned to use in heterogeneity analysis.
- **Data.** Due to time constraints, we left out several variables from our endline survey: information on how much of each planted crop had spoiled, was already consumed, and was stored. We had intended to use these as supplementary outcome measures. We instead focus only on production in this analysis.
- **Outcome variables.** We pre-specified measuring agricultural inputs on a per-acre basis. In the main text, we instead use total expenditure, which we believe better reflects decisions to expand agricultural investment. This is because households ought to make a joint decision to expand land and inputs, maintaining a similar input-to-land ratio. We present results on a per-acre basis in Appendix Table G.10.
- **Outcome variables.** In addition to our pre-specified variables on input expenditure, we add an investment index to Table 4. This is complementary to the q -value approach to dealing with multiple hypotheses, serving as a single summary measure of *ex ante* behavior change. An advantage of the index over the FWER correction is that this index accounts for changes in the *direction* of different measures of investment, while the FWER approach only considers p -values irrespective of sign.
- **Outcome variables.** We pre-specified a comparison between 2022 Kharif crop choice and *planned* 2022 Kharif crop choice (measured at baseline). In the main text, we instead compare 2022 Kharif crop choice to 2021 Kharif crop choice, because this is a revealed preference measure rather than a stated preference measure. We include the stated preference result in Appendix Table G.9 for completeness.
- **Analysis.** For the correlations between WTP and prior beliefs (described in Section 5 and presented in Appendix Tables G.5, G.6, and G.7, we erroneously pre-specified a regression equation that included strata fixed effects and controls chosen by double-selection LASSO.

However, these regressions include only a single experimental group at a time (and do not include the control group), meaning that these control variables remove useful variation rather than adding precision. We therefore omit these controls from the tables.

- **Analysis.** For the correlations between WTP and prior beliefs, we had pre-specified a regression that included standard deviation and squared standard deviation of farmers' prior distributions on the right-hand side to test for possible non-linearity in the relationship between WTP and prior strength. Appendix Table G.5 additionally uses the absolute distance between the share of the prior distribution above an on-time cutoff and an early cutoff and 0.5, because we believe this is easier to interpret. For insurance, our theory predicts that WTP strictly falls with an increase in the farmer's belief that the coming year will be good. We therefore use the simple share before the farmer's on-time cutoff and share before the farmer's early cutoff as regressors in Appendix Table G.7, rather than the difference between the shares and 0.5.
- **Analysis.** For the belief change regressions, we pre-specified heterogeneity with respect to multiple measures of prior strength. Here, we present results with respect to prior SD only, as our outcome measures are all relative to the prior or the forecast (and therefore we do not have specific predictions of movement on the basis of binned prior strength).
- **Analysis.** We pre-specified that we would estimate separate treatment effects for forecast farmers receiving bad news vs. bad news. Because the forecast in 2022 was for an average monsoon, there is a large mass of farmers with priors that are very close to the forecast. We therefore estimate treatment effects by *tercile* of prior, which splits the sample into an optimistic group (who receives bad news), an accurate group (who receives neutral news), and a pessimistic group (who receives good news). Given that the forecast itself gave a date range for the monsoon arrival, and that theoretically we would not expect changes in behavior for neutral news farmers, we believe our current approach is a better representation of the impact of the forecast on farmer decisions. This avoids the attenuation bias that would be created by including the neutral news group in the good news and bad news groups.
- **Analysis.** We pre-specified that we would estimate heterogeneous treatment effects by the change in belief (absolute difference between prior and posterior). However, this is endogenous and therefore difficult to interpret, so we omit it here.
- **Analysis.** We pre-specified that we would estimate treatment effects on crop prices. Because our survey was conducted in early December, many farmers had not yet sold their crops,

leading our individual price data to be extremely noisy and poorly aligned with administrative data on prices. We therefore use district median prices for all outcomes involving crop value, and omit crop price results.

- **Analysis.** We pre-specified that we would estimate treatment effects on crop prices. Because our survey was conducted in early December, many farmers had not yet sold their crops, leading our individual price data to be extremely noisy and poorly aligned with administrative data on prices. We therefore use district median prices for all outcomes involving crop value, and omit crop price results.
- **Analysis.** We did not pre-specify the analysis on crop losses, profits net of losses, and profits for non-flood-affected farmers (Columns (2)–(6) of Table 6).

G Additional pre-specified results

G.1 Pooled forecast treatment results

Table G.1: Effect of the forecast (pooled) and insurance on land use and cropping

	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	-0.119 (0.111)	0.057* (0.032)	0.020 (0.037)	-0.012 (0.039)	0.005 (0.028)
Insurance	0.178 (0.136)	0.062 (0.038)	0.045 (0.046)	0.044 (0.048)	-0.005 (0.037)
q-val Forecast	1.000	1.000	1.000	1.000	1.000
q-val Insurance	0.346	0.279	0.437	0.437	0.963

Notes: This table presents estimates of the treatment effects of forecasts and insurance on farmers' land use and cropping decisions, estimated using Equation (4). 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 an additional crop in 2022 as compared to 2021. Sub Crop is an indicator for removing a crop in 2022 as compared to 2021. Sharpened q -values are adjusted across all outcomes in Tables G.1 and G.2 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.2: Effect of the forecast (pooled) and insurance on inputs

	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast	-1.99 (29.18)	-0.77 (1.54)	1.10 (3.95)	24.38 (50.18)	26.48 (95.40)	0.04 (0.05)
Insurance	97.60** (43.33)	-0.94 (1.34)	0.02 (5.70)	113.49* (64.13)	263.16** (130.22)	0.13** (0.06)
q-val Forecast	1.000	1.000	1.000	1.000	1.000	
q-val Insurance	0.279	0.565	0.996	0.279	0.279	
Control Mean	372.80	7.22	26.81	761.96	1443.49	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 (4). Fert is the amount spent 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. Sharpened q -values are adjusted across all outcomes in Tables G.1 and G.2 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.3: Effect of the forecast (pooled) and insurance on agricultural output

	(1) Prod (Kg)	(2) Value Prod (\$)	(3) Yield
Forecast	-8.55 (5.31)	-186.72 (193.77)	-2.92 (2.73)
Insurance	2.55 (6.79)	134.85 (224.71)	-1.59 (2.59)
q-val Forecast	0.479	0.479	0.479
q-val Insurance	1.000	1.000	1.000
Control Mean	66.91	2419.93	35.37
Observations	1201	1201	1170

Notes: This table presents estimates of the treatment effects of forecasts and insurance on agricultural output, estimated using Equation (4). Prod (Kg) is total agricultural production in kilograms. 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. 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. “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 coefficients. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.4: Effect of the forecast (pooled) and insurance on agricultural profits

	(1) Ag Profit (\$)	(2) Loss (\$)	(3) Profit w/ Loss (\$)	(4) Ag Profit Non-Flood (\$)
Forecast	-213.20 (160.57)	161.83* (90.05)	-31.73 (193.31)	-114.66 (237.23)
Insurance	-146.68 (183.40)	198.41** (91.37)	4.42 (209.56)	402.92 (297.66)
q-val Forecast	0.479	0.479	0.674	
q-val Insurance	1.000	0.221	1.000	
Control Mean	970.62	661.07	1654.24	970.62
Observations	1201	1201	1201	554

Notes: This table presents estimates of the treatment effects of forecasts and insurance on agricultural output, estimated using Equation (4). Ag Profit (\$) is the value of production (evaluated at district-median prices) less total expenditure in USD. Loss (\$) is the value of reported crop losses (evaluated at district-median prices) in USD. Profit w/ loss (\$) is the value of production plus the value of crop losses, less total expenditure in USD. Ag Profit Non-Flood (\$) is agricultural profits for the sample of households that did not report crop losses due to flooding or cyclones 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. “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 coefficients. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

G.2 Correlates of willingness-to-pay

Table G.5: Correlation between willingness-to-pay for the forecast and priors/risk aversion

	Willingness-to-pay for onset forecast					
	(1)	(2)	(3)	(4)	(5)	(6)
Std. Prior	12.172 (26.210)	-17.142 (115.197)				
Std. Prior2		13.634 (48.210)				
Share Before On Time Cutoff – 0.5			-92.939* (50.524)			
Share Before Early Cutoff – 0.5				-31.460 (63.363)		
Prior – Vg. Historical					17.368 (24.260)	
Risk Aversion						-2.722 (1.945)
Mean in Forecast Group	88.84	88.84	88.84	88.84	88.84	88.84
Observations	434	434	434	434	434	434

Notes: This table presents the correlation between forecast treatment group farmers' willingness to pay for the forecast and measures of prior strength and risk aversion. Std. Prior is the standard deviation of the farmer's prior as measured at baseline. Std. Prior2 is this SD squared. The absolute value of the share before on time (and early) cutoff minus 0.5, measures the distance between the likelihood a farmer thinks the monsoon is to arrive (at least) on time and 0.5 such that farmers that are more certain the monsoon either will or will not arrive on time will have higher values, while those who are more uncertain will have low values. The variables' range is between 0 and 0.5. The absolute value of the difference between the farmer's prior and the village's historical average measures the distance between the farmer's belief about this year and the average beliefs of past monsoon arrival within the village. Risk Aversion measures the farmer's choice in an incentivized risk game where higher values indicate that the farmer is more risk averse. The sample includes only farmers in the forecast group. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.6: Correlation between willingness-to-pay for the forecast and prior strength terciles

	Willingness-to-pay for onset forecast
	(1)
Std Prior 2nd Tercile	9.944 (16.442)
Std Prior 3rd Tercile	4.115 (20.843)
Mean in Forecast Group	88.84
Observations	434

Notes: This table presents WTP for the forecast by tercile of the standard deviation of farmers' priors. Std. Prior 2nd / 3rd Tercile is an indicator for the respondent's prior standard deviation being in the 2nd or 3rd tercile as measured at baseline. The omitted group is the 1st tercile. The sample includes only farmers in the forecast group. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.7: Correlation between willingness-to-pay for insurance and priors/risk aversion

	Willingness-to-pay for insurance			
	(1)	(2)	(3)	(4)
Stdv of Prior Distribution	76.008 (69.010)			
Prob mass of beans before individual ontime cutoff		1.055 (37.277)		
Prob mass of beans before individual early cutoff			-75.258 (74.662)	
Risk Preference - higher is more risk averse				-3.867 (4.613)
Mean in Insurance Group	106.02	106.02	106.02	106.02
Observations	221	221	221	221

Notes: This table presents the correlation between insurance treatment group farmers' willingness to pay for insurance and measures of prior strength and risk aversion. Std. Prior is the standard deviation of the farmer's prior as measured at baseline. Prob mass of beans before individual On Time/Early cutoff is the respondent's reported probability that the monsoon will arrive on time or early. Risk Aversion measures the farmer's choice in an incentivized risk game where higher values indicate that the farmer is more risk averse. The sample includes only farmers in the insurance group. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, * $p < 0.05$, ** $p < 0.10$.

G.3 Belief heterogeneity

Table G.8: Effect of the forecast on beliefs by prior strength

	(1) posterior – forecast	(2) posterior – prior	(3) K-S Stat
Forecast	-0.163** (0.083)	-0.215** (0.094)	-0.046* (0.027)
Stdv of Prior × Forecast	-0.211 (0.185)	-0.312 (0.213)	-0.056 (0.070)
Stdv of Prior Distribution	0.239** (0.107)	0.372** (0.155)	0.054 (0.049)
Control Mean	0.70	0.89	0.44
Observations	921	921	921

Notes: This table presents estimates of the treatment effect of forecasts on farmers' beliefs about the onset timing of the Indian Summer Monsoon, estimated using Equation (3). 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. Stdv of Prior is the standard deviation of the respondent's prior belief distribution, where higher values reflect more uncertainty. We exclude households where we were unable to speak to the same respondent when eliciting priors and posteriors. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

G.4 Additional farm input results

Table G.9: Effect of the forecast and insurance on additional inputs

	Panel A: Forecast vs. Insurance		
	(1) Changed plans	(2) Early labor	(3) Late labor
Forecast	-0.020 (0.038)	-40.199* (22.831)	59.135* (33.006)
Insurance	0.024 (0.046)	33.839 (28.866)	78.034* (41.622)
Panel B: Forecast Terciles			
Forecast × Ind Bin 1	-0.067 (0.056)	-72.159** (35.572)	7.779 (58.677)
Forecast × Ind Bin 2	0.001 (0.052)	-61.669** (30.554)	26.606 (42.875)
Forecast × Ind Bin 3	0.056 (0.073)	51.839 (44.387)	207.148*** (69.787)
Test Tercile 1=3	0.158	0.027	0.032
Test Insur. = Ter. 3	0.677	0.708	0.090
Control Mean	0.61	355.10	397.97
Observations	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts and insurance on inputs, estimated using Equations (3, panel A) and (4, panel B). Changed plans is an indicator for whether the farmer said they had changed their plans relative to what they said would do in an “on time” monsoon year. Early labor is total labor expenditure on pre-planting and planting activities in USD. Late labor is total labor expenditure between planting and harvest and during harvest 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, enumerator 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 coefficients. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.10: Effect of the forecast and insurance on inputs per acre

Panel A: Forecast vs. Insurance					
	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total
Forecast	16.26* (9.66)	-1.21 (1.24)	0.11 (2.52)	37.50* (21.34)	89.76*** (34.01)
Insurance	36.21** (14.32)	-1.74 (1.23)	-5.94** (2.54)	23.14 (23.12)	83.10** (39.87)
q-val Forecast	0.141	0.199	0.630	0.141	0.044
q-val Insurance	0.052	0.087	0.052	0.146	0.052

Panel B: Forecast Terciles					
Forecast × Ind Bin 1	19.24 (15.62)	-0.50 (1.63)	-1.70 (4.46)	64.12* (37.51)	120.87** (60.37)
Forecast × Ind Bin 2	9.24 (13.05)	-2.60* (1.38)	0.83 (3.51)	25.98 (27.07)	63.54 (42.84)
Forecast × Ind Bin 3	22.22 (18.10)	0.37 (3.23)	0.80 (4.17)	13.63 (33.55)	90.65 (58.77)
q-val Tercile 1	0.281	0.572	0.572	0.281	0.281
q-val Tercile 2	0.561	0.426	0.952	0.529	0.426
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000
Test Tercile 1=3	0.901	0.804	0.665	0.302	0.717
Test Insur. = Ter. 3	0.489	0.533	0.142	0.771	0.914
Control Mean	182.96	5.17	16.80	400.21	712.92
Observations	1170	1170	1170	1170	1170

Notes: This table presents estimates of the treatment effects of forecasts and insurance on inputs per acre, estimated using Equations (3, panel A) and (4, panel B). 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, all per acre. Total is the total amount spent on all inputs per acre, including all previous outcomes and any other costs reported by farmers. All outcomes are in USD per acre. 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, enumerator 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 coefficients. Sharpened q -values are adjusted across all outcomes in Tables 3 and 4, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

G.5 Heterogeneity

Table G.11: Effect of the forecast on land and crop choice by prior strength

	Panel A: Forecast $\times$ Prior Strength				
	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	-0.127 (0.110)	0.059* (0.032)	0.021 (0.037)	-0.011 (0.039)	0.008 (0.027)
Forecast $\times$ Prior Str.	0.076 (0.338)	0.013 (0.073)	0.035 (0.094)	-0.038 (0.098)	-0.109 (0.078)
	Panel B: Forecast Terciles $\times$ Prior Strength				
	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast $\times$ Ind Bin 1	-0.461*** (0.160)	0.012 (0.050)	-0.053 (0.053)	-0.115* (0.059)	0.011 (0.045)
Forecast $\times$ Ind Bin 2	-0.110 (0.142)	0.050 (0.038)	0.046 (0.051)	0.026 (0.048)	0.023 (0.038)
Forecast $\times$ Ind Bin 3	0.429* (0.241)	0.167*** (0.060)	0.125* (0.065)	0.140** (0.071)	0.013 (0.053)
Forecast $\times$ Bin 1 $\times$ Prior Str.	-0.007 (0.446)	-0.062 (0.107)	0.026 (0.140)	-0.125 (0.155)	-0.035 (0.114)
Forecast $\times$ Bin 2 $\times$ Prior Str.	0.516 (0.512)	-0.057 (0.126)	0.015 (0.162)	-0.200 (0.138)	-0.056 (0.137)
Forecast $\times$ Bin 3 $\times$ Prior Str.	0.369 (0.649)	-0.046 (0.120)	-0.097 (0.159)	0.161 (0.159)	-0.300** (0.130)
Control Mean	2.12	0.51	0.57	0.36	0.39
Observations	1200	1200	1200	1200	1200

Notes: This table presents estimates of the treatment effects of forecasts on farmers' land use and cropping decisions by prior strength. 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 an additional crop in 2022 as compared to 2021. Sub Crop is an indicator for removing a crop in 2022 as 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. Prior Str. is the difference of the respondent's on-time probability from 0.5. It has been de-meaned. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.12: Effect of the forecast on inputs by prior strength

Panel A: Forecast $\times$ Prior Strength						
	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast	-2.98 (29.04)	-0.68 (1.54)	1.12 (3.99)	23.00 (50.07)	23.53 (94.96)	0.04 (0.05)
Forecast $\times$ Prior Str.	15.06 (72.62)	-6.98 (4.45)	-8.89 (11.75)	-71.80 (148.38)	-89.53 (257.03)	-0.02 (0.12)
Panel B: Forecast Terciles $\times$ Prior Strength						
Forecast $\times$ Ind Bin 1	-36.87 (42.66)	-0.40 (2.65)	-0.21 (7.57)	-33.41 (86.83)	-109.11 (162.16)	-0.08 (0.07)
Forecast $\times$ Ind Bin 2	-33.04 (39.13)	-1.84 (1.55)	-0.69 (4.91)	-61.34 (65.74)	-87.62 (121.35)	0.03 (0.06)
Forecast $\times$ Ind Bin 3	88.54 (54.73)	1.92 (3.15)	8.95 (7.19)	264.26** (108.06)	440.63** (185.76)	0.30*** (0.09)
Forecast $\times$ Bin 1 $\times$ Prior Str.	39.93 (104.35)	-4.94 (7.96)	24.83 (19.05)	-229.70 (203.53)	-276.43 (374.79)	-0.16 (0.17)
Forecast $\times$ Bin 2 $\times$ Prior Str.	91.57 (128.82)	-4.68 (5.62)	-6.49 (17.67)	317.96 (235.87)	519.96 (428.40)	0.06 (0.20)
Forecast $\times$ Bin 3 $\times$ Prior Str.	11.86 (126.11)	-6.16 (7.07)	-27.21 (21.08)	122.04 (298.82)	161.91 (481.58)	0.02 (0.23)
Control Mean	372.80	7.22	26.81	761.96	1443.49	0.00
Observations	1200	1200	1200	1200	1200	1200

Notes: This table presents estimates of the treatment effects of forecasts on farmers' input use by prior strength. Fert is the amount spent on fertilizer, Seeds the amount spent on seeds, Irri the amount spent on irrigation, and Labor the amount spent on labor. 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. 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. Prior Str. is the difference of the respondent's on-time probability from 0.5. It has been de-meaned. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.13: Effect of the forecast on land and crop choice by gap between forecast and prior

Panel A: Forecast $\times$ Prior. - Fore.					
	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	-0.137 (0.117)	0.059* (0.032)	0.015 (0.037)	-0.010 (0.039)	0.002 (0.028)
Forecast $\times$ Diff. Prior and Forecast.	0.081 (0.125)	0.010 (0.029)	-0.000 (0.038)	-0.012 (0.039)	0.003 (0.027)
Panel B: Forecast Terciles $\times$ Prior. - Fore.					
Forecast $\times$ Ind Bin 1	-0.341 (0.215)	0.074 (0.061)	-0.019 (0.077)	-0.024 (0.083)	-0.040 (0.056)
Forecast $\times$ Ind Bin 2	0.208 (0.354)	0.019 (0.093)	0.187* (0.112)	0.196* (0.114)	-0.025 (0.085)
Forecast $\times$ Ind Bin 3	0.317 (0.227)	0.152** (0.065)	0.110* (0.064)	0.116 (0.075)	0.036 (0.056)
Forecast $\times$ Bin 1 $\times$ Prior. - Fore.	-0.321 (0.225)	-0.101* (0.061)	-0.061 (0.088)	-0.161* (0.089)	0.088 (0.064)
Forecast $\times$ Bin 2 $\times$ Prior. - Fore.	0.371 (0.371)	-0.031 (0.108)	0.183 (0.132)	0.235* (0.133)	-0.051 (0.095)
Forecast $\times$ Bin 3 $\times$ Prior. - Fore.	0.301 (0.292)	-0.009 (0.043)	0.029 (0.074)	0.035 (0.069)	-0.022 (0.053)
Control Mean	2.12	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 on farmers' land use and cropping decisions by the gap between the forecast and the prior. 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 an additional crop in 2022 as compared to 2021. Sub Crop is an indicator for removing a crop in 2022 as 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. Prior. - Fore. is the standardized absolute difference between the prior and the forecast. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.14: Effect of the forecast on inputs by gap between forecast and prior

Panel A: Forecast $\times$ Prior. - Fore.						
	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast	-4.81 (30.29)	-0.74 (1.60)	0.55 (4.05)	26.01 (54.15)	20.95 (101.90)	0.04 (0.05)
Forecast $\times$ Diff. Prior and Forecast.	17.37 (27.33)	0.13 (1.08)	4.51 (3.74)	41.53 (60.93)	63.63 (99.75)	0.04 (0.05)
Panel B: Forecast Terciles $\times$ Prior. - Fore.						
Forecast $\times$ Ind Bin 1	29.77 (52.34)	3.98 (4.13)	-11.07 (9.76)	112.44 (110.18)	135.61 (208.39)	0.02 (0.09)
Forecast $\times$ Ind Bin 2	30.34 (89.16)	2.62 (2.98)	3.62 (10.03)	10.51 (163.33)	63.37 (309.12)	0.10 (0.15)
Forecast $\times$ Ind Bin 3	79.43 (59.33)	1.94 (3.42)	7.74 (8.64)	222.46** (104.88)	375.66* (192.17)	0.26*** (0.10)
Forecast $\times$ Bin 1 $\times$ Prior. - Fore.	-119.13 (72.49)	-7.63** (3.55)	13.82 (10.72)	-288.84** (123.45)	-526.83** (222.91)	-0.20** (0.08)
Forecast $\times$ Bin 2 $\times$ Prior. - Fore.	77.93 (101.48)	5.77* (3.45)	6.50 (11.68)	74.61 (185.36)	164.96 (350.68)	0.10 (0.16)
Forecast $\times$ Bin 3 $\times$ Prior. - Fore.	19.64 (35.55)	-1.19 (2.27)	3.50 (4.24)	123.51 (125.19)	155.68 (162.47)	0.05 (0.09)
Control Mean	372.80	7.22	26.81	761.96	1443.49	0.00
Observations	1201	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts on farmers' inputs by the gap between the forecast and the prior. 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. 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. 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. Prior. - Fore. is the standardized absolute difference between the prior and the forecast. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.15: Effect of the forecast on land and crop choice by WTP

Panel A: Forecast $\times$ WTP					
	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	-0.262* (0.141)	-0.011 (0.040)	-0.041 (0.044)	-0.050 (0.040)	0.007 (0.034)
Forecast $\times$ WTP	-0.099 (0.122)	-0.019 (0.039)	-0.030 (0.050)	-0.071 (0.045)	0.004 (0.036)
WTP	0.059 (0.102)	0.012 (0.035)	0.027 (0.045)	0.041 (0.038)	0.032 (0.034)

Panel B: Forecast Terciles $\times$ WTP					
Forecast $\times$ Ind Bin 1	-0.776*** (0.253)	-0.005 (0.066)	-0.120 (0.075)	-0.128* (0.072)	0.060 (0.053)
Forecast $\times$ Ind Bin 2	-0.095 (0.194)	-0.075 (0.054)	-0.051 (0.065)	-0.061 (0.057)	0.020 (0.052)
Forecast $\times$ Ind Bin 3	0.371 (0.304)	0.119 (0.084)	0.027 (0.082)	0.071 (0.077)	-0.042 (0.074)
Forecast $\times$ Bin 1 $\times$ WTP	0.019 (0.212)	-0.002 (0.068)	-0.052 (0.089)	-0.159* (0.090)	0.032 (0.070)
Forecast $\times$ Bin 2 $\times$ WTP	-0.202 (0.150)	-0.035 (0.049)	-0.013 (0.052)	-0.045 (0.054)	-0.013 (0.042)
Forecast $\times$ Bin 3 $\times$ WTP	0.105 (0.253)	0.048 (0.089)	0.098 (0.077)	-0.001 (0.074)	0.128* (0.070)
Control Mean	2.12	0.51	0.57	0.36	0.39
Observations	655	655	655	655	655

Notes: This table presents estimates of the treatment effects of forecasts on farmers' land use and cropping decisions by WTP. 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 an additional crop in 2022 as compared to 2021. Sub Crop is an indicator for removing a crop in 2022 as 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. WTP is the willingness-to-pay for the forecast and insurance product. The sample excludes the control group because WTP is undefined for them. The omitted category is insurance. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.16: Effect of the forecast on inputs by WTP

Panel A: Forecast $\times$ WTP						
	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast	-83.97** (40.37)	0.54 (1.48)	1.69 (6.15)	-41.81 (67.91)	-173.16 (125.31)	-0.06 (0.06)
Forecast $\times$ WTP	-37.17 (44.85)	0.57 (1.13)	1.33 (4.01)	21.13 (54.52)	20.26 (106.94)	-0.04 (0.05)
WTP	71.06 (50.90)	-0.21 (0.88)	-1.58 (3.15)	26.56 (54.73)	84.26 (111.88)	0.02 (0.05)

Panel B: Forecast Terciles $\times$ WTP						
Forecast $\times$ Ind Bin 1	-192.75*** (63.92)	-1.04 (2.26)	-5.79 (14.05)	-137.33 (110.53)	-512.87** (227.87)	-0.20** (0.10)
Forecast $\times$ Ind Bin 2	-104.08** (52.38)	-0.55 (1.26)	0.53 (7.63)	-151.49 (100.33)	-236.75 (177.16)	-0.11 (0.09)
Forecast $\times$ Ind Bin 3	69.64 (73.14)	3.65 (3.19)	14.46 (12.08)	233.37 (166.24)	312.29 (275.14)	0.25* (0.14)
Forecast $\times$ Bin 1 $\times$ WTP	-192.03** (80.10)	-0.55 (1.69)	13.89 (8.78)	-72.35 (108.93)	-176.68 (222.45)	0.02 (0.10)
Forecast $\times$ Bin 2 $\times$ WTP	23.05 (42.54)	0.73 (1.06)	0.37 (5.42)	60.87 (69.13)	97.24 (119.79)	-0.07 (0.07)
Forecast $\times$ Bin 3 $\times$ WTP	48.91 (67.11)	2.87 (2.56)	-4.33 (9.37)	143.45 (168.71)	247.85 (259.90)	0.08 (0.12)
Control Mean	372.80	7.22	26.81	761.96	1443.49	0.00
Observations	655	655	655	655	655	655

Notes: This table presents estimates of the treatment effects of forecasts on farmers' land use and cropping decisions by WTP. 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. 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. 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. WTP is the willingness-to-pay for the forecast and insurance product. The sample excludes the control group because WTP is undefined for them. The omitted category is insurance. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.17: Effect of the forecast on land use and crop choice by risk aversion

Panel A: Forecast $\times$ Risk Aversion					
	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast	-0.298** (0.138)	0.037 (0.037)	0.032 (0.045)	-0.030 (0.044)	0.028 (0.034)
Forecast $\times$ Risk Av.	0.456** (0.195)	0.049 (0.059)	-0.040 (0.062)	0.042 (0.068)	-0.065 (0.052)
Panel B: Forecast Terciles $\times$ Risk Aversion					
Forecast $\times$ Ind Bin 1	-0.550*** (0.196)	-0.019 (0.055)	-0.054 (0.067)	-0.122* (0.069)	0.039 (0.052)
Forecast $\times$ Ind Bin 2	-0.112 (0.181)	0.037 (0.048)	0.030 (0.064)	-0.002 (0.058)	0.015 (0.048)
Forecast $\times$ Ind Bin 3	0.079 (0.275)	0.128* (0.077)	0.133* (0.074)	0.064 (0.076)	0.019 (0.068)
Forecast $\times$ Bin 1 $\times$ Risk Av.	0.225 (0.253)	0.084 (0.084)	0.012 (0.094)	0.041 (0.096)	-0.063 (0.071)
Forecast $\times$ Bin 2 $\times$ Risk Av.	0.139 (0.256)	0.019 (0.070)	0.032 (0.086)	0.040 (0.082)	0.000 (0.073)
Forecast $\times$ Bin 3 $\times$ Risk Av.	0.848** (0.422)	0.105 (0.114)	-0.018 (0.109)	0.158 (0.127)	0.013 (0.091)
Control Mean	2.12	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 on farmers' inputs by risk aversion. 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 an additional crop in 2022 as compared to 2021. Sub Crop is an indicator for removing a crop in 2022 as 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. Risk. Av. is the result of an incentivized risk game where higher values indicate the farmer is more risk averse. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.18: Effect of the forecast on inputs by risk aversion

Panel A: Forecast $\times$ Risk Aversion						
	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast	-40.04 (36.36)	0.04 (1.77)	-0.18 (4.35)	-19.22 (61.59)	-77.19 (119.12)	-0.02 (0.06)
Forecast $\times$ Risk Av.	96.15* (52.33)	-2.03 (2.16)	3.08 (9.04)	107.77 (88.57)	258.64 (172.56)	0.16* (0.08)
Panel B: Forecast Terciles $\times$ Risk Aversion						
Forecast $\times$ Ind Bin 1	-9.19 (49.82)	-0.59 (2.90)	8.10 (9.38)	-28.95 (114.18)	-50.22 (217.35)	-0.13 (0.08)
Forecast $\times$ Ind Bin 2	-74.39 (45.97)	-1.99 (1.77)	-6.83 (5.40)	-133.95* (81.07)	-265.44* (145.41)	0.00 (0.07)
Forecast $\times$ Ind Bin 3	-4.80 (61.58)	4.93 (4.85)	-3.84 (4.93)	136.45 (118.52)	141.04 (200.18)	0.18 (0.12)
Forecast $\times$ Bin 1 $\times$ Risk Av.	-46.18 (67.49)	-0.25 (3.09)	-24.57 (17.62)	-30.52 (129.97)	-189.94 (256.07)	0.13 (0.11)
Forecast $\times$ Bin 2 $\times$ Risk Av.	129.31* (67.24)	0.14 (2.58)	16.17 (10.01)	247.32** (114.44)	577.31*** (214.05)	0.08 (0.10)
Forecast $\times$ Bin 3 $\times$ Risk Av.	242.16** (97.67)	-6.59 (5.19)	33.02** (16.38)	306.70 (202.58)	725.92** (340.70)	0.31* (0.17)
Control Mean	372.80	7.22	26.81	761.96	1443.49	0.00
Observations	1201	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecasts on farmers' inputs by the risk aversion. 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. 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. 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. Risk. Av. is the result of an incentivized risk game. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

G.6 Local average treatment effects

Table G.19: Effect of forecast and insurance takeup on beliefs

	(1) posterior – forecast	(2) posterior – prior	(3) K-S Stat
Forecast takeup	-0.204** (0.095)	-0.272** (0.108)	-0.057* (0.030)
Insurance takeup	-0.023 (0.108)	-0.101 (0.125)	-0.021 (0.036)
Control Mean	0.70	0.89	0.44
Observations	921	921	921

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on farmers' beliefs about the onset timing of the Indian Summer Monsoon, estimated using an IV version of Equation (3) where we instrument for forecast and insurance takeup with an indicator for being in a forecast or insurance village. 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. Standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.20: Effect of forecast and insurance takeup on land use and cropping

	(1) Land Ha.	(2) Cash Crop	(3) Changed Crop	(4) Added Crop	(5) Sub Crop
Forecast takeup × Ind Bin 1	-0.588*** (0.205)	0.016 (0.059)	-0.066 (0.064)	-0.140** (0.071)	0.015 (0.054)
Forecast takeup × Ind Bin 2	-0.089 (0.164)	0.048 (0.042)	0.044 (0.057)	0.014 (0.053)	0.012 (0.042)
Forecast takeup × Ind Bin 3	0.454* (0.265)	0.176*** (0.067)	0.130* (0.071)	0.160** (0.078)	0.010 (0.060)
Insurance takeup	0.206 (0.155)	0.071* (0.043)	0.051 (0.053)	0.050 (0.055)	-0.005 (0.042)
q-val Tercile 1	0.043	1.000	1.000	0.286	1.000
q-val Tercile 2	1.000	1.000	1.000	1.000	1.000
q-val Tercile 3	0.095	0.058	0.091	0.075	0.354
Test Tercile 1=3	0.002	0.055	0.037	0.003	0.949
Test Insur. = Ter. 3	0.380	0.144	0.311	0.177	0.823
Control Mean	2.12	0.51	0.57	0.36	0.39
Observations	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on farmers' land use and cropping decisions, estimated using an IV version of Equation (4) where we instrument for forecast and insurance takeup with indicators for being in a forecast or insurance village. 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. 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. “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 coefficients. Sharpened q -values are adjusted across all outcomes in Tables 3 and 4 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.21: Effect of forecast and insurance takeup on inputs

	(1) Fert	(2) Seed	(3) Irri	(4) Labor	(5) Total	(6) Invest Index
Forecast takeup × Ind Bin 1	-40.13 (51.97)	-0.82 (3.18)	-2.46 (9.50)	-55.57 (104.29)	-166.27 (197.03)	-0.10 (0.09)
Forecast takeup × Ind Bin 2	-32.44 (44.04)	-2.28 (1.80)	-1.51 (5.54)	-51.50 (75.96)	-66.09 (139.22)	0.03 (0.06)
Forecast takeup × Ind Bin 3	112.07* (62.04)	2.29 (3.92)	10.26 (8.96)	289.85** (116.78)	495.81** (206.44)	0.32*** (0.11)
Insurance takeup	112.03** (49.71)	-1.06 (1.52)	-0.12 (6.55)	128.06* (73.96)	299.90** (150.00)	0.15** (0.06)
q-val Tercile 1	1.000	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	0.091	0.230	0.141	0.058	0.058	
Test Tercile 1=3	0.053	0.535	0.351	0.031	0.020	0.001
Test Insur. = Ter. 3	1.000	0.410	0.315	0.231	0.403	0.138
Control Mean	372.80	7.22	26.81	761.96	1443.49	0.00
Observations	1201	1201	1201	1201	1201	1201

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on inputs, estimated using an IV version of Equation (4) where we instrument for forecast and insurance takeup with indicators for being in a forecast or insurance village. 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 received good news. “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 coefficients. Sharpened q -values are adjusted across all outcomes in Tables 3 and 4 (except the index), and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.22: Effect of forecast and insurance takeup on agricultural output

	(1) Prod (Kg)	(2) Value Prod (\$)	(3) Yield
Forecast takeup × Ind Bin 1	-20.83** (9.98)	-657.27* (346.04)	-7.94 (5.28)
Forecast takeup × Ind Bin 2	-11.94 (8.53)	-199.74 (264.44)	-0.55 (4.05)
Forecast takeup × Ind Bin 3	16.96 (12.08)	493.54 (450.55)	0.79 (4.48)
Insurance takeup	3.03 (7.79)	156.65 (255.27)	-1.75 (2.88)
q-val Tercile 1	0.209	0.209	0.209
q-val Tercile 2	0.942	1.000	1.000
q-val Tercile 3	1.000	1.000	1.000
Test Tercile 1=3	0.012	0.037	0.163
Test Insur. = Ter. 3	0.262	0.449	0.567
Control Mean	66.91	2419.93	35.37
Observations	1201	1201	1170

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on agricultural output, estimated using an instrumental variables version of Equation (4) where we use an indicator for being in a forecast or insurance offer village as an instrument for forecast or insurance takeup. 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 district median prices for each crop. Yield is kilograms of production per hectare. 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. “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 coefficients. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.23: Effect of forecast and insurance takeup on agricultural profits

	(1) Ag Profit (\$)	(2) Loss (\$)	(3) Profit w/ Loss (\$)	(4) Ag Profit Non-Flood (\$)
Forecast takeup × Ind Bin 1	-484.34* (283.73)	59.45 (164.47)	-360.30 (390.26)	-447.17 (541.64)
Forecast takeup × Ind Bin 2	-124.07 (218.77)	236.75* (137.60)	112.99 (249.83)	-117.11 (277.83)
Forecast takeup × Ind Bin 3	-33.63 (377.11)	208.87 (166.68)	180.70 (411.86)	593.43 (590.37)
Insurance takeup	-164.23 (207.59)	223.52** (104.54)	6.07 (238.66)	561.19 (419.67)
q-val Tercile 1	0.209	0.315	0.249	
q-val Tercile 2	1.000	0.942	1.000	
q-val Tercile 3	1.000	1.000	1.000	
q-val Insurance				
Test Tercile 1=3	0.322	0.496	0.322	0.208
Test Insur. = Ter. 3	0.731	0.933	0.666	0.962
Control Mean	970.62	661.07	1654.24	970.62
Observations	1201	1201	1201	554

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on agricultural profits, estimated using an instrumental variables version of Equation (4) where we use an indicator for being in a forecast or insurance offer village as an instrument for forecast or insurance takeup. Ag Profit (\$) is the value of production (evaluated at district-median prices) less total expenditure in USD. Loss (\$) is the value of reported crop losses (evaluated at district-median prices) in USD. Profit w/ loss (\$) is the value of production plus the value of crop losses, less total expenditure in USD. Ag Profit Non-Flood (\$) is agricultural profits for the sample of households that did not report crop losses due to flooding or cyclones. 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. “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 coefficients. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.24: effect of forecast and insurance takeup on business activity

	(1) Non-Ag Bus.	(2) Non-Ag Invest	(3) Bus Profit
Forecast takeup × Ind Bin 1	0.07 (0.05)	29.96 (87.54)	94.75 (92.38)
Forecast takeup × Ind Bin 2	0.01 (0.04)	9.04 (65.50)	21.71 (53.91)
Forecast takeup × Ind Bin 3	-0.05 (0.05)	-124.09* (72.10)	-33.86 (98.27)
Insurance takeup	0.10*** (0.04)	116.49 (72.93)	119.14* (63.55)
q-val Tercile 1	0.844	0.844	0.844
q-val Tercile 2	1.000	1.000	1.000
q-val Tercile 3	0.395	0.344	0.737
Test Tercile 1=3	0.082	0.180	0.316
Test Insur. = Ter. 3	0.009	0.003	0.181
Control Mean	0.14	157.98	165.51
Observations	1197	1199	1197

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on business activity, estimated using an instrumental variables version of Equation (4) where we use an indicator for being in a forecast or insurance offer village as an instrument for forecast or insurance takeup. 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. 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. “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 coefficients. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table G.25: effect of forecast and insurance takeup on economic wellbeing

	(1) Food cons	(2) Other cons	(3) Asset value	(4) Livestock	(5) Net savings	(6) Welfare index
Forecast takeup × Ind Bin 1	0.38 (0.78)	-0.64 (1.76)	430.70 (322.26)	0.08* (0.05)	421.96 (285.42)	0.17** (0.07)
Forecast takeup × Ind Bin 2	1.34** (0.59)	-0.90 (1.09)	37.70 (172.47)	-0.04 (0.04)	78.00 (247.94)	-0.00 (0.05)
Forecast takeup × Ind Bin 3	1.27 (0.93)	-0.23 (1.57)	-140.86 (160.93)	0.01 (0.06)	159.19 (285.52)	0.05 (0.06)
Insurance takeup	0.51 (0.53)	2.23* (1.14)	-142.81 (180.83)	0.01 (0.03)	-465.43* (245.98)	-0.03 (0.05)
q-val Tercile 1	0.434	0.434	0.434	0.434	0.434	
q-val Tercile 2	0.137	1.000	1.000	0.757	1.000	
q-val Tercile 3	1.000	1.000	1.000	1.000	1.000	
q-val Insurance						
Test Tercile 1=3	0.438	0.853	0.102	0.350	0.545	0.236
Test Insur. = Ter. 3	0.450	0.135	0.992	0.973	0.053	0.282
Control Mean	13.22	9.93	1503.10	0.22	-1031.41	0.00
Observations	1201	1201	1201	1201	1129	1201

Notes: This table presents estimates of the treatment effects of forecast and insurance takeup on economic wellbeing, estimated using an instrumental variables version of Equation (4) where we use an indicator for being in a forecast or insurance offer village as an instrument for forecast or insurance takeup. Food cons is food consumption per household member in USD over the past 30 days. Other cons is non-food consumption (other than medical expenses). Asset value is the value of assets in USD. Livestock is the count of livestock in Tropical Livestock Units. Net savings is savings less outstanding debt in USD. Welfare index is an inverse covariance weighted index of the other five outcomes in the table. It has been excluded from the MHT correction as it is a composite of outcomes already included. 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. “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 coefficients. All regressions include strata fixed effects, enumerator fixed effects, and baseline controls chosen by double-selection LASSO. Sharpened q -values are adjusted for all outcomes in the table, and standard errors are clustered at the village level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

H Seasonal climate forecasts

Forecast time scales Forecasts can be made over a range of time-scales, including short-, medium-, and long-term forecasts. Short-term, or *weather*, forecasts, provide a prediction about precise weather conditions on a particular day, and are issued between one and fourteen days in advance.⁵² Forecasts that provide information beyond this time horizon present information only about average conditions over a longer time period, rather than about an individual day. Medium-range forecasts are issued between 15 and 30 days in advance. Long-range, or *seasonal*, forecasts, which we study in this paper, provide information four or more weeks ahead. These forecasts also tend to provide information on longer time windows, with typical forecasts projecting climatic conditions over a month or entire season. Seasonal forecasts are particularly relevant for agriculture for two reasons. First, with long lead times, farmers can use these forecasts to make meaningful adjustments to key planting decisions, such as amount of land to cultivate and crop choice (Gine et al. (2015)). Second, seasonal forecasts provide information that is highly relevant to agricultural outcomes: climate over the full growing season.


Existing monsoon forecasts Researchers have attempted to produce long-range forecasts of two key features of the Indian Summer Monsoon: rainfall quantity and monsoon onset timing. The Indian Meteorological Department (IMD) produces a statistical forecast of the expected seasonal total rainfall quantity at the beginning of the monsoon each year. These forecasts have traditionally focused on the All-India Rainfall Index (AIRI) (Rajeevan et al. (2007)). One of the most persistent criticisms of the AIRI forecasts is that the AIRI is itself a meaningless spatial average, describing a phenomenon that has both little spatial coherence (Moron et al. (2017)) and little relevance to district- or state-level rainfall amounts. Put differently, an IMD forecast of “normal” monsoon rainfall amounts indicates nothing about rainfall amounts for a specific farmer in a specific location, rendering it useful for climate science but less useful for agriculture. IMD and other agencies have also begun some experiments with dynamical (i.e., physics-based) models of the monsoon, but such forecasts similarly aim to forecast AIRI, rendering them uninformative for local decisions, though they do show some skill nationally (Das et al. (2015)). More recently, the IMD has begun providing region-specific quantity forecasts. However, the accuracy of these forecasts is poor, rendering them of limited usefulness for farmer decision-making: Rosenzweig and Udry (2019) notes that these forecasts have a low (~ 0.2) or even negative correlation with realized rainfall over most of India.

In contrast, seasonal timing forecasts typically deal with the onset of the monsoon. While the monsoon arrives in early–mid June on average, variability in onset timing is high. Appendix

⁵²The 14-day barrier is a physical limit, owing to the variability of the physical weather system.

Figure H.1 plots information about the monsoon onset over India, with Telangana outlined in black. IMD forecasts onset only over the south-western tip of the country — “monsoon onset over Kerala” (MOK)—which is not relevant for most of the country. Though MOK has been the subject of much of the research on onset timing and forecasting (e.g., Preenu et al., 2017), the monsoon does not progress smoothly northwards. Instead, monsoon rainfall frequently halts, and local false starts are common, such that MOK carries no more than a very limited signal for a farmer in parts of India outside of a narrow strip of coastal Kerala. Moron and Robertson (2014) define local agronomic onset and demonstrate the correlation between MOK and local onset over India. In Appendix Figure H.2, they show that there is virtually no signal value of MOK⁵³ in any region in India other than Kerala. Moreover, this forecast typically arrives with only two weeks of advance notice. There is no local IMD monsoon onset forecast, and MOK has been the subject of much of the research on onset timing and forecasting (e.g., Preenu et al., 2017).


Figure H.1: Monsoon onset over India



Notes: The left panel shows the average monsoon onset day (in day-of-year) for the period 1979-2019 across India. The right panel shows the standard deviation of onset for the period 1979-2019. Local onset timing is derived following Moron and Robertson (2014), and captures the timing of the first wet spell of the season that is sufficient to wet the topsoil enough to plant crops and is not immediately followed by a dry spell (in which case it is known as a “false start”). In both panels, grid cells are 0.25 degrees. Telangana, the location of our experiment, is highlighted with a thick black border.

⁵³In the paper, the authors define regional-scale monsoon onset (RSONS) as a summary measure of a number of onset indices over Kerala, which has a correlation of 0.92 with MOK (Moron and Robertson, 2014).

Figure H.2: Monsoon onset over Kerala has limited predictive power elsewhere in India




Notes: (a) Correlations between local-scale onset and the index of regional-scale onset (RSONS) defined in the text. (b) Correlations between local-scale onset and the Niño 3.4 SST index (N34) in May. Crosses indicate statistically significant correlations at the two-sided 95% level (see text). The value in parenthesis gives the fraction of significant grid boxes at the two-sided 95% level of significance according to a random-phase test. *Reproduced from Moron and Robertson (2014).*

PIK forecast We focus on onset forecasts for two main reasons. First, and most importantly, high-quality quantity forecasts are simply not available in our setting. In contrast, there exists such an accurate onset timing forecast. A new model developed by the Potsdam Institute for Climate Impact Research (PIK) (Stolbova et al., 2016) uses observations of climate variables in the months leading up to the beginning of the monsoon to predict the timing of the onset of the monsoon up to one month in advance for a specific region of India and identifies a method for expanding this to other local regions. This forecast substantially outperforms the IMD forecasts that were analysed in Mobarak and Rosenzweig (2014), but is not yet widely available to farmers who might benefit from the information. The output from the PIK model is a probability distribution of potential onset dates of the monsoon for a range of states over the Eastern Ghats with particular accuracy over Telangana. When evaluated for onset dates from 1965-2015, this new scheme was “correct,” defined as local onset falling within ± 7 days of the predicted date, 73% of years in the sample.⁵⁴ Moreover, while MOK date is forecast only two weeks in advance of the average MOK date, the

⁵⁴Stolbova et al. (2016) also predicts withdrawal dates with 8 weeks lead-time and shows 84% of years falling within ± 10 days of the actual withdrawal date.

Figure H.3: The PIK forecast is accurate



Notes: Monsoon OD and prediction based on temperature (green) and relative humidity (orange) and measured (dark blue) (a) Onset date (OD) validated against NCEP/NCAR data. Red and light blue shading indicates positive ENSO (El Niño) and negative ENSO (La Niña) years. (b) Also shown is the difference between the real onset and predicted dates in days. Grey shading indicates range of 7 days, within the prediction is considered accurate (absolute value of the difference between the real onset date in a given year and the predicted onset date). Reproduced from panels A and B of Stolbova et al. (2016).

PIK forecast is issued at least 35 days in advance of the average onset date in Telangana.

Second, farmers demand information on onset timing. Mobarak and Rosenzweig (2014) demonstrate that 40% of farmers purchase insurance against the risk of a *delayed monsoon onset* when randomly offered such a product. In our pilot, more than 60% of farmers stated that they would be willing to pay for a monsoon onset timing forecast. Finally, in a phone survey conducted by the Ministry of Agriculture and Farmers' Welfare in 2024, 88 percent of farmers reported that they would find a monsoon onset forecast, and 88 percent reported that they could use such a forecast for planting decisions.