Memory, Expectations, and Climate Adaptation

Marius Hees* and Frank van Oort

Erasmus School of Economics, Erasmus University Rotterdam

Corresponding author: Marius Hees, hees@ese.eur.nl

September 12, 2026

Abstract. Smallholder farmers face substantial and increasingly variable weather risk, yet take up of insurance and investment in other forms of protection remain low. We show that part of this demand problem begins with how farmers form beliefs about future weather. In a survey experiment with 686 smallholder farmers in eastern Uganda, we elicit memories of past weather and expectations for the coming season, benchmark them against remotely sensed weather records, and experimentally provide accessible information about local weather history. We find that season by season memory is strongly compressed, with large differences in the historical record reduced to only small differences in recall. Remembered weather nevertheless strongly predicts expectations, carrying this compression into beliefs about future risk and limiting farmers' ability to distinguish when protection is most valuable. A complementary measure of salient extreme events reveals a separate, less compressed channel from memory into expectations. These beliefs are economically meaningful, as farmers who expect adverse weather plan substantially more adaptation and protection. We then show that beliefs can be corrected. A simple visual summary of local weather history moves expectations toward the historical record, with the largest corrections among farmers whose initial beliefs are furthest away. Literacy predicts both memory and initial belief accuracy, but not responsiveness to the information. Our results identify belief formation as an important demand side margin in climate adaptation and point to accessible information as a low cost, short run complement to protection policies, with human capital representing a potentially important longer run margin in belief formation and adaptation.

Keywords: climate change; drought; climate adaptation; weather risk; agricultural risk; expectations; memory; smallholder agriculture.

JEL codes: D81, D83, D91, O13, Q12, Q54.

*We are especially grateful to The Ugandan Field Lab, in particular Joshua Balungira and Dr. Arjan Verschoor, and to all enumerators for their exceptional support throughout the project. They provided invaluable advice in designing the survey and experiment, helping tailor the study to local stakeholders, realities, and culture, and played a central role in planning and implementing the fieldwork.

1 Introduction

Smallholder farms produce about a third of the worlds food supply and roughly seventy percent of the calories consumed in the smallholder dominated regions of Africa, Asia, and Latin America (Ricciardi et al., 2018; Samberg et al., 2016; Lowder et al., 2016). For the households that run these farms, weather is a central source of income risk, yet the literature consistently finds low take up of insurance and limited investment in other forms of protection and adaptation (Cole et al., 2013; Carter et al., 2017). As dry spells, floods, and storms become more frequent and weather more erratic, this lack of protection becomes an increasingly important economic concern (IPCC, 2021). How farmers judge the weather they will face, and how those beliefs shape how they prepare for the coming season, is therefore central to understanding climate adaptation and economic resilience.

Much of the literature has approached this problem from the supply side. Over the past two decades, research has designed and tested rainfall insurance, stress tolerant seeds, drought resistant varieties, and other protective technologies, while varying contract design, prices, payment timing, marketing, and trust (Giné et al., 2008; Cole et al., 2013; Karlan et al., 2014; Emerick et al., 2016; Casaburi & Willis, 2018; Carter et al., 2017). Yet demand for protection remains persistently low, even where these products appear economically attractive. Much of this work makes one important simplification: it treats farmers beliefs about future weather as given and broadly correct. But the value of protection depends on farmers beliefs about the weather they expect to face. Farmers will only pay for insurance or other forms of adaptation if they believe adverse weather is sufficiently likely. If those beliefs are distorted or uninformative, better products alone may not raise demand. Understanding how farmers form expectations about future weather is therefore essential to understanding why demand for protection remains low, and what may be needed to change it.

This paper studies how these expectations are formed and quantifies the distortions that shape them. We elicit farmers memories of past seasons, compare them with remotely sensed records of rainfall and vegetation, and trace them into beliefs about the coming season. We then test experimentally whether simple information about local weather history leads farmers to update their beliefs. The central idea is simple. If memory compresses meaningful differences between past seasons, expectations formed from those memories will inherit that compression, potentially weakening the perceived value of protection against unusual weather.

We study this mechanism in the lowland plains of eastern Uganda, west of Mount Elgon, where smallholder farmers grow maize, beans, and cassava on plots averaging under three acres, largely for their own consumption. Production is rainfed, formal weather insurance is largely absent, and access to local historical weather information is limited. Farmers therefore rely heavily on their own experience when forming beliefs about future weather. We surveyed 686 farmers before the 2026 growing season and randomized villages into a control group and two information treatment groups.

This reliance on personal experience creates a measurement challenge. Studying how memory shapes expectations requires eliciting probabilistic beliefs from farmers with limited literacy and formal schooling. We therefore build on the visual probability elicitation approach of Manski (2004) and Delavande et al. (2011). Farmers allocate ten physical tokens across nine images of the same maize field, ranging from extreme drought to severe flooding, thereby expressing a subjective probability distribution without requiring literacy or familiarity with formal probabilities. We repeat this exercise for each of the past five seasons and for the coming season. In a second exercise, farmers allocate the same ten tokens across a visual calendar with ten intervals to elicit beliefs about the onset of the planting rains. We then benchmark recalled seasons against remotely sensed rainfall and vegetation data, allowing us to compare each farmers memory with what actually occurred.

The same visual framework also gives us a natural way to test whether these beliefs can be updated. Farmers in treatment villages receive a simple summary of local weather over the previous five seasons, constructed from the same rainfall and vegetation records and presented using the same images as the elicitation. They then repeat the expectation exercise, while farmers in the control group repeat it without receiving any information. Comparing the two allows us to isolate updating induced by the information from movement caused simply by asking the same question twice. A smaller third arm receives the information before its first expectation elicitation and provides an additional check on the experimental design.

The first result is that farmers remember past weather in a highly compressed way. They broadly remember which seasons were wetter and which were drier, but most of the differences between seasons disappear in memory. A clearly dry and a clearly wet season that are five bins apart in the historical record are remembered as only about half a bin apart. Memory also leans slightly toward drier conditions and is more accurate among literate and more educated farmers.

The second result is that this compression carries directly into expectations. Farmers who remember wetter past seasons also expect wetter conditions in the coming season, showing that past experience is an important input into their forecasts. But because very different seasons are remembered as similar, the expectations built from them vary little as well. Applying the estimated memory distortion to the historical weather record produces implied expectations that contain almost none of the actual variation across places and years. These beliefs also matter economically. Farmers who expect adverse weather report substantially more adaptation spending and are more likely to plan crop insurance, although these relationships are correlational.

The third result is that these beliefs can be updated. Providing farmers with a simple visual summary of local weather history moves expectations toward the historical record and shifts beliefs about the onset of the planting rains toward observed conditions. Farmers whose initial beliefs are furthest from the record update the most, while literacy, education, and other observable characteristics do not predict the response itself. Literacy instead predicts who needs the correction, since farmers who cannot read begin with less accurate memories and beliefs further from the historical record.

Together, these findings make three contributions.

First, we directly measure weather memory as an input into expectation formation. Earlier work infers forgetting from behavior (Gallagher, 2014) or aggregate responses (Moore et al., 2019). We instead observe recalled seasons directly, benchmark them against the historical record, and quantify how much of the variation in experienced weather survives in memory.

Second, we show how memory distortions shape expectations about future weather. Rather than treating subjective expectations as a measured primitive, we trace them back to the experiences from which they are formed. The resulting problem is not simply that farmers hold biased beliefs. Their beliefs preserve remarkably little of the variation in weather they have actually experienced.

Third, we show that these beliefs can be updated with simple information and identify who responds most strongly to that correction. Farmers whose initial beliefs are furthest from the record respond the most, while literacy and education do not predict responsiveness. Yet literacy strongly predicts the quality of the beliefs farmers begin with. Simple visual information can therefore improve beliefs in the short run even among farmers with limited schooling, while the human capital gradient points to a complementary longer run margin in belief formation.

These findings complement rather than replace the supply side literature on climate adaptation. Better contracts, pricing, delivery, and protective technologies remain important, but their effectiveness depends in part on farmers perceiving enough risk to demand them. Understanding how beliefs shape demand can therefore help improve the design, take up, and effectiveness of the supply side policies already available.

The paper proceeds as follows. Section 2 describes the literature, setting, and sample. Section 3 presents the survey and experimental design. Sections 4 and 5 examine memory and its transmission into expectations. Section 6 presents the information experiment, and Section 7 concludes.

2 Context, Setting, and Sample

This section places our study in the literature on climate adaptation, agricultural risk, and expectation formation. We then describe the eastern Ugandan setting, where rainfed production and limited historical weather information make personal experience especially important for forming expectations. Finally, we describe the construction of our sample of 686 farmers used in the analysis.

2.1 Context: Adaptation, Expectations, and Memory

Much of the economic literature on climate adaptation in agriculture has focused on the design and delivery of protective technologies. Research has studied insurance, stress resistant crops, savings, credit, and other forms of protection, while varying prices, contract design, payment timing, marketing, bundling, and access to complementary inputs (Duflo et al., 2011; Emerick et al., 2016; Carter et al., 2017). Rainfall index insurance provides a particularly clear example. By tying payouts to a measured weather index rather than assessed losses, it reduces moral hazard and verification costs, yet take up has remained low outside substantial subsidies (Giné et al., 2008; Cole et al., 2013). Research has therefore focused on improving the product through better index measurement, pricing, payment timing, marketing, and bundling (Carter et al., 2017). These margins matter. Product design and delivery can substantially affect take up, while basis risk can make index insurance unattractive even to risk averse farmers (Casaburi & Willis, 2018; D. J. Clarke, 2016; N. D. Jensen et al., 2016). This work has greatly improved our understanding of these supply side margins, yet low take up remains persistent. Explaining this persistence requires a deeper understanding of how farmers form the beliefs about weather risk that shape demand.

Some work already points in this direction by showing that experience can change demand for insurance. Recent payouts affect subsequent take up (Cole et al., 2014; Stein, 2018), while experimentally induced experience of disaster risk can increase demand for protection (Cai

& Song, 2017). These findings show that demand is not fixed and that recent experience matters. But the experience studied is typically exposure to the product, a payout, or a salient event. Farmers underlying beliefs about the distribution of weather itself are generally not observed directly. Our focus is on this more fundamental stage of demand formation, asking how farmers remember the weather they have experienced and how those memories shape beliefs about what will happen next.

If beliefs about the underlying weather distribution are part of demand formation, the next question is whether they can be measured reliably. A second strand of work shows that subjective probability distributions can be elicited meaningfully even in populations with limited literacy and formal schooling (Manski, 2004; Delavande et al., 2011; Delavande, 2014; Attanasio, 2009). These elicited expectations also predict economically important choices. Beliefs about monsoon onset predict planting decisions beyond measured risk preferences (Giné et al., 2015), and rainfall expectations among pastoralists can be elicited and respond to new information (Lybbert et al., 2007). This work shows that subjective weather beliefs are both measurable and economically meaningful in settings like ours. It largely takes those beliefs as given and asks what they predict or how they respond to information. We instead ask how they are formed from remembered experience.

If expectations are formed from past experience, an important remaining question is how faithfully that experience is remembered. The economics of memory provides a natural mechanism. Building on classic work on availability (Tversky & Kahneman, 1973), recent models emphasize that people do not retrieve past experience perfectly. Some events are more likely to be recalled than others, so the history people remember can differ systematically from the history they actually experienced (Bordalo et al., 2020, 2023). Consistent with this view, lifetime experience shapes risk taking (Malmendier & Nagel, 2011), while elicited probabilities can be compressed toward a mental reference point (Enke & Graeber, 2023). In settings closest to ours, Gallagher (2014) infers forgetting of floods from insurance demand, and Moore et al. (2019) show that perceptions of normal weather place disproportionate weight on recent experience. What this literature rarely observes is memory itself, benchmarked directly against what actually occurred. Agricultural surveys show that farmers can recall recent inputs and harvests with relatively little systematic bias (Beegle et al., 2012), but recalling a recent quantity is different from reconstructing how weather varied across several past seasons.

Taken together, these literatures leave one important link unobserved. We know that experience can affect demand, that subjective expectations can be measured, and that memory need not preserve past experience faithfully. What remains unclear is how farmers remember the weather they experienced, how those memories shape expectations about future weather, and whether information can correct the resulting beliefs. Our design follows this chain directly. We benchmark remembered weather against the historical record, trace those memories into season ahead expectations, and experimentally test whether simple information about local weather history changes the beliefs farmers bring to their decisions.

2.2 Setting: The Plains West of Mount Elgon

The study covers the districts of Bugiri, Butaleja, and Mayuge in eastern Uganda. The area consists of lowland plains on the Lake Victoria basin plateau, at elevations of roughly 1,080 to 1,250 meters, west of the cash crop zones on the slopes of Mount Elgon.[1] Rainfall is bimodal (FEWS NET, 2010). The first and main growing season runs roughly from February to July, followed by a second season from August to December. Farming is predominantly rainfed and oriented toward own consumption. In the 2014 census, between 84 and 95 percent of households in the three districts grew crops, while between 74 and 90 percent reported subsistence farming as their main livelihood (Uganda Bureau of Statistics, 2017). Mean parcel sizes in the surrounding subregions are below half a hectare (Uganda Bureau of Statistics, 2022). In our sample, 95.5 percent of farmers plan to grow maize, 77.3 percent beans, and 52.0 percent cassava.

Three features of the setting make it particularly well suited to studying the demand side of climate adaptation. First, production is rainfed, making beliefs about the coming season an important input into planting, investment, and protection decisions. For households farming close to subsistence, errors in judging future weather can therefore have direct consequences for both production and consumption.

Second, access to local historical weather information is limited. Farmers have few readily available records against which to compare current conditions or their own experience. In such a setting, past seasons provide a natural benchmark for judging what may happen next. Farmers can draw on seasons they remember, together with environmental signals and information from others, to form a view of the conditions they may face. Memory can therefore be an important input into expectations even when it is not their only source of information. The historical summary provided in our experiment supplies a record that is otherwise difficult for farmers to obtain.

This interpretation is also consistent with what we learned during the study run-up and

[1]The 2025 pilot ran at a separate site on the Elgon slopes in Bulambuli district. Online Appendix OA-7 describes the pilot, and no result in this paper uses pilot data.

with the survey itself. During the design phase, the local field team singled out the timing of the planting rains as the dimension of weather farmers now find hardest to judge, which is why onset receives its own elicitation in Section 3.2. We also ask farmers directly which sources they use when forming rainfall expectations. Personal memory of past seasons is the most common answer, named by 61 percent of respondents, followed by discussions with other farmers at 32 percent, while local leaders and government or extension sources are each named by fewer than 13 percent. Formal forecasts do reach the area: 62 percent report access to some rainfall forecast, mostly over the radio, but only a fifth of those receive forecasts at the village level, and when asked whom they trust on weather, farmers most often name experienced farmers (63 percent). Memory and local exchange, rather than formal information, therefore dominate how farmers in this setting think about future weather, which motivates our focus on memory as one input into expectation formation.

Third, formal weather insurance plays little role in current protection decisions. Among the 404 respondents planning any pre-season adaptation, only eight, or 2.0 percent, report plans to purchase crop insurance. The setting therefore allows us to study demand formation before insurance has become a routine part of farmers' adaptation choices.

2.3 Sample: From Villages to Farmers

We constructed the sample in three steps, described in full in Appendix A.1. We began with 594 candidate villages identified with our field partner and restricted the frame to 442 rural villages across 11 non-town subcounties that were reachable by the field team. These were small, predominantly subsistence farming communities suited to the question we study. We also required the surrounding weather history to contain meaningful variation. Using remotely sensed rainfall measured in five-day intervals and satellite vegetation data, we selected areas in which conditions varied both across seasons within a location and across locations. This ensured sufficient variation in experienced weather to study how differences in past conditions are preserved in memory. From the resulting 442 villages, a stratified draw with a fixed seed selected 77 candidate villages.

These 77 villages formed a preselection rather than the final sample. The field team visited each village and sought permission from village leaders and community mobilizers before any treatment was assigned. In total, 60 villages were able to host and organize the survey and experiment. These villages, containing 686 available farmers with between 7 and 14 respondents per village, define the study sample. Because village access was determined before randomization, it affects the population to which our results apply but not treatment comparisons within the experiment. Randomization took place only after the 60 villages had been fixed, as described in Section 3.5. Fieldwork ran from 10 to 31 March 2026, before the 2026 growing season had resolved. Seven enumerators completed all 686 interviews, with no partial or invalid submissions.

Table 1 summarizes the randomized sample. Respondents are 41.4 years old on average, 44.3 percent are women, and mean household size is 7.75. Farms are small, with an average of 2.79 cultivated acres and a median of two acres. Human capital is limited. Just over half report being able to read and write, while 39.2 percent have no formal education. Literacy becomes important later because it predicts both the accuracy of remembered weather and the distance of initial beliefs from the historical record. All 686 respondents identify farming as their occupation. Use of protective inputs is uneven. Some 57.3 percent use fertilizer, 52.6 percent use improved seeds, and 22.6 percent report irrigation. Among respondents planning any adaptation, only 2.0 percent plan to purchase crop insurance.

One feature of the data is important for interpreting the analysis that follows. The objective weather record varies at the subcounty level. The 60 study villages span 42 parishes but only 11 subcounties, giving 55 subcounty-year observations over 2021–2025 rather than the much larger number suggested by the farmer-level panel. The 686 farms also fall within only 17 rainfall grid cells. We therefore do not observe independent farm- or village-level weather histories. Empirical specifications account for this common geographic variation using subcounty fixed effects where appropriate, while standard errors are clustered at the village level. Appendix C.6 discusses the resulting dependence structure in detail.

3 Survey and Experimental Design

Our empirical strategy combines three elements. A broad survey characterizes farmers and their production environment. A visual elicitation measures memories and expectations without relying on literacy or familiarity with formal probabilities. An information experiment then tests whether those beliefs update when farmers observe the local historical weather record. This section describes each element and the randomization that links them.

3.1 The Survey

Each respondent completed a structured interview covering demographics, farming and land use, access to agricultural information, weather memories and expectations, planned adap-

Table 1: Sample characteristics, by experimental arm.

	Arm A	Arm B	Arm C	All
Respondents	274	293	119	686
Villages	24	26	10	60
Age (years)	42.5 (12.9)	40.0 (12.5)	42.1 (12.7)	41.4 (12.7)
Female	0.434	0.457	0.429	0.443
Household size	7.98 (4.38)	7.64 (4.86)	7.50 (5.26)	7.75
Reads and writes	0.504	0.539	0.462	0.512
Cultivated area (acres)	2.84 (2.22)	2.68 (2.42)	2.97 (2.18)	2.79
No formal education	0.383	0.375	0.454	0.392
Primary education	0.427	0.430	0.319	0.410
Secondary or above	0.190	0.195	0.227	0.198
Plans to grow maize	0.960	0.939	0.983	0.955
Plans to grow beans	0.745	0.795	0.782	0.773
Plans to grow cassava	0.529	0.488	0.580	0.520
Plans to grow groundnuts	0.449	0.464	0.445	0.455
Plans to grow rice	0.164	0.119	0.076	0.130
Uses fertilizer	0.580	0.556	0.597	0.573
Uses improved seeds	0.536	0.509	0.546	0.526
Uses irrigation	0.245	0.212	0.218	0.226
Plans crop insurance	0.019	0.024	0.013	0.020

Notes. Arm means, with standard deviations in parentheses for continuous variables. Indicator standard deviations are omitted. Villages are the randomization units. Literacy uses 684 nonmissing responses. The final education category pools secondary, tertiary, and vocational schooling. Crop shares refer to crops planned for the coming season and need not sum to one. Planned insurance is measured among respondents planning any adaptation ($N = 404$). Source: March 2026 survey.

tation and spending, risk perceptions, and the location of the main plot. The survey also records free recall of extreme weather events, the sources farmers use when forming expectations, and the role of social learning. Online Appendix OA-2 reproduces the full instrument. This breadth allows us to characterize who holds more accurate memories, relate expectations to planned protection, and test whether observable farmer characteristics predict updating in the experiment.

3.2 Eliciting Memories and Expectations

Standard probability questions are poorly suited to a setting where many respondents cannot read and formal probabilities are unfamiliar. A common solution in the expectations literature is to represent probabilities physically. Respondents allocate a fixed number of tokens across pictured outcomes, with the allocation representing their subjective probability distribution (Manski, 2004; Delavande et al., 2011; Delavande, 2014). The method requires no literacy and only basic counting, and beliefs elicited in this way have been shown to predict real economic choices (Delavande et al., 2011; Giné et al., 2015). We use this approach to measure two dimensions of weather beliefs, the conditions farmers associate with a season and the timing of the planting rains.

The nine-picture rainfall scale. Our main elicitation uses a card showing the same representative maize field under nine conditions, ranging from extreme drought in bin one, through optimal growing conditions in bin five, to severe flooding in bin nine. The images were constructed from a large set of photographs of agricultural fields in the study region and combined into a synthetic representation designed to look realistic and familiar to respondents, while varying the water conditions depicted. Figure 1 reproduces the card.

Before the elicitation, the enumerator read a fixed description of all nine images and gave the respondent ten physical tokens. Each token represented one vote to be placed on the image or images that best described the season in question. The allocation gives a discrete subjective distribution over the nine ordered outcomes. We repeated the exercise for each of the five past seasons from 2021 to 2025 and for the, then, upcoming 2026 agricultural season.

Respondents who could not remember a past season after two neutral scripted probes were not required to allocate tokens. Recall therefore declines with distance in time, from 99.3 percent of respondents for 2025 to 61.7 percent for 2021, yielding 2,637 respondent year observations. Online Appendix OA-3 reproduces the full elicitation script, image descriptions, and enumerator instructions.

The planting-onset calendar. The second instrument elicits beliefs about when the planting rains begin. Farmers allocate the same ten tokens across ten dated intervals spanning 15 February to 24 April, shown in Figure 2. The allocation therefore gives a subjective distribution over the timing of rainfall onset, using the same physical format as the seasonal weather elicitation.

We measure onset separately because the timing of the first reliable rains is central to planting decisions. Farmers must decide when there is enough moisture to plant without knowing whether the rains will continue, and the local field team described this timing as increasingly difficult to predict. Planting too early can leave germinating seed exposed if the rains stop, while planting too late shortens the effective growing period and can reduce

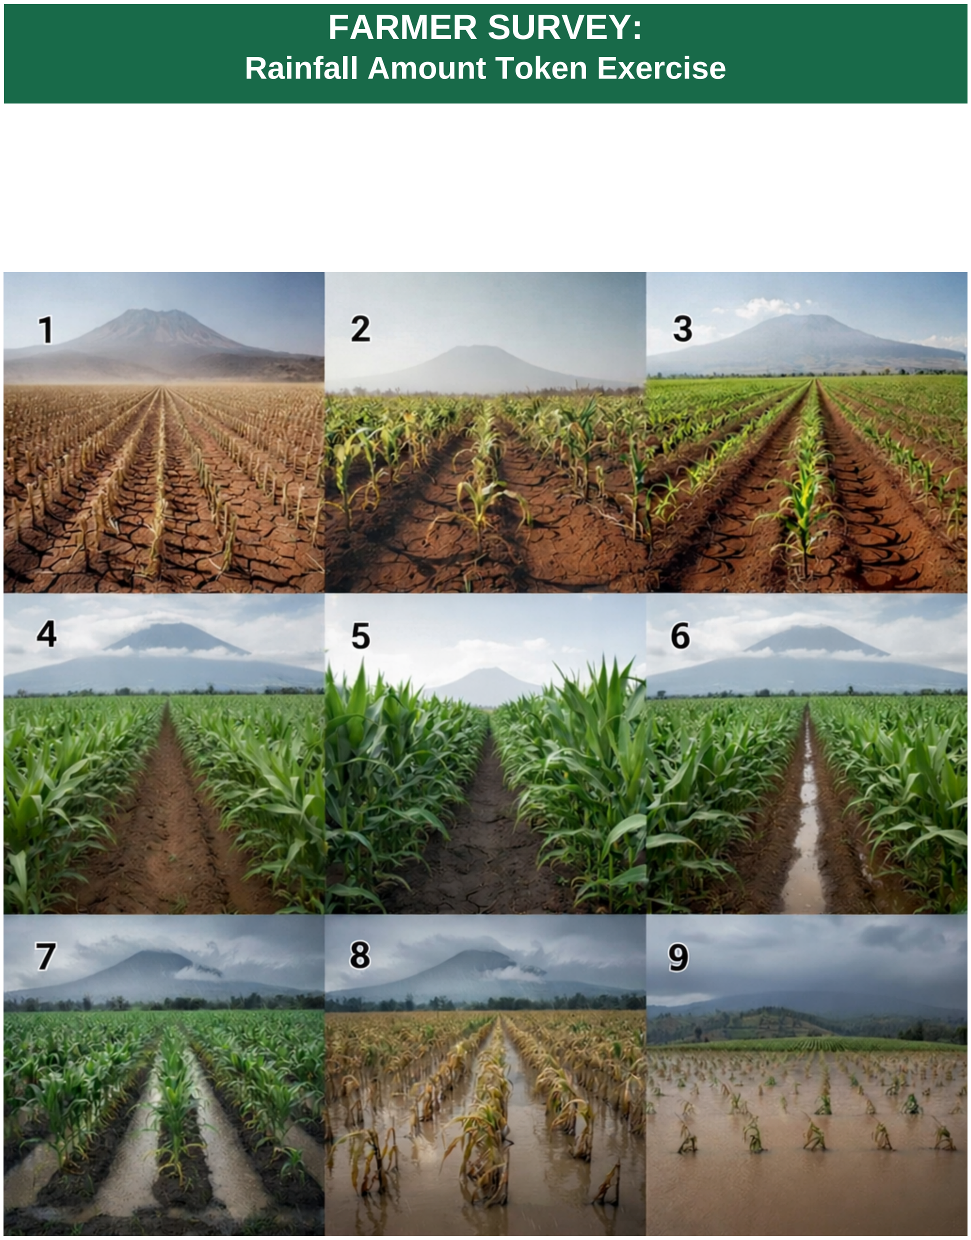

Figure 1: The nine-picture card over which farmers allocated ten tokens.
Notes: Nine-picture rainfall scale used in all survey interviews. Online Appendix OA-3 gives the full elicitation script. Source: authors' field materials.

yields. We repeat the onset exercise for each past season and for 2026.[2]

Separately from the season-by-season elicitation, respondents were asked to recall any extreme weather events without being prompted by year. For each event, they reported the

[2]Elicited onset is consistently 26 to 32 days later than our rainfall based benchmark, while the benchmark varies little across the study area. We therefore use onset only for within person and experimental comparisons, not for cross sectional measures of accuracy. Appendix B.3 provides further details.

hazard type, year, and perceived severity. This free recall exercise produced 550 first-named events, including droughts, floods, hailstorms, and storms, and provides a second measure of weather memory organized around salient events rather than seasons.

We use the same elicitation instruments to measure expectations for the coming season. Farmers allocate tokens across the nine-picture scale for expected 2026 conditions and across the planting-onset calendar for expected timing of the rains, and also report their confidence. Memories and expectations are therefore measured on the same scales, allowing us to trace remembered experience directly into beliefs about the future.

FARMER SURVEY CALENDAR: Start of Rains

Allocate 10 tokens across the intervals to indicate when you remember / expect the seasonal rains started

Interval	Date range	General season	Token(s)
1	15–21 Feb	Very dry; land still hard; few early land preparers start ploughing.	
2	22–28 Feb	First scattered showers in some low-lying areas; usually still unreliable.	
3	1–6 Mar	Land preparation intensifies; some early planters risk maize if showers continue.	
4	7–13 Mar	Transitional period; first rains often occur; seedbed preparation peaks.	
5	14–20 Mar	Some areas record 1–2 good rains; not yet sustained in the highlands.	
6	21–27 Mar	Typical onset for lower slopes; first main planting begins.	
7	28 Mar–3 Apr	Onset consolidates; continuous rain spells; widespread maize and bean planting.	
8	4–10 Apr	Average onset period for mid-slopes; moisture reliable.	
9	11–17 Apr	Late planters begin; upper slopes catch up.	
10	18–24 Apr	Late onset years often start here; second wave of planters.	

Figure 2: The ten-interval planting-onset calendar.

Notes: Planting-onset calendar used in the ten-token elicitation, covering ten intervals from 15 February to 24 April. Respondents could select *No Memory* after two neutral probes. Appendix B.4 describes the scale construction. Source: authors' field materials.

3.3 Benchmarking Memory Against the Record

To assess memory accuracy, we need an objective measure of what occurred on the same scale used in the elicitation. We classify each subcounty-year from 2021 to 2025 into one of the nine weather bins using a fixed rule that combines CHIRPS rainfall (Funk et al., 2015) with Sentinel-2 vegetation data. The classification combines rainfall over three onset-aligned stages of the growing season with peak vegetation greenness measured within cropland, and maps these conditions onto the nine-point scale. Appendix B.1 describes the full classification rule and thresholds.

We benchmark free recall of extreme events separately against the 2001–2025 rainfall record for each farmer's subcounty. Because the seasonal classification produces one objective weather bin for each subcounty-year, the historical record contains 55 independent subcounty-year cells over 2021–2025, as discussed in Section 2.3. The empirical analysis accounts for this geographic structure throughout.

3.4 The Information Treatment

The information treatment presents local weather history in the same visual language used in the elicitation. Each farmer receives a printed packet covering the five seasons from 2021 to 2025. For each season, shown from most to least recent, the packet displays the corresponding picture from the nine-bin scale together with short statements summarizing rainfall conditions, crop conditions, and how the season compared with the other years shown. A companion calendar summarizes rainfall through the season, using one to three water-drop symbols to indicate low, typical, or high rainfall in each ten-day period.

Presenting the historical record in the same format as the elicitation makes the information directly comparable with the memories farmers had just reported and the expectations elicited afterward. Figure 3 shows excerpts of the two central cards of the packet; because every year block repeats the same structure, two to three sample years suffice to convey what respondents were handed. Appendix B.5 reproduces a complete treatment packet at full size.

Three design choices are important for interpreting the results. First, the cards were generated entirely automatically. A single scripted pipeline constructed each subcountys weather history from the satellite data, assigned seasons to bins using the fixed classification rule, generated the accompanying text, and rendered the printed sheets. No researcher judgment therefore entered between the underlying record and the information farmers received. Second, the treatment was designed to be as accessible as the elicitation itself. Information is conveyed through the same pictures and visual symbols respondents had just used, while the short accompanying text could be communicated orally by the enumerator. Reading was therefore not required to understand the card. Third, the treatment contains no forecast. It instead reorganizes recent local weather history into an accessible form, placing the intervention close to Hanna et al. (2014). Farmers receive a usable summary of experience they have already lived through rather than information about what will happen next.

The construction also places limits on what the experiment identifies. Because card content is based on the subcounty classification, the 42 parish cards contain only eight distinct weather histories. The displayed seasons span bins 3 through 7, so the experiment does not identify responses to information about the most extreme conditions. Finally, the second elicitation follows the treatment within minutes, so immediate updating cannot be cleanly separated from anchoring on the displayed information. Appendix B.2 describes these limitations and the resulting interpretation of card content regressions.

3.5 Randomization and Experimental Arms

Randomization. The 60 villages were assigned to three experimental arms at the village level. Village-level assignment limits information spillovers between treated and untreated farmers within the same community. We stratified randomization by subcounty and elevation tercile and used elevation and satellite greenness to form 30 matched pairs, with assignment conducted using a fixed seed (Imai et al., 2009; Bruhn & McKenzie, 2009). The resulting design contains 60 treatment clusters, so standard errors for experimental comparisons are clustered at the village level. Table 1 shows that the arms are similar on baseline respondent characteristics. Appendix A.3 documents the full randomization procedure, pairing, and balance checks. Treatment assignment was hard coded into the survey form five days before fieldwork began, leaving enumerators no discretion over assignment.

Experimental design. Figure 4 summarizes the three arms. All farmers complete the same survey and memory elicitation. Arm A then completes the expectation elicitation twice without receiving information in between. This captures any change caused simply by repeating the exercise. Arm B completes the same two elicitations but receives the information treatment between them, so the difference in changes between Arms B and A identifies the effect of the information. Arm C receives the information before a single expectation elicitation, providing a separate check that the first elicitation itself does not drive the treatment response.

The three arms separate changes caused by repeating the elicitation from changes caused by the information itself. Arm A provides the benchmark. Farmers state their expectations twice with nothing in between, allowing us to measure how much beliefs move simply because the same question is asked again. Arm B follows the same sequence but receives the

Figure 3: The information treatment.

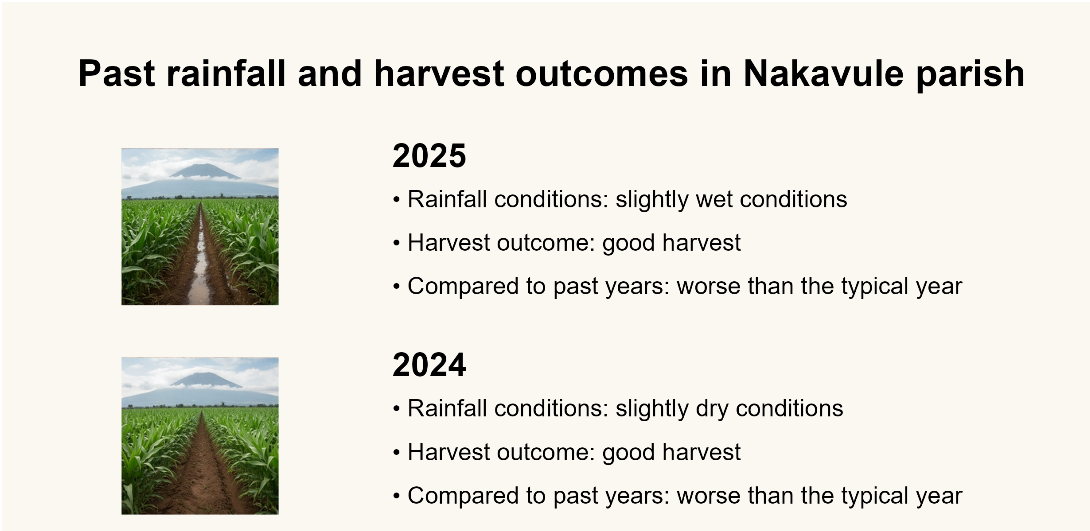

(a) The overview card (excerpt)

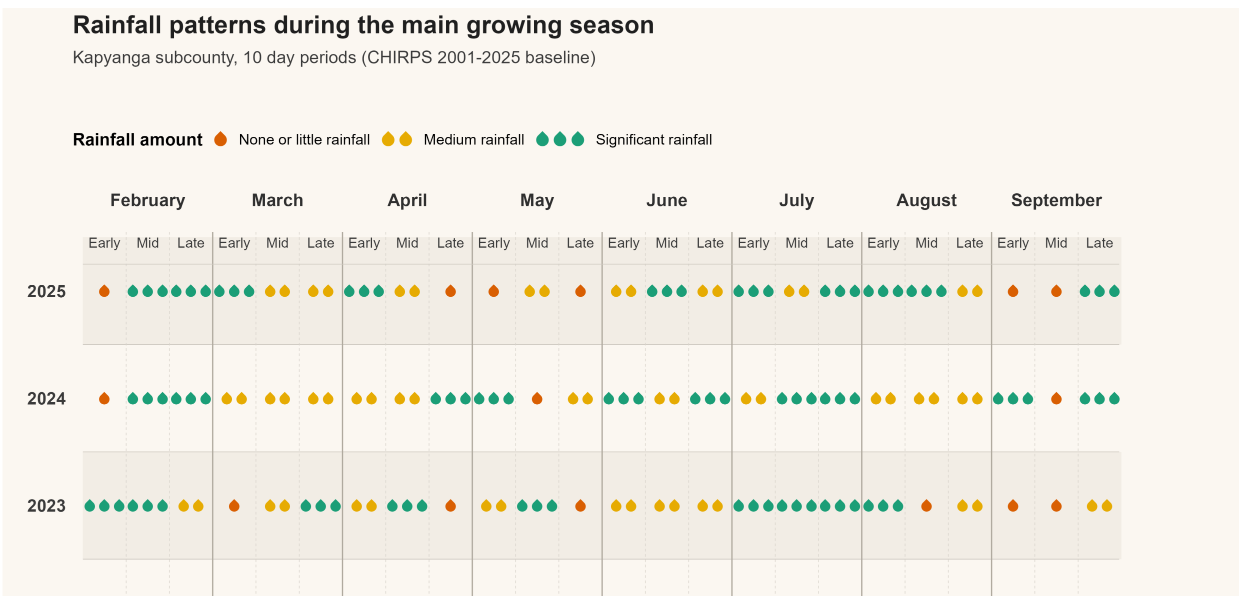

(b) The rainfall calendar (excerpt)

Notes: Excerpts from one parish's information packet, as printed and shown to respondents in Arms B and C. (a) The overview card. Each year block pairs the season's picture bin, the same photograph used in the token exercise, with plain-language lines on rainfall, the harvest, and the comparison with past years; the top two of five blocks are shown. (b) The rainfall calendar. Each row is a season and each cell carries one to three water drops for none-or-little, medium, or significant rainfall against the subcounty's CHIRPS 2001–2025 baseline; the first three of five rows are shown. All remaining blocks and rows repeat the identical structure, and Appendix B.5 reproduces the complete seven-sheet packet. Source: authors' field materials.

information treatment between the two elicitations. The difference in belief revision between Arms B and A therefore identifies the effect of the information beyond any movement caused by repetition. The first elicitation in Arms A and B also provides the untreated expectations

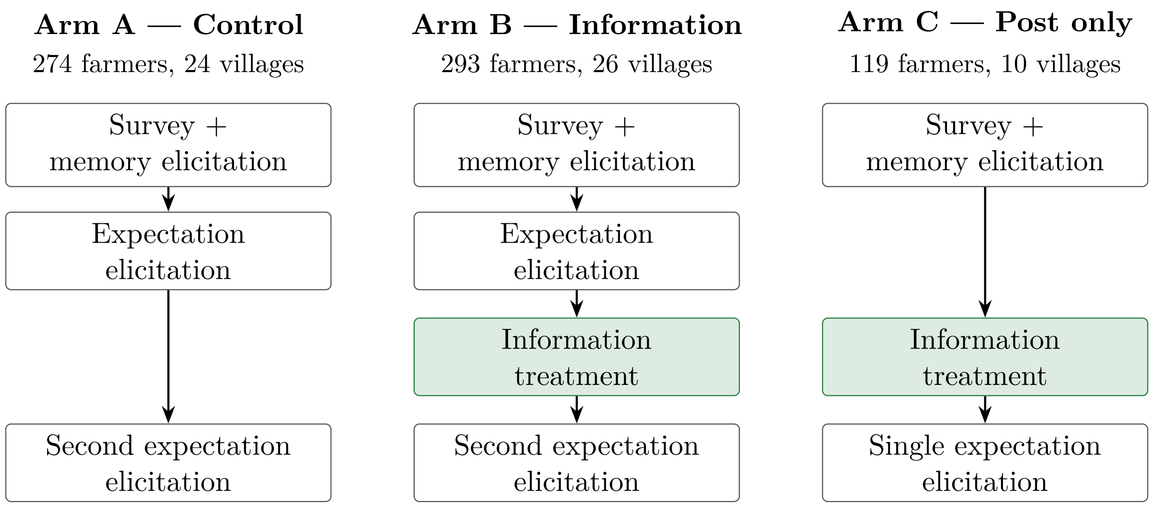

Figure 4: The three experimental arms.

Notes: Arm A repeats the expectation elicitation without information, Arm B receives information between two elicitations, and Arm C receives information before a single elicitation. Green-shaded boxes indicate the information treatment. Villages are the randomization units.

used in Sections 4 and 5.

Arm C provides a separate check on this design. Farmers receive the information before stating their expectations for the first time. Comparing these beliefs with those formed after revision in Arm B allows us to assess whether stating an initial belief before receiving information changes the response to the treatment. Because this is a secondary design check rather than the main experimental comparison, Arm C contains fewer villages. The primary treatment effect throughout the paper is therefore the difference in belief revision between Arms A and B.

One feature of the design qualifies the interpretation of this comparison. After receiving the information and before stating their second expectation, farmers in Arm B answered three short questions about whether the information differed from what they expected, whether it was surprising, and whether it would change their expectations. Arm A answered no comparable questions at that point. The estimated treatment effect therefore captures the information as it was delivered, including any additional effect of these reaction questions. The data cannot separate the two. Section 6 returns to this issue when interpreting the experimental results.

4 Climate Memory

With the measurement and experimental design in place, we first ask what farmers actually remember about past seasons. We compare recalled weather with the independently constructed historical record. The central result is that farmers preserve some of the ordering of seasons but very little of the distance between them. We then characterize this compressed memory, showing where it is centered, how it changes with recall age, and whose memories are more accurate. All results in this section use memories elicited before treatment and benchmarked against the historical record. The experimental treatment plays no role in these estimates.

4.1 Recall and Realized Weather

We begin by asking how much of the variation farmers actually experienced survives in memory. If one season was substantially wetter than another, does the farmer remember a similarly large difference, or do the two seasons look more alike in retrospect? We estimate this relationship by comparing each farmer's recalled value on the nine-picture weather scale with the value indicated by the historical record for that season.

$$\bar{r}_{i,t} = \alpha_i + \beta m_{s(i),t} + \varepsilon_{i,t}, \tag{1}$$

where $\bar{r}_{i,t}$ is the token-weighted mean value of respondent i's recalled distribution for season t, and $m_{s(i),t}$ is the value assigned to that season by the historical record in the respondent's subcounty. Respondent fixed effects absorb persistent differences in how farmers use the scale, so β captures how strongly a farmer's recall responds to differences in the weather that farmer experienced. If differences across seasons were remembered one for one, β would equal one. If recall did not distinguish between seasons at all, it would equal zero. Standard errors are clustered at the village level, with inference based on a wild cluster bootstrap (Cameron et al., 2008).

The estimate is $\beta = 0.0948$ (standard error 0.0342). The historical record varies only across the 11 subcounties, so village-clustered inference understates the uncertainty that the coarse benchmark implies. Clustering on subcounties leaves the point estimate unchanged but raises the standard error to 0.0566, with an asymptotic p-value of 0.094 and a wild bootstrap p-value of 0.176 (Appendix C.6). The point estimate means that only about one tenth of the variation across seasons survives in memory. When two seasons differ by one point on the nine-picture scale in the historical record, their recalled values differ by only about 0.1 points. A clearly dry and a clearly wet season that are five points apart in the historical record are therefore remembered as only about half a point apart. Realized seasons between 2021 and 2025 occupy bins 3 through 7 of the nine-point scale, so the estimate describes compression over moderate variation rather than over the most extreme conditions the scale can represent. Figure 5 makes this compression visible. Farmers retain some sense of which seasons were wetter and which were drier, but most of the difference between those seasons disappears in memory. Recalled values instead cluster around roughly 4.46 on the nine-point scale, slightly below the optimal condition.

This loss of variation matters for judging weather risk. Remembering that one season was wetter than another preserves some information about their ordering. But judging whether a future season is unusually dry or wet also requires remembering how different past seasons were from one another. When those differences are compressed, the weather history available to form that judgment becomes much less informative.

The compression result is not inconsistent with evidence that farmers can recall recent agricultural quantities accurately. Beegle et al. (2012), for example, find little systematic bias in farmers' recall of inputs and harvests from a recent season. Our setting involves a different memory task. Rather than recalling a single quantity from one season, respondents reconstruct how weather differed across several past seasons. Nearly all respondents can report the most recent season, with recall coverage of 99.3 percent, but the central finding is that the differences across seasons are only weakly preserved in memory.

Two mechanisms could generate the compression, and our design cannot distinguish between them. Differences across seasons may simply fade in memory, or recalled seasons may be pulled toward a common reference point, as Godlonton et al. (2018) document for past agricultural outcomes. The interactions that could separate these explanations are not informative in our five-season panel, so we leave the underlying mechanism open.

Where recall is centered. Compression describes how strongly memories respond to differences in realized weather. A separate question is where those memories converge. On average, recalled conditions are 0.324 points drier than the historical record. This average reflects strong compression from both sides of the scale. Seasons that were actually dry are remembered 1.163 points wetter than they were, while wet seasons are remembered 2.226 points drier. Seasons near the middle are remembered 0.379 points drier.

These patterns reinforce the compression result. Dry and wet seasons are both pulled

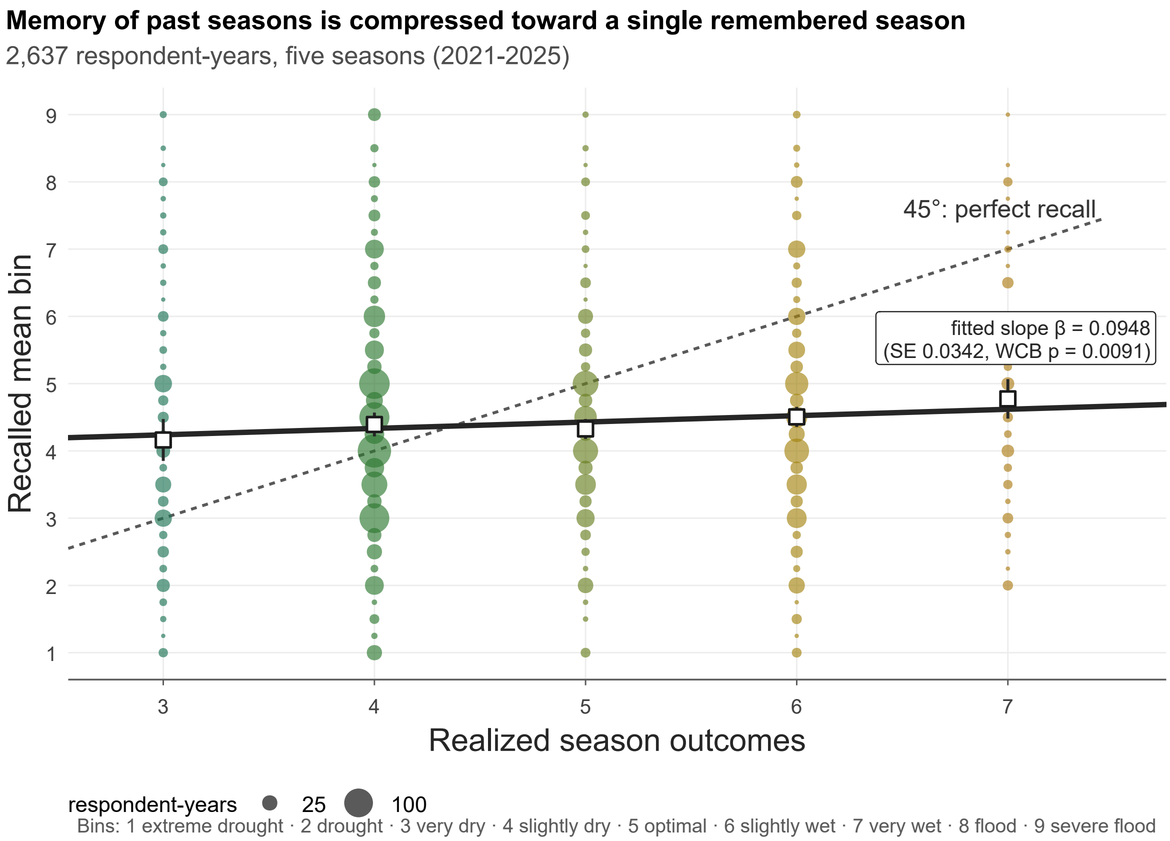

Figure 5: Memory is compressed.

Notes: Discs show combinations of realized and recalled values on the nine-picture weather scale, sized by the number of respondent-year observations and coloured by the realized outcome, from dry (amber) through optimal (green) to wet (blue). White squares show mean recall at each realized value, with 95% village-clustered confidence intervals. The dashed 45-degree line represents perfect recall, while the solid red line shows the estimated relationship, $\beta = 0.0948$ (SE 0.0342). Realized weather varies at the subcounty-year level, giving 55 independent historical observations. Source: authors' computation.

sharply toward the middle of the scale, but the point toward which they converge lies slightly on the dry side of the optimal condition. A related asymmetry appears later when farmers freely recall past droughts and floods, which we examine in Section 5.

How recall changes with time. One possibility is that compression simply reflects memories becoming less accurate as seasons recede into the past. In the raw data, older seasons do appear less accurately recalled, with the accuracy penalty increasing by 0.0624 for each additional year of recall age. But the older seasons in our five-year window also happened to be more extreme, and extreme conditions generate larger recall errors because of compression. Once we account for the realized weather, the relationship with recall age falls to -0.0091 and is statistically indistinguishable from zero. Within this five-year window, we therefore find little evidence that older seasons are less accurately recalled simply because they are older.

What does change is the concentration of the recalled distribution. Memories of older seasons place probability mass over a narrower range of the scale. Entropy falls and sharpness rises, while the probability assigned to the value that actually occurred also falls. One interpretation is that recent seasons retain more of their original nuance, while older seasons gradually settle into a more definite summary of what they were like. Older memories therefore become more concentrated without becoming more accurate.

If farmers draw on these memories when forming expectations, this distinction matters. Recent experience may contribute a broader sense of possible conditions, while more distant experience enters as a narrower and more simplified signal.[3]

Who remembers more accurately. Memory accuracy also differs systematically across farmers. Literate respondents recall past seasons 0.17 standard deviations more accurately on the ranked probability score, and each higher education tier is associated with a further 0.12 standard deviation increase in accuracy. Older respondents and those from larger households recall somewhat less accurately. Taken together, education, literacy, age, and household size strongly predict memory accuracy.

The human capital gradient is important for the results that follow. Farmers with less education and lower literacy begin with less accurate representations of past weather, and Section 6 later shows that literacy also predicts how far initial expectations lie from the historical record. The average tendency to remember conditions as slightly drier, by contrast, does not vary systematically with these characteristics ($p = 0.50$ for their joint significance).

Taken together, these results characterize a strongly compressed memory of past weather. Farmers preserve some ordering across seasons but remarkably little of the distance between them. Recalled conditions converge toward a point slightly on the dry side of optimal, older memories become more concentrated without becoming more accurate, and memory accuracy is systematically related to literacy and education. The next question is whether this compressed picture of the past shapes expectations about the season ahead.

[3]Recalled planting onset shows a related time pattern. Remembered onset moves 1.45 days later for each additional year of recall age, while recalled seasonal conditions show no corresponding drift. Appendix B.3 reports these results.

5 From Memory to Expectations

The setting and survey evidence suggest that farmers draw on remembered experience, alongside other sources of information, when forming expectations about future weather. The previous section showed that this remembered history is highly compressed. We now ask whether that distortion carries into beliefs about the coming season. First, we test directly whether remembered seasons are related to season ahead expectations. Second, we show how much variation survives when past weather is filtered through the estimated memory distortion. Third, we examine a complementary form of memory, free recall of extreme events, and ask how it relates to expectations of future extremes. Finally, we relate these beliefs to planned adaptation and protection. Together, the results trace remembered experience into the expectations farmers bring to economic decisions.

5.1 Memory and Season Ahead Expectations

We begin with the most direct link between memory and expectations. If past experience is an input into beliefs about future weather, farmers who remember wetter recent histories should also expect wetter conditions in the coming season. We test this by relating each farmer's pre-information expectation for 2026 to a weighted average of the five seasons that farmer recalls:

$$\bar{f}_i = \delta_{s(i)} + \lambda \sum_{k=1}^{5} w_k \, \bar{r}_{i,t-k} + \mathbf{x}_i' \boldsymbol{\theta} + u_i, \tag{2}$$

where $\bar{f}_i$ is the mean value of respondent i's expectation for 2026 on the nine-picture weather scale, $\bar{r}_{i,t-k}$ is the corresponding recalled value for a past season, and the w_k are non-negative weights that sum to one. Subcounty fixed effects absorb common differences in local weather conditions, while $\mathbf{x}_i$ contains respondent controls. The coefficient λ measures how strongly differences in remembered weather are reflected in expectations for the coming season. This relationship is descriptive rather than causal.

The estimated λ is 0.4354 (standard error 0.0718) among the 330 respondents in Arms A and B with complete five-season memory histories. A farmer whose remembered history is one point wetter on the nine-picture scale expects the coming season to be about 0.44 points wetter. The relationship is therefore economically substantial. It is also robust to how the five past seasons are weighted. Giving each season equal weight produces an estimate of 0.4092. In Arm A alone, where farmers never receive information, the estimate remains

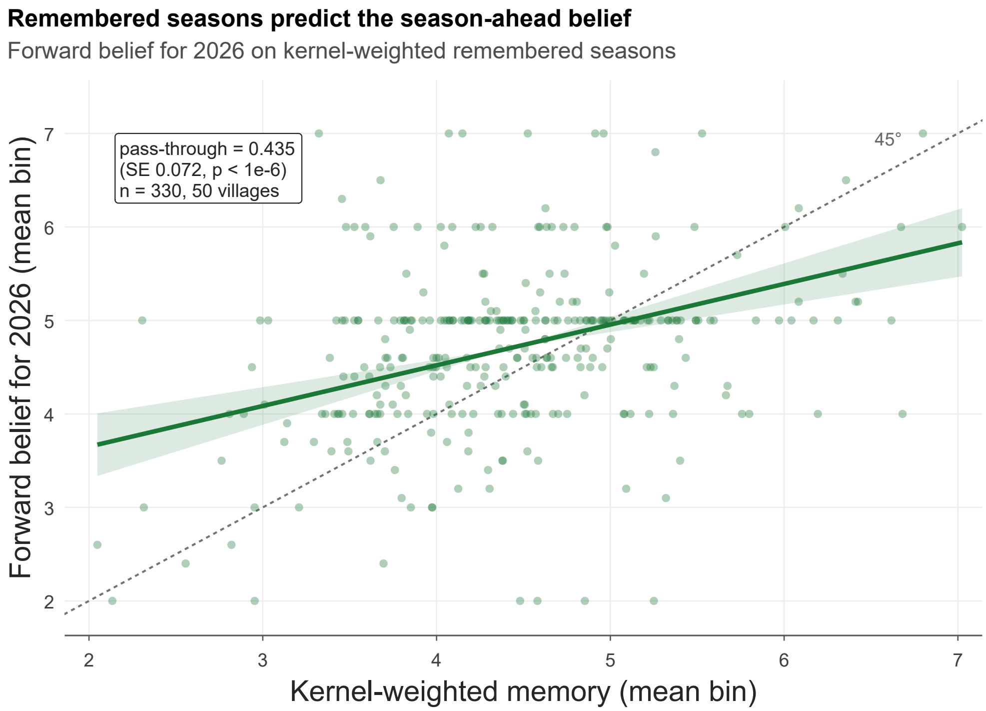

Figure 6: Remembered weather predicts expectations for the coming season.

Notes: Each point represents one of the 330 Arm A and Arm B respondents with a complete five-season memory history. The horizontal axis shows the weighted mean of recalled conditions and the vertical axis the mean pre-information expectation for 2026, both on the nine-picture weather scale. The fitted relationship has slope $\lambda = 0.4354$ (SE 0.0718), with subcounty fixed effects and village-clustered standard errors. Source: authors' computation.

0.3246 (standard error 0.1132). Remembered seasons and expectations are elicited separately, so the relationship does not arise mechanically from the measurement procedure. Given the limited sample, the individual year weights are estimated with too much noise, so we focus on the overall relationship between remembered history and expectations.

The result establishes the link needed for the rest of the argument. Farmers do not form expectations independently of remembered experience. Differences in remembered weather are systematically reflected in what they expect next. The implication of the compression documented in Section 4 therefore depends on how much useful variation remains in the weather history that memory provides.

5.2 The Perceived Weather History

We next ask how much variation survives in that remembered history. We take the actual sequence of weather experienced in each subcounty and apply the memory distortion estimated in Section 4. This gives the historical reference implied by the way farmers remember past seasons.

This object is not a directly observed expectation for each historical year. Rather, it shows what the recent weather history would look like if actual conditions entered memory through the estimated compression and dry lean. For a farmer entering season t, we construct

$$\hat{\mu}_{st} = \frac{1}{5} \sum_{j=t-5}^{t-1} \left[5 + \beta(b_{sj} - 5) + \gamma \right], \qquad \beta = 0.0948, \quad \gamma = -0.324, \tag{3}$$

where b_{sj} is the weather that actually occurred in season j in subcounty s. The comparison is simple. We first calculate the five-season history a farmer would have in mind if past weather were remembered accurately. We then calculate the same history after applying the memory distortion estimated in Section 4. The difference between the two shows how much information about past weather is lost in memory.

The loss is substantial. With accurate memory, recent weather histories differ by 3.58 points on the nine-picture scale across places and years, ranging from clearly dry to clearly wet. After applying the estimated memory distortion, the same histories differ by only 0.34 points. In other words, roughly nine tenths of the variation that distinguishes one recent weather history from another disappears in memory. The standard deviation falls from 0.764 points under accurate memory to just 0.072 points under the estimated memory process.

Why does this matter economically? Protection is most valuable when farmers perceive the coming season as unusually risky. If remembered weather histories look almost the same across places and years, farmers have little basis for judging when the next season is unusually dry, wet, or otherwise risky. This does not imply that compressed memory necessarily lowers average demand for protection, since a farmer who perceives high risk every year could still choose to protect every year. The more precise implication is that compressed memory weakens the ability to adjust protection to differences in perceived risk across seasons.

We also check whether the flatness could be created by the way beliefs are elicited rather than by memory itself. This does not appear to be the case. Nearly half of respondents place all ten tokens on a single image, so the result is not driven by farmers spreading their answers evenly across the scale. The compression is also essentially unchanged when we replace the nine-picture scale with broader dry, normal, and wet categories, and when we account for conditions shared by farmers interviewed in the same village. Appendix D reports these and additional measurement checks. As a further validation, the appendix shows that measuring weather shocks relative to the constructed perceived history adds almost no

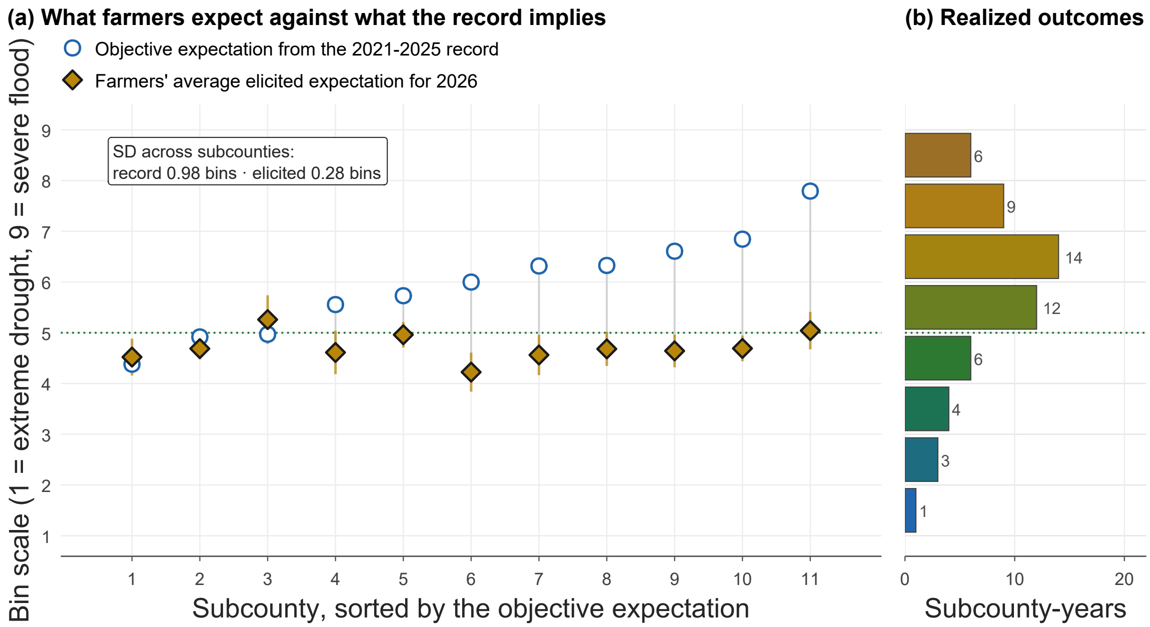

Figure 7: Farmers' expectations barely move while the record does.

Notes: Panel (a) compares, for each of the 11 subcounties, the objective expectation implied by the 2021–2025 record (open blue circles, the mean realized outcome over the five seasons) with farmers' average elicited expectation for 2026 (filled amber diamonds, with 95% confidence intervals), both on the nine-picture scale. The record varies across subcounties with a standard deviation of 0.98 bins; the elicited expectations vary with a standard deviation of 0.28 bins. The dotted green line marks the optimal bin. Panel (b) shows the distribution of realized outcomes over the 55 subcounty-year cells. Applying the estimated memory distortion ($\beta = 0.0948$, $\gamma = -0.324$) to the same record compresses its standard deviation from 0.764 to 0.072 bins, consistent with the elicited pattern shown. Source: authors' computation.

predictive information for farm productivity beyond a conventional rainfall anomaly. This is what we would expect from a perceived history in which most of the original variation has disappeared.

5.3 Extreme Weather Shocks and Expectations of Recurrence

The season by season elicitation asks farmers to reconstruct ordinary variation in weather across several past seasons. But memory need not preserve all experiences equally. Work on availability and selective recall suggests that vivid or distinctive events can remain especially easy to retrieve and can therefore receive disproportionate weight in judgments about future risk (Tversky & Kahneman, 1973; Bordalo et al., 2023). This raises an important possibility in our setting. Even if differences across ordinary seasons are heavily compressed, extreme weather events may remain salient in memory and continue to shape expectations.

We therefore complement the season by season elicitation with a separate free recall exercise focused on extreme weather shocks. Farmers were asked whether they remembered any major weather event from the past and, for each event, reported the type of shock, the year in which it occurred, and its perceived severity. We then asked whether they expected a major weather event in the coming season and, if so, what type of event they expected and how severe they thought it would be.

Empirically, the two memory measures are only weakly related. The correlations between recalling an extreme event and the accuracy, concentration, and dry lean of season by season memory are 0.079, 0.126, and -0.107, respectively. Salient event recall therefore appears to capture information that is largely distinct from the season by season memory measure.

Extreme event memory nevertheless carries strongly into expectations. Among farmers who freely recall a past extreme event, 48.1 percent expect an extreme event in the coming season, compared with 27.8 percent among those who recall no past extreme. With subcounty fixed effects, recalling a past extreme is associated with a 19.4 percentage point higher probability of expecting one in 2026 (standard error 4.5 percentage points). The content of the memory also persists. Farmers who remember more severe events expect more severe events, and those who remember a particular hazard are more likely to expect the same type of hazard again.

The distinction remains when the two forms of memory are considered together. Free recall of an extreme event and the concentration of season by season memory both predict whether a farmer expects an extreme in 2026, while overall season by season memory accuracy does not. Salient events therefore provide a separate channel through which past experience enters beliefs about future extremes.

How well do recalled shocks match the historical record? We next examine whether the events that remain salient in memory also stand out in the historical weather record. For each named event, we locate the reported year in the farmer's subcounty's 2001–2025 rainfall distribution. Recalled floods line up reasonably well with measured rainfall.

The pattern is different for droughts. At the spatial and temporal resolution available in the historical record, years named as droughts are not objectively drier than typical years. They lie around the 61st percentile of seasonal rainfall, while the longest dry spell in recalled drought years lies around the 53rd percentile. Reported drought severity is also unrelated to the measures of realized dryness available to us. Recency predicts which years are named more strongly than these objective drought measures.

This evidence has an important limitation. Our benchmark is measured at the subcounty-

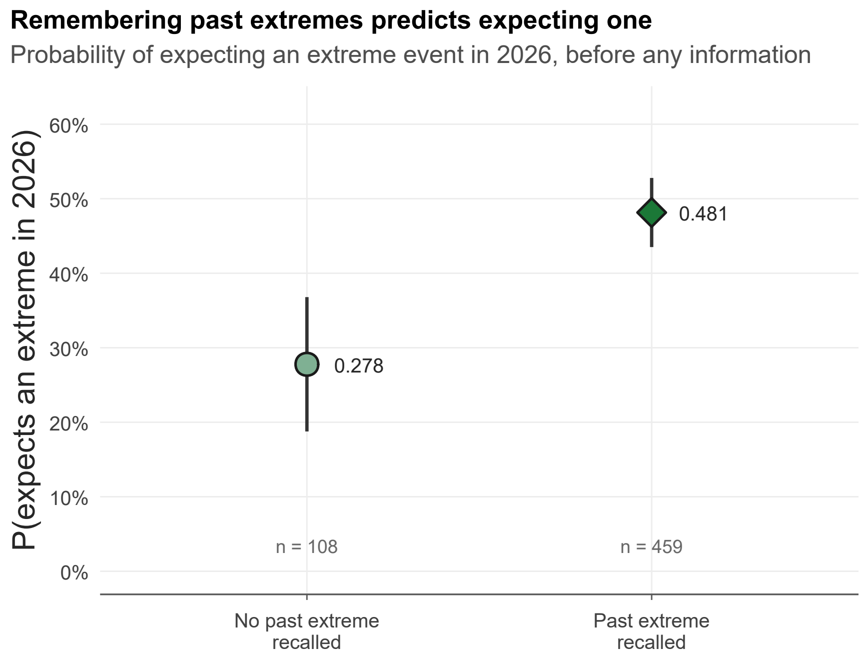

Figure 8: Remembering past extremes predicts expectations of future extremes.
Notes: Bars show the probability of expecting a major weather event in 2026 among farmers who did and did not freely recall one from the past, with 95% village-clustered confidence intervals. With subcounty fixed effects, the difference is 19.4 percentage points (SE 0.045). Source: authors' computation.

season level and can therefore miss dry spells that were highly localized or concentrated in a crop-sensitive part of the season. We therefore do not conclude that farmers did not experience damaging drought conditions. The narrower result is that years recalled as droughts are not systematically drier on any of the measures available at the resolution of our historical record.

One interpretation of the flood-drought asymmetry links it back to the compression result. A flood can be a discrete and visible event, while identifying a drought may require comparing current conditions with a remembered normal and with the variation around that normal. If the distribution of past weather is compressed in memory, that comparison becomes more difficult. Consistent with this interpretation, farmers who recall a drought also show a stronger dry lean in their season by season memories, −0.555 points compared with −0.248 among farmers recalling another type of hazard. The data do not identify this mechanism, but the two elicitation exercises point to a common difficulty in representing the distribution of past weather.

5.4 Expectations and Planned Adaptation

The results so far show that remembered experience is closely related to expectations, but that different forms of memory preserve different information. Season by season memories are strongly compressed, leaving little variation in perceived weather histories, while freely recalled extreme events remain salient and strongly predict expectations of future extremes. We now ask whether these beliefs are reflected in farmers economic decisions. If the expectations we elicit are meaningful for behavior, farmers who perceive greater weather risk should also plan more adaptation and protection. We therefore relate pre-information expectations to planned adaptation choices and spending.

For each outcome, we estimate the relationship between expectations and planned adaptation, controlling for cultivated area, age, household size, and subcounty fixed effects, with standard errors clustered at the village level. Farmers who expect an extreme event in 2026 report substantially greater planned protection. Among those planning any adaptation, expected spending is 0.659 log points higher, equivalent to about 93 percent more spending. These farmers are also 12.5 percentage points more likely to use fertilizer and 4.8 percentage points more likely to plan crop insurance. The insurance result is particularly large relative to its baseline. Among the 404 respondents planning any adaptation, only 2.0 percent report plans to purchase crop insurance.

Memory measures point in a similar direction. Farmers who freely recall a past extreme report greater protective investment. More accurate season by season memory is also associated with greater spending and greater use of protective inputs, while more diffuse memories are associated with higher adaptation spending. Together, these patterns suggest that protection decisions are related not only to expected weather risk, but also to how farmers remember and represent past weather.

These relationships are descriptive and should be interpreted as such. Beliefs, memories, and adaptation plans may all reflect other household characteristics, so we emphasize the overall pattern rather than any single coefficient. Spending is observed only among farmers already planning some form of adaptation, and crop insurance is a stated intention rather than an observed purchase. Beliefs are also only one part of the adaptation decision. Informal risk sharing, contract design, payment timing, liquidity, and present bias can all shape demand (Mobarak & Rosenzweig, 2013; Casaburi & Willis, 2018; Duflo et al., 2011). The result here is that the weather beliefs we elicit are economically meaningful. They are systematically related to the resources farmers plan to commit to protection.

Taken together, the results trace a clear link from remembered experience to economically relevant expectations. Season by season memories strongly predict beliefs about the coming season, even though most of the variation in past weather is lost through compression. Salient extreme events provide a separate channel, strongly predicting expectations of future extremes even though their calibration differs across hazards. These beliefs are in turn closely related to planned adaptation and protection. The next section asks whether providing farmers with a clearer representation of their local weather history changes the expectations they bring to these decisions.

6 Updating Expectations with Weather Information

The previous section showed that farmers' expectations are closely related to remembered experience and to planned adaptation. We now ask whether those beliefs change when farmers are shown a simple summary of local weather history. Relative to farmers who repeat the elicitation without receiving information, treated farmers move their expectations toward the historical benchmark and shift their expected planting onset toward observed rainfall timing. These responses are largest among farmers whose initial beliefs are furthest from the record. Literacy and education do not predict responsiveness to the information. Literacy instead predicts who begins with beliefs further from the record.

6.1 Information Moves Expectations Toward the Historical Record

To isolate the effect of the information, we compare belief revision in Arm B with revision in Arm A. Farmers in both arms state their expectations twice, but only Arm B receives the historical weather information between the two elicitations. This comparison matters because answers can move when the same elicitation is repeated even without new information. Arm A captures that movement, allowing us to distinguish it from revision caused by the information. We estimate

$$\Delta y_i = \alpha + \tau B_i + \mathbf{x}_i'\boldsymbol{\theta} + \eta_i, \tag{4}$$

where Δy_i is the change in respondent i's belief between the two elicitations and B_i indicates assignment to Arm B. The coefficient τ therefore measures how much more beliefs change when farmers receive the historical information than when they simply repeat the elicitation. Standard errors are clustered at the village level over 50 clusters. Arm C does not enter this comparison because farmers in that arm receive the information before their first expectation elicitation. One feature of the design qualifies the interpretation of τ. Arm B also answered three short questions about their reaction to the information before the second elicitation, while Arm A did not. The estimated effect therefore captures the information as delivered, including any effect of answering these questions, and the design cannot separate the two components (Appendix 3.5).

We first ask whether the information shifts farmers' expectations in the direction of the recent weather history they are shown. Relative to Arm A, treated farmers move their beliefs closer to the information on the card, reducing the distance between the two distributions by 0.134 (standard error 0.050). In Arm A, beliefs instead move slightly further away. A second measure gives the same result. When expectations are compared with the realized 2025 season shown on the card, belief accuracy improves by 0.151 points (standard error 0.051). Because 2025 is part of the information treatment, we view these as two measures of the same response rather than separate effects. The two measures are strongly correlated at 0.82. The main takeaway is that the information does not simply cause farmers to revise their answers. It moves their expectations systematically toward the weather history they are shown.

The same pattern appears for a different dimension of expectations. The treatment shifts expected planting onset 2.6 days earlier relative to Arm A (standard error 0.86), which is toward the observed rainfall timing throughout the study area. The information therefore changes both what farmers expect the coming season to look like and when they expect the planting rains to begin. Across both dimensions, beliefs move toward the historical evidence provided by the treatment. Appendix B.3 reports the onset benchmark in detail, while Appendix C.3 reports wild cluster and multiple-testing inference.

The average effects are modest in size, as is common for low-cost agricultural information interventions (Fabregas et al., 2019). More informative for our mechanism is how that average response is generated.

The information changes the magnitude of revision, not its prevalence. The treatment does not substantially increase the share of farmers who revise toward the record, which rises only from 27.7 to 31.4 percent. Instead, it changes the direction and size of revision. Farmers who update move further toward the historical benchmark, while large movements away from it become less common. The information therefore improves the quality of updating among a relatively stable group of farmers willing to revise. This matters because its value lies not only in encouraging correction, but also in preventing beliefs from drifting further from the historical benchmark. Appendix C.2 reports the full distribution of belief changes.

6.2 Who Updates and Who Has Most to Gain?

The response depends strongly on the initial beliefs of each farmer. Farmers who already expect an early planting onset, close to the historical timing, barely change their beliefs. Those starting with later expectations move about four days earlier after receiving the information. The same pattern holds across the full range of initial beliefs. The further a farmer starts from the record, the more that farmer moves toward it.

This is not simply the movement we see when the elicitation is repeated, since the comparison is made relative to Arm A. The information therefore acts selectively. It leaves beliefs largely unchanged when they are already close to the historical benchmark and produces larger corrections when they are further away

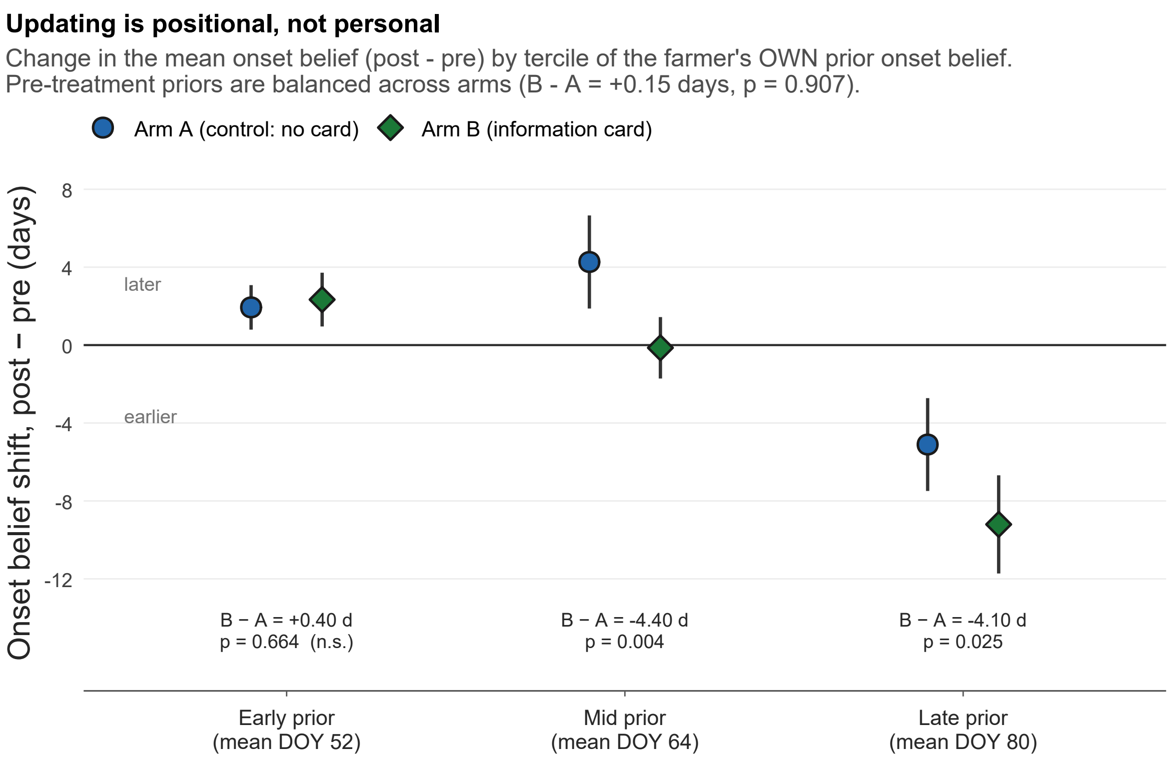

Figure 9: Who revises depends on where the initial belief sits.

Notes: Bars show the mean change in expected planting onset between the first and second elicitations by tercile of the farmer's initial belief, separately for Arms A and B. Negative values indicate revision toward an earlier onset. The treatment effect is +0.40 days for farmers with early initial beliefs, −4.41 days for those in the middle tercile, and −4.10 days for those in the late tercile. Bars show 95% village-clustered confidence intervals. Source: authors' computation.

Observable farmer characteristics, by contrast, do not meaningfully predict the response.

Across 22 pre-treatment characteristics, including education, literacy, age, gender, farm size, information access, and risk attitudes, and four measures of updating, no characteristic survives adjustment for multiple testing. The design also rules out relatively large differences in responsiveness, with differential effects larger than about one quarter of a standard deviation excluded. What predicts updating is therefore primarily where the farmer begins relative to the information, rather than observable characteristics of the farmer.

This result differs from settings in which information interventions generate larger responses among more educated recipients (Cole et al., 2017; R. Jensen, 2010). One possible explanation is the format of the information. The treatment presents a simple visual summary using the same images farmers already use to express their beliefs, so interpreting the message does not require reading or familiarity with formal probabilities. More generally, previous work emphasizes that learning depends on how well information is matched to the recipient and the decision problem (Hanna et al., 2014; BenYishay & Mobarak, 2019; Cole & Fernando, 2021). Our results are consistent with that view. Once the information is made directly accessible, literacy and schooling do not appear to constrain short-run belief revision.

Initial Belief Accuracy and Literacy. While literacy does not predict how strongly farmers update, it does predict the accuracy of the beliefs they begin with. Before receiving the information, literate farmers' beliefs are about 0.28 standard deviations closer to the historical record. This pattern is robust across alternative measures of belief distance, while years of schooling add little once literacy is held fixed.

This result connects directly to the memory patterns documented earlier in the paper. Farmers who cannot read recall past seasons less accurately, and their season ahead beliefs also lie further from the historical record. Yet they respond just as strongly when shown the visual information. The group that begins with less accurate memories and beliefs is therefore no less able to correct them when the information is made accessible. A visual format can reach precisely those farmers for whom the initial information gap is greatest.

6.3 Updating and Farm Outcomes

The 2026 season provided little scope to test whether improved beliefs translate into better farm outcomes. Although rainfall was below its long-run mean and included a 25-day dry spell, no study location was classified as dry, vegetation showed no detectable damage, and farm productivity remained close to its historical average. Consistent with this relatively benign realization, we find no detectable treatment effects on satellite-measured farm outcomes. Appendix D.8 reports these results and their statistical power.

Our experiment therefore establishes that weather beliefs can be moved toward the historical record, but not whether these corrections improve outcomes when weather risk materializes. There is good reason to expect this final link. Subjective weather beliefs predict consequential agricultural choices such as planting decisions (Giné et al., 2015), while agricultural information interventions have been shown to improve farm outcomes (Fabregas et al., 2019). Our results identify an intermediate step in this chain by showing that simple historical information can improve the beliefs farmers bring to these decisions. The 2026 season does not allow us to test whether those improved beliefs translate into better choices and greater resilience when adverse weather actually occurs. Following similar interventions through droughts, floods, or other extreme seasons would provide an important test of this final link.

Taken together, the experiment shows that distorted expectations are not fixed. Simple visual information moves beliefs toward the historical record, with the largest responses among farmers who begin furthest away, while literacy predicts initial belief accuracy rather than the ability to update. The final section considers what these findings imply for climate adaptation and the design of information and protection policies.

7 Conclusion

Smallholder farmers face substantial weather risk but invest relatively little in insurance and other forms of protection. Much of the literature has explained this through features of the products available to farmers, including prices, contract design, payment timing, basis risk, and trust (Giné et al., 2008; Cole et al., 2013; Karlan et al., 2014; Casaburi & Willis, 2018; Carter et al., 2017). This paper studies an earlier part of the decision. Before farmers choose whether protection is worth buying, they must form beliefs about the weather they expect to face. We show that these beliefs are shaped by memories that preserve remarkably little of the variation in past weather.

Our first contribution is to measure this memory directly. Section 4 shows that farmers broadly remember which seasons were wetter and which were drier, but most of the differences between seasons disappear in recall. Two seasons that are five points apart on our nine-picture weather scale are remembered as only about half a point apart. This extends work that infers forgetting from insurance behavior or aggregate perceptions (Gallagher, 2014; Moore et al., 2019) and complements models in which selective retrieval causes remembered experience to differ systematically from actual experience (Bordalo et al., 2020, 2023). Our contribution is to observe the memory itself and benchmark it season by season against the historical record. The complementary extreme-event exercise in Section 5 also shows that memory does not treat all experiences alike. Salient hazards remain strongly related to expectations of future extremes, although their correspondence with the historical record differs across floods and droughts. Memory therefore preserves some features of past weather while losing much of the variation needed to judge how extreme a future season is likely to be.

Our second contribution is to trace remembered experience into expectations. A large literature has established that subjective expectations can be elicited meaningfully and that weather beliefs predict agricultural decisions (Manski, 2004; Delavande et al., 2011; Giné et al., 2015; Lybbert et al., 2007). Section 5 traces these beliefs back to one of their underlying inputs and asks how remembered experience shapes them. Farmers who remember wetter histories also expect wetter conditions in the coming season. But because different weather histories become similar in memory, much of their informational content disappears before expectations are formed. Applying the estimated memory distortion to the historical record removes roughly nine tenths of the variation across recent weather histories. The problem is therefore not simply that expectations are biased in one direction. Farmers have a much less informative representation of the weather variation they have experienced. This can make limited demand for protection consistent with farmers own perceived risk, even when it appears puzzling relative to objective weather patterns. These beliefs are economically meaningful as well. Farmers who expect adverse weather plan substantially more adaptation spending and are more likely to plan crop insurance, although these relationships remain descriptive rather than causal.

Our third contribution is to show that these beliefs can be updated. Section 6 shows that a simple visual summary of recent local weather history moves expectations toward the historical record and shifts expected planting onset toward observed rainfall timing. The largest corrections occur among farmers whose initial beliefs are furthest from the record, consistent with a broad information-experiment literature in which revision depends on the distance between prior beliefs and new information (Haaland et al., 2023; Cavallo et al., 2017; Nguyen, 2008). Just as importantly, literacy, education, and other observable characteristics do not predict responsiveness. This differs from settings in which information reaches more educated recipients first (Cole et al., 2017; R. Jensen, 2010). One reason may be the format. The same visual language allows us both to elicit probabilistic beliefs without requiring literacy and to communicate historical information without requiring farmers to read or interpret formal probabilities. This is consistent with evidence that learning depends importantly on how information is matched to its recipient (Hanna et al., 2014; BenYishay & Mobarak, 2019; Cole & Fernando, 2021).

Historical weather information is increasingly available from satellite and other remote sensing sources, but having the data is different from presenting it in a form farmers can readily access and interpret. Our treatment reorganizes recent local weather history into a simple visual format that farmers can process regardless of literacy. It can be generated from existing weather data and requires little more than a printed summary. The intervention does not make every farmer revise, but among those who do, it produces larger movements toward the historical benchmark and limits large movements away from it. It also produces the largest corrections where initial beliefs are furthest from the record. Accessible historical information may therefore offer a low-cost complement to forecasts, extension services, insurance, and other adaptation policies, particularly in settings where farmers otherwise have few records against which to assess their own experience. This complements evidence that well-designed agricultural information can improve decisions and outcomes even when sophisticated forecasting remains difficult (Fabregas et al., 2019; Rosenzweig & Udry, 2019).

The literacy gradient points to a complementary longer run margin. Literacy does not predict the ability to respond to our information treatment, but Sections 4 and 6 show that it is strongly associated with the quality of the memories and beliefs farmers begin with. Farmers who cannot read recall past weather less accurately and enter the season with beliefs further from the historical record, yet respond no less strongly once the information is made accessible. These relationships suggest a broader an indirect potential return to human capital. Literacy and education may matter for agricultural adaptation not only because they provide access to written information or formal technologies, but also because they may help people organize past experience and form more informative beliefs from it. Short run information provision and longer run human capital investment may therefore operate on different parts of the same process. The former can improve the information available for a current decision, while the latter may strengthen how experience is remembered and translated into beliefs over time.

One important link remains open. Section 6.3 shows that the relatively benign 2026 season provided little scope to test whether corrected expectations ultimately improve farm outcomes when weather risk materializes. Existing evidence gives good reason to take this link seriously. Subjective weather beliefs predict consequential agricultural choices such as planting decisions (Giné et al., 2015), while agricultural information interventions can improve farm outcomes (Fabregas et al., 2019). Our experiment identifies an intermediate step in this chain by showing that simple historical information can change the beliefs farmers bring to those decisions. Following similar interventions through droughts, floods, and other extreme seasons would provide an important test of whether better calibrated expectations translate into better adaptation and greater resilience.

The demand side emphasized here complements rather than replaces the supply side of climate adaptation. Better contracts, pricing, payment timing, index design, and protective technologies remain important (Casaburi & Willis, 2018; D. J. Clarke, 2016; N. D. Jensen et al., 2016; Carter et al., 2017). But farmers will only value these forms of protection if they perceive the underlying risks as sufficiently likely. Our results identify a belief formation process that precedes the choice of any particular product, show how distorted memory can make those beliefs less informative and weaken the perceived case for protection in unusually risky seasons, and demonstrate that simple information can improve the beliefs underlying these decisions. Climate adaptation therefore depends not only on the protection farmers can access, but also on whether they can form sufficiently accurate and informative beliefs about the risks they are preparing for.

Declarations

Competing interests. The authors declare no competing financial or non-financial interests.

Funding. This study was funded by Erasmus University Rotterdam and the Erasmus Research Institute of Management (ERIM) through departmental research support, as disclosed in the consent form reproduced in Online Appendix OA-5. The funded proposal anticipated satellite-based rather than survey-based outcomes. This difference is documented in Appendix E.

Ethics. The research was reviewed and approved by the Internal Review Board of the Erasmus School of Economics (IRB-E), under approval number ETH2425-0832, granted on 20 May 2025 and valid for three years. The approval documentation is reproduced in Online Appendix OA-5.

Pre-registration. The analysis plan was pre-registered with the AEA RCT Registry in March 2026, before data collection began. Section 6 and Appendix E document deviations from the pre-registered analysis plan.

Data availability. Replication data and code will be made publicly available upon publication. CHIRPS rainfall data are publicly available (Funk et al., 2015).

Generative AI. During the preparation of this work, the authors used Claude (Anthropic) to assist with drafting and editing the manuscript and with preparing and checking analysis code. All AI-assisted output was reviewed and edited by the authors, who take full responsibility for the final content.

References

Attanasio, O. P. (2009). Expectations and perceptions in developing countries: Their measurement and their use. *American Economic Review: Papers and Proceedings*, *99*(2), 87–92.

Beegle, K., Carletto, C., & Himelein, K. (2012). Reliability of recall in agricultural data. *Journal of Development Economics*, *98*(1), 34–41.

Belloni, A., Chernozhukov, V., & Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, *81*(2), 608–650.

BenYishay, A., & Mobarak, A. M. (2019). Social learning and incentives for experimentation and communication. *Review of Economic Studies*, *86*(3), 976–1009.

Bordalo, P., Conlon, J. J., Gennaioli, N., Kwon, S. Y., & Shleifer, A. (2023). Memory and probability. *Quarterly Journal of Economics*, *138*(1), 265–311.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2020). Memory, attention, and choice. *Quarterly Journal of Economics*, *135*(3), 1399–1442.

Bruhn, M., & McKenzie, D. (2009). In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics*, *1*(4), 200–232.

Burke, M., & Lobell, D. B. (2017). Satellite-based assessment of yield variation and its determinants in smallholder african systems. *Proceedings of the National Academy of Sciences*, *114*(9), 2189–2194.

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

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, *90*(3), 414–427.

Carter, M., de Janvry, A., Sadoulet, E., & Sarris, A. (2017). Index insurance for developing country agriculture: A reassessment. *Annual Review of Resource Economics*, *9*, 421–438.

Casaburi, L., & Willis, J. (2018). Time versus state in insurance: Experimental evidence from contract farming in kenya. *American Economic Review*, *108*(12), 3778–3813.

Cavallo, A., Cruces, G., & Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, *9*(3), 1–35.

Clarke, D., Romano, J. P., & Wolf, M. (2020). The Romano–Wolf multiple-hypothesis correction in Stata. *Stata Journal*, *20*(4), 812–843.

Clarke, D. J. (2016). A theory of rational demand for index insurance. *American Economic Journal: Microeconomics*, *8*(1), 283–306.

Cole, S., & Fernando, A. N. (2021). 'Mobile'izing agricultural advice technology adoption diffusion and sustainability. *Economic Journal*, *131*(633), 192–219.

Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., & Vickery, J. (2013). Barriers to household risk management: Evidence from india. *American Economic Journal: Applied Economics*, *5*(1), 104–135.

Cole, S., Giné, X., & Vickery, J. (2017). How does risk management influence production decisions? evidence from a field experiment. *Review of Financial Studies*, *30*(6), 1935–1970.

Cole, S., Stein, D., & Tobacman, J. (2014). Dynamics of demand for index insurance: Evidence from a long-run field experiment. *American Economic Review: Papers and Proceedings*, *104*(5), 284–290.

Conley, T. G., & Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *American Economic Review*, *100*(1), 35–69.

Delavande, A. (2014). Probabilistic expectations in developing countries. *Annual Review of Economics*, *6*, 1–20.

Delavande, A., Giné, X., & McKenzie, D. (2011). Measuring subjective expectations in developing countries: A critical review and new evidence. *Journal of Development Economics*, *94*(2), 151–163.

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American Economic Review*, *101*(6), 2350–2390.

Emerick, K., de Janvry, A., Sadoulet, E., & Dar, M. H. (2016). Technological innovations, downside risk, and the modernization of agriculture. *American Economic Review*, *106*(6), 1537–1561.

Enke, B., & Graeber, T. (2023). Cognitive uncertainty. *Quarterly Journal of Economics*, *138*(4), 2021–2067.

Fabregas, R., Kremer, M., & Schilbach, F. (2019). Realizing the potential of digital development: The case of agricultural advice. *Science*, *366*(6471), eaay3038.

FEWS NET. (2010). *Livelihood mapping and zoning exercise: Uganda* (Tech. Rep.). Washington, DC: Famine Early Warning Systems Network, United States Agency for International Development.

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., ... Michaelsen, J. (2015). The climate hazards infrared precipitation with stations – a new environmental record for monitoring extremes. *Scientific Data*, *2*, 150066.

Gallagher, J. (2014). Learning about an infrequent event: Evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics*, *6*(3), 206–233.

Giné, X., Townsend, R., & Vickery, J. (2008). Patterns of rainfall insurance participation in rural india. *World Bank Economic Review*, *22*(3), 539–566.

Giné, X., Townsend, R., & Vickery, J. (2015). *Forecasting when it matters: Evidence from semi-arid india.* (Unpublished working paper)

Godlonton, S., Hernandez, M. A., & Murphy, M. (2018). Anchoring bias in recall data: Evidence from central america. *American Journal of Agricultural Economics*, *100*(2), 479–501.

Haaland, I., Roth, C., & Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, *61*(1), 3–40.

Hanna, R., Mullainathan, S., & Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics*, *129*(3), 1311–1353.

Imai, K., King, G., & Nall, C. (2009). The essential role of pair matching in cluster-randomized experiments. *Statistical Science*, *24*(1), 29–53.

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.* Cambridge: Cambridge University Press.

Jensen, N. D., Barrett, C. B., & Mude, A. G. (2016). Index insurance quality and basis risk: Evidence from northern kenya. *American Journal of Agricultural Economics*, *98*(5), 1450–1469.

Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *Quarterly Journal of Economics*, *125*(2), 515–548.

Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics*, *129*(2), 597–652.

Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Annals of Applied Statistics*, *7*(1), 295–318.

Lowder, S. K., Skoet, J., & Raney, T. (2016). The number, size, and distribution of farms, smallholder farms, and family farms worldwide. *World Development*, *87*, 16–29.

Lybbert, T. J., Barrett, C. B., McPeak, J. G., & Luseno, W. K. (2007). Bayesian herders: Updating of rainfall beliefs in response to external forecasts. *World Development*, *35*(3), 480–497.

MacKinnon, J. G., & Webb, M. D. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, *32*(2), 233–254.

Malmendier, U., & Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, *126*(1), 373–416.

Manski, C. F. (2004). Measuring expectations. *Econometrica*, *72*(5), 1329–1376.

Mayuge District Local Government. (2015). *District development plan II, 2015/2016–2019/2020* (Tech. Rep.). Mayuge, Uganda: Mayuge District Local Government.

Mobarak, A. M., & Rosenzweig, M. (2013). Informal risk sharing, index insurance, and risk taking in developing countries. *American Economic Review: Papers and Proceedings*, *103*(3), 375–380.

Moore, F. C., Obradovich, N., Lehner, F., & Baylis, P. (2019). Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change. *Proceedings of the National Academy of Sciences*, *116*(11), 4905–4910.

Nguyen, T. (2008). *Information, role models and perceived returns to education: Experimental evidence from madagascar.* (Unpublished working paper, MIT)

Ricciardi, V., Ramankutty, N., Mehrabi, Z., Jarvis, L., & Chookolingo, B. (2018). How much of the world's food do smallholders produce? *Global Food Security*, *17*, 64–72.

Romano, J. P., & Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica*, *73*(4), 1237–1282.

Rosenzweig, M. R., & Udry, C. R. (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). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w25894

Samberg, L. H., Gerber, J. S., Ramankutty, N., Herrero, M., & West, P. C. (2016). Subnational distribution of average farm size and smallholder contributions to global food production. *Environmental Research Letters*, *11*(12), 124010.

Stein, D. (2018). Dynamics of demand for rainfall index insurance: Evidence from a commercial product in india. *World Bank Economic Review*, *32*(3), 692–708.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207–232.

Uganda Bureau of Statistics. (2017). *National population and housing census 2014 — area specific profiles: Bugiri, butaleja and mayuge districts* (Tech. Rep.). Kampala, Uganda: Uganda Bureau of Statistics. (Separate district profile volumes)

Uganda Bureau of Statistics. (2022). *The annual agricultural survey 2019 report* (Tech. Rep.). Kampala, Uganda: Uganda Bureau of Statistics.

A Sample, randomization, and balance

This appendix documents how the study sample was constructed and how villages were assigned to the three experimental arms. It describes the sampling frame and village selection, access to participating communities, respondent recruitment within villages, the matched-pair randomization and rerandomization procedure, balance across arms, and the fieldwork ultimately completed. Appendix B describes the construction of the objective weather benchmark and information treatment, while the survey and field materials are reproduced in the Online Appendix.

A.1 The sampling frame and village selection

Village selection proceeded in three stages. We first screened candidate districts, then constructed a rural village sampling frame within the selected districts, and finally drew a set of candidate villages using a stratified procedure.

Stage one: district selection. Ten districts were considered before the study area was fixed: Bugiri, Bulambuli, Butaleja, Gomba, Kamuli, Kiboga, Kyankwanzi, Mayuge, Nakasongola, and Namutumba. The screening exercise compared rainfall characteristics across 94 subcounties using CHIRPS pentad data, including seasonal rainfall moments, the frequency of dry pentads, and rainfall-onset timing. Two overlapping data vintages covering 2018–2022 and 2021–2025 were examined. The surviving project documentation does not record the exact criterion by which Bugiri, Butaleja, and Mayuge were selected from the ten candidate districts. We therefore report the screen as documented and do not reconstruct an ex post selection rule.

Stage two: the village sampling frame. The initial village listing contains 594 rows covering 483 distinct village names across 85 parishes and 14 subcounties in Bugiri, Butaleja, and Mayuge. Of these, 511 carry a field-partner assessment, with 327 classified as "Good" and 184 as "Very good." The remaining 83 have no assessment and are excluded. We then remove settlements identified as town councils, town boards, townships, markets, or towns. This removes three subcounties and leaves an effective rural sampling frame of 442 villages across 11 subcounties and 70 parishes. The population from which the study villages were drawn is therefore rural communities in the 11 non-town subcounties of Bugiri, Butaleja, and Mayuge that the local field partner considered suitable for the study.

Stage three: the candidate-village draw. Villages rated "Very good" received higher sampling priority than those rated "Good." Within each parish, up to five villages were initially sampled, with the higher-priority villages receiving three times the sampling weight. Where village names differed only by a trailing "A" or "B," at most one was retained for a given parish. The procedure then capped each subcounty at seven villages and sampled up to 80 villages from the remaining pool using the same three-to-one priority weighting. With a fixed random seed of 321, this procedure selected 77 candidate villages.

From 77 candidates to 60 study villages. The 77 villages were a preselection rather than the final experimental sample. Before any treatment was assigned, the field team visited the candidate communities and sought permission to conduct the survey, meeting local council leadership and community mobilizers who would later help convene respondents. Access was secured in 60 villages. These 60 villages constitute the experimental sample.

This access stage restricts the population to communities whose local leadership agreed to host the study. It therefore matters for external validity. Participating villages may differ from villages that did not participate in local institutional capacity, prior contact with researchers, or other unobserved characteristics. We cannot determine the direction or magnitude of this selection because no survey data were collected in the 17 candidate villages that did not enter the final sample.

The access process does not affect the internal treatment comparison because it was completed before randomization. The 60-village sample was fixed first, and treatment assignment was conducted only afterward, on 5 March 2026, five days before the first interview on 10 March. No village entered or left the sample after treatment assignment, and treatment status did not exist when village access was negotiated. Selection into the study therefore determines the population to which the experiment applies but not the validity of treatment comparisons within that population.

A limitation of the field records is that the identities of the 17 candidate villages that did not enter the final sample, and the reasons for their exclusion, were not recorded systematically. We therefore cannot compare participating and nonparticipating candidate villages on the weather characteristics used in the original sampling frame. The archived sampling documentation also contains one description error: it refers to a final-sample category labeled "Yes," although no such assessment category was used. The final sample is instead defined by the 60 villages in which access was secured.

The realized sample. Table 2 summarizes the geographic structure of the final sample. Twenty-eight of the 42 participating parishes contribute one village, ten contribute two villages, and four contribute three.

Geographic coordinates. Exact village coordinates were not collected. Villages were instead assigned the centroid of their parish polygon using the Uganda Bureau of Statistics 2016 parish boundaries. Coordinates were matched by district and parish, with subcounty used to resolve repeated parish names. Two parish-name spellings were standardized to match the official geographic boundaries.

This means that the 60 study villages are represented by only 42 distinct coordinate pairs because villages in the same parish share a centroid. Geographic variables used in the randomization are consequently measured at parish or subcounty level rather than at the

Table 2: The realized village sample, by district.

District	Subcounties	Parishes	Villages	Respondents
Bugiri	4	16	19	215
Butaleja	3	11	18	197
Mayuge	4	15	23	274
Total	11	42	60	686

Notes: The final sample contains 60 villages across 42 parishes and 11 subcounties in Bugiri, Butaleja, and Mayuge. Respondent locations are assigned using the standardized village-to-district mapping used in the survey rather than the free-text district field, which contains multiple spelling variants.

village itself. Maps in this paper therefore show parish-centroid locations and should not be interpreted as exact village coordinates.

A.2 Respondent selection within villages

The surviving field documentation does not record a formal within-village sampling rule. In particular, it does not establish whether respondents were drawn from a household listing, selected using a spatial rule or quota, or recruited from lists maintained by local leaders. We therefore do not characterize within-village recruitment as random and do not make claims about refusals or replacement that are not documented.

The realized sample contains 686 respondents across all 60 villages, with between 7 and 14 respondents per village, a median of 12, and a mean of 11.4. Only two villages contain fewer than ten respondents, with 7 in Buhwaya and 9 in Mawanga, while three villages contain 14. This distribution is consistent with an approximate field target of twelve respondents per village, but such a target is not documented in the surviving protocol and should therefore be interpreted only as an inference from the realized sample.

Coverage is similar across experimental arms. Arm A contains 274 respondents in 24 villages, with 7 to 14 respondents per village. Arm B contains 293 respondents in 26 villages, with 9 to 14 per village. Arm C contains 119 respondents in 10 villages, with 11 to 14 per village. Mean village sample sizes are 11.4, 11.3, and 11.9 in Arms A, B, and C, respectively.

Respondent identifiers were designed to incorporate enumerator and village codes, providing an additional check that interviews were associated with the intended field assignment. All 686 respondent identifiers are unique. Of these, 683 follow the intended coding structure and identify one of the seven enumerators. Three contain malformed codes. These appear to be data-entry errors rather than evidence of additional enumerators and do not affect treatment assignment or respondent uniqueness.

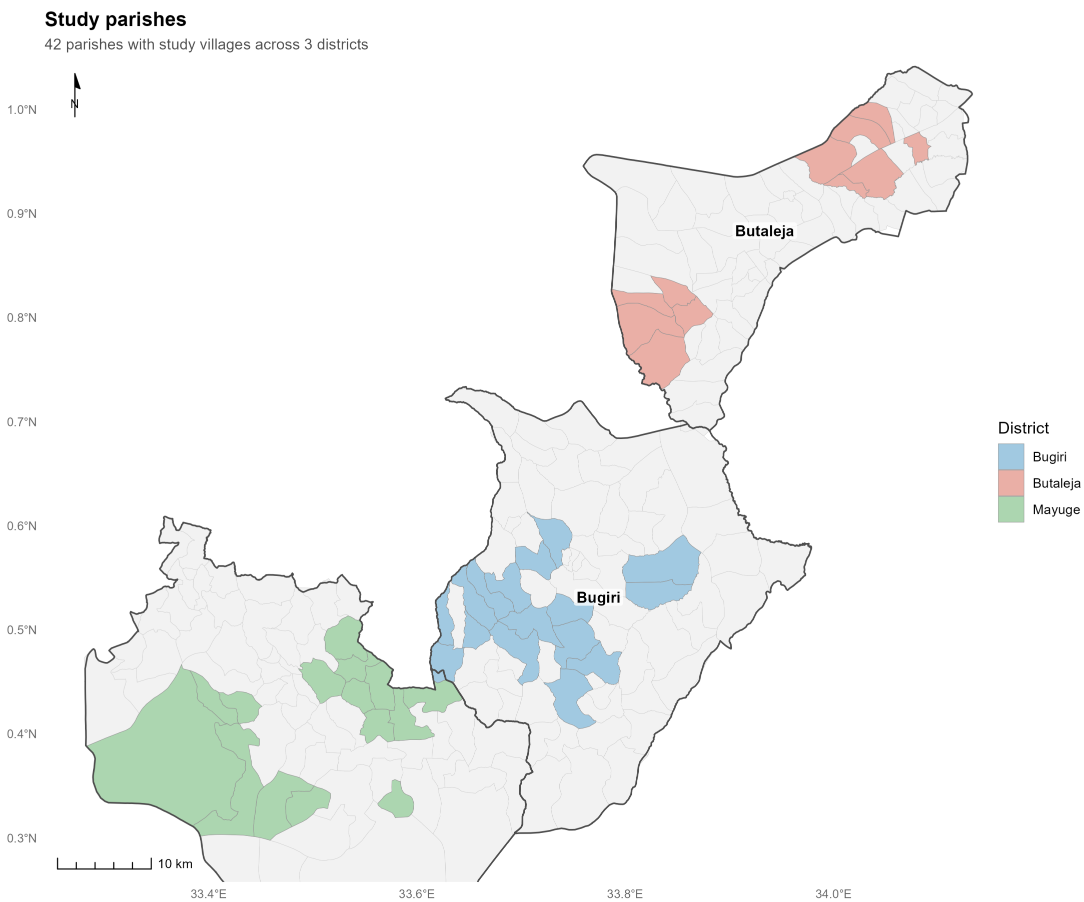

Figure 10: The study area: 42 parishes containing study villages across the three districts.
Notes: Parishes containing at least one of the 60 study villages are shown by district. Unshaded polygons are other parishes within Bugiri, Butaleja, and Mayuge. Villages are represented geographically by parish centroids, so parish is the finest mapped unit. Figure 11 shows the corresponding locations separately by district. Source: authors' construction using Uganda Bureau of Statistics 2016 parish boundaries.

A.3 Randomization

Why randomization occurs at the village level. Treatment is assigned at village rather than household level because the intervention consists of information that can readily diffuse between nearby farmers. Respondents receive a printed summary of local weather history, while social interactions are themselves an important channel through which agricultural information can spread (Conley & Udry, 2010; BenYishay & Mobarak, 2019). Household-level treatment assignment within the same village would therefore create a substantial risk of information spillovers into the control group. The corresponding cost is statistical power. The experimental treatment varies across 60 villages rather than 686 individuals, and inference is clustered accordingly.

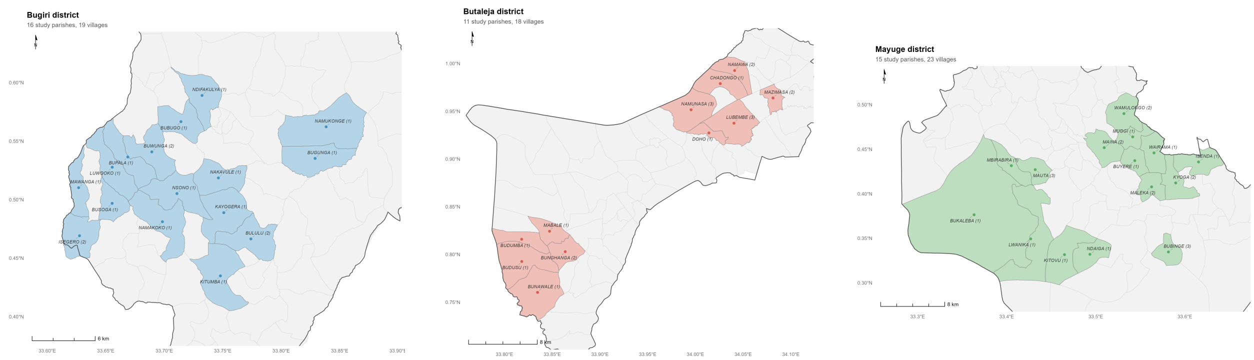

(a) Bugiri (b) Butaleja (c) Mayuge

Figure 11: Study villages by district.

Notes: The study contains 19 villages in Bugiri, 18 in Butaleja, and 23 in Mayuge. Locations are plotted at parish centroids, so villages within the same parish share a coordinate and may overlap on the map. Source: authors' construction.

Randomization covariates and strata. Five geographic covariates were used in the randomization. The first is mean elevation within a 500-meter radius of the parish centroid, measured using SRTM. The remaining four are subcounty-level characteristics calculated over the 2001–2025 historical weather record: mean seasonal rainfall, the coefficient of variation of seasonal rainfall, mean peak EVI, and the coefficient of variation of peak EVI. Only elevation varies within subcounty.

Villages were stratified by subcounty and elevation tercile. Elevation terciles are divided at 1,118 and 1,171 meters, placing exactly 20 villages in each tercile. Crossing these groups with the 11 subcounties produces 17 nonempty strata containing between one and six villages.

Pair formation. Within each stratum, villages were paired by nearest-neighbor matching on standardized elevation and mean peak EVI. At each step, the two closest remaining villages were paired until fewer than two villages remained. This produced 25 within-stratum pairs. Ten villages remained unmatched because they belonged to strata with odd numbers of villages. These villages were pooled and matched using the same covariates, with a strong preference for matches within the same subcounty. This produced five additional pairs. Three remain within subcounty, while two necessarily cross subcounty boundaries: Kityerera with Mpungwe and Buwunga with Kigandalo. All 60 villages are therefore contained in 30 matched pairs. Pair-matched cluster randomization follows Imai et al. (2009).

Matching was close for most pairs. Fifteen of the 30 pairs have no distance on the standardized matching variables, nine have distances at or below 0.32 standard-deviation units, and four have distances between 0.62 and 1.65. The remaining two are the forced cross-subcounty pairs and were deliberately penalized in the matching procedure to ensure that within-subcounty matches were selected whenever possible.

Assignment of the three experimental arms. Arm C is smaller than Arms A and B, so treatment assignment uses ten of the 30 matched pairs as donor pairs for Arm C. Within each donor pair, one village is randomly assigned to Arm C and the second is randomly assigned to either Arm A or Arm B. In each of the remaining 20 pairs, one village is assigned to Arm A and the other to Arm B. Arm A and Arm B therefore each receive 20 villages from the regular matched pairs plus the villages assigned to them from the ten Arm C donor pairs. In the accepted assignment, four donor-pair villages were allocated to Arm A and six to Arm B, yielding 24 villages in Arm A, 26 in Arm B, and 10 in Arm C.

Rerandomization. We used rerandomization to improve balance on the five geographic covariates. A candidate assignment was accepted only if three conditions were satisfied simultaneously. The largest absolute standardized difference between Arms A and B across the five covariates had to be no greater than 0.10. The corresponding maximum difference for either the A–C or B–C comparison had to be no greater than 0.25. Finally, the difference in the number of villages assigned to Arms A and B could not exceed four.

The tighter balance requirement for Arms A and B reflects their role as the primary experimental comparison, while Arm C is a smaller robustness arm containing only ten villages. Candidate assignments were generated sequentially from a fixed randomization procedure. The accepted assignment was the 844th candidate draw, corresponding to random seed 966, and was finalized on 5 March 2026. Balance is assessed following Bruhn & McKenzie (2009).

Balance achieved by the design. Table 3 reports balance on the five randomization covariates. The largest standardized difference between Arms A and B is 0.0847, for mean peak EVI. The largest difference involving Arm C is 0.2331, for elevation between Arms B and C. Mahalanobis distances between the five-dimensional arm means are 0.0164 for A–B, 0.3334 for A–C, and 0.3003 for B–C. Treatment assignment is also distributed across all three districts. Bugiri contributes 7, 9, and 3 villages to Arms A, B, and C; Butaleja contributes 8, 8, and 2; and Mayuge contributes 9, 9, and 5.

Assignment in the field. Treatment assignment was fixed before fieldwork and preloaded for each of the 60 study villages. Enumerators selected the village from the study list, after which the corresponding experimental arm was assigned automatically and displayed for confirmation. Enumerators therefore had no discretion over treatment assignment. This procedure differed from the pilot, in which treatment status could be selected manually, and was adopted specifically to eliminate that source of discretion.

Realized assignment matches the randomized assignment exactly. Arm A contains 274 respondents in 24 villages, Arm B contains 293 respondents in 26 villages, and Arm C contains 119 respondents in 10 villages. No respondent is associated with a treatment status different from the assignment of the reported study village.

Figure 12 shows the geographic distribution of the randomized assignment.

Table 3: Design-stage balance across arms on the five randomization covariates.

	Arm A ($n = 24$)		Arm B ($n = 26$)		Arm C ($n = 10$)	
Village covariate	Mean	SD	Mean	SD	Mean	SD
Elevation (m)	1147.167	51.908	1149.538	47.171	1138.600	35.616
Season rainfall (mm)	709.922	9.962	710.715	9.724	710.673	10.022
Rainfall CV	0.137	0.017	0.137	0.016	0.138	0.021
Peak EVI mean	0.593	0.029	0.591	0.027	0.591	0.027
Peak EVI CV	0.076	0.026	0.075	0.024	0.075	0.027

	A − B	A − C	B − C
Elevation (m)	−0.0505	0.1826	0.2331
Season rainfall (mm)	−0.0817	−0.0773	0.0044
Rainfall CV	−0.0148	−0.0899	−0.0751
Peak EVI mean	0.0847	0.0801	−0.0047
Peak EVI CV	0.0342	0.0290	−0.0052
Mahalanobis D^2	0.0164	0.3334	0.3003

Notes: The upper panel reports arm means and standard deviations for the five geographic covariates used in the matched-pair rerandomization. The lower panel reports pairwise mean differences divided by the pooled standard deviation across all 60 villages and the Mahalanobis distance between arm centroids over the five covariates jointly. Elevation is measured within a 500-meter buffer around the parish centroid. The rainfall and vegetation measures are subcounty-level quantities constructed from the 2001–2025 CHIRPS and Sentinel-2 historical record and are therefore constant within subcounty. Respondent-level balance is reported in Appendix A.4.

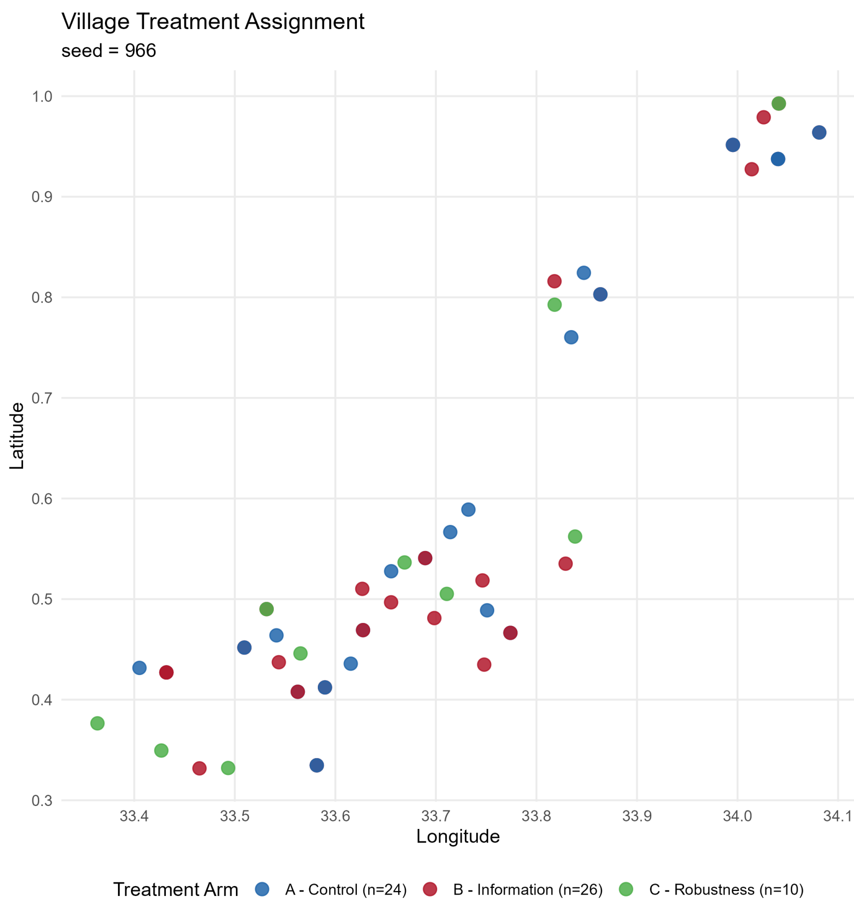

Figure 12: Village treatment assignment.

Notes: The 60 study villages are shown by experimental arm: 24 control villages in Arm A, 26 information villages in Arm B, and 10 villages in Arm C. Treatment assignment is distributed throughout the 11 study subcounties rather than concentrated geographically. Locations are represented by parish centroids, so villages within the same parish overlap and the 60 villages appear at 42 distinct coordinates. The accepted rerandomization draw used seed 966.

A.4 Respondent characteristics and balance

Table 1 in the main text reports respondent characteristics by experimental arm. This section provides additional balance evidence, distinguishing the geographic variables explicitly balanced by the design from respondent characteristics that were not used in the randomization.

Design-stage balance. Table 3 shows that the rerandomization achieved the balance restrictions imposed on the five geographic covariates. The largest standardized A–B difference is 0.0847 and the largest standardized difference involving Arm C is 0.2331, both within the pre-specified acceptance thresholds. Pair-matched cluster randomization follows Imai et al. (2009), balance is assessed following Bruhn & McKenzie (2009), and covariate adjustment where used follows Lin (2013).

Respondent-level balance. Respondent characteristics were not included in the rerandomization because they were measured only after village assignment. We assess ex post balance across 74 pre-treatment respondent characteristics using village-clustered comparisons between experimental arms. Seven of the 74 characteristics differ between Arms A and B at the ten percent level, compared with 7.4 differences expected by chance under independent tests. Eleven differ between Arms A and C, consistent with the smaller size of Arm C and the wider balance tolerance applied to comparisons involving that arm.

Among the respondent characteristics displayed in Table 1, the largest A–B difference is age. Respondents average 42.5 years in Arm A and 40.0 years in Arm B, corresponding to a standardized difference of 0.200 and a village-clustered p-value of 0.066. No other characteristic shown in the main descriptive table differs between Arms A and B at the ten percent level. Age is included in the pre-specified covariate set used in the treatment-effect regressions.

Two initial belief measures are especially relevant for the updating analysis and are also balanced across the main experimental arms. Before treatment, expected planting onset differs between Arms A and B by 0.149 days on the fielded calendar ($p = 0.91$), while the mean weather expectation differs by -0.051 bins. Neither difference is statistically distinguishable from zero.

A.5 Realized fieldwork

Fieldwork dates. Fieldwork took place from 10 through 31 March 2026 over 20 interviewing days. All 686 interviews were completed during March, with the first recorded submission on 10 March and the final submission on 31 March. The first 2026 growing season had not yet resolved, so all season-ahead beliefs analyzed in the paper were elicited before the corresponding weather outcome was known. No interviews were recorded on 15, 17, or 22 March.

Daily interviewing volume increased over the field period. The first five interviewing days contained 32, 31, 31, 35, and 34 interviews. The final full week contained 46, 45, 44, 46, and 48 interviews, followed by 32 interviews on the closing day. The smallest daily total was 25 interviews on 19 March.

Field team. Seven enumerators conducted the survey. They completed 94, 92, 101, 97, 97, 108, and 94 interviews, respectively, giving a range of 92 to 108 interviews per enumerator. Three respondent identifiers contain malformed enumerator codes, but the field records provide no evidence of an additional enumerator. Each enumerator worked across 56 to 59 of the 60 villages, indicating that the field team moved through the study area together rather than assigning separate sets of villages to individual enumerators. All survey submissions were made through the same project account.

Completion. All 60 randomized villages were visited and 686 interviews were completed, with 7 to 14 respondents per village. The realized data contain 686 distinct respondent records and no partial or invalid submissions. Every respondent is associated with one of the 60 sampled villages. The analysis sample therefore coincides with the completed fieldwork sample, with no attrition between data collection and the main respondent-level analysis.

Recall and expectation elicitation. The nine-picture weather exercise shown in Figure 1 was administered for each season from 2021 through 2025 and for the upcoming 2026 season, with the forward expectation elicited a second time where required by the experimental arm. The complete instrument is reproduced in Online Appendix OA-2.

For each past season, respondents were first asked whether they could remember the season. Those unable to recall it after two neutral prompts were not required to allocate tokens. Recall coverage falls from 99.3 percent for 2025 to 61.7 percent for 2021, producing the 2,637 respondent-year observations used in Section 4. Because the instrument records whether a respondent reports being unable to remember a season, rather than directly observing whether the underlying memory has been lost, this decline should be interpreted as declining elicited recall coverage rather than as a direct measure of forgetting.

Extreme-weather memory was elicited separately without prompting respondents with particular years. Respondents named the hazard, year, and perceived severity of remembered extreme events. The resulting data contain 550 first-mentioned events. Across all reported events, respondents name 210 droughts, 134 floods, 103 hailstorms, and 36 storms.

Interview duration. Actual interview duration cannot be recovered from the available survey records because interview start and end times were not retained. We therefore do not report an empirical measure of interview duration.

Data vintage. All fieldwork counts and descriptive statistics in this appendix are based on the frozen survey export obtained on 1 April 2026 at 08:28:56, immediately after completion of fieldwork.

B The objective benchmark and the information cards

This appendix documents how the objective weather benchmark and information treatment were constructed. It describes the rule used to map observed rainfall and vegetation into the nine-picture weather scale, the geographic resolution of the resulting information, the limits this places on analyses of card content, the construction of the planting-onset measure, and one complete information packet as respondents saw it.

B.1 Construction of the information cards

A central feature of the treatment is its geographic resolution. Weather conditions were constructed at the subcounty level and then presented on cards labeled by parish. The study covers 11 subcounties and 42 parishes, so the 42 parish-labeled cards are generated from 11 underlying subcounty weather histories. Across the five displayed seasons, the objective record therefore contains 55 independent subcounty-year cells rather than 210 independent parish-year observations. This structure follows directly from how the treatment was constructed and determines the effective geographic variation in card content.

Geographic unit and data sources. Each subcounty is represented by a centroid surrounded by a two-kilometer extraction buffer. Rainfall is measured using CHIRPS (Funk et al., 2015), with a historical baseline covering the 25 seasons from 2001 through 2025. The information cards display the five seasons from 2021 through 2025. Vegetation conditions for these display years are measured using Sentinel-2 L2A surface reflectance at 10-meter resolution.

Season and rainfall onset. Rainfall onset is identified from daily CHIRPS observations. Beginning on 10 February, onset is defined as the first date for which rainfall over the preceding 21 days reaches at least 25 millimeters and rainfall over the subsequent seven days reaches at least a further 10 millimeters. The search window ends on 15 May, with 30 April used as a fallback if the threshold is never met. The growing season then runs from the detected onset date until the earlier of 180 days after onset and 30 September. A separate fixed window from 10 February through 30 September is used for the rainfall-calendar component of the information treatment so that calendar coverage remains identical across years.

An important feature of this onset definition is that it produces almost no geographic or interannual variation in the study area. Across the 55 subcounty-year cells, detected onset falls on 10 February in 48 cases and on 11 February in the remaining seven. The accumulation threshold is therefore met at or within one day of the beginning of the search window in every case. By contrast, the onset calendar shown to respondents begins on 15 February and ends on 24 April. The rainfall-based benchmark consequently falls before the first elicitation interval in every subcounty-year. This is why the benchmark maps to the earliest onset interval everywhere and why we do not use the instrument to make cross-sectional statements about onset accuracy. Appendix B.3 discusses this issue in detail.

Vegetation and the cropland mask. Vegetation is measured using EVI2, calculated for each Sentinel-2 observation as

$$\mathrm{EVI2} = 2.5\frac{\mathrm{NIR} - \mathrm{Red}}{\mathrm{NIR} + 2.4\,\mathrm{Red} + 1}.$$

Pixels classified as cloud, shadow, snow, or missing are excluded. For each subcounty-year, the cropland mask is based on four conditions. The 10th percentile of EVI2 between 15 January and 31 March must be below 0.20, the 90th percentile between 1 May and 15 September must exceed 0.55, the difference between these two measures must be at least 0.25, and the seasonal vegetation peak must occur between 1 May and 15 September. The resulting mask is smoothed using a local majority rule.

Daily mean vegetation is retained only when at least 30 percent of cropland pixels contain valid observations. Isolated daily outliers are removed using a rolling median and median absolute deviation criterion, after which the time series is smoothed using a three-day rolling median. Peak EVI is defined as the maximum of this smoothed series between 1 May and 15 September. Across the 55 subcounty-year cells, peak EVI ranges from 0.468 to 0.748, and every seasonal vegetation peak falls within the expected April-to-September growing-period window.

Mapping observed weather into the nine-picture scale. Each subcounty-year is assigned to one of the nine weather pictures by combining vegetation performance with the direction of rainfall conditions.

The vegetation component places peak EVI into six fixed bands. Values of 0.65 or above are classified as optimal, values from 0.62 to below 0.65 as near-optimal, values from 0.55 to below 0.62 as good, values from 0.45 to below 0.55 as average, values from 0.35 to below 0.45 as bad, and values below 0.35 as failed.

The rainfall component is constructed over three stages relative to detected onset. Establishment covers days 0 through 30, flowering covers days 31 through 75, and the pre-harvest stage begins on day 76. Within each stage and subcounty, rainfall amount is classified relative to the 2001–2025 distribution. Stage rainfall below the 25th percentile is classified as dry, rainfall between the 25th and 75th percentiles as balanced, and rainfall above the 75th percentile as wet.

We separately characterize the distribution of rainfall within each stage. Pentad rainfall is classified as even, moderate, or lumpy using both the coefficient of variation across pentads and the share of stage rainfall falling in the wettest pentad, benchmarked against the median and 75th percentile of the corresponding historical distributions. Rainfall amount and within-stage distribution are combined into a stage score ranging from -2 for dry conditions to $+2$ for wet and evenly distributed conditions. Establishment receives a weight of 0.25, flowering a weight of 0.40, and pre-harvest a weight of 0.35. The resulting weighted score is divided at -1.0, -0.3, 0.3, and 1.0 to classify the season as dry, slightly dry, balanced, slightly wet, or wet.

Two additional rules address cases in which this stage-based classification may miss important rainfall patterns. First, a dry-spell rule classifies a season as dry when the longest sequence of days with rainfall below one millimeter reaches 44 days. A balanced season is reclassified as slightly dry when the longest such sequence reaches 28 days. This rule does not bind in the realized sample and changes none of the 55 subcounty-year classifications.

Second, when the stage-based rainfall classification is balanced but vegetation is below the optimal range, an additional measure determines whether the season is placed on the dry or wet side of the scale. This measure equally weights seasonal rainfall relative to its 25-year mean and the complement of the share of ten critical ten-day periods from early March through early June with little or no rainfall. Values of 1.10 or above are assigned to the wet side, values of 0.90 or below to the dry side, and intermediate values are assigned to the slightly dry category. This rule applies in 17 of the 55 subcounty-year cells, all of which are classified as slightly dry.

The final picture bin is determined by crossing the vegetation band with the five rainfall categories. For optimal and near-optimal vegetation, dry through wet rainfall conditions map respectively to bins 4, 5, 5, 5, and 6. For good vegetation, they map to bins 4, 4, 4, 6, and 6. For average vegetation, they map to bins 3, 3, 3, 7, and 7. For bad vegetation, they map to bins 2, 2, 2, 8, and 8. For failed vegetation, they map to bins 1, 1, 1, 9, and 9. These assignments follow a fixed lookup rule and are not smoothed or estimated from the realized survey outcomes.

Variation in the realized weather classifications. Among the 55 subcounty-year observations, seven are assigned to bin 3, twenty-one to bin 4, eleven to bin 5, thirteen to bin 6, and three to bin 7. No realized subcounty-year falls into bins 1, 2, 8, or 9. Across the 11 subcounties, the five displayed seasons generate eight distinct unordered sets of weather bins and nine distinct ordered sequences. Buwunga, Muterere, and Nankoma each display the set $\{3, 4, 4, 5, 5\}$. Kigandalo and Kityerera share the ordered sequence $4, 6, 4, 6, 3$, from 2025 back through 2021, while Buwunga and Nankoma share $4, 5, 5, 4, 3$. These limited differences in card content determine what can be identified from regressions using characteristics of the displayed weather history, as discussed in Appendix B.2.

The printed information packet. The cards were produced as high-resolution printed sheets at 300 dpi. Each parish-labeled overview carries the heading "Past rainfall and harvest outcomes in [parish] parish" and presents one block for each of the five displayed years, ordered from most recent to least recent. Each block places the assigned weather picture beside three short statements constructed mechanically from the underlying weather measures.

The first statement describes rainfall conditions using the language associated with the nine weather bins, ranging from "extreme drought" and "drought conditions" through "well balanced rainfall" to "flood conditions" and "severe flooding." The second describes the harvest outcome using peak EVI thresholds at 0.35, 0.43, 0.50, 0.55, 0.60, and 0.65, with descriptions ranging from “crop failure or very little harvest” to “excellent harvest.” The third compares vegetation performance with the other five displayed seasons using the mean and standard deviation of peak EVI within that card. This final statement is therefore relative to the five displayed years rather than to the longer historical distribution.

A separate rainfall calendar divides each month into three approximately ten-day periods. Rainfall in each period is classified relative to the 2001–2025 distribution for the same subcounty and calendar period, using terciles. When fewer than eight historical observations are available for a particular period, month-level terciles are used instead. The three rainfall categories are displayed using one, two, or three water-drop symbols.

Each parish packet contains seven sheets. Five show rainfall and vegetation for the individual seasons from 2021 to 2025, one summarizes rainfall over ten-day periods across all five seasons, and one provides the parish-labeled overview. Across the 42 parishes, this produced 294 printed sheets. Because the annual panels and rainfall calendars are constructed at subcounty level, only 108 of these sheets contain distinct underlying content. Information packets were assembled by village and shown only in the 36 treated villages, comprising 26 villages in Arm B and 10 in Arm C. Appendix B.5 reproduces one complete packet.

B.2 What card content can and cannot support

The construction above places clear limits on what can be learned from regressions that use characteristics of the information cards themselves. Three features are particularly important.

There are few independent information histories. Although 42 parish-labeled cards were printed, their weather content is generated from 11 subcounty histories and reduces to only eight distinct card distributions and nine distinct ordered sequences over the five displayed years. Regressions using the distribution of weather shown on the card therefore rest on only eight distinct information profiles across 11 subcounties. Village-level clustering does not create additional independent variation in these card-level regressors.

The experiment contains no extreme card content. Across the 210 parish-year displays, realized weather bins range only from 3 through 7. No respondent is shown a season in bin 1, 2, 8, or 9. An indicator for exposure to an extreme card is therefore constant at zero and cannot be estimated. The nearest available proxy, the distance of the most extreme displayed bin from the center of the scale, varies only between one and two bins and produces no detectable relationship with updating. The complete information packet reproduced in Appendix B.5 illustrates the limited support of the treatment distribution.

Card shape explains little of the variation in revision. Measures of the distribution shown on the information card, including normalized entropy, standard deviation, and range, do not meaningfully predict the magnitude of belief revision. A joint Wald test over six card characteristics is statistically indistinguishable from zero, although the available variation is too limited for this null to be especially informative. A variance decomposition gives a clearer sense of magnitudes. The eight distinct card distributions, using seven degrees of freedom, explain $R^2 = 0.026$ of the variation in revision magnitude. By comparison, the respondent's absolute initial gap explains $R^2 = 0.081$ with one degree of freedom, while village indicators explain $R^2 = 0.090$ with 25 degrees of freedom.

One exploratory decomposition does reveal an association with information content. We separate the respondent's initial gap from the card into two components. The first is the distance between the card and the leave-one-out consensus initial belief in the respondent's village. The second is the distance between that village consensus and the respondent's own initial belief. Because the first component does not contain the respondent's own elicited belief, it is not mechanically generated by measurement error in that individual's initial elicitation.

Pooling the 567 respondents in Arms A and B across 50 village clusters, the relationship between this village-consensus component and revision has a slope of 0.141 in Arm A. The Arm B interaction is +0.2282 with a standard error of 0.1095, giving an Arm B slope of 0.3692. We interpret this result cautiously for two reasons. First, the regressor combines variation from eight distinct information cards with variation in 50 village-level consensus beliefs, so the interaction cannot be attributed to card variation alone. Second, the interpretation assumes that the village consensus does not itself contain a common elicitation shock arising from enumerator, session, or other village-level influences. Because this analysis was exploratory and the underlying card variation is limited, we report it here rather than in the main text.

B.3 The onset calendar and two measurement caveats

Planting onset is elicited using the same ten-token format as the weather-amount exercise. Respondents allocate tokens across ten dated intervals running from 15 February through 24 April. Nine intervals cover seven days and the third covers six days. Figure 2 reproduces the fielded calendar, and Online Appendix OA-2 transcribes the interval labels. Two features of the instrument constrain how onset magnitudes can be interpreted.

Elicited onset differs systematically from the rainfall benchmark. Reported onset falls substantially later than the CHIRPS-based benchmark in every displayed season. On the fielded day scale, mean recalled onset and the corresponding rainfall benchmark are 73.0 and 41.2 in 2021, 70.5 and 41.1 in 2022, 71.3 and 41.3 in 2023, 69.9 and 41.1 in 2024, and 67.4 and 41.0 in 2025. The resulting gap is approximately 26 to 32 days.

Part of this difference is built into the measurement scales. The earliest interval available to respondents has a midpoint of day 49, whereas the objective benchmark is day 41 or 42 in every subcounty-year. The two measures may also represent different concepts. The rainfall algorithm identifies the first sustained rainfall event satisfying its threshold, whereas farmers may interpret onset as the beginning of sufficiently reliable planting rains. A further concern is that some descriptions on the onset calendar originated in the Elgon pilot and refer to highlands and slopes, as documented in Online Appendix OA-7. The data cannot distinguish among these explanations. We therefore make no cross-sectional claims about the absolute accuracy of elicited onset. The analysis uses the onset measure only for within-person comparisons and randomized treatment contrasts.

The objective benchmark maps to the first interval everywhere. Because CHIRPS-based onset occurs before the first elicitation interval in every subcounty-year, the objective benchmark maps to interval one across all 210 parish-year displays. Cross-sectional measures of onset accuracy therefore contain no independent variation in the benchmark. For the randomized comparison, however, this feature makes the direction of revision unambiguous. Since the historical benchmark lies earlier than the elicited support everywhere, movement toward an earlier expected onset is movement toward the historical rainfall timing. This is the interpretation used for the treatment effect in Section 6.1.

B.4 The day scale of the fielded calendar

Every onset result reported in days is calculated directly from respondents' ten-token allocations using the midpoints of the calendar intervals actually fielded. The calendar covers 15 February through 24 April. Its interval midpoints are days 49, 56, 62.5, 69, 76, 83, 90, 97, 104, and 111 of the year. A respondent's expected or recalled onset is the token-weighted mean of these values and must therefore lie between 49 and 111.

Because the third interval contains six rather than seven days, these midpoints are not exactly affine in the interval number. The day-scale measure is therefore not simply a linear rescaling of the interval-scale measure. We calculate and compare both representations. Results expressed directly on the interval scale, slopes that are invariant to affine transformations, and standardized ratios are substantively unchanged. The main onset treatment effect on the interval scale is -0.3776 intervals (SE 0.1270, $p = 0.0046$). As a support check, every day-scale statistic reported in the paper lies within the feasible range of 49 to 111 days. The initial-belief tercile means of 51.7, 63.7, and 80.4 days all satisfy this restriction.

B.5 One complete information card set

Respondents in Arms B and C were shown a printed information packet corresponding to their parish and subcounty. Figures 13 through 15 reproduce the complete packet used in Nakavule village, Nakavule parish, Kapyanga subcounty, Bugiri district, which was assigned to Arm B. The packet contains seven sheets. One provides an overview of the five previous seasons, one presents rainfall in ten-day intervals for all five seasons, and five present rainfall and vegetation separately for each season. The overview is labeled at parish level, while the rainfall calendar and annual rainfall-and-vegetation panels contain subcounty-level weather information.

Treatment administration. Enumerators were instructed to show respondents the local historical rainfall charts after the first expectation elicitation. After viewing the packet and before repeating the expectation exercise, respondents in Arm B answered a short series of questions about their reaction to the information. They were asked whether the information differed from what they remembered about past rainfall, how surprising they found it on a five-point scale, whether they wished to adjust their expectations, and, if so, how confident they were in those updated expectations. They were also asked whether the information changed their planned adaptations and, if so, what changes they intended to make, whether historical rainfall information was useful for planning, and whether they would like to receive similar information in the future. Respondents then repeated the same nine-picture weather exercise and ten-interval onset exercise for the upcoming season.

As discussed in the main text, these reaction questions were asked only in Arm B between the information treatment and the second elicitation. The treatment effect therefore identifies the information intervention as delivered, including any additional effect of these questions, and the two components cannot be separated. The order in which the different sheets were physically presented is not recorded. We therefore make no claim about sequencing effects within the information packet.

What this packet illustrates. The overview is the only sheet labeled at parish level, although its weather classification is generated from the same underlying subcounty record. The rainfall calendar and five annual panels are identical for every treated village within Kapyanga subcounty. This illustrates why the 42 parish-labeled packets reduce to a much smaller number of distinct displayed weather distributions.

The Nakavule overview also illustrates the limited support of the treatment. It classifies 2021, 2022, and 2025 as "slightly wet conditions," 2024 as "slightly dry conditions," and 2023 as "well balanced rainfall." These correspond to bins 4 through 6 of the nine-picture scale. No season in this packet is close to either extreme, and the same is true of the treatment as a whole. No respondent was shown a season classified as drought, extreme drought, flood, or severe flooding.

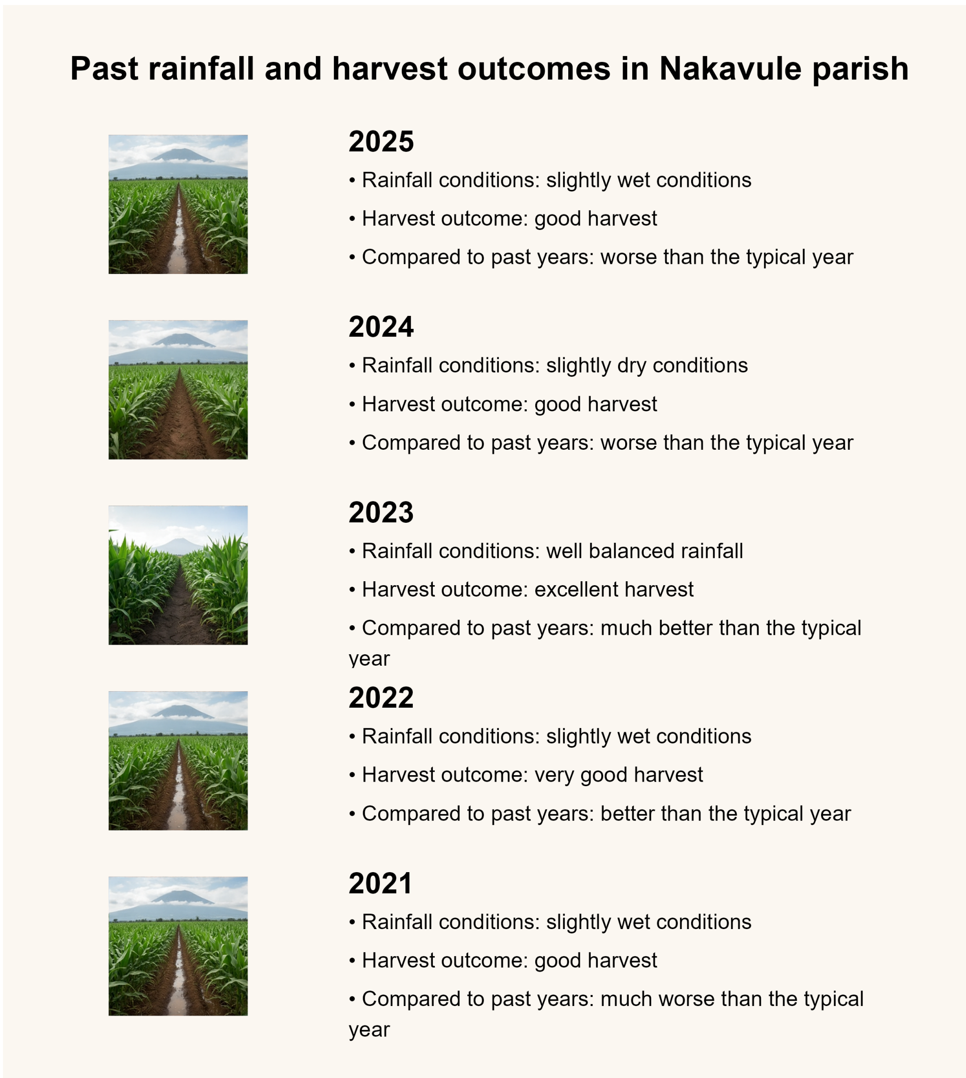

Figure 13: Information card, sheet one: overview of the five previous seasons. Nakavule parish, Kapyanga subcounty, Arm B.

Notes: The overview presents the seasons from 2025 back through 2021. For each season, it displays the same picture used in the token elicitation together with short descriptions of rainfall conditions, harvest conditions, and performance relative to the other displayed years. The sheet is labeled at parish level, while the underlying weather classification is constructed at the subcounty level as described in Section B.1.

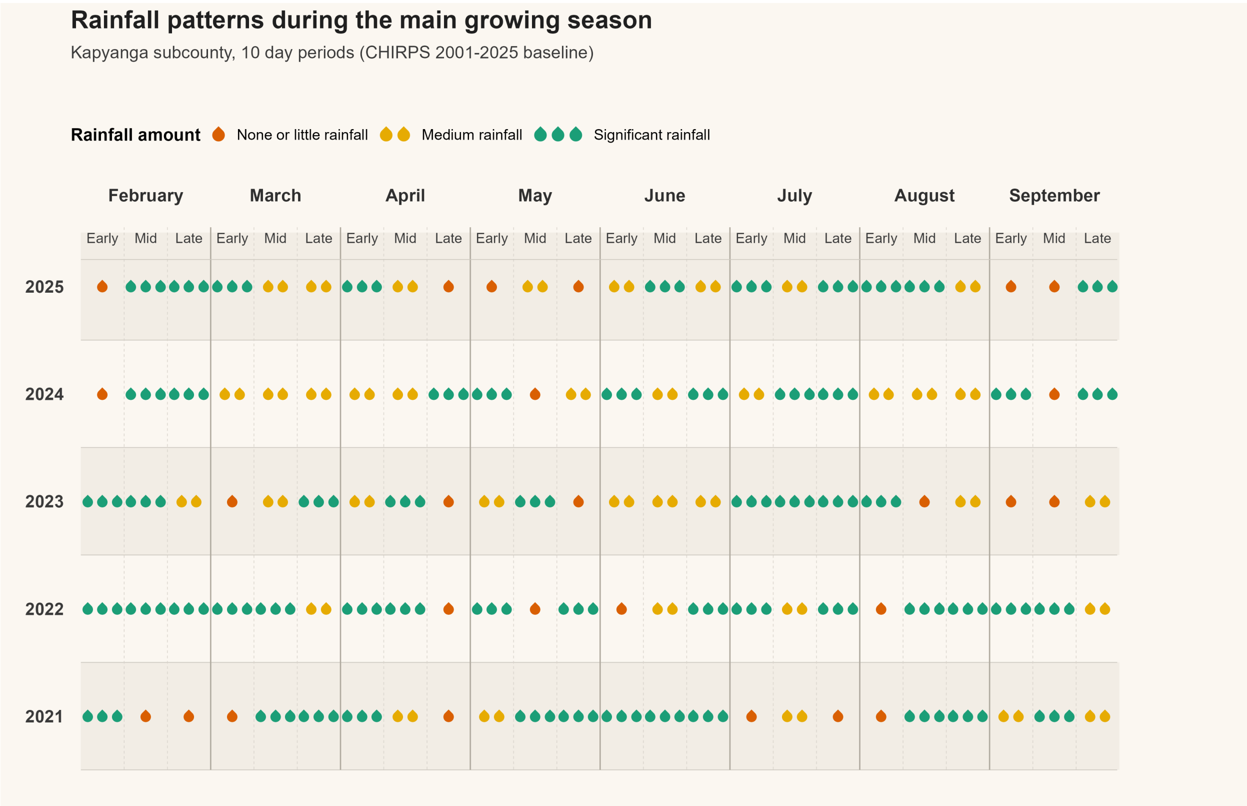

Figure 14: Information card, sheet two: ten-day rainfall calendar for the five previous seasons. Kapyanga subcounty.

Notes: Each row represents one season and each column an approximately ten-day period from February through September. Cells display one, two, or three water-drop symbols according to rainfall relative to the 2001–2025 historical distribution for the same subcounty and calendar period (Funk et al., 2015). This sheet is constructed at subcounty level and is identical for all treated villages within Kapyanga.

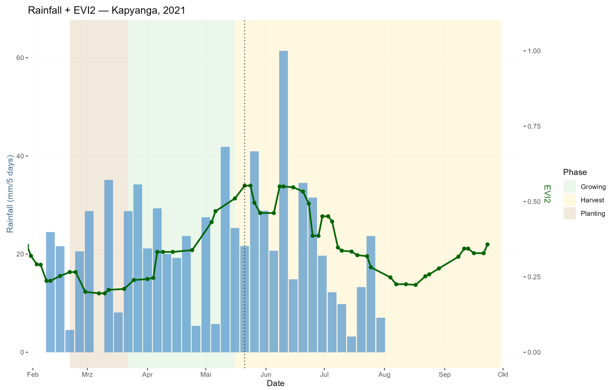

(a) 2021

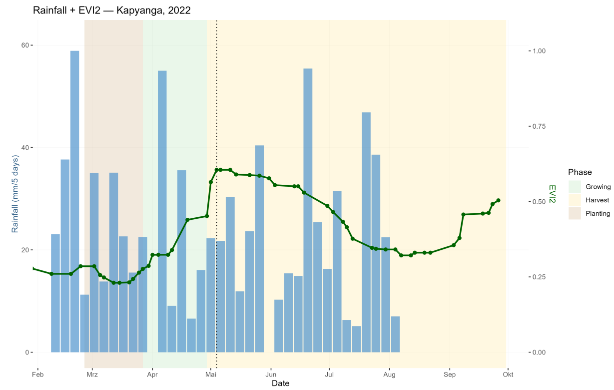

(b) 2022

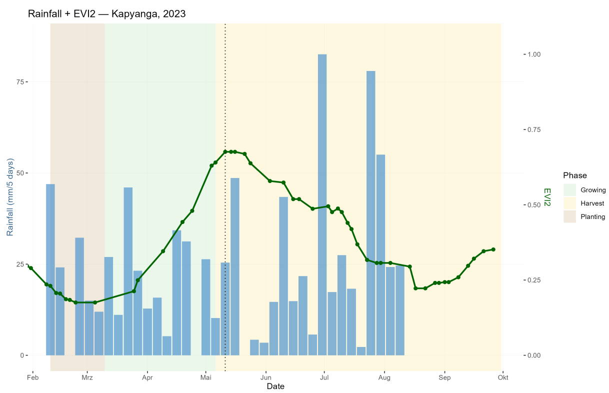

(c) 2023

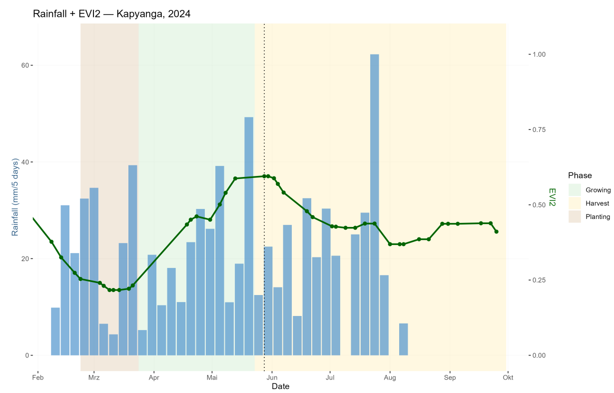

(d) 2024

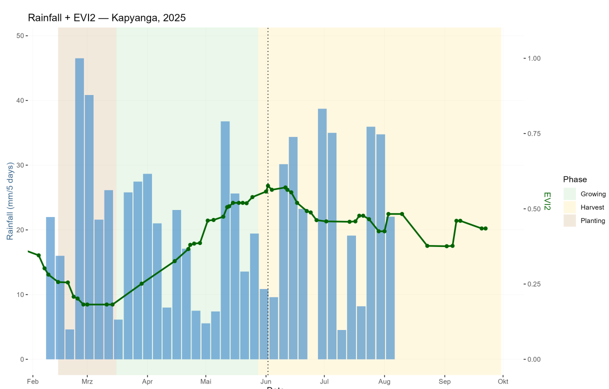

(e) 2025

Figure 15: Information card, sheets three to seven: rainfall and vegetation by season. Kapyanga subcounty, 2021–2025.

Notes: The five annual sheets display CHIRPS rainfall in millimeters per five-day period on the left axis (Funk et al., 2015) and EVI2 vegetation on the right axis. Shaded regions indicate the planting, growing, and harvest phases, and the vertical dotted line marks the end of the growing window. These sheets are constructed at subcounty level and are identical across treated villages within Kapyanga.

C Statistical inference

This appendix documents the inference procedures used throughout the paper. We describe the restricted wild cluster bootstrap, the use of Arm A to benchmark repeated elicitation, Romano–Wolf stepdown adjustment, double selection, the treatment of multiple hypothesis testing, and the clustering structure that determines effective sample sizes. Computational details needed for exact reproduction but not for evaluating the statistical claims are provided in Online Appendix OA-8.

Three features of the design are particularly important for inference. First, the number of independent clusters is limited. The full experiment contains 60 villages, the main two-arm comparison contains 50 villages, Arm B contains 26 treated villages, Arm C contains 10 treated villages, and the objective weather record varies across only 11 subcounties. Second, Arm A repeats the same expectation elicitation without receiving information, allowing us to measure the amount of apparent updating generated by repeated elicitation alone. Third, substantial parts of the analysis are exploratory. We therefore distinguish explicitly between confirmatory, exploratory, and descriptive results. Sections C.1, C.2, and C.5 address these issues in turn.

C.1 The wild cluster bootstrap

Why it is needed here. Conventional cluster-robust inference relies on the number of clusters becoming large. Randomization in this study occurs at the village level, and the relevant comparisons contain 60, 50, 34, or 26 village clusters, with only 10 treated villages in Arm C. The objective weather record itself varies across only 11 subcounties. At these cluster counts, asymptotic approximations can over-reject. We therefore use the restricted wild cluster bootstrap of Cameron et al. (2008), which constructs the reference distribution under the null using the same clustering structure as the observed data.

Bootstrap procedure. For each hypothesis test, we first estimate the unrestricted specification with cluster-robust standard errors and calculate the studentized statistic $t = \hat{\beta}/\widehat{\text{se}}$. We then estimate the restricted specification on exactly the same estimation sample, imposing the null by excluding the tested regressor. Let $\hat{y}_i^0$ and $\hat{u}_i^0$ denote the fitted values and residuals from this restricted model. For each bootstrap replication b, one random weight $w_{g(i),b}$ is assigned to each cluster, so all observations within a village receive the same weight, and the bootstrap outcome is constructed as

$$y_{ib}^* = \hat{y}_i^0 + w_{g(i),b}\hat{u}_i^0.$$

We then re-estimate the unrestricted specification on each bootstrap sample, recompute the cluster-robust standard error using the same clusters, and store the studentized statistic t_b^*. Studentization allows the bootstrap statistics to remain comparable when residual variances differ across replications. The two-sided bootstrap p-value is

$$p = \frac{1 + \sum_{b=1}^{B_{\text{eff}}} \mathbf{1}\{|t_b^*| \geq |t|\}}{B_{\text{eff}} + 1},$$

with analogous signed comparisons for one-sided tests. Replications that do not produce a finite statistic are excluded, and B_{eff} denotes the number of valid replications. Inference is not reported if fewer than 99 valid replications remain.

Bootstrap settings. We use $B = 9{,}999$ replications and a fixed random seed of 20260411. The default weights are Rademacher, taking values -1 and $+1$ with equal probability. When the number of treated clusters falls below twelve, we instead use Webb six-point weights (MacKinnon & Webb, 2017),

$$\{-\sqrt{3/2},\ -1,\ -\sqrt{1/2},\ \sqrt{1/2},\ 1,\ \sqrt{3/2}\}.$$

This rule applies to all contrasts involving Arm C, which contains 10 treated villages, and to inference for the compression regression when clustering at the 11-subcounty level. It does not apply to the main Arm B versus Arm A comparison, which contains 26 treated villages. The finite-sample adjustment in the p-value formula implies a minimum reportable value of 1/10,000, so no bootstrap p-value below 0.0001 is reported.

Results using bootstrap inference. Table 4 summarizes the results for which wild cluster bootstrap inference is reported.

Result	Cluster	G	Weights	B
Recall-on-truth slope, $\beta = 0.0948$	village	60	Webb	9,999
same estimate, clustered at the level at which truth varies	subcounty	11	Webb	9,998
Updating family, 11 outcomes	village	50	Rademacher	9,999
Learning elasticity ϕ_1 and arm interaction	village	50	Rademacher	9,999
Receptivity factor, within-arm and interacted	village	50	Rademacher	9,999
Farm ITT, Arm B versus Arm A	village	50	Rademacher	9,999
Farm ITT, contrasts involving Arm C	village	34, 60	Webb	9,999
Memory-predictor survivors, 37 cells	village	50	Webb	9,999
Predictor-block joint tests, 56 cells	village	50	—	9,999
Adaptation composition and financing families	village	50	Rademacher	9,999
Index regressions at the truth resolution	subcounty	11	—	1,999

Table 4: Results carrying wild cluster bootstrap p-values.

Notes: G is the number of clusters in the estimation sample. Webb six-point weights are used for contrasts involving Arm C and for the compression regression when inference is clustered on the 11 subcounties. The mixture-weight comparison in Table 6 instead uses a village cluster bootstrap with 1,999 draws and is the one inference result in the paper that does not use the restricted wild cluster procedure described above.

For the compression slope, the restricted bootstrap gives a two-sided p-value of 0.0091 using Webb weights over the 60 village clusters. An independent computational implementation gives $p = 0.0106$, while the asymptotic village-clustered p-value is 0.0056. The statement in Section 4.1 that the slope is significant at the one percent level under bootstrap inference refers to the value $p = 0.0091$. When the same regression is instead clustered at the subcounty level, where the objective weather record varies, the restricted bootstrap p-value rises to 0.1759 over 11 clusters. Section C.6 discusses the implications of this coarser level of independent weather variation.

C.2 Placebo testing against Arm A

The repeated elicitation in Arm A provides an empirical benchmark for changes in reported beliefs that occur even when no information is supplied. This design feature is particularly important for statistics constructed from pre-to-post belief changes.

C.2.1 The design feature that makes the comparison possible

Arm A completes the same forward expectation elicitation twice, with no information card and no debrief between the two elicitations. All 274 Arm A respondents completed both rounds. Every within-Arm-B updating statistic therefore has a direct counterpart among respondents who received no information, including movement toward the card, movement toward the realized record, changes in expected planting onset, the magnitude of revision, changes in stated confidence, and parameters constructed from pre-to-post change scores.

We use this repeated control elicitation as the empirical benchmark for the amount of movement generated by the measurement process itself. In Arm A, 57.3 percent of respondents meaningfully change their ten-token allocation between elicitations. The mean absolute shift is 7.71 tokens out of ten, stated confidence rises by 0.208 Likert points, and the distance between the elicited belief and the card that these respondents never saw increases by 0.037 in earth-mover distance. Repeating the elicitation within the same interview can therefore generate substantial apparent updating even in the absence of information.

C.2.2 The three treatment effects that remain after the Arm A comparison

Table 5 reports the three treatment effects emphasized in Section 6.1, each expressed relative to the corresponding change in Arm A.

The first two rows of Table 5 are estimated as differences in the raw pre-to-post change between Arms B and A, without fixed effects, so the treatment contrast decomposes exactly into the two arm means shown in the table. These specifications give an earth-mover-distance effect of -0.1338 (SE 0.0497, clustered $p = 0.0097$, restricted bootstrap $p = 0.0123$) and a ranked-probability-score effect of -0.1509 (SE 0.0506, clustered $p = 0.0044$, restricted bootstrap $p = 0.0047$). Specifications including subcounty fixed effects give closely similar estimates of -0.1399 (SE 0.0438) and -0.1497 (SE 0.0475), with bootstrap p-values close to

Pre-to-post change	Arm A	Arm B	DiD	SE	p
Earth-mover distance to the card	+0.0369	−0.0969	−0.1338	0.0497	0.0097
Ranked probability score against realized 2025	+0.0346	−0.1163	−0.1509	0.0506	0.0044
Believed planting onset, days[†]	+0.552	−2.039	−2.591	0.862	0.0042

Table 5: Treatment effects benchmarked against repeated elicitation in Arm A.

Notes: $N = 567$, of whom 274 are in Arm A and 293 in Arm B, across 50 villages and 11 subcounties. Lower values indicate improvement for both distance measures, so a negative difference-in-differences corresponds to movement toward the card or the realized record. For planting onset, the change within Arm A is statistically indistinguishable from zero ($p = 0.452$), while the change within Arm B is not ($p = 0.00031$). Because the first two outcomes drift away from the card in Arm A, the within-Arm-B changes of −0.0969 and −0.1163 would understate the corresponding treatment effects by roughly 28 percent if repeated elicitation were ignored. For planting onset, the treatment contrast is −2.5913 days with a village-clustered standard error of 0.8624 and $t = -3.0049$, giving $p = 0.0042$ under the finite-sample clustered correction and $p = 0.0027$ under the normal approximation. [†]Onset days are measured on the fielded calendar with interval midpoints from 49 to 111. Appendix B.4 documents the construction of the day scale. The corresponding effect on the interval scale is −0.3776 (SE 0.1270, $p = 0.0046$).

0.005. The two specifications therefore agree in sign, magnitude, and inference. Romano–Wolf adjustment in Section C.3 is applied to the fixed-effects family.

For expected planting onset, the village-clustered p-value is 0.0042 for the −2.591-day treatment effect on the fielded calendar and 0.0046 for the corresponding −0.3776-interval effect, based on 50 village clusters and 567 respondents. The normal approximation gives p-values of 0.0027 and 0.0029, respectively. Both specifications therefore place the onset effect below one percent under either convention.

C.2.3 What a design without repeated control elicitation would imply

Several parameters that appear informative when estimated within Arm B are also reproduced in Arm A. Table 6 reports three examples.

C.2.4 A common mechanical explanation

The three patterns in Table 6 can arise from the same arithmetic feature of change-score regressions that condition on the elicited initial belief.

Write the elicited initial probability for bin b as $p^{\text{pre}} = \pi + \varepsilon$ and the second elicitation as $p^{\text{post}} = \pi + \nu$, where π is the respondent's underlying belief and ε and ν are independent elicitation errors with variance σ_ε^2. Under the null of no information, $\mathbb{E}[\nu \mid \varepsilon] = 0$, so the second elicitation is a second noisy measurement of the same underlying belief. The change score and the initial distance from the card are then

$$\delta = p^{\text{post}} - p^{\text{pre}} = \nu - \varepsilon, \qquad \text{gap} = I_v - p^{\text{pre}} = (I_v - \pi) - \varepsilon, \tag{5}$$

Quantity	Arm B	Arm A	Arm comparison
Learning elasticity ϕ_1	0.419	0.374	+0.045 (SE 0.0543), $p = 0.41$
Mixture weight on the card, ξ, OLS reconstruction	0.470	0.186	—
Mixture weight on the card, ξ, cluster-bootstrap reconstruction	0.470	0.412	+0.063 (SE 0.0679), $p = 0.35$
Receptivity factor on meaningful revision	+0.0585	+0.0727	−0.0164 (SE 0.0212), $p = 0.43$

Table 6: Parameters reproduced by repeated elicitation in Arm A.

Notes: The learning elasticity is the coefficient on the card-minus-prior gap in a bin-level regression of the change in bin probability with respondent fixed effects, estimated on 2,637 bin-cells in Arm B and 2,466 in Arm A. Arm A reproduces 89.3 percent of the Arm B estimate and has a t statistic of 9.59. The mixture-weight estimates are reported under two reconstructions. The ordinary least-squares reconstruction gives an Arm A coefficient of 0.186, while the village cluster-bootstrap reconstruction gives 0.412 in Arm A against 0.470 in Arm B and an arm difference statistically indistinguishable from zero. The receptivity factor has $\alpha = 0.748$ and $\omega = 0.901$, is orthogonal to education, and predicts meaningful revision in both arms, with a somewhat larger coefficient in Arm A. All arm comparisons use 50 village clusters.

where I_v denotes the information shown on the card. Both constructed variables contain the same term, $-\varepsilon$. It follows that

$$\mathrm{Cov}(\delta, \mathrm{gap}) = \sigma_\varepsilon^2$$

and

$$\mathrm{Var}(\mathrm{gap}) = \mathrm{Var}(I_v - \pi) + \sigma_\varepsilon^2.$$

The corresponding population regression coefficient is therefore

$$\phi = \frac{\sigma_\varepsilon^2}{\mathrm{Var}(I_v - \pi) + \sigma_\varepsilon^2} > 0. \tag{6}$$

Under these assumptions, a positive relationship between revision and the initial card gap arises mechanically even when no information is received, and the relationship becomes stronger as elicitation noise increases. The mixture-weight specification captures a closely related object. Regressing p^{post} on p^{pre} and I_v assigns weight to I_v whenever measurement error attenuates the coefficient on the initial elicitation. That attenuation depends on σ_ε^2, not necessarily on learning. The receptivity result can arise through the outcome rather than the regressor. Any characteristic associated with greater token reshuffling will increase $|\delta|$ and can therefore predict measured revision even without treatment-specific learning.

Equation (6) also shows why the repeated control elicitation is useful. The mechanical component is not a fixed additive bias whose magnitude can be known in advance. It depends on elicitation noise relative to the dispersion of the true distance between beliefs and the signal. Arm A provides an empirical estimate of this component. In this experiment, it reproduces about ninety percent of the within-Arm-B learning elasticity.

Expected planting onset requires a different benchmark because the information card contains no parish-specific onset signal and therefore provides no analogous onset gap. The relevant mechanical force is regression to the mean in the initial elicitation. In Arm A, the slope of revision on initial expected onset is -0.2835 (SE 0.0381, $p = 1.0 \times 10^{-13}$). The corresponding slope in Arm B is -0.4317 (SE 0.0595, $p = 4.0 \times 10^{-13}$), giving an arm difference of -0.1482 (SE 0.0700, $t = -2.117$, $p = 0.034$). Recomputed on the fielded calendar, the arm difference is -0.1486 (SE 0.0705, $t = -2.107$, $p = 0.040$). The treatment-relevant object is therefore the difference between arms rather than the within-arm slope itself. This is the continuous counterpart of the tercile analysis in Section 6.2. Because a regression slope of revision in days on an initial belief in days is invariant to affine transformations of the day scale, the two versions agree to three decimal places even though the levels of the calendar measures differ.

C.2.5 Why the positional result is different

Regression to the mean pulls extreme initial measurements inward from both directions. The pattern in Section 6.2 is instead asymmetric. Farmers whose initial expected onset is already early, with a mean of day 51.7, show essentially no treatment response, while those in the middle and late terciles move 4.41 and 4.10 days earlier. Moreover, each treatment effect is estimated relative to Arm A within the same initial-belief tercile. Mechanical mean reversion is therefore differenced out within each part of the initial-belief distribution rather than only in the aggregate.

C.2.6 Netting out mechanical revision and residual contamination

As an additional diagnostic, we estimate the relationship between initial beliefs and mechanical revision using Arm A alone, both at the bin level and at the individual level. These relationships are then used to predict the revision that would be expected mechanically for each respondent in Arm B. Excess revision is defined as observed revision minus this predicted mechanical component.

This adjustment has an important asymmetry. Because the mechanical relationship is estimated within Arm A, the resulting Arm A excess measure is orthogonal by construction to the regressors used to generate the prediction, whereas the Arm B excess measure is not. For predictors correlated with these mechanical regressors, this construction tends to shrink the Arm A coefficient and can therefore increase the apparent difference between arms. The resulting bias works against finding equal coefficients across arms, making the null arm differences reported below conservative. We also examine the remaining relationship between each excess measure and the initial elicitation, since removing the conditional mean estimated in Arm A does not eliminate all first-elicitation noise. The corresponding diagnostics are reported in Online Appendix OA-8.

C.2.7 Implications for information experiments

More generally, information experiments that relate belief revision to the initial belief-to-signal gap, estimate mixture weights on a signal, or ask which characteristics predict the magnitude of revision face a mechanical component whenever the same noisy elicitation enters both sides of the estimating equation. Without a group that repeats the same elicitation in the absence of information, this component cannot be separately observed. A repeated elicitation in the control group provides a direct empirical benchmark for it.

C.3 Romano–Wolf stepdown

We use Romano–Wolf stepdown adjustment for pre-defined families of related hypotheses. Shared cluster draws are used across hypotheses so that dependence among the test statistics is preserved even when estimation samples differ. The computational construction and validation of the procedure are documented in Online Appendix OA-8. Here we report the hypothesis families and the resulting adjusted inference.

Families and results. Table 7 lists every family to which Romano–Wolf adjustment is applied.

Family	J	G	B	Weights	Survivors at $q < 0.10$
Updating difference-in-differences	11	50	9,999	Rademacher	0
Updating screen, amount, Arm B	4×22	26	9,999	Rademacher	0 in each of the four
Updating screen, onset, Arm B	3×22	26	9,999	Rademacher	1 in one of the three
Updating screen, Arm A placebo	3×22	24	9,999	Rademacher	1 in one of the three
Memory-quality predictor screen	420	50	4,999	Rademacher	37
Adaptation composition	10	50	9,999	Rademacher	0
Adaptation financing channels	8	50	9,999	Rademacher	0
Demographic selection into recall	5	60	999	Rademacher	2
Prior gap on adaptation spending and channels	2	50	9,999	Rademacher	2

Table 7: Romano–Wolf hypothesis families and adjusted results.

Notes: J is the number of hypotheses in the family, G the number of village clusters, and B the number of bootstrap replications using shared cluster draws. The memory-quality screen consists of 105 candidate predictors crossed with four pre-specified outcomes, and the full 420-cell grid defines the family. The updating screen crosses 22 pre-treatment predictors with four amount outcomes and three onset outcomes, with each of the seven outcomes corrected as a separate family. The Arm A analyses repeat the corresponding screens within the group that received no information.

The updating family. Among the eleven pre-to-post belief changes, three have unadjusted p-values below 0.05: earth-mover distance at $p = 0.0061$, ranked probability score at $p = 0.0060$, and stated rainfall confidence at $p = 0.0179$. None remains below five percent after Romano–Wolf adjustment. The corresponding adjusted values are $q = 0.0597$, 0.0625, and 0.1318. The main result in Section 6.1 instead relies on the earth-mover distance having been designated as the primary outcome before estimation. The paper therefore does not interpret the stepdown procedure as selecting a significant outcome ex post.

The 22-predictor updating screen. Across the four weather-amount outcomes, none of the 22 predictors survives at $q < 0.10$ in Arm B in any family. The smallest adjusted values are $q = 0.1068$ for memory entropy when predicting bin-level excess revision and $q = 0.1218$ for the absolute initial gap when predicting excess revision magnitude. The corresponding Arm A placebo screen is informative about these associations. The absolute initial gap reaches $q = 0.061$ in Arm A, compared with $q = 0.122$ in Arm B. The stronger association in the untreated arm indicates that this predictor captures a feature of repeated elicitation rather than a treatment-specific response.

The only adjusted survivor elsewhere in the screen produces the same qualitative conclusion. For excess earliness of the onset belief in Arm B, the receptivity factor has a coefficient of -1.25 days per standard deviation (SE 0.34, $t = -3.72$) and $q = 0.0519$. In the corresponding Arm A placebo family, the coefficient is -1.36 days (SE 0.47, $t = -2.92$), which is larger in magnitude, although its adjusted value is $q = 0.2033$ because another statistic leads that family. Both coefficients use the fielded calendar. Together with Table 6, this pattern is consistent with the receptivity factor capturing a general propensity to revise or reshuffle the elicited distribution rather than a treatment-specific propensity to absorb information. Across the 154 predictor-by-outcome arm interactions, which directly test differential response to the card, no adjusted value falls below 0.10.

How we interpret the stepdown adjustment. Romano–Wolf adjustment is most naturally interpreted for hypothesis families defined before the corresponding estimates are examined. We therefore use it as a confirmatory device for pre-registered and otherwise pre-committed families. We also report it for the 22-predictor updating screen, where the adjusted results are useful for comparing the treated and placebo arms, but we do not treat an ex post exploratory family as though it had been pre-specified.

The descriptive belief-investment relationships in Section 5.4 are not Romano–Wolf adjusted and are not presented as confirmatory evidence. Roughly thirty specifications are estimated in that analysis. The object of interest is the pattern across outcomes, as Appendix D.5 explains, and individual p-values close to conventional thresholds should be interpreted as suggestive rather than as independently confirmatory findings.

C.4 Double selection

Purpose. We use double selection (Belloni et al., 2014) to reduce discretion over the control set in the heterogeneity analyses. In both applications, the question is whether a pre-treatment characteristic predicts an outcome after allowing the remaining candidate characteristics to enter flexibly through the selection procedure. Stepwise selection is not used. The candidate sets, penalty choices, and estimability rules are documented in Online

Appendix OA-8.

Results. Double selection changes little and selects few additional controls. For excess revision magnitude in Arm B, the absolute initial gap is estimated at +0.1487 with $p = 0.0067$ and no selected controls. Memory entropy is estimated at +0.1427 with $p = 0.00047$ and two selected controls, while the receptivity factor is estimated at +0.1124 with $p = 0.0146$ and three selected controls. For excess earliness of expected onset, the receptivity factor is the leading target, with an estimate of -1.13 days (SE 0.38), $p = 0.0031$, and a Benjamini–Hochberg adjusted value of $q = 0.068$, with three controls selected.

The placebo comparisons determine how these estimates should be interpreted. The variables selected most clearly by double selection are the absolute initial gap and the receptivity factor, but both also predict revision in Arm A. The absolute initial gap is more statistically pronounced in the untreated arm, while the receptivity coefficient is larger in magnitude there. Double selection therefore addresses discretion in the choice of controls but cannot remove mechanical dependence built into the construction of the revision outcome. In this setting, the repeated Arm A elicitation provides the more relevant discipline for distinguishing treatment-specific updating from general re-elicitation effects.

C.5 Multiple-hypothesis accounting

Table 8 provides a complete inventory of the main hypothesis families, the multiplicity adjustments applied to them, and the inferential status of the corresponding claims.

The general rule is to apply multiplicity adjustment when a hypothesis family was defined before estimation and to use the repeated Arm A elicitation whenever inference is based on a change score. Where neither form of discipline applies, as in Section 5.4, the analysis is explicitly labeled descriptive and individual p-values are not asked to support confirmatory claims.

C.6 Clustering and effective sample size

The most important constraint on statistical precision comes from the geographic resolution at which the objective weather record varies.

The objective record varies at the subcounty level. The weather panel contains 42 parish rows, but there is only one distinct picture classification and one peak-greenness value for each subcounty-year. Across the five seasons from 2021 to 2025, the nominal 210 parish-year observations therefore contain only **55 independent weather observations**, corresponding to 11 subcounties observed for five years. Treating the 42 parishes as independent weather realizations would overstate precision, so the paper reports 11 subcounties as the effective number of independent geographic units on the weather side.

Family	Tests	Correction	Status of the claim
Registered family $\{H1, H3, H4\}$	3	Romano–Wolf, registered	Confirmatory in intent but not executable as written. $H2$ is suspended, $H1$ is degenerate, and $H3$ has an effective n of 11. See Appendix E.
Farm-outcome ITT	4×3	Wild cluster bootstrap; no stepdown across outcomes	Confirmatory in design and null in result. Reported with confidence intervals rather than as evidence of absence.
Updating difference-in-differences	11	Romano–Wolf on shared draws	Exploratory. The primary outcome was designated a priori. No outcome survives the stepdown at five percent, and the paper reports this explicitly.
Updating predictor screen	22×7	Romano–Wolf by outcome family, Benjamini–Hochberg, and Arm A placebo families	Exploratory and reported as a bounded null with a stated minimum detectable effect.
Memory-quality predictor screen	420	Romano–Wolf over the full grid and Holm adjustment over 56 block tests	Exploratory. The family is the complete search rather than the subset of surviving coefficients.
Belief–investment associations	≈ 30	*None*	Descriptive and correlational. The pattern across outcomes is the object of interest, and individual p-values are interpreted as suggestive.
Adaptation composition and financing	$10 + 8$	Romano–Wolf by family	Exploratory and null.

Table 8: Confirmatory, exploratory, and descriptive inference across the paper.
Notes: "Tests" counts coefficients of interest rather than regressions. The updating predictor screen estimates each of its 154 predictor-by-outcome cells in Arm B, in Arm A as a placebo, and as an interaction between the two arms. The belief–investment analysis is the only family for which no multiplicity correction is applied and is correspondingly the only one that makes no confirmatory inferential claim. Sections 4 through 6 are exploratory except where a registered hypothesis is explicitly identified.

The implication is visible in the main compression estimate. The recall-on-truth slope in Section 4.1 is 0.0948. With respondent fixed effects and standard errors clustered over 60 villages, the standard error is 0.0342 and $t = 2.77$. When inference is instead clustered on the 11 subcounties, which is the level at which the objective regressor varies, the point estimate is unchanged but the standard error rises to 0.0566, giving $t = 1.67$, an asymptotic p-value of 0.094, and a Webb-weighted restricted bootstrap p-value of 0.176. Restricting the analysis to the clean pre-information sample in Arms A and B gives a slope of 0.0756 with a village-clustered standard error of 0.0389. The point estimate is therefore reasonably stable across these specifications, while its precision depends strongly on recognizing that only 11 independent geographic weather histories underlie the objective record.

The rainfall data provide little independent farm-level variation. The 686 farms occupy **17 distinct CHIRPS grid cells**, with a median of 36 farms per cell. Only five of the 11 subcounties contain two or more grid cells, while the remaining six contain a single cell. Within those six subcounties, every farm is therefore assigned the same rainfall measure by construction. This is why the pre-registered farm-pixel fallback described in Appendix E could not be implemented. It also explains why the registered primary specification is degenerate: a rainfall z-score that is constant within subcounty is fully absorbed by subcounty fixed effects.

Consequences for effective degrees of freedom. On the weather side, the effective sample contains 55 subcounty-year cells rather than 3,430 farm-year observations. In the backtest of Section 5.2, which is evaluated on a nominal 220-row panel, the intraclass correlation by year is 0.469 for the leading model comparison. This implies a design effect of 5.69 and an effective sample size of 38.6 rather than 220. Diebold–Mariano comparisons are therefore reported using both subcounty clustering with $G = 11$ and year clustering, while leave-one-subcounty-out cross-validation uses 11 folds. Bootstrap inference for index regressions at this resolution uses 1,999 replications over 11 clusters, and all corresponding coefficients are statistically null.

Which results are affected. This geographic constraint applies whenever the regressor of interest is the objective weather record. It therefore affects the compression slope, the signed dry lean, the flatness statistics constructed from the 55 subcounty-year cells, the backtest and its cross-validation, and the registered hypotheses comparing a belief index with a rainfall anomaly.

The same constraint does not govern the randomized information experiment. Treatment is assigned at the village level, and the specifications in Section 6 regress within-respondent belief changes on village-level assignment. The relevant cluster count is therefore the 50 villages represented in Arms A and B, of which 26 are treated, and the geographic resolution of the weather record does not determine treatment-effect inference. The same reasoning applies to the farm-outcome ITT. The pooled comparison uses 60 villages and contrasts involving Arm C use 34 villages. The limited power of those farm-outcome tests instead reflects the resolution of the outcome data: farm identity explains only 0.35 percent of the variation in peak greenness, while the CHIRPS cell explains 63.4 percent, leaving little independent farm-level signal for a village-randomized intervention to affect. Finally, the geographic weather constraint does not determine inference for the belief–investment relationships in Section 5.4 or the extreme-event memory results in Section 5.3, because the regressors in those analyses vary at the individual level. Those specifications cluster standard errors at the village level and include subcounty fixed effects to account for the common geographic environment.

D Additional robustness

This appendix reports additional analyses supporting the main results. We present the out-of-sample backtest underlying Section 5.2, power and equivalence bounds for null results that matter for the interpretation of the paper, measurement checks for the compression and flatness results, the full specification inventory underlying Section 5.4, supplementary results on dry-spell memory, two additional figures, and a detailed account of the 2026 growing season and farm-outcome analysis.

D.1 Inference

The inference procedures used throughout the paper are documented in Appendix C. Standard errors are clustered at the village level, using 50 clusters for the main comparison between Arms A and B and 60 clusters where the full experimental sample is used. Wild cluster bootstrap inference uses 9,999 replications with Rademacher weights and the finite-sample p-value adjustment described in Appendix C.1 (Cameron et al., 2008). For contrasts involving Arm C, which contains only 10 villages, we use Webb six-point weights and matched-pair randomization inference because conventional Rademacher weights can be size-distorted with very few treated clusters (MacKinnon & Webb, 2017). Multiple testing is addressed using Romano–Wolf stepdown within defined outcome families (Romano & Wolf, 2005), while variable selection where used follows the double-selection procedure of Belloni et al. (2014) rather than stepwise selection.

D.2 The backtest of the perceived weather history

Section 5.2 shows that applying the estimated memory distortion removes most of the variation across recent weather histories. A natural validation question is whether measuring weather shocks relative to this constructed perceived history adds predictive information for farm productivity beyond a conventional rainfall anomaly. We test this out of sample.

We define a perceived shock as the deviation of realized weather from the perceived historical reference in Equation (3). This is the belief-relative analogue of a conventional rainfall anomaly. We then ask whether this measure predicts farm productivity out of sample. The analysis covers 3,430 farm-year observations, corresponding to 686 farms over the five seasons from 2021 through 2025. Farm productivity is measured using peak vegetation greenness relative to each farm's own observed history. This outcome does not enter the construction of the perceived-weather index.

Table 9 reports leave-one-subcounty-out prediction. None of the alternative weather-shock measures improves on predicting the sample mean, and the memory-distorted measure performs slightly worse than the conventional rainfall anomaly. This is consistent with the flatness result in Section 5.2. Once the historical reference varies very little across places and years, subtracting that reference from realized weather produces an index that is close to an absolute weather anomaly. Correspondingly, the distorted index has a correlation of 0.985

with the conventional rainfall anomaly. The undistorted historical reference, which preserves the variation that would remain under accurate memory, performs marginally better than the other weather measures but still does not outperform the sample mean.

Table 9: Out-of-sample prediction of farm productivity across alternative weather-shock measures.

Shock measure	Leave-one-subcounty-out R^2	Reference
Predicting the sample mean	0.0000	benchmark
Undistorted expectation ($\beta = 1$)	-0.0027	Eq. (3)
Conventional rainfall anomaly	-0.0078	conventional measure
Distorted memory index	-0.0116	Eq. (3)
Distorted drought surprise	-0.0273	Eq. (3)

Notes: Out-of-sample R^2 for farm productivity over 3,430 farm-year observations, corresponding to 686 farms in 55 subcounty-year cells over 2021–2025, using leave-one-subcounty-out cross-validation. Farm productivity is peak vegetation greenness relative to the farm's own observed history and does not enter the construction of the shock measures. Negative values indicate worse out-of-sample prediction than the sample mean. Source: authors' computation.

The conclusion is unchanged under alternative cross-validation schemes. Table 10 compares leave-one-year-out prediction over five folds, leave-one-subcounty-out prediction over 11 folds, and a leakage-free leave-one-year-out procedure in which the farm-specific historical productivity benchmark is recomputed within each training sample so that the held-out year does not enter the benchmark. Under all three procedures, the distorted memory index performs slightly worse than the conventional rainfall anomaly, with differences in out-of-sample R^2 of -0.0016, -0.0038, and -0.0008, respectively. No specification outperforms prediction by the sample mean. We also estimate leave-one-subcounty-out models including year fixed effects, with the same qualitative conclusion.

D.3 Power and equivalence bounds for the null results

Several null results play a substantive role in the paper. For these results, a failure to reject zero is not by itself sufficient to establish that economically meaningful differences are absent. We therefore report minimum detectable effects, confidence intervals, and equivalence bounds where appropriate.

Observable characteristics and updating. Across 22 pre-treatment characteristics and four measures of weather-belief updating, the minimum detectable differential effect ranges from approximately 0.22 to 0.25 outcome standard deviations for a one-standard-deviation change in the predictor. The largest estimated differential effects range from approximately 0.16 to 0.21 standard deviations. The experiment therefore rules out relatively large differences in responsiveness across observable farmer characteristics, while remaining compatible with smaller differences.

Table 10: Out-of-sample R^2 by shock measure and cross-validation scheme.

Folds	Leave-one-year-out 5	Leave-one-subcounty-out 11	Leakage-free leave-one-year-out 5
Predicting the sample mean	0.0000	0.0000	0.0000
Undistorted expectation ($\beta = 1$)	-0.0368	-0.0027	-0.0311
Conventional rainfall anomaly	-0.0412	-0.0078	-0.0351
Distorted memory index	-0.0429	-0.0116	-0.0359
Distorted drought surprise	-0.0406	-0.0273	-0.0375
Distorted minus conventional	-0.0016	-0.0038	-0.0008

Notes: Out-of-sample R^2 for farm productivity over 3,430 farm-year observations from 2021 through 2025. Farm productivity is peak vegetation greenness relative to each farm's own historical performance. In the leakage-free specification, that benchmark is recomputed separately within each training fold. Negative values indicate worse out-of-sample prediction than the sample mean. The final row reports the difference between the distorted-memory index and the conventional rainfall anomaly. Source: authors' computation.

Demographics, updating, and memory quality. The contrast between updating and initial memory quality is particularly informative. For the three belief-updating outcomes, joint tests of education, literacy, age, and household size give p-values of 0.82, 0.31, and 0.30. Across the corresponding cells, minimum detectable effects range from 0.13 to 0.19 standard deviations and equivalence bounds from 0.09 to 0.21 standard deviations. One of the twelve individual coefficients, household size on the change in mean rainfall belief, has $p = 0.028$, but this isolated result is not unusual in a family of twelve tests.

The same demographic characteristics strongly predict memory quality. Joint tests reject zero for recall accuracy measured by the ranked probability score and by a linear scoring rule, with p-values below 0.0001 under alternative fixed-effect specifications, and for memory entropy at $p = 0.007$. Literacy is the strongest individual predictor, with literate respondents recalling past seasons approximately 0.17 standard deviations more accurately. By contrast, the signed recall bias toward drier conditions is jointly unrelated to these characteristics, with $p = 0.50$. Minimum detectable effects across the 16 memory-quality cells range from 0.08 to 0.17 standard deviations, with equivalence bounds between 0.05 and 0.23. The observed education and literacy gradient in recall accuracy lies within this detectable range and is observed in the data. The contrast in the main text is therefore supported by the power calculations: human capital is systematically associated with initial memory accuracy, while neither the average dry lean nor short-run responsiveness to information shows a comparable demographic gradient. This is also consistent with the broader memory-quality predictor screen in Table 7, where 37 of 420 cells survive Romano–Wolf adjustment.

Farm-outcome treatment effects. The intention-to-treat estimates for the four primary Arm B farm outcomes are statistically indistinguishable from zero, but the precision differs substantially across outcomes. For standardized peak greenness, the estimate is -0.140 with a 95 percent confidence interval of $[-0.456, +0.176]$, already expressed in standard-deviation units. For the own-history benchmark, the estimate is -0.0033 with an interval of $[-0.0111, +0.0045]$. Relative to the outcome standard deviation of 0.0401, this corresponds to approximately $[-0.28, +0.11]$ standard deviations. For the village-mean benchmark, the estimate is $+0.0003$ with an interval of $[-0.0004, +0.0010]$. For the double-differenced benchmark, the estimate is $+0.0002$ with an interval of $[-0.0002, +0.0007]$, corresponding to approximately $[-0.02, +0.06]$ standard deviations given an outcome standard deviation of 0.0116. Interactions between treatment and the pre-treatment belief gap have bootstrap p-values between 0.29 and 0.75 across the four outcomes. The tightest interval, for the double-differenced benchmark, therefore excludes favorable Arm B effects larger than approximately six percent of an outcome standard deviation. The other outcomes leave considerably wider ranges of economically relevant effects open.

D.4 Measurement checks on the compression and flatness results

The main compression result is not sensitive to the use of a nine-picture elicitation scale. Reclassifying the nine bins into three broader categories, dry, normal, and wet, produces a recall-on-truth slope of 0.0537 with a standard error of 0.0313 and a bootstrap p-value of 0.0982 over 2,559 respondent-year observations. Perfect preservation of differences across seasons would still imply a slope of one after this recoding. The lower precision therefore reflects the variation lost through coarsening rather than the disappearance of compression.

The estimate is also insensitive to alternative fixed-effect structures. Replacing respondent fixed effects with village fixed effects gives a slope of 0.0997 (SE 0.0337), while including both respondent and village effects gives 0.0948 (SE 0.0391). These estimates are close to the preferred coefficient of 0.0948 and provide little indication that common village-level survey conditions generate the compression pattern.

A dedicated enumerator field was not retained in the survey export. As Appendix A notes, however, enumerator codes can be recovered from the respondent identifier for 683 of the 686 interviews. We therefore do not interpret the absence of a dedicated field as evidence that enumerator effects are impossible to examine. Any robustness claim based on enumerator fixed effects should use these recoverable codes and is separate from the village-level checks reported above.

A second concern is that apparent flatness could simply reflect respondents spreading their tokens across many pictures. The elicited distributions show the opposite. Before treatment, forward expectations occupy 1.76 bins on average, with a median of two, and 48.68 percent of the 567 Arm A and Arm B distributions place all ten tokens on a single picture. Recalled distributions occupy between 1.93 and 2.12 bins on average, each with a median of two. Across the 2,637 respondent-year recall observations, 35.0 percent are point masses. Half of respondents, 50.2 percent, give a point-mass recall distribution for at least one season, while 21.3 percent give one for every season they report. Planting-onset beliefs are still more concentrated, occupying between 1.63 and 1.76 intervals on average with medians between one and two. The elicited distributions are therefore generally concentrated. The flatness result in Section 5 concerns the limited variation in their means across objective weather histories, not excessive dispersion within individual responses.

The numerical flatness result is calculated over the 55 subcounty-year cells, giving each cell equal weight. The constructed perceived historical reference has a standard deviation of 0.0724 bins and ranges from 4.653 to 4.992. Under accurate memory, the corresponding historical reference has a standard deviation of 0.7644 bins and ranges from 4.760 to 8.339. These are the values reported in Section 5.2 and Figure 7. Weighting subcounty-year cells by the number of sampled farms would produce different summary statistics because heavily sampled locations would receive greater influence. The paper uses the unweighted distribution because the object of interest is variation across independent subcounty-year weather histories.

D.5 Specification inventory for the belief–investment associations

Section 5.4 presents selected associations between weather beliefs, memory, and planned adaptation. These relationships are descriptive rather than causal. To make the underlying specification search transparent, Table 11 reports the complete grid of five belief and memory measures crossed with five adaptation outcomes. Each specification includes subcounty fixed effects and controls for cultivated acreage, age, and household size, with standard errors clustered at the village level. The pattern across outcomes is the object of interest. Individual p-values, particularly those close to conventional thresholds, should be interpreted as suggestive.

Two associations in this grid are not emphasized in the main text. The relationship between memory accuracy and irrigation changes sign across alternative specifications and is therefore not interpreted as a robust pattern. Recalled event severity is also negatively associated with fertilizer use, but the magnitude is small and the result does not align with the broader pattern across severity outcomes. We report both coefficients in Table 11 for completeness without attaching a substantive interpretation to them.

D.6 Dry-spell salience in memory

We also ask whether seasons containing longer dry spells are remembered as drier after conditioning on the overall realized seasonal weather classification. In the cross section, the relationship goes in the opposite direction, with years containing longer dry spells recalled as somewhat wetter. This association disappears when respondent fixed effects are introduced and is also statistically indistinguishable from zero when inference is clustered at the subcounty level. The cross-sectional pattern therefore appears to reflect differences across locations rather than within-person memory encoding of dry spells.

The support of this test is limited. Longest dry spells in the recalled seasons range only from 5 to 13 days, whereas the 2026 season contains a 25-day dry spell. The null result therefore does not show that unusually long dry spells have no effect on memory. The historical sample simply contains no events close to the magnitude observed in 2026.

Table 11: Belief and memory measures and planned adaptation outcomes.

Regressor	Outcome	Coef.	SE	p	n
Expects future extreme	Adaptation spend (log)	0.659	0.175	0.0004	323
	Fertilizer	0.125	0.042	0.0043	567
	Improved seeds	0.066	0.051	0.206	567
	Irrigation	−0.051	0.030	0.089	567
	Plans insurance	0.048	0.020	0.020	327
Recalled past extreme	Adaptation spend (log)	0.781	0.151	< 0.0001	397
	Fertilizer	−0.014	0.045	0.758	686
	Improved seeds	−0.056	0.046	0.221	686
	Irrigation	0.035	0.033	0.294	686
	Plans insurance	0.018	0.007	0.018	404
Recalled severity (1–10)	Adaptation spend (log)	−0.108	0.036	0.0037	333
	Fertilizer	−0.022	0.010	0.038	550
	Improved seeds	−0.011	0.009	0.200	550
	Irrigation	0.010	0.009	0.276	550
	Plans insurance	−0.003	0.002	0.196	339
Memory RPS (higher = worse)	Adaptation spend (log)	−0.351	0.128	0.0081	396
	Fertilizer	−0.055	0.027	0.046	685
	Improved seeds	−0.172	0.026	< 0.0001	685
	Irrigation	0.065	0.026	0.017	685
	Plans insurance	−0.005	0.009	0.555	403
Memory entropy	Adaptation spend (log)	1.356	0.359	0.0004	396
	Fertilizer	−0.051	0.094	0.593	685
	Improved seeds	0.198	0.100	0.052	685
	Irrigation	−0.090	0.075	0.234	685
	Plans insurance	0.000	0.027	0.986	403

Notes: Each row reports a separate regression of the indicated adaptation outcome on the belief or memory measure. Specifications include subcounty fixed effects and controls for cultivated acreage, age, and household size, with village-clustered standard errors. There are 50 clusters when the regressor requires the pre-treatment forward elicitation and 60 otherwise. Adaptation spending and planned-insurance outcomes are defined among respondents planning any adaptation. Recalled severity is the maximum severity assigned to a named past event. Memory RPS is the ranked probability score of recalled weather and increases as recall becomes less accurate. All relationships in this table are descriptive and correlational.

D.7 Supplementary figures

The following figures provide additional detail for the extreme-event memory results in Section 5.3 and the planned-adaptation results in Section 5.4.

D.8 The 2026 season and the farm-outcome test in detail

The realized season. Cumulative rainfall through 30 June 2026 was 76.5 percent of the full-season long-run mean. Even if the remainder of the season had matched the lowest historical July–August rainfall observed at each location, the resulting full-season rainfall

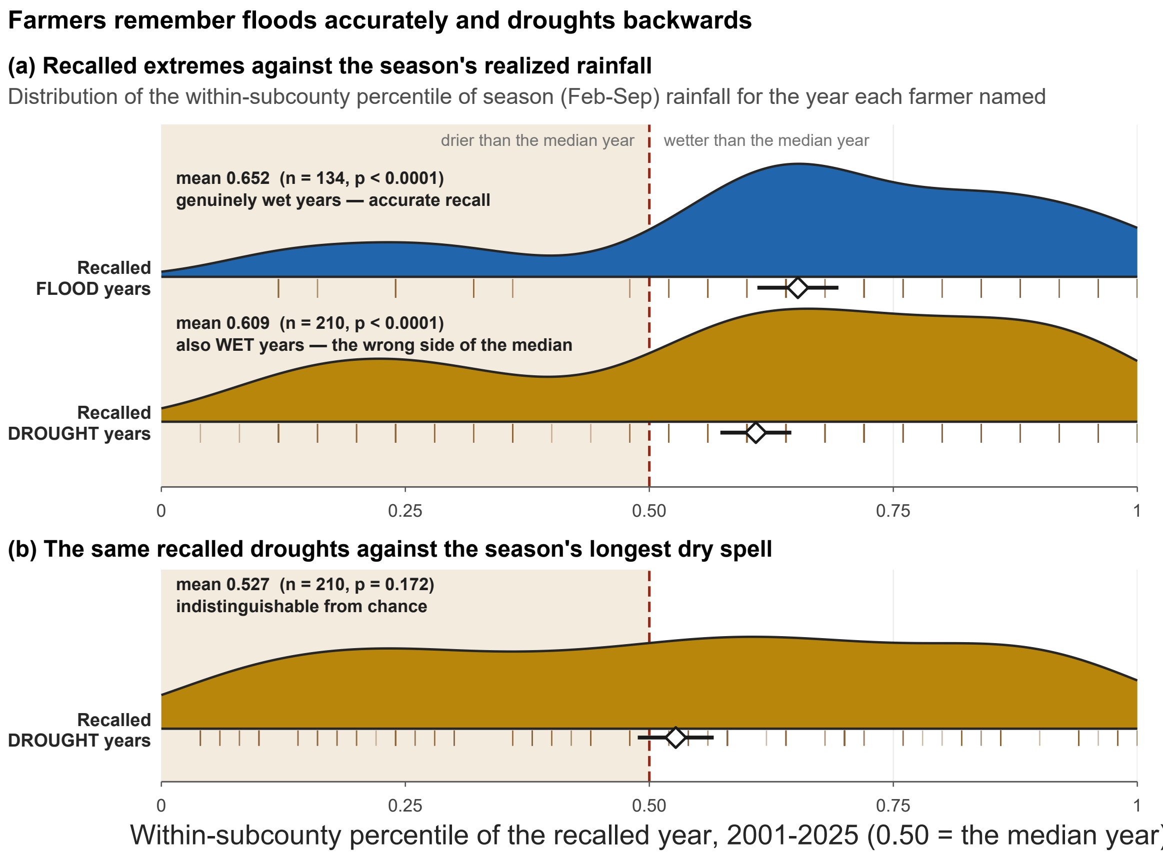

Figure 16: Recalled floods and droughts differ in their correspondence with the historical record.

Notes: For each farmer who freely named a past extreme event and reported its year, the figure locates that year within the 2001–2025 rainfall distribution of the respondent's subcounty. Panel (a) shows that recalled flood years are relatively wet, with a mean seasonal-rainfall percentile of 0.652 ($n = 134$, $p < 0.0001$). Other flood measures show the same pattern, with mean percentiles of 0.613 for heavy-rain days and 0.601 for maximum daily rainfall, both $p < 0.0001$. Years recalled as droughts have a mean seasonal-rainfall percentile of 0.609 ($n = 210$, $p < 0.0001$), placing them on the wetter rather than drier side of the local distribution. Panel (b) shows that recalled drought years are also not unusually dry when measured by longest dry spell, with a mean percentile of 0.527 ($p = 0.172$), or by the number of rain days, with a mean percentile of 0.519 ($p = 0.354$). Some 37.1 percent of recalled drought years fall in the upper dry-spell tercile, compared with 33.3 percent mechanically. Reported drought severity is essentially unrelated to the available measures of realized dryness ($r = 0.002$ and $r = 0.063$), while reported flood severity is positively related to realized rainfall. Diamonds show means with 95 percent confidence intervals and ticks show individual recalled events. The benchmark is measured at subcounty-season resolution and may therefore miss drought conditions that were highly localized or concentrated in a crop-sensitive part of the season.

ratio would have been 0.885, above the pre-defined dry-season threshold of 0.85. No study location could therefore be classified as dry under the fixed seasonal rule.

The season nevertheless contained an unusually long dry spell. The longest observed spell with daily rainfall below the relevant threshold lasted 25 consecutive days. Across 275 historical subcounty-year observations, the previous maximum was 20 days and the median was nine. The 2026 episode remained three days below the 28-day threshold that would

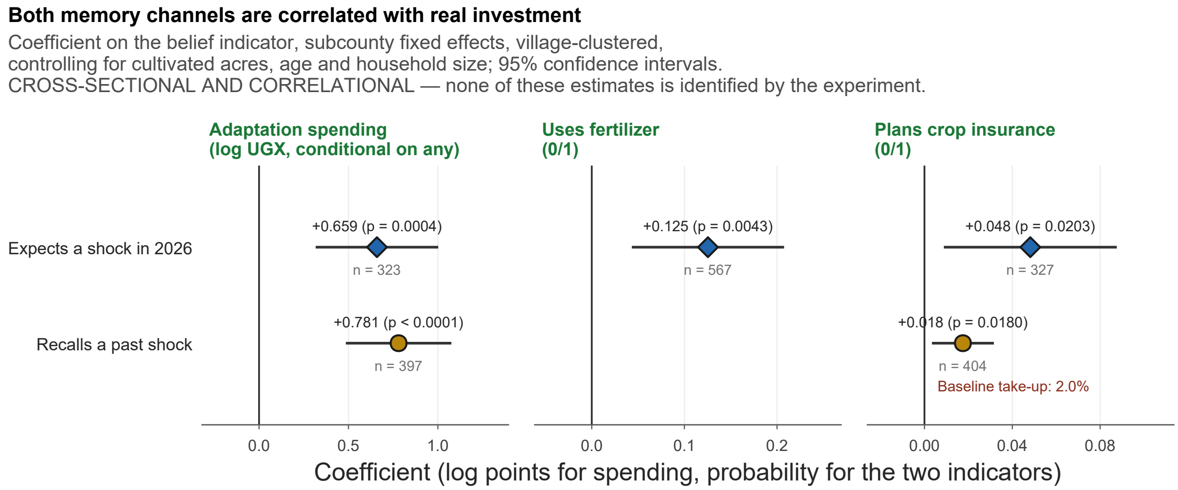

Figure 17: Weather beliefs are associated with planned adaptation and protection.

Notes: The figure reports coefficients on two belief measures, whether the farmer expects an extreme event in the upcoming 2026 season and whether the farmer freely recalled a past extreme event, estimated separately for planned adaptation outcomes. Specifications include subcounty fixed effects and controls for cultivated acreage, age, and household size, with village-clustered standard errors. Whiskers show 95 percent confidence intervals. Expecting an extreme event is associated with 0.659 log points greater planned adaptation spending ($p = 0.0004$, $n = 323$), a 12.5 percentage point higher probability of fertilizer use ($p = 0.0043$, $n = 567$), and a 4.8 percentage point higher probability of planning crop insurance ($p = 0.020$, $n = 327$). Among respondents planning any adaptation, 2.0 percent plan to purchase crop insurance. Recalling a past extreme is associated with 0.781 log points greater adaptation spending ($p < 0.0001$, $n = 397$) and a 1.8 percentage point higher probability of planned insurance ($p = 0.018$). A higher ranked probability score, indicating less accurate recall, is associated with lower planned use of improved seeds (-0.172, $p < 0.0001$) and lower adaptation spending (-0.351, $p = 0.0081$), while greater memory entropy is associated with greater spending ($+1.356$, $p = 0.0004$). These estimates are cross-sectional and descriptive rather than causal.

trigger the slightly-dry override described in Appendix B.

Despite its length, the dry spell is not associated with detectable vegetation damage in the available satellite measures. Peak greenness changes by approximately -0.0031 per additional dry-spell day, a relationship statistically indistinguishable from zero. Across the 17 rainfall grid cells, the rank correlation between dry-spell length and vegetation performance is 0.051 and the relationship is not monotonic. The same relationship is also null over the historical range of dry spells from 5 to 13 days. We therefore retain the pre-specified weather classification rather than recalibrating the dry-spell threshold ex post.

By the time farm outcomes were measured, the vegetation cycle was sufficiently advanced to characterize the season. All 686 farms were at least 30 days beyond their observed vegetation peak. The median peak occurred on 13 May, day 133 of the year, median peak EVI2 was 0.555, and by 5 August median vegetation senescence had covered 57 percent of the decline from peak greenness back toward pre-planting levels.

Farm-level vegetation performance was close to its recent historical average. The median farm's 2026 peak greenness was -0.0012 relative to its own 2021–2025 history, 47.5 percent of farms performed above their own historical benchmark, and the mean within-farm rank of 2026 among the six seasons from 2021 through 2026 was 3.65, compared with 3.5 for an exactly middle-ranked season. The aggregate realization was therefore broadly typical, although the lower tail was somewhat heavier than in an exactly average year.

Treatment estimates. Interactions between treatment assignment and the pre-treatment belief gap have bootstrap p-values between 0.29 and 0.75 across the four farm outcomes. The corresponding confidence intervals for the main Arm B effects are reported in Section D.3. One coefficient involving Arm C reaches conventional significance for raw peak greenness, with an estimate of -0.610, but the relationship disappears for each benchmarked outcome. Given that Arm C contains only 10 villages and the pattern is not robust across outcome definitions, we do not interpret this coefficient substantively.

Reported adaptation intentions also decline between the two survey measurements, but the corresponding pattern in Arm A shows that this decline cannot be attributed to the information treatment. The share reporting plans for any adaptation falls by 8.03 percentage points in Arm A, which receives no information, compared with 5.14 points in Arm B and 8.40 points in Arm C. The fact that the largest decline occurs in the untreated arm is consistent with a repeated-question effect rather than a treatment response and provides another example of why repeated elicitation matters for interpreting within-interview changes.

Why the farm-outcome test has limited power. Farm outcomes are measured using satellite vegetation as a proxy for agricultural productivity, following Burke & Lobell (2017). At the spatial resolution available here, most of the variation in this measure occurs across locations rather than individual farms. In the 2021–2026 farm panel, the CHIRPS rainfall cell explains 63.4 percent of the variation in peak greenness, village explains 6.2 percent, year explains 2.4 percent, farm identity explains 0.35 percent, and reported crop explains essentially none.

This leaves little independent farm-level variation for a village-randomized intervention to affect. Benchmarking each farm against its own history and then differencing this change from the corresponding village-level historical gain substantially reduces common geographic variation. Under this double-differenced measure, the rainfall-cell share falls from 0.92 to 0.005. The remaining variation is correspondingly small, with a standard deviation of 0.0116 compared with 0.0401 for the simpler own-history difference. The farm-outcome nulls should therefore be interpreted together with their confidence intervals rather than as evidence that improved beliefs could not affect outcomes in a season with more consequential weather variation.

E Pre-registration and deviations

This appendix summarizes the pre-registration record and deviations between the registered analysis and the paper. The record consists of the AEA RCT Registry entry filed in March 2026 and a more detailed operational pre-analysis plan dated 11 April 2026. A post-harvest addendum dated 5 August 2026 reconciled the two under an explicit precedence rule before any 2026 outcome data were processed.

What was registered. The AEA registry specifies six hypotheses. H1 concerns systematic divergence between recalled weather and the satellite-derived historical record; H2 links lower memory divergence to more accurate pre-season forecasts; H3 tests whether localized historical weather information moves beliefs toward the objective record and increases their precision; H4 tests effects on end-of-season vegetation; H5 compares belief-relative weather shocks with conventional rainfall anomalies in predicting agronomic outcomes; and H6 tests whether receiving information before elicitation anchors initial beliefs closer to the objective record. The 11 April operational plan defines four confirmatory tests using different numbering: H1 is the registry's H5, comparing a belief-based shock measure with a conventional rainfall anomaly; H2 uses self-reported yield for the same comparison; H3 is an out-of-sample kernel test; and H4 is a confidence-calibration test related to the registry's H3. After the suspension of operational H2 below, the confirmatory family is $\{H1, H3, H4\}$ in the operational-plan numbering. The registry identifier is reported in the Declarations.

Operational H2 is suspended. No post-harvest survey was conducted, and there was no corresponding instrument, fieldwork contract, budget allocation, or ethics coverage extending beyond May 2026. This is consistent with the funded proposal, which anticipated remotely sensed vegetation rather than survey-based yield outcomes. We therefore treat the self-reported-yield hypothesis as infeasible rather than reconstructing an outcome after observing the satellite results.

Operational H1 required a pre-outcome specification change. The registered specification compares a belief-based index with a rainfall z-score while including subcounty fixed effects. Because the rainfall measure is constant within subcounty, it is fully absorbed by those fixed effects. The registered farm-pixel fallback is also uninformative at the available resolution: the 686 farms occupy only 17 CHIRPS cells, with a median of 36 farms per cell, and only five of 11 subcounties contain more than one cell. The 5 August addendum therefore specifies a comparison between the individual belief-based shock measure and the subcounty rainfall anomaly without subcounty fixed effects, with inference clustered over the 60 villages. This change was fixed before the 2026 outcome data were processed.

Operational H3 has 11 independent geographic units. The objective weather classification varies at subcounty rather than parish level. Although the registered comparison refers to 42 parishes, the weather side contains only 11 independent geographic units.

Treating parishes as independent would therefore overstate precision. The paper uses the subcounty as the relevant level of independent variation whenever inference depends on the objective weather record; Appendix C.6 provides the corresponding effective sample sizes and alternative inference.

Status of the remaining analyses. Analyses in Sections 4 through 6 that fall outside the confirmatory family are labeled exploratory unless a registered hypothesis is explicitly identified. The treatment-effect family in the AEA registry is reported in full, but the farm-outcome results are interpreted cautiously because the main information arm contains only 26 treated villages and is underpowered relative to the registered minimum detectable effect. We therefore report these estimates with confidence intervals and power calculations rather than interpreting null significance as evidence of no effect. A complete record of intermediate analyses and numerical reconciliations between the registered documents, subsequent analysis plans, and final specifications is provided in the Online Appendix.

F Variable definitions

This appendix defines the variables used in the paper, grouped by the object they measure.

F.1 Belief variables

Variable	Definition
Recalled distribution	Allocation of ten tokens across the nine ordered weather pictures for a named past season from 2021 through 2025. Bin 1 represents extreme drought, bin 5 the center of the scale, and bin 9 severe flooding.
Mean recalled bin	Token-weighted mean of the recalled distribution, $$\bar{r}_{it} = \sum_b b \frac{n_{itb}}{10},$$ where n_{itb} is the number of tokens respondent i assigns to bin b for season t.
Forward belief	The same nine-picture elicitation for the upcoming 2026 season. It is elicited before the information treatment in Arms A and B and after the information treatment in Arm C.
Mean forward bin	Token-weighted mean of the forward-belief distribution.
Memory entropy	Shannon entropy of the recalled token allocation, averaged across the seasons recalled by the respondent. Higher values indicate a more diffuse recalled distribution.
Sharpness	Concentration of the recalled distribution, defined as the complement of entropy under the normalization used in the paper. Higher values indicate a more concentrated allocation.
Probability assigned to the realized bin	Share of the respondent's ten tokens assigned to the objectively classified weather bin for that season.
Local concentration	Share of tokens assigned to the realized weather bin or one of its two immediately adjacent bins.
Memory accuracy	Ranked probability score of the recalled distribution relative to the realized weather bin. Lower values indicate more accurate recall, so coefficient signs are interpreted accordingly.
Signed bias	Mean recalled bin minus the realized bin. Negative values indicate recall that is drier than the objective record, while positive values indicate recall that is wetter.

continued from previous page

Variable	**Definition**
Recall distance	Number of years between the survey year and the season being recalled.
Onset belief	Token-weighted mean of the ten planting-onset intervals, expressed as day of year using the midpoints of the fielded calendar. The calendar runs from 15 February through 24 April, with midpoints at days 49, 56, 62.5, 69, 76, 83, 90, 97, 104, and 111. Appendix B.4 describes the construction in detail.
Recalled extreme event	Indicator and associated characteristics of an extreme weather event freely recalled by the respondent, including hazard type, year, and reported severity. Respondents are not prompted with particular years.
Expects a future extreme	Indicator equal to one if the respondent reports expecting any extreme weather event during the upcoming 2026 season.
Confidence	Self-reported confidence in the elicited weather expectation, measured on the survey Likert scale.

F.2 Objective weather and outcome variables

Variable	**Definition**
Realized weather bin	Classification of each season into the same nine-picture scale used in the belief elicitation, based on CHIRPS rainfall and satellite vegetation. The classification varies at the subcounty-year level, yielding 55 independent cells across 11 subcounties and five seasons. Appendix B.1 gives the complete classification rule.
Season rainfall	Total CHIRPS rainfall over the defined growing-season window.
Heavy-rain days	Number of days during the season with at least 30 mm of rainfall.
Maximum daily rainfall	Largest single-day rainfall total during the season.
Longest dry spell	Longest sequence of consecutive days during the season with rainfall below 1 mm. The same threshold enters the dry-spell rule used in the weather classification.
Within-subcounty percentile	Rank of a season's weather measure within the corresponding subcounty's 2001–2025 historical distribution, rescaled to the unit interval.
Rainfall anomaly (z)	Seasonal rainfall relative to the subcounty's historical mean, expressed in standard-deviation units.

continued from previous page

Variable	Definition
Peak EVI2	Maximum two-band enhanced vegetation index during the relevant growing period, measured within a 110-meter buffer around the respondent's recorded farm location.
Own-history benchmark	Peak EVI2 relative to the same farm's peak-EVI2 history over 2021–2025. Construction requires at least three usable historical seasons.
Double-differenced benchmark	Change in a farm's peak EVI2 relative to its own history, net of the corresponding village mean change. This transformation reduces the share of variance attributable to the rainfall grid cell from 0.92 to 0.005. Its standard deviation is 0.0116, compared with 0.0401 for the unadjusted own-history difference.

NDVI is not used in any classification threshold, outcome, or reported result in the paper. Some values from the available farm-level NDVI extraction exceed one because of instability in the underlying ratio calculation. We therefore do not report or interpret NDVI-based quantities. EVI2, which is used for the farm-outcome analysis, is unaffected.

F.3 Adaptation and investment variables

Variable	Definition
Adaptation spending	Log planned expenditure on adaptation in Ugandan shillings. The amount is observed only among respondents who report planning some form of adaptation, so spending regressions condition on this group.
Uses fertilizer	Indicator for planned fertilizer use. The sample mean is 57.3 percent.
Uses improved seeds	Indicator for planned use of improved seeds. The sample mean is 52.6 percent.
Uses irrigation	Indicator for planned irrigation use. The sample mean is 22.6 percent.
Plans insurance	Indicator for planning to purchase crop insurance for the upcoming season. Among respondents planning any adaptation, the sample mean is 2.0 percent. The survey does not establish whether an index or indemnity insurance product was commercially available to each respondent in the study area during 2026.

F.4 Treatment and updating variables

Variable	Definition
Arm A	Control arm. Respondents complete the same forward-belief elicitation twice with no information treatment between the two measurements. The arm contains 274 respondents in 24 villages.
Arm B	Information arm. Respondents complete the forward-belief elicitation, receive the historical weather information treatment, and then complete the forward elicitation again. The arm contains 293 respondents in 26 villages.
Arm C	Post-only information arm. Respondents receive the historical weather information before completing a single forward-belief elicitation. No pre-treatment belief is observed. The arm contains 119 respondents in 10 villages.
Earth-mover distance to the historical record	Wasserstein distance between the respondent's elicited forward distribution and the weather distribution summarized by the information card. Lower values indicate greater similarity to the historical record shown on the card.
Ranked probability score against realized 2025 weather	Ranked probability score of the elicited forward distribution evaluated against the realized 2025 weather bin. Lower values indicate a distribution closer to the realized benchmark.
Accuracy gain	Pre-treatment distance or scoring-rule value minus its post-treatment value. Positive values therefore indicate movement toward the relevant historical benchmark.
Absolute prior gap	Absolute distance between the mean of the information-card distribution and the respondent's mean pre-treatment belief. Because the initial elicitation enters both this gap and any subsequent change score, regressions using this quantity are subject to the mechanical relationship discussed in Appendix C.2.
Prior onset tercile	Tercile of the respondent's pre-treatment expected planting onset. Mean initial beliefs are days 51.7, 63.7, and 80.4 in the early, middle, and late terciles, respectively. All lie within the feasible support of the fielded calendar, $[49, 111]$; see Appendix B.4.

Online Appendix. The accompanying Online Appendix contains the survey instrument and elicitation materials, treatment materials, additional field documentation, supplementary analyses, pre-registration records, and computational and replication details.

Online Appendix

The Online Appendix collects the documentation of the study: the study site in detail, the survey instrument item by item, the enumerator script for the two token elicitations, the complete enumerator manual as fielded, the consent form, the fieldwork protocol, the pilot and what it changed, and the implementation detail behind the inference procedures. The material a reader needs to evaluate the paper's claims is in Appendices A–F above.

OA-1 Study setting

The study covers 60 villages in 42 parishes and 11 subcounties across Bugiri, Butaleja, and Mayuge districts in eastern Uganda. The study subcounties are Buwunga, Kapyanga, Muterere, and Nankoma in Bugiri; Budumba, Kachonga, and Mazimasa in Butaleja; and Bukatube, Kigandalo, Kityerera, and Mpungwe in Mayuge. As Appendix A of the paper explains, exact village coordinates were not collected, so villages are represented geographically by their parish centroids. These centroids span 0.332 to 0.993 degrees north and 33.363 to 34.081 degrees east. Mean elevation within a 500-meter buffer ranges from 1,076 to 1,248 meters, with a mean of 1,147 meters, placing the study area on the Lake Victoria basin plateau rather than on the slopes of Mount Elgon. This distinction is relevant because the 2025 pilot was conducted in Bulambuli district on the Elgon slopes; Section OA-7 describes the pilot and the changes made before the main study.

Rainfall is bimodal. The survey focuses on the first agricultural season, described in the instrument as running approximately from February through July. The objective-weather construction described in Appendix B of the paper allows the first-season growing window to extend through 30 September depending on detected rainfall onset. The second agricultural season, which runs approximately from August through December, is not elicited. Over the 25 seasons from 2001 through 2025, CHIRPS seasonal rainfall averages range from 697.8 to 730.3 millimeters across the 11 subcounties, with within-subcounty coefficients of variation between 0.113 and 0.173 (Funk et al., 2015). Across the five seasons displayed in the information treatment, subcounty rainfall totals range from 610.6 to 885.3 millimeters. Long-run mean rainfall differs by only 32.5 millimeters across the 11 subcounties, consistent with the relatively limited spatial variation in rainfall documented in the paper.

Agriculture in the study area is predominantly smallholder and rainfed. All 686 respondents identify farming as their occupation. Among crops planned for the upcoming season, 95.5 percent of respondents name maize, 77.3 percent beans, 52.0 percent cassava, 45.5 percent groundnuts, 15.3 percent soybeans, 13.0 percent rice, and 10.8 percent sweet potatoes; no other crop is named by more than seven percent of respondents. Mean respondent age is 41.4 years, 44.3 percent of respondents are women, mean household size is 7.75, and 350 of 684 respondents with nonmissing literacy responses report being able to read and write. Mean cultivated area is 2.79 acres and the median is two acres. Of the 686 respondents, 269 report no formal education, 281 primary education, 119 secondary education, 13 tertiary education, and four vocational training.

District-level statistics describe a similarly agriculture-dependent setting. In the 2014 National Population and Housing Census, 88.8 percent of households in Bugiri, 95.3 percent in Butaleja, and 83.7 percent in Mayuge report growing crops. Subsistence farming is reported as the main livelihood by 80.8, 89.8, and 74.2 percent of households in the three districts, respectively, compared with 69 percent nationally (Uganda Bureau of Statistics, 2017). The Annual Agricultural Survey reports that the average agricultural household nationally operates 1.3 hectares over 2.1 parcels, while mean parcel size is 0.47 hectares in Busoga, which includes Bugiri and Mayuge, and 0.37 hectares in Bukedi, which includes Butaleja (Uganda Bureau of Statistics, 2022). The bimodal rainfall regime and irrigated rice production in the wider region are documented by FEWS NET (FEWS NET, 2010). The Mayuge District Development Plan reports annual rainfall ranging from approximately 900 millimeters in the north to 1,200 millimeters in the south (Mayuge District Local Government, 2015).

OA-2 Survey instrument

This section reproduces the pre-season survey instrument in field order. The instrument contains 328 numbered entries, comprising 311 respondent-facing questions, 15 stored calculations, and two enumerator-facing notes. The entries are organized into 19 modules following the structure of the fielded questionnaire. For each item, we report the fielded variable name, question wording, response type, answer options, routing conditions, and relevant validation constraints. Question wording and answer options are reproduced as fielded rather than edited for presentation.

Language. The machine-readable questionnaire contains a single English-language label for each item and no stored written translations. Interviews were administered face to face. Where respondents did not use English, enumerators translated orally during administration. Because no written translation or back-translation is preserved in the study records, we reproduce only the English fielded instrument.

Interview timing. The fielded questionnaire does not retain interview start and end timestamps, so realized interview duration cannot be recovered from the available survey data.

Experimental routing. Experimental arm is assigned at the village level and fixed by the village selected at the beginning of the questionnaire. Arms A and B complete the pre-information forward-belief modules, while Arm C is routed directly to the information treatment and subsequent forward elicitation. All three arms complete the final forward elicitation. Arm A therefore repeats the same elicitation without receiving information in between, providing the repeated-elicitation benchmark used in the experimental analysis.

Token elicitations. Two visual elicitation instruments recur throughout the questionnaire. The seasonal-weather exercise asks respondents to allocate ten physical tokens across nine ordered pictures ranging from extreme drought to severe flooding. The planting-onset exercise asks respondents to allocate ten tokens across ten dated intervals. Section OA-3 reproduces the standardized instructions and enumerator protocol for both exercises.

Reading the arm gating. Assignment is at the village level and is computed on the device by the `calculate` field `arm` (Module OA-2.6), which maps each of the 60 sampled villages to one of three arms: 24 villages to Arm A, 26 to Arm B, and 10 to Arm C. Two features of the ordering carry most of the identification weight in this paper, and both are properties of the form itself rather than of the analysis.

First, the pre-card adaptation block `adapt_group` carries the relevance condition `${arm} = 'A' or ${arm} = 'B' or ${arm} = 'C'`, which is exhaustive over the three arms, and it sits at form position 324–329, ahead of the information block `info_treatment_group` at position 336. Every respondent in every arm therefore answers the pre-card adaptation questions before any information is shown, and `adapt_any_pre` is defined for all three arms.

Second, the pre-card belief blocks `belief_amount_future` and `onset_future_group` carry the relevance condition `${arm} = 'A' or ${arm} = 'B'`. Arm C respondents are routed past them entirely, so Arm C has no pre-card belief elicitation at all and `fut_pre_total_tokens` is zero for every one of the 119 Arm C respondents. Arm C's single forward elicitation is the `post_amount` and `post_onset_group` block, which is administered after the card.

The second forward elicitation itself, `post_amount` and `post_onset_group`, carries *no* relevance condition. All three arms answer it. For Arm A it is a second reading of an unchanged instrument with nothing in between, which is what makes Arm A a within-respondent placebo for the elicitation itself rather than a pure no-contact control.

Table OA1: The `village_list` choice list and the village-to-arm assignment encoded in the `arm` calculate field

Option code	Option label (subcounty, district)	Arm
`bubalule_b`	Bubalule B (Kityerera, Mayuge)	C
`bubinge`	Bubinge (Kityerera, Mayuge)	A
`bufuta_a`	Bufuta A (Bukatube, Mayuge)	A
`bugalo`	Bugalo (Mazimasa, Butaleja)	B
`bugodo_b`	Bugodo B (Kapyanga, Bugiri)	C
`bugolola_a`	Bugolola A (Budumba, Butaleja)	B
`buguluma`	Buguluma (Bukatube, Mayuge)	A
`buhwaya`	Buhwaya (Budumba, Butaleja)	A
`bukerewe`	Bukerewe (Buwunga, Bugiri)	A
`bulomi`	Bulomi (Mazimasa, Butaleja)	B
`bulumba`	Bulumba (Kachonga, Butaleja)	B
`bunghanga`	Bunghanga (Budumba, Butaleja)	A
`busimo`	Busimo (Kityerera, Mayuge)	B
`busolo`	Busolo (Kapyanga, Bugiri)	B
`buswikira`	Buswikira (Mpungwe, Mayuge)	B
`butegwa_b`	Butegwa B (Nankoma, Bugiri)	B
`buwembula`	Buwembula (Muterere, Bugiri)	A
`buwuge`	Buwuge (Buwunga, Bugiri)	B
`bwalula_a`	Bwalula A (Nankoma, Bugiri)	A
`dumbu`	Dumbu (Budumba, Butaleja)	C
`kasokwe`	Kasokwe (Mpungwe, Mayuge)	A
`kazinga`	Kazinga (Kigandalo, Mayuge)	B
`kigulamo`	Kigulamo (Kigandalo, Mayuge)	A
`kikandwa`	Kikandwa (Kityerera, Mayuge)	A
`kiteigalwa`	Kiteigalwa (Buwunga, Bugiri)	A

Table OA1, continued

Option code	Option label (subcounty, district)	Arm
kyaluya	Kyaluya (Buwunga, Bugiri)	C
kyando	Kyando (Bukatube, Mayuge)	B
kyete	Kyete (Mpungwe, Mayuge)	B
lubanga	Lubanga (Kachonga, Butaleja)	B
lubanyi	Lubanyi (Muterere, Bugiri)	A
lubembe	Lubembe (Mazimasa, Butaleja)	A
lukunhu	Lukunhu (Kigandalo, Mayuge)	A
lwaba	Lwaba (Kapyanga, Bugiri)	A
maleka	Maleka (Kigandalo, Mayuge)	B
matiki	Matiki (Nankoma, Bugiri)	B
mauta	Mauta (Bukatube, Mayuge)	B
mawanga	Mawanga (Buwunga, Bugiri)	B
mitimito	Mitimito (Kityerera, Mayuge)	B
mpumu	Mpumu (Mpungwe, Mayuge)	C
muhula_kagondo	Muhula Kagondo (Kachonga, Butaleja)	A
musima	Musima (Mpungwe, Mayuge)	B
mwezi	Mwezi (Mpungwe, Mayuge)	A
nabuyanja	Nabuyanja (Budumba, Butaleja)	A
nahanja	Nahanja (Mazimasa, Butaleja)	A
nahigogola	Nahigogola (Kachonga, Butaleja)	B
nakabale	Nakabale (Mazimasa, Butaleja)	A
nakalanga	Nakalanga (Bukatube, Mayuge)	C
nakavule	Nakavule (Kapyanga, Bugiri)	B
namalenha	Namalenha (Buwunga, Bugiri)	A
namatooke	Namatooke (Mpungwe, Mayuge)	C
namawa_a	Namawa A (Kachonga, Butaleja)	C
namujugume	Namujugume (Kachonga, Butaleja)	A
namunasa	Namunasa (Mazimasa, Butaleja)	B
napindo	Napindo (Budumba, Butaleja)	B
ndhokero_b	Ndhokero B (Bukatube, Mayuge)	C
ngunga	Ngunga (Muterere, Bugiri)	B
nongo	Nongo (Muterere, Bugiri)	B
nsono	Nsono (Nankoma, Bugiri)	C
nvunwa	Nvunwa (Kigandalo, Mayuge)	A
wandegeire	Wandegeire (Buwunga, Bugiri)	B

Notes: This is the complete `village_list` choice list of the XLSForm, given verbatim: the option code as stored, and the option label as it appeared on the enumerator's screen, with subcounty and district in parentheses. The Arm column is not part of the choice list; it is the value the `calculate` field `arm` assigns to a respondent who selects that village. Arm A: 24 villages. Arm B: 26 villages. Arm C: 10 villages. The mapping is hard-coded in the form and is not re-drawn on the device, so assignment is fixed before the enumerator opens the questionnaire.

The two token instruments. Two elicitation devices recur throughout the instrument and are described once here. The *rainfall amount* exercise asks the respondent to place ten physical tokens across nine ordered pictures of a field, running from image 1 (extreme drought) through image 5 (optimal) to image 9 (severe flood); the nine labels are given verbatim as the question text of the nine token cells in Module OA-2.7. The *onset timing* exercise asks for ten tokens across ten dated intervals of the planting window. The interval edges are printed on the enumerator's calendar card and are not stored in the XLSForm, which labels the cells only as "onset interval 1" through "onset interval 10"; they are reproduced in Table OA2.

Table OA2: Onset intervals shown on the enumerator's calendar card

Interval	Date range	Text printed alongside
1	15–21 Feb	Very dry; land still hard; few early land preparers start ploughing.
2	22–28 Feb	First scattered showers in some low-lying areas; usually still unreliable.
3	1–6 Mar	Land preparation intensifies; some early planters risk maize if showers continue.
4	7–13 Mar	Transitional period; first rains often occur; seedbed preparation peaks.
5	14–20 Mar	Some areas record 1–2 good rains; not yet sustained in the highlands.
6	21–27 Mar	Typical onset for lower slopes; first main planting begins.
7	28 Mar–3 Apr	Onset consolidates; continuous rain spells; widespread maize and bean planting.
8	4–10 Apr	Average onset period for mid-slopes; moisture reliable.
9	11–17 Apr	Late planters begin; upper slopes catch up.
10	18–24 Apr	Late onset years often start here; second wave of planters.

Notes: Reproduced from the enumerator calendar card "FARMER SURVEY CALENDAR: Start of Rains", whose instruction line reads: "Allocate 10 tokens across the intervals to indicate when you remember / expect the seasonal rains started." These edges are not stored in the XLSForm and must be read from the card to interpret `y2025_int1`–`y2021_int10`, `fut_pre_int1`–`fut_pre_int10`, and `fut_post_int1`–`fut_post_int10`. The intervals are not uniform: nine of the ten are exactly seven days wide and the third, 1–6 March, is six, so the calendar spans 15 February to 24 April and nothing outside it, with midpoints at days 49, 56, 62.5, 69, 76, 83, 90, 97, 104 and 111 of the year. The physical card is reproduced as the calendar card reproduced in the paper's design section. Every day-scale magnitude in this draft is computed against those midpoints; Appendix B of the paper documents the construction.

OA-2.1 Identification and location

Arms: All arms. Administered before randomization is revealed to the enumerator.

No.	Variable	Question, type, options, and routing
1	`resp_id` → `respondent_id` *text*	Respondent ID *Required.*
2	`Date` → `interview_date` *datetime*	Date
3	`village_name` → `village_name_label` *select_one village_list*	Village name (select from list) *Required.* *Options:* the 60 study villages. The list is too long to sit in this cell; all 60 option codes and labels, with the arm each village is assigned to, are given in Table OA1. *Appearance:* `minimal`
4	`village_confirm` → *(echo column, dropped)* *select_one yn*	*Dynamic text:* `concat("Is the village you selected: ", jr:choice-name(${village_name}, "${village_name}"), "?")` *Required.* *Options* (code — label as read out), `yn`: `yes` — Yes `no` — No *Constraint:* `. = 'yes'` *Message:* "Please go back and reselect the correct village before continuing"
5	`Village_ID` → `village_id` *integer*	Village ID
6	`district` *text*	District or region

OA-2.2 Demographic information

XLSForm group(s): `demo_group`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `demo_group` — *"Demographic information".* *Block relevance:* none; every respondent who reaches this point answers it.		

continued from previous page

No.	Variable	Question, type, options, and routing
7	`age` *integer*	Age in years *Required.* *Constraint:* `. >= 10 and . <= 99` *Message:* "Please enter an age between 10 and 99"
8	`gender` *select_one gender*	Gender *Required.* *Options* (code — label as read out), `gender`: `male` — Male `female` — Female `other` — Other
9	`marital_status` *select_one marital*	Marital status *Options* (code — label as read out), `marital`: `single` — Single `married` — Married `widowed` — Widowed `divorced` — Divorced
10	`hh_head` *select_one yn*	Are you the head of the household *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
11	`hh_size` *integer*	Household size (including Children) *Required.* *Constraint:* `. > 0` *Message:* "Household size must be greater than zero"
12	`literate` → `literacy` *select_one yn*	Can you read and write *Options:* the 2-item `yn` list, given in full at Q4.
13	`educ_level` → `education` *select_one educ*	Highest level of education attended *Options* (code — label as read out), `educ`: `none` — No formal education `primary` — Primary education `secondary` — Secondary education `vocational` — Vocational training `tertiary` — Tertiary education

OA-2.3 Farming background, land tenure, and decision rights

XLSForm group(s): `Farming_Background`.
Arms: All arms.

No.	Variable	Question, type, options, and routing

Block `Farming_Background` — *"Farming Background"*.
Block relevance: none; every respondent who reaches this point answers it.

14	`farmer_status` → `is_farmer` *select_one yn*	Are you a farmer *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
15	`land_ownership` → `tenure_own, tenure_rent, tenure_sharecrop, tenure_communal, tenure_other` *select_multiple landown*	Which types of land tenure apply to the land you farm (tick all that apply) *Required.* *Options* (code — label as read out), `landown`: `own` — Own `rent` — Rent `sharecrop` — Sharecrop `communal` — Communal `other` — Other
16	`land_own_acres` → `acres_owned` *decimal*	How many acres are owned by your household or close family (Grand-Parents, Parents, Children or Siblings) *Relevance:* Asked only if the option `own` was ticked at `land_ownership`. `selected(${land_ownership}, 'own')` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"
17	`land_rent_acres` → `acres_rented` *decimal*	How many acres are rented *Relevance:* Asked only if the option `rent` was ticked at `land_ownership`. `selected(${land_ownership}, 'rent')` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"
18	`land_sharecrop_acres` → `acres_sharecrop` *decimal*	How many acres are farmed under sharecropping arrangements *Relevance:* Asked only if the option `sharecrop` was ticked at `land_ownership`. `selected(${land_ownership}, 'sharecrop')` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"
19	`land_communal_acres` → `acres_communal` *decimal*	How many acres are communal land that you farm *Relevance:* Asked only if the option `communal` was ticked at `land_ownership`. `selected(${land_ownership}, 'communal')` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"

continued from previous page

No.	Variable	Question, type, options, and routing
20	`land_other_acres` → `acres_other` *decimal*	How many acres are under other land tenure arrangements *Relevance:* Asked only if the option `other` was ticked at `land_ownership`. `selected(${land_ownership}, 'other')` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"
21	`decision_role` *select_one yn*	Are you involved in farm related decision making *Options:* the 2-item `yn` list, given in full at Q4.
22	`top_decision_1` *select_one dec_parties*	Most important decision maker for crop and planting decisions *Required.* *Options* (code — label as read out), `dec_parties`: `Individual` — Survey Taker (Individual); `aid` — Aid organizations `hh_head` — Household head (if other than individual); `spouse` — Spouse `elders` — Elders; `other_dm` — Other Decision Maker `parents` — Parents; `na_none` — NA / No one `neighbors` — Neighbors; `other_dm_2` — Other Decision Maker `chief` — Village official or chief; `other_dm_3` — Other Decision Maker `official_info` — Official information source; `na_none_2` — NA / No one *Displayed subset:* the nine substantive parties plus `other_dm`; the "NA / No one" codes are hidden here.
23	`top_decision_1_other` → `other_decision_1` *text*	Specify other decision maker (first) *Relevance:* Asked only if "Other Decision Maker" (`other_dm`) was chosen at `top_decision_1`. `${top_decision_1} = 'other_dm'`
24	`top_decision_2` *select_one dec_parties*	Second most important decision maker *Required.* *Options:* the 14-item `dec_parties` list, given in full at Q22. *Displayed subset:* the nine substantive parties plus `na_none` and `other_dm_2`. *Constraint:* `. != ${top_decision_1}` *Message:* "Please select a different decision maker from the first choice"

continued from previous page

No.	Variable	Question, type, options, and routing
25	`top_decision_2_other` → `other_decision_2` *text*	Specify other decision maker (second) *Relevance:* Asked only if "Other Decision Maker" (`other_dm_2`) was chosen at `top_decision_2`. `${top_decision_2} = 'other_dm_2'`
26	`top_decision_3` *select_one dec_parties*	Third most important decision maker *Required.* *Options:* the 14-item `dec_parties` list, given in full at Q22. *Displayed subset:* the nine substantive parties plus `other_dm_3` and `na_none_2`. *Constraint:* `. != ${top_decision_1} and . != ${top_decision_2}` *Message:* "Please select a different decision maker from the first two choices"
27	`top_decision_3_other` → `other_decision_3` *text*	Specify other decision maker (third) *Relevance:* Asked only if "Other Decision Maker" (`other_dm_3`) was chosen at `top_decision_3`. `${top_decision_3} = 'other_dm_3'`
28	`parents_farmers` *select_one yn*	Were your parents farmers *Options:* the 2-item `yn` list, given in full at Q4.
29	`grandparents_farmers` *select_one yn*	Were your grandparents farmers *Options:* the 2-item `yn` list, given in full at Q4.
30	`farm_size_acres` → `total_cultivated_acres` *decimal*	How much land does your household cultivate in total in acres *Required.* *Constraint:* `. > 0` *Message:* "Total cultivated area must be greater than zero"
31	`owned_area_acres` → `owned_cultivated_acres` *decimal*	How much of this cultivated land is owned by your household or direct relatives *Constraint:* `. >= 0 and . <= ${farm_size_acres}` *Message:* "Value must be between 0 and total cultivated acres"
32	`irrigation_access` *select_one yn*	Is there a water (irrigation) source (rivers, wells, pond, canals or similar sources) near or at your property *Options:* the 2-item `yn` list, given in full at Q4.
33	`irrigation_use` *select_one yn*	Do you use irrigation *Options:* the 2-item `yn` list, given in full at Q4. *Relevance:* Asked only if `irrigation_access` was answered Yes. `${irrigation_access} = 'yes'`

continued from previous page

No.	Variable	Question, type, options, and routing
34	`irrigated_area_acres` → `irrigated_acres` *decimal*	How many acres of your cultivated land are irrigated *Relevance:* Asked only if `irrigation_use` was answered Yes. `${irrigation_use} = 'yes'` *Constraint:* `. >= 0 and . <= ${farm_size_acres}` *Message:* "Irrigated acres must not exceed total cultivated acres"

OA-2.4 Crop production and agricultural practices

XLSForm group(s): `crop_practices_group`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `crop_practices_group` — *"Crop production and agricultural practices"*. *Block relevance:* none; every respondent who reaches this point answers it.		
35	`planned_crops` → `crop_maize ... crop_other (22 dummies)` *select_multiple crop_list*	Which crops do you plan to grow this season *Required.* *Options* (code — label as read out), `crop_list`: `maize` — Maize; `soybeans` — Soybeans `beans` — Beans; `sunflower` — Sunflower `cassava` — Cassava; `coffee` — Coffee `sweet_potatoes` — Sweet potatoes; `tomatoes` — Tomatoes `groundnuts` — Groundnuts; `onions` — Onions `sorghum` — Sorghum; `cabbages` — Cabbages `millet` — Millet; `pumpkin` — Pumpkin `rice` — Rice; `vegetables` — Mixed vegetables `bananas` — Bananas or plantains; `sugar_cane` — Sugar Cane `simsim` — Sesame Simsım; `coconut` — Coconut `cowpeas` — Cowpeas; `other` — Other
36	`planned_crops_other` → `crop_other_specify` *text*	Please specify the other crops *Relevance:* Asked only if the option `other` was ticked at `planned_crops`. `selected(${planned_crops}, 'other')`

continued from previous page

No.	Variable	Question, type, options, and routing
37	`top_crop1` *select_one crop_list*	Biggest crop by area *Required.* *Options:* the 22-item `crop_list` list, given in full at Q35. *Displayed subset:* restricted to the crops ticked at `planned_crops`.
38	`top_crop2` *select_one crop_list*	Second biggest crop by area *Required.* *Options:* the 22-item `crop_list` list, given in full at Q35. *Displayed subset:* restricted to the crops ticked at `planned_crops`. *Constraint:* `. != ${top_crop1}` *Message:* "Please select a different crop from the first choice"
39	`top_crop3` *select_one crop_list*	Third biggest crop by area *Required.* *Options:* the 22-item `crop_list` list, given in full at Q35. *Displayed subset:* restricted to the crops ticked at `planned_crops`. *Constraint:* `${top_crop1} != ${top_crop2} and ${top_crop1} != ${top_crop3} and ${top_crop2} != ${top_crop3}` *Message:* "Please select three different crops"

Block `crop1_details` — *"Details for your biggest crop"*.
Block relevance: Asked only if a crop was named at `top_crop1`. `${top_crop1} != ''`

No.	Variable	Question, type, options, and routing
40	`crop1_area` *decimal*	How many acres will you plant with ${top_crop1}? *Required.* *Constraint:* `. >= 0 and . <= ${farm_size_acres}` *Message:* "Must not exceed total cultivated acres"
41	`crop1_plant_weeks` *→ (not renamed; 11 week dummies in the export)* *select_multiple plant_weeks*	In which weeks between 1 February and 15 April do you plan to plant ${top_crop1}? (Select at most two.) *Options* (code — label as read out), `plant_weeks`: `w1` — 1 to 7 February `w2` — 8 to 14 February `w3` — 15 to 21 February `w4` — 22 to 28 February `w5` — 1 to 7 March `w6` — 8 to 14 March `w7` — 15 to 21 March `w8` — 22 to 28 March `w9` — 29 March to 4 April `w10` — 5 to 11 April `w11` — 12 to 18 April *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"

continued from previous page

No.	Variable	Question, type, options, and routing
42	`crop1_harvest_weeks` → `crop1_harvest_weeks_label` *select_multiple harvest_weeks*	In which weeks between 1 June and 31 August do you expect to harvest ${top_crop1}? (Select at most two.) *Options* (code — label as read out), `harvest_weeks`: `h1` — 1 to 7 June; `h7` — 15 to 21 July `h2` — 8 to 14 June; `h8` — 22 to 31 July `h3` — 15 to 21 June; `h9` — 1 to 7 August `h4` — 22 to 30 June; `h10` — 8 to 14 August `h5` — 1 to 7 July; `h11` — 15 to 21 August `h6` — 8 to 14 July; `h12` — 22 to 31 August *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
43	`crop1_critical_months` → `crop1_critical_months_label` *select_multiple month_list*	Months when rainfall is most critical for ${top_crop1} (Select at most two.) *Options* (code — label as read out), `month_list`: `jan` — January; `jul` — July `feb` — February; `aug` — August `mar` — March; `sep` — September `apr` — April; `oct` — October `may` — May; `nov` — November `jun` — June; `dec` — December *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
44	`irrigated_area_acres_crop1` → `crop1_irrigated_acres` *decimal*	How many acres of your cultivated land are irrigated for ${top_crop1}? *Relevance:* Asked only if `irrigation_use` was answered Yes. `${irrigation_use} = 'yes'` *Constraint:* `. >= 0 and . <= ${crop1_area}` *Message:* "Irrigated acres must not exceed cultivated acres for this crop"

Block `crop2_details` — *"Details for your second biggest crop"*.
Block relevance: Asked only if a crop was named at `top_crop2`. `${top_crop2} != ''`

45	`crop2_area` *decimal*	How many acres will you plant with ${top_crop2}? *Required.* *Constraint:* `. >= 0 and . + ${crop1_area} <= ${farm_size_acres}` *Message:* "Sum of 1st and 2nd crop areas must not exceed total cultivated acres"

continued from previous page

No.	Variable	Question, type, options, and routing
46	`crop2_plant_weeks` → *(not renamed)* *select_multiple plant_weeks*	In which weeks between 1 February and 15 April do you plan to plant ${top_crop2}? (Select at most two.) *Options:* the 11-item `plant_weeks` list, given in full at Q41. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
47	`crop2_harvest_weeks` → *(not renamed)* *select_multiple harvest_weeks*	In which weeks between 1 June and 31 August do you expect to harvest ${top_crop2}? (Select at most two.) *Options:* the 12-item `harvest_weeks` list, given in full at Q42. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
48	`crop2_critical_months` → *(not renamed)* *select_multiple month_list*	Months when rainfall is most critical for ${top_crop2} (Select at most two.) *Options:* the 12-item `month_list` list, given in full at Q43. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
49	`irrigated_area_acres_crop2` → `crop2_irrigated_acres` *decimal*	How many acres of your cultivated land are irrigated for ${top_crop2}? *Relevance:* Asked only if `irrigation_use` was answered Yes. `${irrigation_use} = 'yes'` *Constraint:* `. >= 0 and . <= ${crop2_area}` *Message:* "Irrigated acres must not exceed cultivated acres for this crop"

Block `crop3_details` — *"Details for your third biggest crop"*.
Block relevance: Asked only if a crop was named at `top_crop3`. `${top_crop3} != ''`

No.	Variable	Question, type, options, and routing
50	`crop3_area` *decimal*	How many acres will you plant with ${top_crop3}? *Required.* *Constraint:* `. >= 0 and . + ${crop1_area} + ${crop2_area} <= ${farm_size_acres}` *Message:* "Sum of all three crop areas must not exceed total cultivated acres"
51	`crop3_plant_weeks` → *(not renamed)* *select_multiple plant_weeks*	In which weeks between 1 February and 15 April do you plan to plant ${top_crop3}? (Select at most two.) *Options:* the 11-item `plant_weeks` list, given in full at Q41. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"

continued from previous page

No.	Variable	Question, type, options, and routing
52	`crop3_harvest_weeks` → *(not renamed)* *select_multiple harvest_weeks*	In which weeks between 1 June and 31 August do you expect to harvest ${top_crop3}? (Select at most two.) *Options:* the 12-item `harvest_weeks` list, given in full at Q42. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
53	`crop3_critical_months` → *(not renamed)* *select_multiple month_list*	Months when rainfall is most critical for ${top_crop3} (Select at most two.) *Options:* the 12-item `month_list` list, given in full at Q43. *Constraint:* `count-selected(.) >= 1 and count-selected(.) <= 2` *Message:* "Please select one or two weeks"
54	`irrigated_area_acres_crop3` → `crop3_irrigated_acres` *decimal*	How many acres of your cultivated land are irrigated for ${top_crop3}? *Relevance:* Asked only if `irrigation_use` was answered Yes. `${irrigation_use} = 'yes'` *Constraint:* `. >= 0 and . <= ${crop3_area}` *Message:* "Irrigated acres must not exceed cultivated acres for this crop"
55	`farm_practices_used` → `prac_mechanized`, `prac_fertilizer`, `prac_pest_control`, `prac_improved_seeds`, `prac_crop_rotation`, `prac_other` *select_multiple farm_practices*	Does your household use any of the following for farming *Options* (code — label as read out), `farm_practices`: `mechanized` — Mechanized equipment electric or petrol powered `improved_seeds` — Use of improved seeds `fertilizer` — Use of fertilizers `rotation` — Crop rotation `pest_control` — Pest control methods `other` — Other
56	`farm_practices_other` → `prac_other_specify` *text*	Please specify other farming practices *Relevance:* Asked only if the option `other` was ticked at `farm_practices_used`. `selected(${farm_practices_used}, 'other')`

OA-2.5 Access to agricultural information and forecasts

XLSForm group(s): `info_access_group`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `info_access_group` — *"Access to agricultural information"*. *Block relevance:* none; every respondent who reaches this point answers it.		
57	`ext_access` → `extension_access` *select_one yn*	Do you have access to agricultural extension services *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
58	`weather_info_any` *select_one yn*	Do you receive any information on past present or future weather *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
59	`weather_info_level` *select_one info_level*	At what level is this weather information provided (If you receive information at multiple levels, select the most detailed one, e.g. select "District" rather than "National" if you receive both.) *Options* (code — label as read out), `info_level`: `national` — National `regional` — Regional `district` — District `village` — Village `other` — Other *Relevance:* Asked only if `weather_info_any` was answered Yes. `${weather_info_any} = 'yes'`
60	`weather_info_level_other` *text*	Please specify the other level *Relevance:* Asked only if "Other" was chosen at `weather_info_level`. `${weather_info_level} = 'other'`
61	`weather_info_specific` → `weather_info_subcounty_specific` *select_one yn*	Is this weather information specific to your location (sub county level or below) *Options:* the 2-item `yn` list, given in full at Q4. *Relevance:* Asked only if `weather_info_any` was answered Yes. `${weather_info_any} = 'yes'`
62	`weather_info_past` *select_one yn*	Is this information about past weather *Options:* the 2-item `yn` list, given in full at Q4. *Relevance:* Asked only if `weather_info_any` was answered Yes. `${weather_info_any} = 'yes'`

continued from previous page

No.	Variable	Question, type, options, and routing
63	`weather_info_past_seasons` → `weather_info_n_seasons` *integer*	How many past seasons are reported *Relevance:* Asked only if `weather_info_past` was answered Yes. `${weather_info_past} = 'yes'` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"
64	`weather_info_content` → `info_content_rain, info_content_temp, info_content_soil, info_content_storms, info_content_other` *select_multiple info_content*	What information is reported *Options* (code — label as read out), `info_content`: `rain` — Rainfall `temp` — Temperature `soil` — Soil moisture `storms` — Storms `other` — Other *Relevance:* Asked only if `weather_info_any` was answered Yes. `${weather_info_any} = 'yes'`
65	`weather_info_content_other` → `info_content_other_specify` *text*	Please specify the other reported information *Relevance:* Asked only if the option `other` was ticked at `weather_info_content`. `selected(${weather_info_content}, 'other')`
66	`weather_info_sources` → `src_past_radio, src_past_tv, src_past_mobile, src_past_meetings, src_past_extension, src_past_neighbors, src_past_internet, src_past_other` *select_multiple info_source*	What are your main sources of weather information *Options* (code — label as read out), `info_source`: `radio` — Radio `tv` — Television `phone` — Mobile phone alerts `meetings` — Community meetings `extension` — Agricultural extension agents `neighbors` — Neighbors or other farmers `internet` — Internet via smartphone or device `other` — Other *Relevance:* Asked only if `weather_info_any` was answered Yes. `${weather_info_any} = 'yes'`
67	`weather_info_sources_other` → `src_past_other_specify` *text*	Please specify the other weather information source *Relevance:* Asked only if the option `other` was ticked at `weather_info_sources`. `selected(${weather_info_sources}, 'other')`
68	`forecast_any` *select_one yn*	Do you receive weather forecasts *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
69	`forecast_level` *select_one info_level*	At what level is the weather forecast provided (If you receive information at multiple levels, select the most detailed one, e.g. select "District" rather than "National" if you receive both.) *Options:* the 5-item `info_level` list, given in full at Q59. *Relevance:* Asked only if `forecast_any` was answered Yes. `${forecast_any} = 'yes'`
70	`forecast_level_other` *text*	Please specify the other level *Relevance:* Asked only if “Other” was chosen at `forecast_level`. `${forecast_level} = 'other'`
71	`forecast_specific` *select_one yn*	Is this forecast specific to your location (sub county level or below) *Options:* the 2-item `yn` list, given in full at Q4. *Relevance:* Asked only if `forecast_any` was answered Yes. `${forecast_any} = 'yes'`
72	`forecast_content` → `forecast_content_rain, forecast_content_temp, forecast_content_soil, forecast_content_storms, forecast_content_other` *select_multiple info_content*	What is forecasted *Options:* the 5-item `info_content` list, given in full at Q64. *Relevance:* Asked only if `forecast_any` was answered Yes. `${forecast_any} = 'yes'`
73	`forecast_content_other` → `forecast_content_other_specify` *text*	Please specify the other forecasted information *Relevance:* Asked only if the option `other` was ticked at `forecast_content`. `selected(${forecast_content}, 'other')`
74	`forecast_sources` → `src_fcst_radio, src_fcst_tv, src_fcst_mobile, src_fcst_meetings, src_fcst_extension, src_fcst_neighbors, src_fcst_internet, src_fcst_other` *select_multiple info_source*	What are your main sources of weather forecasts *Options:* the 8-item `info_source` list, given in full at Q66. *Relevance:* Asked only if `forecast_any` was answered Yes. `${forecast_any} = 'yes'`

continued from previous page

No.	Variable	Question, type, options, and routing
75	`forecast_sources_other` → `src_fcst_other_specify` *text*	Please specify the other weather forecast source *Relevance:* Asked only if the option `other` was ticked at `forecast_sources`. `selected(${forecast_sources}, 'other')`

OA-2.6 Experimental assignment

XLSForm group(s): `arm_group`.
Arms: All arms. The arm is computed from the village, not drawn on the device.

No.	Variable	Question, type, options, and routing
Block `arm_group` — *"Experimental assignment"*. *Block relevance:* none; every respondent who reaches this point answers it.		
76	`arm` *calculate*	*Not asked. Stored calculation.* *Calculation:* a 60-branch nested `if()` on `village_name`; the full village-to-arm mapping it encodes is reproduced in Table OA1.
77	`arm_note` → *(echo column, dropped)* *note*	*Dynamic text:* `concat("✓ This village is assigned to Treatment Arm ", ${arm})`

OA-2.7 Recalled rainfall amount, 2021–2025

XLSForm group(s): `belief_amount_2025 ... belief_amount_2021`.
Arms: All arms. This is the memory instrument and it runs before any treatment.

No.	Variable	Question, type, options, and routing
Block `belief_amount_2025` — *"Memory of rainfall amount in the first agricultural season (February-July) 2025"*. *Block relevance:* none; every respondent who reaches this point answers it.		
78	`rem_2025` *select_one yn*	Do you remember the rainfall conditions in the first agricultural season (February-July) 2025 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
79	`y2025_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
80	`y2025_img2` *integer*	Tokens on image 2 drought *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
81	`y2025_img3` *integer*	Tokens on image 3 very dry *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
82	`y2025_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
83	`y2025_img5` *integer*	Tokens on image 5 optimal *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
84	`y2025_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
85	`y2025_img7` *integer*	Tokens on image 7 very wet *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
86	`y2025_img8` *integer*	Tokens on image 8 flood *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
87	`y2025_img9` *integer*	Tokens on image 9 severe flood *Required.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2025_img1}+${y2025_img2}+${y2025_img3}+${y2025_img4}+${y2025_img5}+${y2025_img6}+${y2025_img7}+${y2025_img8}) = 10` *Message:* “Total number of tokens for this year must be 10”
88	`y2025_total_tokens` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'` *Calculation:* `${y2025_img1}+${y2025_img2}+${y2025_img3}+${y2025_img4}+${y2025_img5}+${y2025_img6}+${y2025_img7}+${y2025_img8}+${y2025_img9}`
89	`comment_belief_amount_2025` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_2025` was answered Yes. `${rem_2025} = 'yes'`

Block `belief_amount_2024` — *“Memory of rainfall amount in the first agricultural season (February-July) 2024”.*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
90	`rem_2024` *select_one yn*	Do you remember the rainfall conditions in the first agricultural season (February-July) 2024 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
91	`y2024_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
92	`y2024_img2` *integer*	Tokens on image 2 drought *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
93	`y2024_img3` *integer*	Tokens on image 3 very dry *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
94	`y2024_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
95	`y2024_img5` *integer*	Tokens on image 5 optimal *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
96	`y2024_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
97	`y2024_img7` *integer*	Tokens on image 7 very wet *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
98	`y2024_img8` *integer*	Tokens on image 8 flood *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
99	`y2024_img9` *integer*	Tokens on image 9 severe flood *Required.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2024_img1}+${y2024_img2}+${y2024_img3}+${y2024_img4}+${y2024_img5}+${y2024_img6}+${y2024_img7}+${y2024_img8}) = 10` *Message:* "Total number of tokens for this year must be 10"
100	`y2024_total_tokens` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'` *Calculation:* `${y2024_img1}+${y2024_img2}+${y2024_img3}+${y2024_img4}+${y2024_img5}+${y2024_img6}+${y2024_img7}+${y2024_img8}+${y2024_img9}`
101	`comment_belief_amount_2024` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_2024` was answered Yes. `${rem_2024} = 'yes'`

Block `belief_amount_2023` — *"Memory of rainfall amount in the first agricultural season (February-July) 2023".*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
102	`rem_2023` *select_one yn*	Do you remember the rainfall conditions in the first agricultural season (February-July) 2023 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

No.	Variable	Question, type, options, and routing
103	y2023_img1 *integer*	Tokens on image 1 extreme drought *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"
104	y2023_img2 *integer*	Tokens on image 2 drought *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"
105	y2023_img3 *integer*	Tokens on image 3 very dry *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"
106	y2023_img4 *integer*	Tokens on image 4 slightly dry *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"
107	y2023_img5 *integer*	Tokens on image 5 optimal *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"
108	y2023_img6 *integer*	Tokens on image 6 slightly wet *Required.* *Relevance:* Asked only if rem_2023 was answered Yes. ${rem_2023} = 'yes' *Constraint:* . >= 0 and . <= 10 *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
109	`y2023_img7` *integer*	Tokens on image 7 very wet *Required.* *Relevance:* Asked only if `rem_2023` was answered Yes. `${rem_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
110	`y2023_img8` *integer*	Tokens on image 8 flood *Required.* *Relevance:* Asked only if `rem_2023` was answered Yes. `${rem_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
111	`y2023_img9` *integer*	Tokens on image 9 severe flood *Required.* *Relevance:* Asked only if `rem_2023` was answered Yes. `${rem_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2023_img1}+${y2023_img2}+${y2023_img3}+${y2023_img4}+${y2023_img5}+${y2023_img6}+${y2023_img7}+${y2023_img8}) = 10` *Message:* "Total number of tokens for this year must be 10"
112	`y2023_total_tokens` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_2023` was answered Yes. `${rem_2023} = 'yes'` *Calculation:* `${y2023_img1}+${y2023_img2}+${y2023_img3}+${y2023_img4}+${y2023_img5}+${y2023_img6}+${y2023_img7}+${y2023_img8}+${y2023_img9}`
113	`comment_belief_amount_2023` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_2023` was answered Yes. `${rem_2023} = 'yes'`

Block `belief_amount_2022` — *"Memory of rainfall amount in the first agricultural season (February-July) 2022".*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
114	`rem_2022` *select_one yn*	Do you remember the rainfall conditions in the first agricultural season (February-July) 2022 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
115	`y2022_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
116	`y2022_img2` *integer*	Tokens on image 2 drought *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
117	`y2022_img3` *integer*	Tokens on image 3 very dry *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
118	`y2022_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
119	`y2022_img5` *integer*	Tokens on image 5 optimal *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
120	`y2022_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
121	`y2022_img7` *integer*	Tokens on image 7 very wet *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
122	`y2022_img8` *integer*	Tokens on image 8 flood *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
123	`y2022_img9` *integer*	Tokens on image 9 severe flood *Required.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10 and (.  + ${y2022_img1}+${y2022_img2}+${y2022_img3}+${y2022_img4}+${y2022_img5}+${y2022_img6}+${y2022_img7}+${y2022_img8}) = 10` *Message:* "Total number of tokens for this year must be 10"
124	`y2022_total_tokens` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'` *Calculation:* `${y2022_img1}+${y2022_img2}+${y2022_img3}+${y2022_img4}+${y2022_img5}+${y2022_img6}+${y2022_img7}+${y2022_img8}+${y2022_img9}`
125	`comment_belief_amount_2022` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_2022` was answered Yes. `${rem_2022} = 'yes'`

Block `belief_amount_2021` — *"Memory of rainfall amount in the first agricultural season (February-July) 2021".*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
126	`rem_2021` *select_one yn*	Do you remember the rainfall conditions in the first agricultural season (February-July) 2021 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	**Variable**	**Question, type, options, and routing**
127	`y2021_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
128	`y2021_img2` *integer*	Tokens on image 2 drought *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
129	`y2021_img3` *integer*	Tokens on image 3 very dry *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
130	`y2021_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
131	`y2021_img5` *integer*	Tokens on image 5 optimal *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
132	`y2021_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
133	`y2021_img7` *integer*	Tokens on image 7 very wet *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
134	`y2021_img8` *integer*	Tokens on image 8 flood *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
135	`y2021_img9` *integer*	Tokens on image 9 severe flood *Required.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2021_img1}+${y2021_img2}+${y2021_img3}+${y2021_img4}+${y2021_img5}+${y2021_img6}+${y2021_img7}+${y2021_img8}) = 10` *Message:* "Total number of tokens for this year must be 10"
136	`y2021_total_tokens` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'` *Calculation:* `${y2021_img1}+${y2021_img2}+${y2021_img3}+${y2021_img4}+${y2021_img5}+${y2021_img6}+${y2021_img7}+${y2021_img8}+${y2021_img9}`
137	`comment_belief_amount_2021` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_2021` was answered Yes. `${rem_2021} = 'yes'`

OA-2.8 Recalled onset timing, 2021–2025

XLSForm group(s): `onset_2025 ... onset_2021`.
Arms: All arms. Runs before any treatment.

No.	Variable	Question, type, options, and routing
Block `onset_2025` — *"Memory of onset timing in the first agricultural season (February-July) 2025".* *Block relevance:* none; every respondent who reaches this point answers it.		
138	`rem_onset_2025` *select_one yn*	Do you remember the onset timing in the first agricultural season (February-July) 2025 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
139	`y2025_int1` *integer*	Tokens for onset interval 1 *Required.* *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
140	`y2025_int2` *integer*	Tokens for onset interval 2 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
141	`y2025_int3` *integer*	Tokens for onset interval 3 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
142	`y2025_int4` *integer*	Tokens for onset interval 4 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
143	`y2025_int5` *integer*	Tokens for onset interval 5 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
144	`y2025_int6` *integer*	Tokens for onset interval 6 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
145	`y2025_int7` *integer*	Tokens for onset interval 7 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
146	`y2025_int8` *integer*	Tokens for onset interval 8 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
147	`y2025_int9` *integer*	Tokens for onset interval 9 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
148	`y2025_int10` *integer*	Tokens for onset interval 10 *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2025_int1}+${y2025_int2}+${y2025_int3}+${y2025_int4}+${y2025_int5}+${y2025_int6}+${y2025_int7}+${y2025_int8}+${y2025_int9}) = 10` *Message:* "Total number of tokens for this year must be 10"
149	`y2025_onset_total` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'` *Calculation:* `${y2025_int1}+${y2025_int2}+${y2025_int3}+${y2025_int4}+${y2025_int5}+${y2025_int6}+${y2025_int7}+${y2025_int8}+${y2025_int9}+${y2025_int10}`
150	`comment_onset_2025` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_onset_2025` was answered Yes. `${rem_onset_2025} = 'yes'`

Block `onset_2024` — *"Memory of onset timing in the first agricultural season (February-July) 2024".*
Block relevance: none; every respondent who reaches this point answers it.

continued from previous page

No.	Variable	Question, type, options, and routing
151	`rem_onset_2024` *select_one yn*	Do you remember the onset timing in the first agricultural season (February-July) 2024 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.
152	`y2024_int1` *integer*	Tokens for onset interval 1 *Required.* *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
153	`y2024_int2` *integer*	Tokens for onset interval 2 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
154	`y2024_int3` *integer*	Tokens for onset interval 3 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
155	`y2024_int4` *integer*	Tokens for onset interval 4 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
156	`y2024_int5` *integer*	Tokens for onset interval 5 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
157	`y2024_int6` *integer*	Tokens for onset interval 6 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
158	`y2024_int7` *integer*	Tokens for onset interval 7 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
159	`y2024_int8` *integer*	Tokens for onset interval 8 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
160	`y2024_int9` *integer*	Tokens for onset interval 9 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
161	`y2024_int10` *integer*	Tokens for onset interval 10 *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2024_int1}+${y2024_int2}+${y2024_int3}+${y2024_int4}+${y2024_int5}+${y2024_int6}+${y2024_int7}+${y2024_int8}+${y2024_int9}) = 10` *Message:* "Total number of tokens for this year must be 10"
162	`y2024_onset_total` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'` *Calculation:* `${y2024_int1}+${y2024_int2}+${y2024_int3}+${y2024_int4}+${y2024_int5}+${y2024_int6}+${y2024_int7}+${y2024_int8}+${y2024_int9}+${y2024_int10}`
163	`comment_onset_2024` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_onset_2024` was answered Yes. `${rem_onset_2024} = 'yes'`

Block `onset_2023` — *"Memory of onset timing in the first agricultural season (February-July) 2023"*.
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
164	`rem_onset_2023` *select_one yn*	Do you remember the onset timing in the first agricultural season (February-July) 2023 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
165	`y2023_int1` *integer*	Tokens for onset interval 1 *Required.* *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
166	`y2023_int2` *integer*	Tokens for onset interval 2 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
167	`y2023_int3` *integer*	Tokens for onset interval 3 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
168	`y2023_int4` *integer*	Tokens for onset interval 4 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
169	`y2023_int5` *integer*	Tokens for onset interval 5 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
170	`y2023_int6` *integer*	Tokens for onset interval 6 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
171	`y2023_int7` *integer*	Tokens for onset interval 7 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
172	`y2023_int8` *integer*	Tokens for onset interval 8 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
173	`y2023_int9` *integer*	Tokens for onset interval 9 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10` *Message:* "Enter a number between 0 and 10"
174	`y2023_int10` *integer*	Tokens for onset interval 10 *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Constraint:* `.  >= 0 and .  <= 10 and (.  + ${y2023_int1}+${y2023_int2}+${y2023_int3}+${y2023_int4}+${y2023_int5}+${y2023_int6}+${y2023_int7}+${y2023_int8}+${y2023_int9}) = 10` *Message:* "Total number of tokens for this year must be 10"
175	`y2023_onset_total` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'` *Calculation:* `${y2023_int1}+${y2023_int2}+${y2023_int3}+${y2023_int4}+${y2023_int5}+${y2023_int6}+${y2023_int7}+${y2023_int8}+${y2023_int9}+${y2023_int10}`
176	`comment_onset_2023` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_onset_2023` was answered Yes. `${rem_onset_2023} = 'yes'`

Block `onset_2022` — *"Memory of onset timing in the first agricultural season (February-July) 2022".*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
177	`rem_onset_2022` *select_one yn*	Do you remember the onset timing in the first agricultural season (February-July) 2022 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
178	`y2022_int1` *integer*	Tokens for onset interval 1 *Required.* *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
179	`y2022_int2` *integer*	Tokens for onset interval 2 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
180	`y2022_int3` *integer*	Tokens for onset interval 3 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
181	`y2022_int4` *integer*	Tokens for onset interval 4 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
182	`y2022_int5` *integer*	Tokens for onset interval 5 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
183	`y2022_int6` *integer*	Tokens for onset interval 6 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
184	`y2022_int7` *integer*	Tokens for onset interval 7 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
185	`y2022_int8` *integer*	Tokens for onset interval 8 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
186	`y2022_int9` *integer*	Tokens for onset interval 9 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
187	`y2022_int10` *integer*	Tokens for onset interval 10 *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2022_int1}+${y2022_int2}+${y2022_int3}+${y2022_int4}+${y2022_int5}+${y2022_int6}+${y2022_int7}+${y2022_int8}+${y2022_int9}) = 10` *Message:* "Total number of tokens for this year must be 10"
188	`y2022_onset_total` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'` *Calculation:* `${y2022_int1}+${y2022_int2}+${y2022_int3}+${y2022_int4}+${y2022_int5}+${y2022_int6}+${y2022_int7}+${y2022_int8}+${y2022_int9}+${y2022_int10}`
189	`comment_onset_2022` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_onset_2022` was answered Yes. `${rem_onset_2022} = 'yes'`

Block `onset_2021` — *"Memory of onset timing in the first agricultural season (February-July) 2021"*.
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
190	`rem_onset_2021` *select_one yn*	Do you remember the onset timing in the first agricultural season (February-July) 2021 *Required.* *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	**Variable**	**Question, type, options, and routing**
191	`y2021_int1` *integer*	Tokens for onset interval 1 *Required.* *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
192	`y2021_int2` *integer*	Tokens for onset interval 2 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
193	`y2021_int3` *integer*	Tokens for onset interval 3 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
194	`y2021_int4` *integer*	Tokens for onset interval 4 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
195	`y2021_int5` *integer*	Tokens for onset interval 5 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
196	`y2021_int6` *integer*	Tokens for onset interval 6 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
197	`y2021_int7` *integer*	Tokens for onset interval 7 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
198	`y2021_int8` *integer*	Tokens for onset interval 8 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
199	`y2021_int9` *integer*	Tokens for onset interval 9 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
200	`y2021_int10` *integer*	Tokens for onset interval 10 *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Constraint:* `. >= 0 and . <= 10 and (. + ${y2021_int1}+${y2021_int2}+${y2021_int3}+${y2021_int4}+${y2021_int5}+${y2021_int6}+${y2021_int7}+${y2021_int8}+${y2021_int9}) = 10` *Message:* "Total number of tokens for this year must be 10"
201	`y2021_onset_total` *calculate*	*Not asked. Stored calculation.* *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'` *Calculation:* `${y2021_int1}+${y2021_int2}+${y2021_int3}+${y2021_int4}+${y2021_int5}+${y2021_int6}+${y2021_int7}+${y2021_int8}+${y2021_int9}+${y2021_int10}`
202	`comment_onset_2021` *text*	Enumerator: note any comments or remarks made by the participant during this section *Relevance:* Asked only if `rem_onset_2021` was answered Yes. `${rem_onset_2021} = 'yes'`

OA-2.9 Pre-card forward expectations and confidence

XLSForm group(s): `belief_amount_future`, `onset_future_group`, `prediction_confidence`.
Arms: **Arms A and B only.** Arm C has no pre-card belief elicitation.

No.	Variable	Question, type, options, and routing

Block `belief_amount_future` — *"Expectations about rainfall amount upcoming season"*. *Block relevance:* Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`

No.	Variable	Question, type, options, and routing
203	`fut_img1` → `fut_pre_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
204	`fut_img2` → `fut_pre_img2` *integer*	Tokens on image 2 drought *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
205	`fut_img3` → `fut_pre_img3` *integer*	Tokens on image 3 very dry *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
206	`fut_img4` → `fut_pre_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
207	`fut_img5` → `fut_pre_img5` *integer*	Tokens on image 5 optimal *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
208	`fut_img6` → `fut_pre_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
209	`fut_img7` → `fut_pre_img7` *integer*	Tokens on image 7 very wet *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
210	`fut_img8` → `fut_pre_img8` *integer*	Tokens on image 8 flood *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
211	`fut_img9` → `fut_pre_img9` *integer*	Tokens on image 9 severe flood *Required.* *Constraint:* `. >= 0 and . <= 10 and (. + ${fut_img1}+${fut_img2}+${fut_img3}+${fut_img4}+${fut_img5}+${fut_img6}+${fut_img7}+${fut_img8}) = 10` *Message:* "Total number of tokens must be 10"
212	`fut_total_tokens` → `fut_pre_total_tokens` *calculate*	*Not asked. Stored calculation.* *Calculation:* `${fut_img1}+${fut_img2}+${fut_img3}+${fut_img4}+${fut_img5}+${fut_img6}+${fut_img7}+${fut_img8}+${fut_img9}`
213	`comment_belief_amount_future` → `comment_fut_pre_amount` *text*	Enumerator: note any comments or remarks made by the participant during this section

Block `onset_future_group` — *"Expectations about onset timing upcoming season".*
Block relevance: Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`

No.	Variable	Question, type, options, and routing
214	`fut_int1` → `fut_pre_int1` *integer*	Tokens for onset interval 1 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
215	`fut_int2` → `fut_pre_int2` *integer*	Tokens for onset interval 2 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
216	`fut_int3` → `fut_pre_int3` *integer*	Tokens for onset interval 3 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
217	`fut_int4` → `fut_pre_int4` *integer*	Tokens for onset interval 4 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
218	`fut_int5` → `fut_pre_int5` *integer*	Tokens for onset interval 5 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
219	`fut_int6` → `fut_pre_int6` *integer*	Tokens for onset interval 6 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
220	`fut_int7` → `fut_pre_int7` *integer*	Tokens for onset interval 7 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
221	`fut_int8` → `fut_pre_int8` *integer*	Tokens for onset interval 8 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
222	`fut_int9` → `fut_pre_int9` *integer*	Tokens for onset interval 9 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
223	`fut_int10` → `fut_pre_int10` *integer*	Tokens for onset interval 10 *Required.* *Constraint:* `. >= 0 and . <= 10 and (. + ${fut_int1}+${fut_int2}+${fut_int3}+${fut_int4}+${fut_int5}+${fut_int6}+${fut_int7}+${fut_int8}+${fut_int9}) = 10` *Message:* “Total number of tokens must be 10”
224	`fut_onset_total` → `fut_pre_onset_total` *calculate*	*Not asked. Stored calculation.* *Calculation:* `${fut_int1}+${fut_int2}+${fut_int3}+${fut_int4}+${fut_int5}+${fut_int6}+${fut_int7}+${fut_int8}+${fut_int9}+${fut_int10}`
225	`comment_onset_future_group` → `comment_fut_pre_onset` *text*	Enumerator: note any comments or remarks made by the participant during this section

Block `prediction_confidence` — *“Prediction Confidence”.*
Block relevance: Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`

continued from previous page

No.	Variable	Question, type, options, and routing
226	`conf_amount` → `conf_rain_pre` *select_one conf_scale*	Confidence in rainfall amount expectations *Required.* *Options* (code — label as read out), `conf_scale`: `5` — 5 Very High Confidence `4` — 4 High Confidence `3` — 3 Moderate Confidence `2` — 2 Low Confidence `1` — 1 Very Low Confidence
227	`conf_onset` → `conf_onset_pre` *select_one conf_scale*	Confidence in onset expectations *Required.* *Options:* the 5-item `conf_scale` list, given in full at Q226.

OA-2.10 Memory of extreme events

XLSForm group(s): `extreme_memory_grp`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `extreme_memory_grp` — *"Memory of extreme events"*. *Block relevance:* none; every respondent who reaches this point answers it.		
228	`extreme_past_any` *select_one yn*	Do you remember any years with major weather events (e.g. drought, flood, hailstorm, storm, etc.) *Options:* the 2-item `yn` list, given in full at Q4.
229	`extreme_past_n` *integer*	How many major weather events do you remember (0 to 5) *Required.* *Relevance:* Asked only if `extreme_past_any` was answered Yes. `${extreme_past_any} = 'yes'` *Constraint:* `. >= 0 and . <= 5` *Message:* "Enter a number between 0 and 5"
Block `extreme_event1` — *"Details of event 1"*. *Block relevance:* Asked only if the respondent said they remember at least 1 major weather event(s). `${extreme_past_n} >= 1`		

continued from previous page

No.	Variable	Question, type, options, and routing
230	`extreme1_type` *select_one extreme_type*	Type of event *Options* (code — label as read out), `extreme_type`: `drought` — Drought `flood` — Flood `hailstorm` — Hailstorm `storm` — Storm `other` — Other *Relevance:* Asked only if the respondent said they remember at least 1 major weather event(s). `${extreme_past_n} >= 1`
231	`extreme1_type_other` *text*	If other, specify the type of event *Relevance:* Asked only if "Other" was chosen at `extreme1_type`. `${extreme1_type} = 'other'`
232	`extreme1_severity` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Required.* *Relevance:* Asked only if the respondent said they remember at least 1 major weather event(s). `${extreme_past_n} >= 1` *Constraint:* `. >= 1 and . <= 10` *Message:* "Enter a value from 1 to 10"
233	`extreme1_year` *integer*	In which year did this event happen *Required.* *Relevance:* Asked only if the respondent said they remember at least 1 major weather event(s). `${extreme_past_n} >= 1` *Constraint:* `. >= 1900 and . <= 2100` *Message:* "Enter a valid year"

Block `extreme_event2` — *"Details of event 2".*
Block relevance: Asked only if the respondent said they remember at least 2 major weather event(s). `${extreme_past_n} >= 2`

No.	Variable	Question, type, options, and routing
234	`extreme2_type` *select_one extreme_type*	Type of event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if the respondent said they remember at least 2 major weather event(s). `${extreme_past_n} >= 2`
235	`extreme2_type_other` *text*	If other, specify the type of event *Relevance:* Asked only if "Other" was chosen at `extreme2_type`. `${extreme2_type} = 'other'`

continued from previous page

No.	Variable	Question, type, options, and routing
236	`extreme2_severity` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Required.* *Relevance:* Asked only if the respondent said they remember at least 2 major weather event(s). `${extreme_past_n} >= 2` *Constraint:* `. >= 1 and . <= 10` *Message:* "Enter a value from 1 to 10"
237	`extreme2_year` *integer*	In which year did this event happen *Required.* *Relevance:* Asked only if the respondent said they remember at least 2 major weather event(s). `${extreme_past_n} >= 2` *Constraint:* `. >= 1900 and . <= 2100` *Message:* "Enter a valid year"

Block `extreme_event3` — *"Details of event 3"*.
Block relevance: Asked only if the respondent said they remember at least 3 major weather event(s). `${extreme_past_n} >= 3`

No.	Variable	Question, type, options, and routing
238	`extreme3_type` *select_one extreme_type*	Type of event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if the respondent said they remember at least 3 major weather event(s). `${extreme_past_n} >= 3`
239	`extreme3_type_other` *text*	If other, specify the type of event *Relevance:* Asked only if "Other" was chosen at `extreme3_type`. `${extreme3_type} = 'other'`
240	`extreme3_severity` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Required.* *Relevance:* Asked only if the respondent said they remember at least 3 major weather event(s). `${extreme_past_n} >= 3` *Constraint:* `. >= 1 and . <= 10` *Message:* "Enter a value from 1 to 10"

continued from previous page

No.	Variable	Question, type, options, and routing
241	`extreme3_year` *integer*	In which year did this event happen *Required.* *Relevance:* Asked only if the respondent said they remember at least 3 major weather event(s). `${extreme_past_n} >= 3` *Constraint:* `. >= 1900 and . <= 2100` *Message:* "Enter a valid year"

Block `extreme_event4` — *"Details of event 4"*.
Block relevance: Asked only if the respondent said they remember at least 4 major weather event(s). `${extreme_past_n} >= 4`

No.	Variable	Question, type, options, and routing
242	`extreme4_type` *select_one extreme_type*	Type of event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if the respondent said they remember at least 4 major weather event(s). `${extreme_past_n} >= 4`
243	`extreme4_type_other` *text*	If other, specify the type of event *Relevance:* Asked only if "Other" was chosen at `extreme4_type`. `${extreme4_type} = 'other'`
244	`extreme4_severity` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Required.* *Relevance:* Asked only if the respondent said they remember at least 4 major weather event(s). `${extreme_past_n} >= 4` *Constraint:* `. >= 1 and . <= 10` *Message:* "Enter a value from 1 to 10"
245	`extreme4_year` *integer*	In which year did this event happen *Required.* *Relevance:* Asked only if the respondent said they remember at least 4 major weather event(s). `${extreme_past_n} >= 4` *Constraint:* `. >= 1900 and . <= 2100` *Message:* "Enter a valid year"

Block `extreme_event5` — *"Details of event 5"*.
Block relevance: Asked only if the respondent said they remember at least 5 major weather event(s). `${extreme_past_n} >= 5`

continued from previous page

No.	Variable	Question, type, options, and routing
246	`extreme5_type` *select_one extreme_type*	Type of event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if the respondent said they remember at least 5 major weather event(s). `${extreme_past_n} >= 5`
247	`extreme5_type_other` *text*	If other, specify the type of event *Relevance:* Asked only if "Other" was chosen at `extreme5_type`. `${extreme5_type} = 'other'`
248	`extreme5_severity` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Required.* *Relevance:* Asked only if the respondent said they remember at least 5 major weather event(s). `${extreme_past_n} >= 5` *Constraint:* `. >= 1 and . <= 10` *Message:* "Enter a value from 1 to 10"
249	`extreme5_year` *integer*	In which year did this event happen *Required.* *Relevance:* Asked only if the respondent said they remember at least 5 major weather event(s). `${extreme_past_n} >= 5` *Constraint:* `. >= 1900 and . <= 2100` *Message:* "Enter a valid year"

OA-2.11 Pre-card expectations of extreme events

XLSForm group(s): `extreme_expect_group`.
Arms: **Arms A and B only.**

No.	Variable	Question, type, options, and routing
Block `extreme_expect_group` — *"Expectations of extreme events"*. *Block relevance:* Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`		
250	`extreme_future_any` → `extreme_future_any_pre` *select_one yn*	Do you expect a major weather event in the upcoming season *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
251	`extreme_future_type` → `extreme_future_type_pre` *select_one extreme_type*	Type of expected event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if `extreme_future_any` was answered Yes. `${extreme_future_any} = 'yes'`
252	`extreme_future_type_other` → `extreme_future_type_pre_other` *text*	If other, specify *Relevance:* Asked only if "Other" was chosen at `extreme_future_type`. `${extreme_future_type} = 'other'`
253	`extreme_future_severity` → `extreme_future_severity_pre` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Relevance:* Asked only if `extreme_future_any` was answered Yes. `${extreme_future_any} = 'yes'` *Constraint:* `. >= 1 and . <= 10` *Message:* "Value must be between 1 and 10"

OA-2.12 Drivers of expectations and social learning

XLSForm group(s): `factors_group, social_group`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `factors_group` — *"Factors influencing expectations"*. *Block relevance:* All three arms (the condition is exhaustive and gates nothing out). `${arm} = 'A' or ${arm} = 'B' or ${arm} = 'C'`		

continued from previous page

No.	Variable	Question, type, options, and routing
254	`factors_considered` → `driver_memory, driver_farmers, driver_leaders, driver_govt, driver_tradition, driver_religion, driver_env, driver_media, driver_ext, driver_other` *select_multiple factors*	What factors do you consider when forming rainfall expectations *Options* (code — label as read out), `factors`: `memory` — Personal memory of past seasons `discuss` — Discussions with other farmers `leaders` — Information from community leaders `gov` — Government weather forecasts `trad` — Traditional knowledge or beliefs `religious` — Religious or spiritual guidance `env` — Environmental indicators `media` — Media sources `extension` — Agricultural extension services `other` — Other
255	`factors_considered_other` → `driver_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `factors_considered`. `selected(${factors_considered}, 'other')`

Block `social_group` — *"Social interactions and information sharing"*.
Block relevance: All three arms (the condition is exhaustive and gates nothing out).
`${arm} = 'A' or ${arm} = 'B' or ${arm} = 'C'`

No.	Variable	Question, type, options, and routing
256	`discuss_expectations` *select_one yn*	Do you discuss your rainfall expectations with others *Options:* the 2-item `yn` list, given in full at Q4.
257	`discuss_with` → `discuss_family, discuss_neighbors, discuss_leaders, discuss_extension, discuss_friends, discuss_religious, discuss_others` *select_multiple discuss_who*	With whom do you usually discuss *Options* (code — label as read out), `discuss_who`: `family` — Family members `neighbors` — Neighboring farmers `leaders` — Community leaders `extension` — Agricultural extension agents `friends` — Friends `religious` — Religious leaders `other` — Others *Relevance:* Asked only if `discuss_expectations` was answered Yes. `${discuss_expectations} = 'yes'`
258	`discuss_with_other` → `discuss_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `discuss_with`. `selected(${discuss_with}, 'other')`

No.	Variable	Question, type, options, and routing
259	`discuss_freq` *select_one freq*	How often do you have these discussions *Options* (code — label as read out), `freq`: `daily` — Daily `weekly` — Weekly `monthly` — Monthly `occas` — Occasionally *Relevance:* Asked only if `discuss_expectations` was answered Yes. `${discuss_expectations} = 'yes'`
260	`trusted_opinions` → `trust_experienced_farmers`, `trust_elders`, `trust_experts`, `trust_religious`, `trust_media`, `trust_govt`, `trust_others` *select_multiple trust_who*	Whose opinions do you trust most *Options* (code — label as read out), `trust_who`: `farmers` — Experienced farmers `elders` — Elders in the community `experts` — Agricultural experts `religious` — Religious leaders `media` — Media sources `gov` — Government officials `other` — Others
261	`trusted_opinions_other` → `trust_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `trusted_opinions`. `selected(${trusted_opinions}, 'other')`
262	`trust_reasons` → `trust_quality_experience`, `trust_quality_accuracy`, `trust_quality_education`, `trust_quality_age`, `trust_quality_wealth`, `trust_quality_family`, `trust_quality_status`, `trust_quality_other` *select_multiple trust_why*	What qualities make you trust someones knowledge *Options* (code — label as read out), `trust_why`: `exper` — Farming experience `accuracy` — Past accuracy in predictions `educ` — Education level `age` — Age `wealth` — Wealth `family` — Family relation `status` — Social status `other` — Other
263	`trust_reasons_other` → `trust_quality_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `trust_reasons`. `selected(${trust_reasons}, 'other')`

OA-2.13 Pre-card planned adaptation

XLSForm group(s): `adapt_group`.
Arms: **All three arms.** This is the block that yields `adapt_any_pre` for every respondent.

No.	Variable	Question, type, options, and routing
		Block `adapt_group` — *"Planned adaptation strategies based on expectations"*. *Block relevance:* All three arms (the condition is exhaustive and gates nothing out). `${arm} = 'A' or ${arm} = 'B' or ${arm} = 'C'`
264	`adapt_any` → `adapt_any_pre` *select_one yn*	Are you planning to change your farming practices based on your rainfall expectations *Options:* the 2-item `yn` list, given in full at Q4.
265	`adapt_types` → `adapt_pre_variety, adapt_pre_planting_date, adapt_pre_irrigation, adapt_pre_fertilizer, adapt_pre_diversify, adapt_pre_insurance, adapt_pre_off_farm, adapt_pre_credit, adapt_pre_training, adapt_pre_other` *select_multiple adapt_ops*	What changes are you planning to implement *Options* (code — label as read out), `adapt_ops`: `variety` — Changing crop varieties; `dates` — Adjusting planting dates; `irrigation` — Investing in irrigation; `fert` — Modifying fertiliser use; `diversify` — Diversifying crops; `insurance` — Purchasing insurance; `offfarm` — Seeking off farm employment; `credit` — Accessing credit or loans; `training` — Attending training or workshops; `other` — Other *Relevance:* Asked only if `adapt_any` was answered Yes. `${adapt_any} = 'yes'`
266	`adapt_types_other` → `adapt_pre_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `adapt_types`. `selected(${adapt_types}, 'other')`
267	`adapt_cost` → `adapt_pre_cost_ugx` *decimal*	Estimated total cost of these adaptations (in Ugandan Shillings) *Relevance:* Asked only if `adapt_any` was answered Yes. `${adapt_any} = 'yes'` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"

continued from previous page

No.	Variable	Question, type, options, and routing
268	`adapt_finance` → `adapt_pre_fin_savings`, `adapt_pre_fin_family_loans`, `adapt_pre_fin_bank`, `adapt_pre_fin_microfinance`, `adapt_pre_fin_govt`, `adapt_pre_fin_assets`, `adapt_pre_fin_community`, `adapt_pre_fin_other` *select_multiple finance*	How do you plan to finance these adaptations *Options* (code — label as read out), `finance`: `savings` — Own savings `gov` — Government subsidies or grants `family` — Loans from family or friends `assets` — Sale of assets `bank` — Bank loans `group` — Community savings groups `micro` — Microfinance institutions `other` — Other *Relevance:* Asked only if `adapt_any` was answered Yes. `${adapt_any} = 'yes'`
269	`adapt_finance_other` → `adapt_pre_fin_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `adapt_finance`. `selected(${adapt_finance}, 'other')`

OA-2.14 Risk perceptions and attitudes

XLSForm group(s): `risk_group`.
Arms: All arms.

No.	Variable	Question, type, options, and routing
Block `risk_group` — *"Risk perceptions and attitudes"*. *Block relevance:* none; every respondent who reaches this point answers it.		
270	`risk_willingness` *select_one conf_scale3*	How willing are you to take risks in farming *Required.* *Options* (code — label as read out), `conf_scale3`: `5` — 5 Extremely Willing `4` — 4 Very Willing `3` — 3 Somewhat Willing `2` — 2 A Little Willing `1` — 1 Not Willing At All

continued from previous page

No.	Variable	Question, type, options, and routing
271	`weather_concern` *select_one conf_scale4*	How concerned are you about weather impact on your income *Required.* *Options* (code — label as read out), `conf_scale4`: `5` — 5 Extremely Concerned `4` — 4 Very Concerned `3` — 3 Somewhat Concerned `2` — 2 A Little Concerned `1` — 1 Not Concerned At All

OA-2.15 Historical rainfall information card and immediate reaction

XLSForm group(s): `info_treatment_group`.
Arms: **Arms B and C only.** Arm A never sees the card.

No.	Variable	Question, type, options, and routing
Block `info_treatment_group` — *"Historical rainfall information block"*. *Block relevance:* Arms B and C only; Arm A never sees the information card. `${arm} = 'B' or ${arm} = 'C'`		
272	`info_given_note` → `info_card_shown_note` *note*	Show the respondent the local historical rainfall charts now
273	`info_differs` *select_one yn*	Does this information differ from what you remembered about past rainfall *Options:* the 2-item `yn` list, given in full at Q4.
274	`info_surprise` *select_one conf_scale2*	How surprised are you by this information *Options* (code — label as read out), `conf_scale2`: `5` — 5 Extremely Surprised `4` — 4 Very Surprised `3` — 3 Somewhat Surprised `2` — 2 A Little Surprised `1` — 1 Not Surprised At All
275	`adjust_expectations` *select_one yn*	After seeing this information would you like to adjust your expectations *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
276	`conf_updated` *select_one conf_scale*	Confidence in updated expectations *Options:* the 5-item `conf_scale` list, given in full at Q226. *Relevance:* Asked only if `adjust_expectations` was answered Yes. `${adjust_expectations} = 'yes'`
277	`info_useful` *select_one yn*	Is historical rainfall information helpful for planning *Options:* the 2-item `yn` list, given in full at Q4.
278	`info_future` *select_one yn*	Would you like to receive similar information in the future *Options:* the 2-item `yn` list, given in full at Q4.

OA-2.16 Second forward elicitation

XLSForm group(s): `post_amount, post_onset_group`.
Arms: **All three arms.** Neither group carries a relevance condition.

No.	Variable	Question, type, options, and routing
Block `post_amount` — *"Expectations about rainfall amount upcoming season"*. *Block relevance:* none; every respondent who reaches this point answers it.		
279	`post_img1` → `fut_post_img1` *integer*	Tokens on image 1 extreme drought *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
280	`post_img2` → `fut_post_img2` *integer*	Tokens on image 2 drought *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
281	`post_img3` → `fut_post_img3` *integer*	Tokens on image 3 very dry *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"
282	`post_img4` → `fut_post_img4` *integer*	Tokens on image 4 slightly dry *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* "Enter a number between 0 and 10"

continued from previous page

No.	Variable	Question, type, options, and routing
283	`post_img5` → `fut_post_img5` *integer*	Tokens on image 5 optimal *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
284	`post_img6` → `fut_post_img6` *integer*	Tokens on image 6 slightly wet *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
285	`post_img7` → `fut_post_img7` *integer*	Tokens on image 7 very wet *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
286	`post_img8` → `fut_post_img8` *integer*	Tokens on image 8 flood *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
287	`post_img9` → `fut_post_img9` *integer*	Tokens on image 9 severe flood *Required.* *Constraint:* `. >= 0 and . <= 10 and (. + ${post_img1}+${post_img2}+${post_img3}+${post_img4}+${post_img5}+${post_img6}+${post_img7}+${post_img8}) = 10` *Message:* “Total number of tokens must be 10”
288	`post_total_tokens` → `fut_post_total_tokens` *calculate*	*Not asked. Stored calculation.* *Calculation:* `${post_img1}+${post_img2}+${post_img3}+${post_img4}+${post_img5}+${post_img6}+${post_img7}+${post_img8}+${post_img9}`
289	`comment_post_amount` → `comment_fut_post_amount` *text*	Enumerator: note any comments or remarks made by the participant during this section

Block `post_onset_group` — *“Expectations about onset timing upcoming season”.*
Block relevance: none; every respondent who reaches this point answers it.

No.	Variable	Question, type, options, and routing
290	`post_int1` → `fut_post_int1` *integer*	Tokens for onset interval 1 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”

continued from previous page

No.	Variable	Question, type, options, and routing
291	`post_int2` → `fut_post_int2` *integer*	Tokens for onset interval 2 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
292	`post_int3` → `fut_post_int3` *integer*	Tokens for onset interval 3 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
293	`post_int4` → `fut_post_int4` *integer*	Tokens for onset interval 4 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
294	`post_int5` → `fut_post_int5` *integer*	Tokens for onset interval 5 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
295	`post_int6` → `fut_post_int6` *integer*	Tokens for onset interval 6 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
296	`post_int7` → `fut_post_int7` *integer*	Tokens for onset interval 7 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
297	`post_int8` → `fut_post_int8` *integer*	Tokens for onset interval 8 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
298	`post_int9` → `fut_post_int9` *integer*	Tokens for onset interval 9 *Required.* *Constraint:* `. >= 0 and . <= 10` *Message:* “Enter a number between 0 and 10”
299	`post_int10` → `fut_post_int10` *integer*	Tokens for onset interval 10 *Required.* *Constraint:* `. >= 0 and . <= 10 and (. + ${post_int1}+${post_int2}+${post_int3}+${post_int4}+${post_int5}+${post_int6}+${post_int7}+${post_int8}+${post_int9}) = 10` *Message:* “Total number of tokens must be 10”

continued from previous page

No.	Variable	Question, type, options, and routing
300	`post_onset_total` → `fut_post_onset_total` *calculate*	*Not asked. Stored calculation.* *Calculation:* `${post_int1}+${post_int2}+${post_int3}+${post_int4}+${post_int5}+${post_int6}+${post_int7}+${post_int8}+${post_int9}+${post_int10}`
301	`comment_post_onset_group` → `comment_fut_post_onset` *text*	Enumerator: note any comments or remarks made by the participant during this section

OA-2.17 Post-card confidence, extremes, and adaptation (Arms A and B)

XLSForm group(s): `prediction_confidence2`, `extreme_expect_group2`, `adapt_group2`.
Arms: **Arms A and B only.**

No.	Variable	Question, type, options, and routing
Block `prediction_confidence2` — *"Prediction Confidence"*. *Block relevance:* Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`		
302	`conf_amount2` → `conf_rain_post` *select_one conf_scale*	Confidence in updated rainfall amount expectations *Required.* *Options:* the 5-item `conf_scale` list, given in full at Q226.
303	`conf_onset2` → `conf_onset_post` *select_one conf_scale*	Confidence in updated onset expectations *Required.* *Options:* the 5-item `conf_scale` list, given in full at Q226.
Block `extreme_expect_group2` — *"Expectations of extreme events"*. *Block relevance:* Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`		
304	`extreme_future_any2` → `extreme_future_any_post` *select_one yn*	Do you expect a major weather event in the upcoming season *Options:* the 2-item `yn` list, given in full at Q4.

continued from previous page

No.	Variable	Question, type, options, and routing
305	`extreme_future_type2` → `extreme_future_type_post` *select_one extreme_type*	Type of expected event *Options:* the 5-item `extreme_type` list, given in full at Q230. *Relevance:* Asked only if `extreme_future_any2` was answered Yes. `${extreme_future_any2} = 'yes'`
306	`extreme_future_type2_other` → `extreme_future_type_post_other` *text*	If other, specify *Relevance:* Asked only if "Other" was chosen at `extreme_future_type2`. `${extreme_future_type2} = 'other'`
307	`extreme_future_severity2` → `extreme_future_severity_post` *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Relevance:* Asked only if `extreme_future_any2` was answered Yes. `${extreme_future_any2} = 'yes'` *Constraint:* `. >= 1 and . <= 10` *Message:* "Value must be between 1 and 10"

Block `adapt_group2` — *"Planned adaptation strategies based on expectations"*.
Block relevance: Arms A and B only; Arm C respondents skip this item. `${arm} = 'A' or ${arm} = 'B'`

No.	Variable	Question, type, options, and routing
308	`adapt_any2` → `adapt_any_post` *select_one yn*	Are you planning to change your farming practices based on your updated rainfall expectations *Options:* the 2-item `yn` list, given in full at Q4.
309	`adapt_types2` → `adapt_post_variety ... adapt_post_other (10 dummies)` *select_multiple adapt_ops*	What changes are you planning to implement *Options:* the 10-item `adapt_ops` list, given in full at Q265. *Relevance:* Asked only if `adapt_any2` was answered Yes. `${adapt_any2} = 'yes'`
310	`adapt_types2_other` → `adapt_post_other_specify` *text*	If other, specify *Relevance:* Asked only if the option `other` was ticked at `adapt_types2`. `selected(${adapt_types2}, 'other')`
311	`adapt_cost2` → `adapt_post_cost_ugx` *decimal*	Estimated total cost of these adaptations (in Ugandan Shillings) *Relevance:* Asked only if `adapt_any2` was answered Yes. `${adapt_any2} = 'yes'` *Constraint:* `. >= 0` *Message:* "Value must be zero or positive"

continued from previous page

No.	Variable	Question, type, options, and routing
312	adapt_finance2 → adapt_post_fin_savings ... adapt_post_fin_other (8 dummies) *select_multiple finance*	How do you plan to finance these adaptations *Options:* the 8-item finance list, given in full at Q268. *Relevance:* Asked only if adapt_any2 was answered Yes. ${adapt_any2} = 'yes'
313	adapt_finance2_other → adapt_post_fin_other_specify *text*	If other, specify *Relevance:* Asked only if the option other was ticked at adapt_finance2. selected(${adapt_finance2}, 'other')

OA-2.18 Confidence, extremes, and adaptation (Arm C)

XLSForm group(s): prediction_confidence3, extreme_expect_group3, adapt_group3.
Arms: **Arm C only.**

No.	Variable	Question, type, options, and routing
Block prediction_confidence3 — *"Prediction Confidence"*. *Block relevance:* Arm C only. ${arm} = 'C'		
314	conf_amount3 → conf_rain_C *select_one conf_scale*	Confidence in rainfall amount expectations *Required.* *Options:* the 5-item conf_scale list, given in full at Q226.
315	conf_onset3 → conf_onset_C *select_one conf_scale*	Confidence in onset expectations *Required.* *Options:* the 5-item conf_scale list, given in full at Q226.
Block extreme_expect_group3 — *"Expectations of extreme events"*. *Block relevance:* Arm C only. ${arm} = 'C'		
316	extreme_future_any3 → extreme_future_any_C *select_one yn*	Do you expect a major weather event in the upcoming season *Options:* the 2-item yn list, given in full at Q4.
317	extreme_future_type3 → extreme_future_type_C *select_one extreme_type*	Type of expected event *Options:* the 5-item extreme_type list, given in full at Q230. *Relevance:* Asked only if extreme_future_any3 was answered Yes. ${extreme_future_any3} = 'yes'

continued from previous page

No.	Variable	Question, type, options, and routing
318	extreme_future_type3_other → extreme_future_type_C_other *text*	If other, specify *Relevance:* Asked only if "Other" was chosen at extreme_future_type3. ${extreme_future_type3} = 'other'
319	extreme_future_severity3 → extreme_future_severity_C *integer*	Severity of this event on a scale from 1 to 10 (1= Not Severe at All, 5= Moderately Severe, 10 = Extremely Severe) *Relevance:* Asked only if extreme_future_any3 was answered Yes. ${extreme_future_any3} = 'yes' *Constraint:* . >= 1 and . <= 10 *Message:* "Value must be between 1 and 10"

Block adapt_group3 — *"Planned adaptation strategies based on expectations"*.
Block relevance: Arm C only. ${arm} = 'C'

No.	Variable	Question, type, options, and routing
320	adapt_any3 → adapt_any_C *select_one yn*	Are you planning to change your farming practices based on your rainfall expectations *Options:* the 2-item yn list, given in full at Q4.
321	adapt_types3 → adapt_C_variety ... adapt_C_other (10 dummies) *select_multiple adapt_ops*	What changes are you planning to implement *Options:* the 10-item adapt_ops list, given in full at Q265. *Relevance:* Asked only if adapt_any3 was answered Yes. ${adapt_any3} = 'yes'
322	adapt_types3_other → adapt_C_other_specify *text*	If other, specify *Relevance:* Asked only if the option other was ticked at adapt_types3. selected(${adapt_types3}, 'other')
323	adapt_cost3 → adapt_C_cost_ugx *decimal*	Estimated total cost of these adaptations (in Ugandan Shillings) *Relevance:* Asked only if adapt_any3 was answered Yes. ${adapt_any3} = 'yes' *Constraint:* . >= 0 *Message:* "Value must be zero or positive"
324	adapt_finance3 → adapt_C_fin_savings ... adapt_C_fin_other (8 dummies) *select_multiple finance*	How do you plan to finance these adaptations *Options:* the 8-item finance list, given in full at Q268. *Relevance:* Asked only if adapt_any3 was answered Yes. ${adapt_any3} = 'yes'

continued from previous page

No.	Variable	Question, type, options, and routing
325	adapt_finance3_ other → adapt_C_fin_other_ specify *text*	If other, specify *Relevance:* Asked only if the option other was ticked at adapt_finance3. selected(${adapt_finance3}, 'other')

OA-2.19 Plot location

Arms: All arms.

No.	Variable	Question, type, options, and routing
326	latitude_farm_ manual *text*	Latitude coordinate of interviewees farm (typical range for eastern Uganda: 0.5 to 2.5) *Required.* *Appearance:* numbers
327	longitude_farm_ manua → longitude_farm_manual *text*	Longitude coordinate of interviewees farm (typical range for eastern Uganda: 33 to 35) *Required.* *Appearance:* numbers
328	gps_main_plot → gps_main_plot_label, gps_lat, gps_lon, gps_ alt, gps_precision *geopoint*	GPS of main farming plot *Required.*

OA-2.20 Items in the form that we could not fully resolve

Six features of the XLSForm are worth flagging for anyone replicating from it, none of which changes what was asked but all of which will trip up a naive read of the workbook.

First, the choices sheet contains five blank spacer rows with an empty list_name and an empty name: four of them separate the four five-point scales conf_scale, conf_scale2, conf_scale3, and conf_scale4 from one another and from extreme_type, and one sits between finance and village_list. They define no options and are inert, but a script that groups the sheet by list_name will report a phantom, unnamed, five-item choice list.

Second, the choice list arm, with the three labels "Group A", "Group B", and "Group C", is declared in the choices sheet but is never referenced by any question. Arm membership is stored by the `calculate` field `arm`, which writes the bare letters `A`, `B`, and `C`, so those labels never reach the data.

Third, the variable `longitude_farm_manua` is missing its final letter in the form. The analysis dictionary repairs it to `longitude_farm_manual`, and because that dictionary addresses columns by position rather than by header text the repair is silent; the two names must be matched by column position.

Fourth, the plant-week, harvest-week, and critical-month multi-selects for the second and third crops are collected by the form but are assigned no analysis name in `survey_dictionary.R`, so they are present in the raw export and absent from the analysis frame. For the first crop, only the harvest-week and critical-month selections are named; the plant-week selection is not.

Fifth, the constraint message on all nine of the "at most two" crop-calendar multi-selects reads "Please select one or two weeks", including on the three *months* questions, where the unit is a month and not a week. This is a copy-and-paste artifact in the message string only. The constraint itself, `count-selected(.)  >= 1 and count-selected(.)  <= 2`, is correct in every case.

Sixth, two option labels contain spelling slips that we reproduce verbatim above because they are what the enumerator read from the screen: "5 Extremely Surpised" in the surprise scale `conf_scale2`, and "Sesame Simsım" in the crop list, where the intended word is the Ugandan term *simsim* for sesame.

OA-3 Token elicitation protocol

Both belief elicitations followed standardized enumerator instructions. Respondents were given exactly ten physical tokens, counted in front of them before each exercise. Enumerators were instructed not to suggest an answer, characterize any outcome as normal or desirable, or allocate tokens on the respondent's behalf. Respondents could move tokens until satisfied with their allocation, after which the enumerator confirmed that exactly ten tokens had been used before recording the response. The survey form independently enforced the ten-token total.

Seasonal-weather exercise. Before respondents allocated tokens, enumerators described the nine pictures in order. Picture 1 represented extreme drought, picture 5 optimal rainfall and growing conditions, and picture 9 severe flooding, with intermediate pictures representing progressively wetter conditions. The standardized descriptions were:

> Picture 1: "An extreme drought. The soil is cracked and plants are dry or dead and almost nothing can grow."
>
> Picture 2: "A drought. There is very little rain and plants are badly stressed."
>
> Picture 3: "Very dry conditions. Plants are weak and growing slowly."
>
> Picture 4: "Slightly dry conditions. Plants grow but they are smaller and the soil is drying."
>
> Picture 5: "Optimal rainfall. The soil is moist and plants are healthy."
>
> Picture 6: "Slightly wet conditions. There is more water than usual but plants still grow normally."
>
> Picture 7: "Very wet conditions. Water stays for some days and plants begin to weaken."
>
> Picture 8: "A flood. Water stands on the field and the crop is damaged."
>
> Picture 9: "A severe flood. The whole field is under water and the crop is destroyed."

Enumerators then explained the token allocation using the following standardized instruction:

> "You have ten tokens. Each token is like a vote. You can put your tokens on any of the pictures. If you think one picture shows the season very well you can put all ten tokens on that picture. If you think the season was between two pictures you can divide the tokens between them. You must use all ten tokens."

The same picture card was used in all study villages. The pictures are generic representations of the ordered weather scale rather than photographs of the respondent's own village. Treatment variation therefore does not arise from the elicitation images themselves.

Order and examples. The seasonal-weather exercise was introduced before the planting-onset exercise. Respondents completed it for each remembered season from 2021 through 2025 and for the upcoming 2026 season, with an additional forward elicitation where required by experimental assignment. Worked examples were framed around farmers in other locations and were explicitly identified as examples. Their purpose was to demonstrate that respondents could either concentrate all tokens on one outcome or distribute them across outcomes.

Planting-onset exercise. The onset exercise used the same ten-token procedure over ten dated intervals from 15 February through 24 April. Enumerators introduced the task by explaining that the intervals represented possible dates at which the main rains could begin. Respondents could concentrate or spread their tokens in the same manner as in the seasonal-weather exercise. The fielded calendar is reproduced in the paper and its interval construction is documented in Appendix B.

Respondents unable to recall a season. Enumerators could use up to two neutral probes before recording that a respondent could not remember a past season. Respondents who continued to report no memory were not required to allocate tokens. The resulting recall indicator therefore measures whether a season could be elicited under this protocol, not whether the underlying memory objectively existed.

Standardization rules. Respondents who requested more or fewer than ten tokens were reminded that ten tokens were used for every participant. Respondents reluctant to use extreme pictures were reminded that any picture was permissible. These instructions standardized the elicitation without directing respondents toward the center or tails of the scale.

OA-4 Enumerator manual

Enumerators worked from a 59-page field manual dated 19 December 2025. The manual presents the questionnaire screen by screen and distinguishes respondent-facing text from enumerator instructions. It documents the permitted probes, routing rules, consistency checks, and common administration errors. The manual was finalized before the main fieldwork and was used throughout the March 2026 survey.

For completeness, the fielded manual is provided as a separate supplementary file, *Pilot Enumerator Manual, 19 December 2025*. It is reproduced without editorial revision.

OA-5 Consent and ethics documentation

Written informed consent was obtained from all respondents before the survey began. Consent was recorded on paper using a standardized information-and-consent form and retained separately from the tablet questionnaire. The form identifies Erasmus University Rotterdam and the Erasmus Research Institute of Management (ERIM) as the sources of research support and records ethics approval under reference ETH2425-0832. The approval was granted by the Internal Review Board of the Erasmus School of Economics on 20 May 2025 and was valid for three years.

The consent form explains the purpose of the study, the voluntary nature of participation, the types of personal and research data collected, confidentiality and data-retention procedures, possible discomfort associated with recalling adverse weather events, and respondents' right to stop the interview or decline individual questions. It also states that anonymized research data may be retained and reused for research and educational purposes.

Consent was recorded on paper rather than through the electronic questionnaire. The fielded survey form contains no digital consent item, while the study documentation records signed and scanned consent forms among the retained project materials. The fielded consent form was written in English. No written translated version or back-translation is preserved in the study records.

The original fielded consent form is provided as a separate supplementary file, *Consent Form – Fielded Version*. We reproduce the original document rather than an edited transcription so that handwritten amendments, strike-throughs, formatting, and unresolved template language remain visible exactly as they appeared in the field.

OA-6 Fieldwork procedures

Fieldwork took place from 10 through 31 March 2026 across the 60 study villages. Seven enumerators administered the survey face to face using the fielded electronic questionnaire, printed elicitation materials, and the enumerator manual described in Section OA-4.

Interview administration. The manual instructed enumerators to conduct interviews in a private setting where possible, record only answers given by the respondent, follow the questionnaire routing as programmed, and consult the field team rather than improvise when an item was unclear. Respondent-facing wording was distinguished from enumerator instructions in the field materials. The two token exercises followed the standardized protocol described in Section OA-3. Enumerators also received a one-page field aid summarizing the two token tasks.

Built-in data checks. The electronic questionnaire imposed a series of validation rules during administration. Token allocations were required to sum to ten, individual token entries were restricted to the feasible range, acreage and household variables were constrained to nonnegative values where appropriate, and event severity and event year were restricted to their permitted ranges. Villages were selected from the fixed study list and subsequently confirmed, while experimental arm was assigned automatically from village identity rather than entered by the enumerator. Free-text comment fields allowed enumerators to record unusual circumstances during the belief modules.

The district variable remained a free-text field and consequently contains multiple spellings of the three study districts. All geographic assignment in the analysis therefore uses the standardized village-to-district mapping rather than this free-text field.

Consent. Written consent was obtained before interviewing using the form reproduced in Section OA-5. Consent was recorded on paper and retained separately from the tablet questionnaire.

Language. The electronic questionnaire and enumerator manual preserve the field instrument in English and contain no stored written translations. Interviews were conducted face to face, with oral translation by enumerators where needed. Because the study records do not preserve a written translation protocol or identify the language used in each interview, we do not assign a specific language of administration at the respondent level.

Training and monitoring documentation. The surviving study records do not contain a sufficiently detailed training log to reconstruct the duration or content of enumerator training beyond the field manuals and pilot materials described here. They likewise do not contain a separate back-check dataset or systematic back-check log. We therefore restrict claims about field quality control to the procedures documented in the instrument and field materials.

OA-7 Pilot and changes to the field protocol

The pilot was conducted on 24 and 29 December 2025 in Namikelo, Sango, and Burukuru villages in Bulambuli district on the slopes of Mount Elgon. Thirty interviews were completed, ten in each village, by five enumerators. All pilot respondents followed the same experimental sequence. No pilot observation is included in the analysis reported in the paper.

Comparing the pilot questionnaire and materials with those fielded in March 2026 identifies eleven changes.

1. **Recall years were changed to a consecutive five-season window.** The pilot elicited 2025, 2024, 2023, 2021, and 2019. The final instrument instead elicited each season from 2021 through 2025. This aligned the recall window with the five seasons displayed in the information treatment.

2. **Village identification was converted from free text to a closed list.** The final questionnaire required selection from the 60 study villages followed by explicit confirmation, eliminating spelling variation in the village identifier.

3. **Experimental assignment was automated.** In the pilot, enumerators selected the assigned arm from the allocation sheet. In the final instrument, experimental arm was determined automatically from village identity. Enumerators therefore had no discretion over assignment once the village was selected correctly.

4. **Enumerator comment fields were added.** The final questionnaire introduced comment fields after the recall and forward-belief modules so that unusual responses or administration circumstances could be recorded.

5. **Interview date and village identifier fields were added.** The final questionnaire records an interview date and numeric village identifier alongside the respondent identifier.

6. **A manual coordinate fallback was added.** Numeric latitude and longitude fields were introduced as a fallback when the primary map-based location capture could not be used.

7. **Separate "other" fields were added to decision-maker questions.** The final instrument records a separate text response for each ranked decision-maker item when "other" is selected.

8. **A crop-ranking item was removed.** A separate rank-order question used in the pilot was dropped because the final questionnaire already collected information on the respondent's three principal crops through dedicated crop modules.

9. **The post-information reaction block was shortened.** Three reaction items used in the pilot were removed before the main fieldwork. The final instrument retained questions about whether the information differed from memory, surprise, willingness

to revise expectations, confidence after revision, usefulness of historical information, and demand for similar information in the future. As discussed in the paper, reaction questions still occur between the information treatment and the second forward elicitation in Arm B but not in Arm A.

10. **The planting-onset calendar was corrected.** The pilot materials attached year labels to the calendar and included an erroneous February date in the 2026 calendar. The final field card removed year labels so that the same instrument could be used for both recalled and expected onset and correctly defined the third interval as 1–6 March.

11. **Questionnaire metadata were standardized.** Form labels and version information were updated for the final field instrument.

Elements retained from the pilot. The nine-picture weather scale, ten-token format, ten planting-onset intervals, experimental-arm structure, and core read-aloud instructions were retained for the main study. One site-specific feature of the onset calendar also remained. Several interval descriptions refer to highlands and slopes, reflecting the Elgon pilot location, although the main study was conducted on the lower-elevation plateau. We therefore do not interpret absolute cross-sectional differences between reported and satellite-based onset as measurement of forecast accuracy. As Appendix B of the paper explains, onset is used only for within-person and randomized comparisons.

The available study records do not contain a separate written pilot debrief. We therefore restrict this account to changes that can be established by comparing the pilot and fielded instruments and materials.

OA-8 Computational details for statistical inference

Appendix C of the paper describes the inferential procedures and reports the resulting estimates. This section provides additional computational detail on the accelerated wild cluster bootstrap, the Romano–Wolf stepdown procedure, and the double-selection specifications. File-level replication instructions are provided separately in the replication package.

OA-8.1 Wild cluster bootstrap implementation

The restricted wild cluster bootstrap is computationally demanding because each hypothesis requires repeated estimation of the model and its cluster-robust variance. We therefore use a matrix implementation that partials out fixed effects before the bootstrap replications, precomputes the relevant cross-products, and subsequently updates only the bootstrap outcome and studentized statistic. For specifications in which villages are nested within a fixed-effect unit, the weighted restricted residual is itself residualized within each bootstrap replication before the cluster-robust variance is computed. This preserves the residual transformation that would result from re-estimating the complete fixed-effect model.

The accelerated implementation was validated against direct repeated estimation on smaller bootstrap samples. Coefficients and test statistics coincide numerically, while differences in simulated p-values at full replication counts are within the Monte Carlo variation implied by the bootstrap.

OA-8.2 Romano–Wolf stepdown

The stepdown procedure follows Romano & Wolf (2005) and the resampling implementation described by D. Clarke et al. (2020). A single set of cluster-level bootstrap weights is shared across all hypotheses in a family. This is important because individual hypotheses can be estimated on different subsets of respondents while sharing villages. Whenever a village enters more than one hypothesis, it therefore receives the same bootstrap weight in the same replication, preserving dependence among the test statistics.

For each hypothesis, unrestricted and null-restricted specifications are estimated as described in Appendix C. Bootstrap outcomes are generated under the null and studentized statistics are calculated for each replication. Statistics are transformed according to the direction of the registered test, using the signed statistic for one-sided hypotheses and its absolute value for two-sided hypotheses. Hypotheses are ordered by their observed transformed statistic. At each step, the bootstrap reference statistic is the maximum over the hypotheses remaining in the family. Adjusted p-values are calculated from the corresponding exceedance frequency and made weakly increasing along the stepdown sequence. A bootstrap replication is retained for the family only when every hypothesis produces a finite statistic, preserving the common-draw structure.

Accelerated implementations are used for families sharing a common regression structure. These implementations were validated against the general procedure on smaller bootstrap samples before being used for the full set of replications.

OA-8.3 Double selection

Double selection is used in the heterogeneity analyses to reduce discretion over which correlated pre-treatment characteristics enter as controls (Belloni et al., 2014). For each target predictor, selection is conducted twice: once with the outcome as the dependent variable and once with the target predictor as the dependent variable. The final control set is the union of variables selected in these two equations. The reported coefficient is then obtained by re-estimating the outcome equation on the target predictor and the selected controls, with standard errors clustered at the village level.

Updating screen. The candidate set contains 22 pre-treatment characteristics. Continuous characteristics are standardized over the pooled Arms A and B sample so that coefficients are comparable across arms. Binary characteristics enter without standardization. Penalization uses the theoretically motivated heteroskedasticity-robust penalty with post-selection refitting rather than cross-validation. The final regressions use village-clustered inference.

Memory-quality screen. The same procedure is applied to the candidate predictors of memory accuracy, concentration, signed bias, and related memory outcomes. Subcounty indicators and experimental-arm assignment are always retained rather than subjected to selection. They are partialled out of the outcome and candidate predictors before the selection step. Candidate controls that are perfectly collinear after this transformation are excluded. Specifications with insufficient cluster-level degrees of freedom relative to the number of free parameters are classified as not estimable rather than retained with unstable cluster-robust inference.

The replication package reports the complete candidate sets, penalty parameters, estimability rules, random seeds, and numerical validation checks.