Seeds of (climate) change: Private adaptation and subsidized insurance in West Bengal

Pre-analysis plan

Fiona Burlig, Manzoor H. Dar, Xavier Giné, Erin M. Kelley, and Gregory Lane*

May 2024

Abstract

Adaptation is required to cope with climate change. Theory predicts that agricultural insurance, which protects farmers against climate damages, and enjoys billions of subsidy dollars, may either increase *or* decrease private adaptation. Subsidized insurance may crowd in adaptation by limiting farmers' risks from experimenting with new technologies. However, it may instead crowd out adaptation, by insulating farmers against climate risk. We test which of these effects dominates with an RCT in West Bengal, India. We randomize 300 villages into a control group and index insurance arms, where farmers receive payouts if floods occur. We estimate the impact of insurance on farmer willingness-to-pay for both flood-tolerant and high-yield-variety seeds, providing a direct test of the impacts of insurance on demand for adaptation. We also induce random variation in seed take-up, in order to estimate the effects of specialty seeds, insurance, and their interaction on agricultural inputs and *ex post* welfare outcomes.

Keywords: Climate; adaptation; insurance; agriculture
JEL Codes: D81; D25; O12; O13; Q12; Q54

*Burlig: Harris School of Public Policy and Energy Policy Institute (EPIC), University of Chicago and NBER. Email: burlig@uchicago.edu. Dar: International Crops Research Institute for the Semi-Arid Tropics. Email: m.dar@cgiar.org. Giné: Development Research Group, World Bank. Email: xgine@worldbank.org. Kelley: Harris School of Public Policy, University of Chicago, and Development Impact Evaluation Unit, World Bank. Email: erinmkelley@uchicago.edu. Lane: Harris School of Public Policy, University of Chicago and NBER. Email: laneg@uchicago.edu. We are grateful to Susanna Berkouwer, Josh Dean, Kyle Emerick, Michael Kremer, and Andrew Plantinga for helpful comments and suggestions. We thank Sam Hsu, Mohit Negi, and Amrita Pal for excellent research assistance. We thank the Development Economics Center at the University of Chicago and the World Bank for generously funding this project. This research has IRB approval from the University of Chicago (Protocol No. IRB24-0719), and is registered on the AEA RCT registry (Identification No. AEARCTR-XXX). All remaining errors are our own.

1 Introduction

Climate shocks reduce welfare, particularly among the poor. The World Bank estimates that extreme weather events are responsible for pushing 26 million people into poverty every year, generating consumption losses of approximately 520 billion US dollars annually (Hallegatte et al. (2017)). As the planet warms, the frequency and severity of these extreme events is rising (Hirabayashi et al. (2013); Burgess et al. (2017)). Adaptation—measures taken to directly reduce climate damages, such as air conditioning use, sea wall construction, or planting climate-resilient seeds—will be essential for coping with climate change (Barreca et al. (2016)). However, adaptation remains constrained, particularly in low-income countries (Carleton et al. (2022), Lane (2024)). To what extent will widespread, substantial, and rapidly-growing subsidies for insurance that protects agents against climate damages (e.g., FEMA's National Flood Insurance Program; the USDA's Federal Crop Insurance Program; India's Pradhan Mantri Fasal Bima Yojana, and products offered by the People's Insurance Company of China) limit private climate adaptation?

We use a simple model of decision-making under risk to investigate how subsidized insurance impacts the demand for private climate adaptation. Under actuarily fair insurance pricing, and absent other market failures, farmers would choose the private (and also socially) optimal level of investment in both insurance and other adaptation technologies. However, we can then show that subsidized insurance has two counterveiling effects on adaptation investment, in the spirit of Ehrlich and Becker (1972).[1] On one hand, because climate insurance and private climate adaptation both protect agents against the same risks, subsidies that lower the price of insurance may lead to a "crowd-out" effect. On the other hand, because climate insurance protects agents from facing the full brunt of the bad state of the world if a costly adaptation measure fails, these subsidies may instead lead to a "crowd-in" effect. Which of these effects ultimately dominates is ambiguous.

We set up and implement a sharp empirical test of the impacts of subsidized flood insurance on demand for one particular technology – flood-tolerant seeds – in order to shed important light on the effect of these subsidies more broadly. We take the predictions of the model to data by implementing a cluster-randomized trial in 300 villages in West Bengal in conjunction with the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT), a well-respected international organization known for developing advanced crop varieties. We randomize villages into three groups (sampling 6 farmers per village): a control group; a group that receives a fully-subsidized index insurance product that pays out a low

[1]Note that in our model, we abstract away from a third channel: moral hazard, whereby individuals manipulate their outcomes in order to receive an insurance payout. In our experiment, we use an index insurance product whose payouts do not depend on farmer behavior to exclude this possibility.

amount in the event of a flood; and a group that receives a fully-subsidized index insurance product that pays out a high amount in the event of a flood. After offering (treated) farmers the insurance product, we elicit every sampled farmers' willingness-to-pay for both a flood-tolerant seed minikit and a high-yield-variety seed minikit using a Becker, DeGroot, and Marschak (1964) mechanism. We test for crowd-in vs. crowd-out by estimating the impact of insurance on demand for flood-tolerant seeds, high-yield-variety seeds, and the difference between the two. We further measure the extent to which these impacts vary with insurance payout levels.

In addition, through the BDM process, we induce experimental variation in seed take-up: in each village, at random, one farmer will be offered the high-yield-variety minikit for the market price and one farmer will be offered the flood-tolerant seed minikit for the market price (together making up the "seed control group"); two farmers will receive the high-yield-variety seed minikit for free; and two farmers will receive the flood-tolerant seed minikit for free. This design enables us to also study how specialty seeds, insurance, and their interaction impact *ex ante* agricultural input decisions and *ex post* welfare outcomes. Finally, we test how these impacts change with flood realizations.

The main contribution of this study is to experimentally estimate the extent to which climate insurance crowds in or crowds out demand for private climate adaptation. In doing so, we contribute to a literature on the economics of insurance in both rich and poor countries. Recent work in the US suggests that public protection against climate shocks such as hurricanes (e.g., Kousky, Luttmer, and Zeckhauser (2006); Deryugina (2017)), wildfires (e.g., Baylis and Boomhower (2021)), and adverse growing season weather (e.g., Annan and Schlenker (2015)) may limit private adaptation.[2]

In contrast, the development literature on climate insurance has broadly argued for expanding the provision of climate insurance, arguing that insurance can meaningfully improve farmer welfare by enabling risky but profitable investments (e.g., Karlan et al. (2014)), but demand remains limited (Jensen and Barrett (2017); Carter et al. (2017); Cole and Xiong (2017)), potentially due to market failures such as liquidity constraints (e.g., Casaburi and Willis (2018)). Our experiment provides evidence—within a single, unified setting—on both the crowd-in vs. crowd-out impacts of crop insurance on adaptation and the direct impacts of such insurance (at two different coverage levels) on farmer welfare, allowing us to comment on the efficiency of such subsidies.

Our research stands to be particularly important in light of the massive growth in subsidized climate insurance currently taking place in low-income countries. Since the 2007

[2]In a related literature, Boomhower et al. (2023) and Ostriker and Russo (2024) consider the economics of adaptation mandates for fires and floods, respectively.

launch of China's agricultural insurance subsidy program, these subsidies have totaled over 41.4 billion USD, with year-over-year growth of 22% (Peoples' Republic of China (2024)). In India, the location of this study, the 40 million farmers enrolled in the central PMFBY crop insurance scheme as of 2023-24—up 27% from 2022-23—are only charged up to 2% of the actuarily fair premium. Measuring the extent to which these insurance programs are ushering in or hindering climate adaptation is therefore critical for understanding their welfare consequences.

In addition, we contribute important *in-situ* estimates of the effectiveness of a widely-promoted climate adaptation technology—specialty seeds—to a climate change economics literature which has largely focused on mitigation (e.g., Nordhaus (1993); Pindyck (2013)) or on the costs of climate change (e.g., Carleton and Hsiang (2016); Hsiang et al. (2017)). A more recent literature estimates climate damages accounting for adaptation (Auffhammer (2022); Carleton et al. (2022); Rode et al. (2021) Hultgren et al. (2022)), but does not directly evaluate explicit adaptation approaches. In contrast, we use an experiment to build on a small development literature which explicitly tests individual climate adaptation approaches (e.g., Lane (2024); Aker and Jack (2023); Burlig et al. (2024)). In doing so, our work is most closely related to Emerick et al. (2016), who measures the benefits of Swarna-Sub1 in Odisha, and Boucher et al. (2022), who test the effectiveness of drought-tolerant seeds in Africa. We build on this research by estimating demand both for high-yield-variety and flood-tolerant seeds, and measuring how both demand for and use of these seeds varies with more vs. less generous insurance.

The remainder of this pre-analysis plan proceeds as follows. Section 2 provides relevant details about the research setting. Section 3 presents a simple theoretical model of farmer decision-making under risk. Section 4 describes our experimental design. Section 5 discusses the data we will collect over the course of our experiment. Section 6 presents our planned analysis, including a detailed outline of our regression specifications. Section 7 concludes.

2 Research context

2.1 Flood-prone agriculture in West Bengal

West Bengal is an ideal setting for this study. The fourth-most populous state in India, West Bengal is home to more than 7.1 million farming households. Small-holder agriculture is common: 96% of farmers are small and marginal, with an average land-holding of only 0.77 hectares per household (Government of West Bengal (2024)). At the same time, flooding is extremely common: more than 42 percent of the state's land is deemed flood-prone. In the 41 years between 1960 and 2000, severe flooding occurred in all but five (Irrigation and Waterways Department (2024)).

Paddy rice is an extremely important crop in West Bengal: of the state's 5.5 million hectares of land under cultivation, 5.8 million are planted to rice, which has an average productivity of 2.6 tonnes per hectare (Indian Council of Agricultural Research (2024)).[3] Of this paddy land, 40% is prone to submergence (Raghu, Veettil, and Das (2022)). If traditional-variety rice crops are submerged for more than 14 days, crop losses will be 100% (Ismail et al. (2013)). According to a 2016 survey of nearly 5,000 farmers across Assam, Odisha, and West Bengal, however, adoption of the promising flood-tolerant variety Swarna-Sub1 was only 9%, suggesting that private adoption of specialty seeds for risk mitigation is low (Raghu, Veettil, and Das (2022)).

West Bengal does not participate in the national Pradhan Mantri Fasal Bima Yojana crop insurance scheme. Instead, the government operates its own program, Bangla Sashya Bima. In conversations with farmers and sarpanches, we understand that payouts under this scheme are relatively rare.[4] In 2023, despite the India Meterological Department identifying seven of West Bengal's 23 districts as anomalously "moderately dry" or "severely dry," (India Meterological Department (2023)), payouts for deficit rainfall were distributed to fewer than 250,000 farmers across the state.[5] Per internal ICRISAT survey data from 2022, only approximately 34% of farmers in West Bengal insured their crops, further suggesting limited engagement with both public and private insurance schemes.

2.2 Implementing partner

We partner with ICRISAT. ICRISAT is an international organization, headquartered in Hyderabad, Telangana. They have over 50 years of experience in India, and are known across the country for breeding and disseminating high-performance crops. ICRISAT is a trusted partner for farmers and local extension services. The organization has a broad network of local partners, enabling them to act as a trustworthy supplier of seeds and insurance. By working with ICRISAT, we give this project the best possible chance of success, increasing farmer trust in the products.[6]

[3] West Bengal engages in significant double- and triple-cropping, with a cropping intensity of 176% across agricultural land.

[4] To the extent that farmers already have access to some crop insurance that protects against floods, this will attenuate our treatment effects relative to a world with no existing insurance.

[5] https://www.deccanherald.com/india/west-bengal/bengal-govt-offers-crop-insurance-money-to-farmers-affected-by-inadequate-rainfall-2730865

[6] Members of the research team also partnered with ICRISAT in Burlig et al. (2024), where we were able to successfully distribute monsoon onset forecasts and insurance against a late monsoon. Though measuring farmer trust is difficult, we find that farmers substantially changed their agricultural inputs in response to

3 Model

To formalize the main argument that access to subsidized weather index insurance can influence climate change adaptation, imagine a farmer living in a flood-prone area that can grow the main crop (i.e. rice) using three types of seeds: flood-tolerant, high-yielding or traditional. Both flood-tolerant and high-yielding seeds cost roughly the same and are more expensive than traditional seeds. Flood-tolerant seeds outperform the other seeds in case of floods, and so can be thought of as a climate resilient technology. In contrast, high-yielding seeds outperform the other seeds only in the absence of floods.

Our goal is to estimate the willingness to pay (WTP) for high-yield and flood-tolerant seeds and traditional seeds depending on whether or not the farmer has access to subsidized insurance. In particular, if WTP for a seed increases under insurance, then access to insurance has "crowd in" effect on demand for that seed. In contrast, if WTP lowers, then insurance has a "crowd-out" effect.[7]

Output from high-yielding, flood tolerant and traditional seeds in the high state H (i.e. no flood) is Y_{HH}, Y_{HF} and Y_{HT}, respectively, which occurs with probability p and Y_{LH}, Y_{LF} and Y_{LT} in the low state L (i.e. a flood) with probability $1-p$. We assume that high-yielding seeds outperform the other seeds in state H, $Y_{HH} > Y_{HF} = Y_{HT} > 0$ and that flood tolerant seeds outperform the other seeds in state L, $Y_{LF} > Y_{LH} = Y_{LT} = 0$. The cost of high-yield, flood tolerant and traditional seeds is C_H, C_F and 0, respectively. When farmers are offered flood insurance for free, the insurance pays out Y^I (either a high, $\bar{Y}^I$, or low, $\underline{Y}^I$, amount) in the low state L. We also assume that farmers have illiquid assets W.

Farmers, however, believe that high-yielding or flood-tolerant seeds may not perform as intended. In particular, we assume that the farmer believes that with probability $(1-\pi_F)$ the flood-tolerant seeds will perform as traditional seeds in the low state L, and that with probability $(1-\pi_H)$ the high-yielding seeds will perform as traditional seeds in the high state H. For simplicity, we assume that $\pi_F = \pi_H = \pi$.[8]

The WTP for flood-tolerant and high-yield seeds will depend on the probability π that the seeds perform as intended and whether the farmer is insured. We denote by B_F^I (B_F^U) the WTP which refers to the amount of money that the farmer is willing to pay for seeds F when insured (uninsured) such that the farmer is indifferent between using seeds F and the traditional seeds. Thus, if the actual cost of seed F satisfies $C_F \leq B_F^I$, then the farmer will

both products, providing revealed-preference evidence that farmers found these products credible.

[7]Note, we could also conceptualize "crowd-in" and "crowd-out" effects as the difference in WTP between the high-yielding and flood-tolerant seed. This difference, rather than WTP above the traditional variety, would determine actual crop choice if both types of seeds were freely available in the market at the same price.

[8]This belief could arise due to fears of counterfeit seed or skepticism of the technology itself.

adopt seeds F when the farmer is insured, and the traditional seeds otherwise. The WTP B_F^I satisfies $U_F^I\left(B_F^I\right) = U_T^I$, and B_F^U satisfies $U_F^U\left(B_F^U\right) = U_T^U$, where U_i^j refers to the expected utility of planting seed i ($i = F, H$), under insurance coverage j, where $j = U$ for uninsured farmers and $j = I$ for insured farmers. In particular, B_F^I and B_F^U solve, respectively:

$$pu(Y_{HF} + W - B_F^I) + (1-p)[\pi u(Y_{LF} + Y^I + W - B_F^I) + (1-\pi)u(Y_{LT} + Y^I + W - B_F^I)] = \\ pu(Y_{HT} + W) + (1-p)u(Y_{LT} + Y^I + W) \tag{1}$$

and

$$pu(Y_{HF} + W - B_F^U) + (1-p)\left[\pi u\left(Y_{LF} + W - B_F^U\right) + (1-\pi)\,u\left(Y_{LT} + W - B_F^U\right)\right] = \\ pu\left(Y_{HT} + W\right) + (1-p)\,u\left(Y_{LT} + W\right). \tag{2}$$

Analogously, B_H^I and B_H^U solve

$$p[\pi u\left(Y_{HH} + W - B_H^I\right) + (1-\pi)\,u\left(Y_{HT} + W - B_H^I\right)] + (1-p)\ u\left(Y_{LH} + Y^I + W - B_H^I\right) = \\ pu\left(Y_{HT} + W\right) + (1-p)\,u\left(Y_{LT} + Y^I + W\right) \tag{3}$$

and

$$p\left[\pi u\left(Y_{HH} + W - C_H\right) + (1-\pi)\,u\left(Y_{HT} + W - C_H\right) + (1-p)\ u\left(Y_{LH} + W - C_H\right) = \\ pu\left(Y_{HT} + W\right) + (1-p)\,u\left(Y_{LT} + W\right). \tag{4}$$

We parameterize this model in order to show predictions of changes in WTP for each seed type both with and without insurance.[9] Panel A of Figure 1 plots the change in WTP for flood-tolerant seed when an uninsured farmer is offered insurance. In particular, it plots the change in WTP as a function of farmer beliefs that flood-tolerant and high-yield seed will perform as expected (π) when the farmer is insured with high coverage: $\Delta^F\left(\pi; Y^I\right) = B_F^I\left(\pi; Y^I\right) - B_F^U\left(\pi\right)$. The figure plots this relation ship for the high level of insurance coverage in light blue, and the low level of insurance coverage in navy blue.

Panel A illustrates the ambiguous effect of crop insurance on adaptation. When farmers have relatively low beliefs about the probability of seed success, then flood insurance crowds-in demand for flood-tolerant seeds. This arises because flood insurance offers some protection in the case when a flood occurs and the seeds are revealed to not be, in fact, flood-tolerant.

[9] We assume for this exercise that the utility function $u\left(c\right)$ is CRRA, $u\left(c\right) = \frac{c^{1-\sigma}-1}{1-\sigma}, 0 < \sigma < 1$.

However, if farmers are more confident in the quality of the seed, then insurance coverage begins to crowd-out demand. Higher levels of insurance coverage make crowd-out more likely.

Panel B in Figure 1 plots the difference in WTP between the flood-tolerant seed and high-yield seed as a function of π for different levels of insurance coverage: $\Delta^I(\pi) = B_F^I(\pi) - B_H^I(\pi)$. The gray line shows the difference without insurance, the navy blue line for a low coverage insurance, and the light blue line for high coverage insurance. As expected, without insurance, farmers are willing to pay more for flood-tolerant seeds than the high-yielding seed due to the fact that farmers are risk-averse and only the flood-tolerant seeds increase income in the flood state of the world. However, as flood insurance coverage increases, farmers are willing to pay less for flood-tolerant seeds relative to high-yield seeds. With a high level of flood insurance, whether WTP for the flood-tolerant seed is higher or lower than the high-yielding seed depends on farmers beliefs about seed quality π.

In sum, Figure 1 illustrates the different potential impacts of subsidized insurance on demand for climate adaptation. If farmers are unsure the effectiveness of the flood-tolerant seed, access to insurance may crowd-in demand. However, the seed are known to be effective with little risk of failure if a flood occurs, then insurance will unambiguously crowd-out demand, and thus hinders adaptation.

Figure 1: Impact of insurance on willingness-to-pay for private adaptation (model)

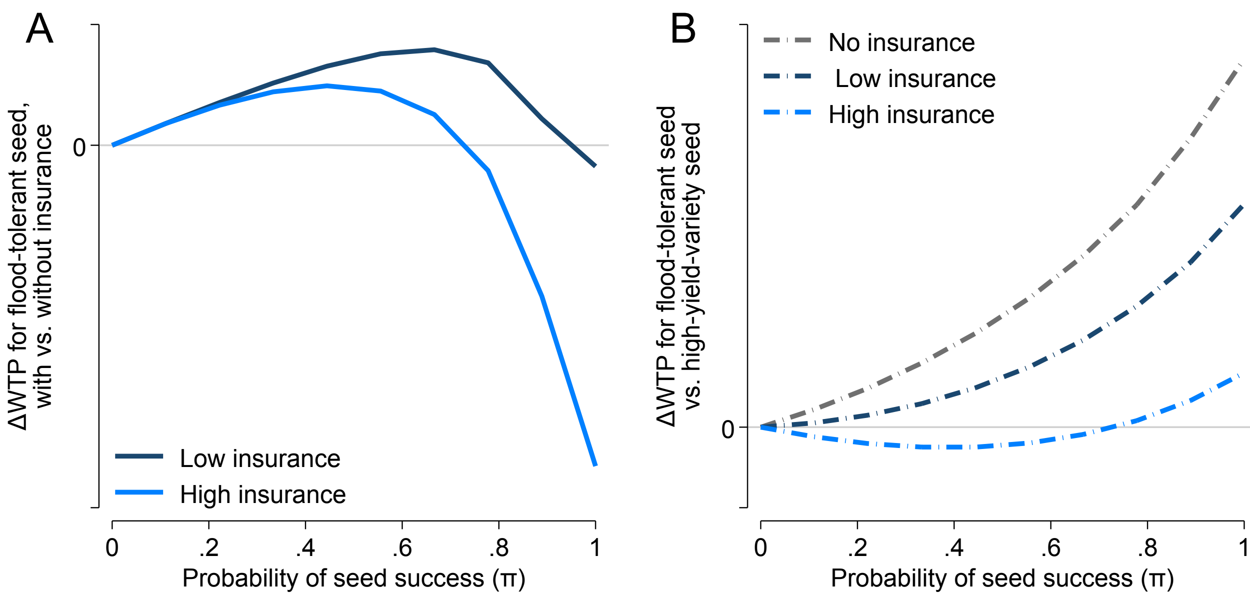

Notes: This figure plots the main results of our model. Panel A shows the effect of insurance on willingness-to-pay (WTP) for flood-tolerant seeds, as the probability of seed success (π) increases. The navy line plots the difference in WTP between low-payout insurance and no insurance, while the light blue line plots the difference between high-payout insurance and no insurance. When the curves lie above zero, insurance crowds in demand for flood-tolerant seeds. When the curves fall below zero, insurance instead crowds out demand for flood-tolerant seeds. Panel B shows the difference between WTP for flood-tolerant seeds and high-yield-variety seeds at varying levels of π. We show results without insurance (gray), with low-payout insurance (navy), and with high-payout insurance (light blue). With no insurance, farmers have higher demand for flood-tolerant seeds than for high-yield-variety seeds. As insurance coverage rises, however, demand for flood-tolerant seeds falls relative to high-yield-variety seeds. This figure is generated with the following values: $Y_{HF} = Y_{HT} = 3$, $Y_{HH} = 5$, $Y_{LF} = 0.5, Y_{LH} = Y_{LT} = 0, \underline{Y}^I = 0.25, \bar{Y}^I = 0.5, W = 0.4, p = 0.5, \sigma = 0.9$.

4 Experimental design

Informed by our model, we will implement a cluster-randomized trial in order to answer three main research questions:

1. How does flood insurance impact the demand for flood-tolerant seeds vs. demand for high-yield-variety seeds? (i.e., does the "crowd-out" effect or the "crowd-in" effect dominate?)

2. How does flood insurance affect farming practices with vs. without specialty seeds?

3. What is the impact of flood insurance, specialty seeds, and their interaction on farmers who experience flooding?

4. How do these effects vary with the level of insurance coverage?

Sampling To answer these questions, we first construct a sample of 300 villages, drawn from three particularly flood-prone districts in West Bengal: Murshidabad, Paschim Medinipur, and Hooghly. We then sample blocks that are (i) relatively close to river gauges operated by the Government of West Bengal's Irrigation and Waterways Department, which we use to trigger our index insurance product, described below; (ii) most frequently flooded in the district, per the Government's Flood Hazard Atlas (Department of Space (2021)); (iii) suitable for Swarna-Sub1, a flood-tolerant seed (described further below), per ICRISAT and the International Rice Research Institute; and (iv) with limited access to existing crop insurance, per a qualitative survey conducted by the research team. Within these blocks, we randomly sample 300 villages, using a geospatial algorithm to ensure that we do not include any neighboring villages in the sample, limiting the possibility of risk-sharing across villages in different treatment arms.

Randomization We randomize in two steps. First, we randomly assign each of our 300 villages into an insurance control group (125 villages who receive no insurance product); a high-payout insurance group (125 villages); and a low-payout insurance group (50 villages).[10] We sample six farming households per village. We exclude households that did not cultivate any rice in the last 3 years, since our research design is focused on improved rice seeds. Every sample household in a given village will receive the same insurance treatment. To ensure balance and increase statistical power, we stratify our village-level randomization by block.

[10] We divide villages into treatment arms unequally because our primary comparison of interest is between control villages and high-payout insurance villages. We therefore design the experiment to maximize power on this comparison, while still allowing us to compare point estimates between high-payout insurance and low-payout insurance to estimate the extent to which treatment effects vary with insurance coverage levels.

Second, we randomize individuals to receive offers to purchase either a flood-tolerant seed minikit or a high-yield-variety seed minikit at either market price or for free. This randomization yields four groups: market-price flood-tolerant seed offer (1/6 of households); free flood-tolerant seed offer (1/3 of households); market-price high-yield-variety seed offer (1/6 of households); and free high-yield-variety seed offer (1/3 of households). Because these seeds are already available outside of the experiment, we treat the market price groups as a "seed control" group, which we pool for the purposes of analysis.[11] In this second step, we stratify the randomization by village.

Figure 2 presents a diagram of the experimental design. The first layer shows the village-level randomization into the three insurance arms. The second layer shows the household-level randomization into seed offers.

[11] One concern with randomizing seed offers at the individual level is spillovers of seeds from seed-treatment farmers to seed-control farmers. We believe this is not a major cause for concern for three reasons. First, we are only interacting with 6 households per village. The probability that, when sharing seeds from a 5kg minikit, the two households in the high-yield-variety seed treatment group or the two households in the flood-tolerant seed treatment group chooses to share seeds with the two households in the seed control group is very low. In 2011 (the latest Census of India), West Bengal had approximately 540 households per village on average. Second, Emerick et al. (2016) provided Swarna-Sub1 seeds to households in Odisha, finding no evidence of within-village spillovers. Finally, any sharing of seeds between treatment and control farmers will attenuate our estimated impacts of these seeds towards zero. As discussed in Section 6, we will estimate LATEs using take-up of seeds as the endogenous variable, instrumenting with free seed offers.

Figure 2: Experimental design

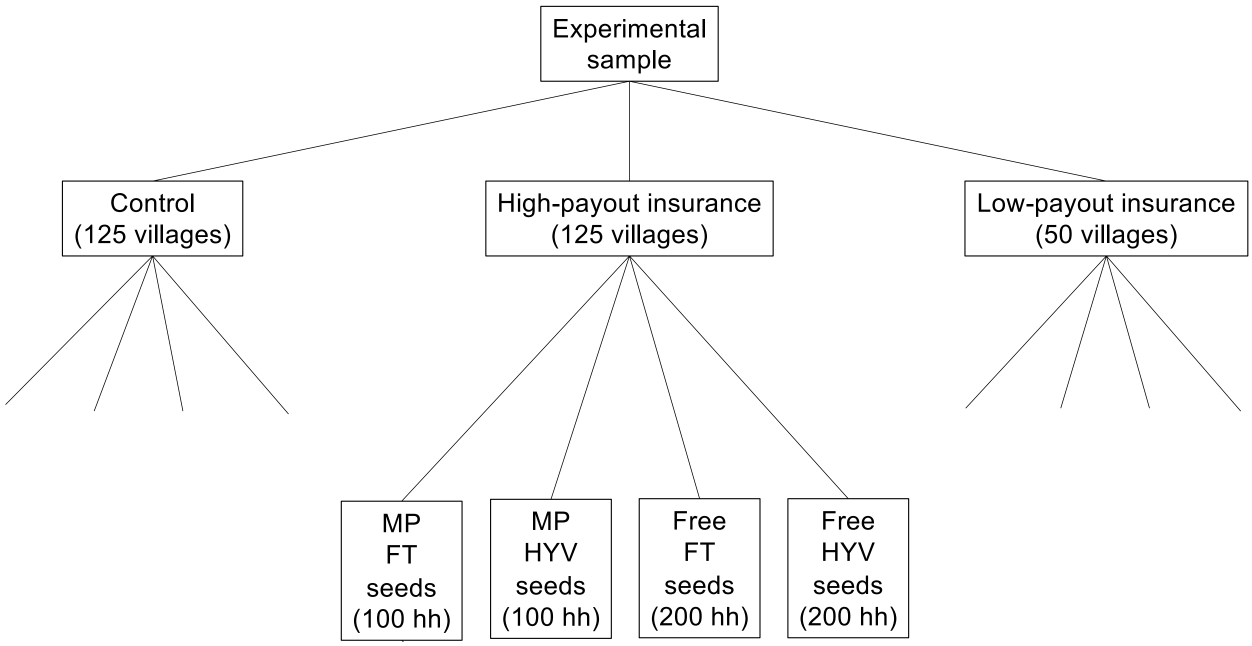

Notes: This figure shows the design for our cluster-randomized experiment. We randomly divide the 300 villages that make up the experimental sample into an insurance control group, a low-payout insurance group, and a high-payout insurance group. We then randomize the six sampled households per village into a group that is offered flood-tolerant (FT) seeds at market price (MP), a group that is offered high-yield-variety seeds (HYV) at market price (MP), a group that is offered FT seeds for free, and a group that is offered HYV seeds for free. The two MP groups together make up the "seed control" group.

Insurance Our implementing partner, ICRISAT, will provide treatment farmers with an index insurance product, which pays out in the event of a flood. Though insurance products being offered in low-income countries vary, we use index insurance to ensure that the type of risk addressed by our flood-tolerant seeds is the same as that addressed by our insurance product, as well as to avoid traditional issues of moral hazard and adverse selection. To avoid low take-up of insurance (since this is well-documented, and measuring insurance demand is not central to this study) we provide farmers with insurance for free. All treatment farmers will be shown an info-sheet describing the insurance product , and will confirm their participation in the insurance product by signing a slip.

Following Lane (2024), we assign each village in our sample to its closest river gauge, maintained by the Government of West Bengal's Irrigation and Waterways Department. The insurance product pays out if and only if the relevant river gauge reaches a trigger height at least one day between June 25th and October 31st. We set the gauge-specific trigger height such that it was reached 30% of the years from 2014 to 2023.[12] If a village's gauge is triggered, we will notify farmers via SMS.

In order to measure the elasticity of our outcomes with respect to insurance coverage, we randomize insurance villages into a low-payout product and a high-payout product. Both products have the same trigger. If triggered, the high-payout product will pay out INR 10,000 (approximately USD 120), while the low-payout product will pay out INR 5,000 (approximately USD 60). Farmers will receive payouts in November.

Specialty seeds In order to test the predictions of our model, we will measure each farmer's willingness-to-pay for both a 5kg flood-tolerant seed mini-kit and a 5kg high-yield-variety seed minikit using a BDM process, which we describe in further detail below. Both seeds were selected to be high-quality varieties, suitable for flood-prone areas of West Bengal, available at local markets and therefore likely known to farmers, but not fully adopted at the time of our study. Both seeds also command a similar price on the market.

Flood-tolerant rice seed: We use Swarna-Sub1 as our flood-tolerant seed variety. Swarna-Sub1 is a medium long rice grain that measures 46-48 inches, and matures in 135-140 days. This seed performs similarly to the standard Swarna variety under non-flood conditions, but substantially better under flood conditions (Singh, Mackill, and Ismail (2009); Singh, Mackill, and Ismail (2011); Dar et al. (2013); Emerick et al. (2016)). More specifically, it produces yields of 2200-2400 kg/hectare in both flood and non-flood years.

High-yield-variety rice seed: We use Pratiksha as our high-yield variety. Pratiskha is a medium slender rice grain that measures 40-42 inches, and matures in 140-145 days. It outperforms traditional rice varieties in the area (by almost 1000kg/ha under normal conditions), producing yields comparable to Swarna, another common high-yield variety in West Bengal.

Seed offers We will elicit farmers' willingness to pay (WTP) for Swarna-Sub1 and Pratiksha minikits through a Becker-Degroot-Marshack (BDM) mechanism. We model the BDM procesure closely after Berkouwer and Dean (2022) and Burlig et al. (2024). To measure farmers' WTP for each seed, the enumerator will explain that they will be offering the farmer an opportunity to buy seeds through a two-step procedure. In the first step, they will elicit the farmer's stated WTP for the seed through a grid-search process. They will repeat the grid search for both 5kg seed mini-kits. After these WTP amounts are confirmed

[12] We include only gauges that have data during at least 9 of these 10 years.

for each seed, the enumerator will reveal a screen on their tablet which shows a randomized draw, containing (A) which seed the farmer may purchase during the interview, and (B) the price at which the farmer may buy this seed. If the price on the tablet is greater than the farmer's WTP for that seed, they may not purchase the product, and no cash changes hands. If the price on the tablet is below or equal to the farmer's WTP, the enumerator will provide the farmer with the seed mini-kit in exchange for the price in cash.

Because it is vital that this procedure is thoroughly understood by households before they begin, the enumerator will play a "practice" round with a common household product (either a bar of soap or a small bottle of shampoo). Therefore, any misunderstanding about the process will be resolved before the BDM procedure for the seeds is started.[13]

We will set the distribution of BDM prices to create a wedge in seed take-up between the "seed control" groups, whose BDM price will be set at the market price, and the "seed treatment" groups, whose BDM price will be set to zero, to induce take-up of the seeds. The underlying distribution will be unknown to both farmers and to the enumerators, following Burchardi et al. (2021), meaning that farmers will not have any incentive to bid strategically. In this way, we will maximize power by ensuring a large gap in take-up of each seed vs. the control group without compromising incentive compatibility. Prices will be randomly assigned to each participant prior to the baseline visit. The enumerator will not be aware of the seed or price draw prior to conducting the survey, and will bring minikits of both seeds to the interview.

Using the experiment for identification As described in our theoretical model, insurance has the potential to influence farmers' WTP for climate adaptation through two main channels:

1. *Crowd-out effect:* By covering the same risk, insurance and climate-resilient seeds may serve as substitutes

2. *Crowd-in effect:* By providing a backstop in the event of a poor realization, insurance lowers the risks associated with experimenting with a new technology

Our experiment allows us to tease these channels apart. Specifically, comparing WTP for the high-yield-variety seed between the insurance control group and the low- and high-payout insurance groups measures the crowd-in effect, while comparing WTP for the flood-tolerant seed between the insurance control group and the low- and high-payout insurance groups

[13] We will also conduct a take-it-or-leave-it (TIOLI) round with all households, where they receive a randomized price for either the soap or the shampoo. By doing this, we will be able to map out demand for the practice products using both BDM and TIOLI, allowing us to validate our BDM procedure.

measures the combination of the crowd-in and crowd-out effects. Therefore, the difference between WTP for the flood-tolerant seed and high-yield-variety seed across these groups identifies the crowd-out effect alone.[14] In addition, we can estimate the elasticity of these effects with respect to insurance coverage by comparing the low-payout insurance group to the control group and the high-payout insurance group.

Our design also allows us to measure four additional effects, going beyond WTP and seed take-up. First, we can measure the effects of insurance on agricultural practices (*ex ante* outcomes) and realizations (*ex post* outcomes) by comparing the insurance control group to the insurance treatment groups. Second, we can estimate the effects of flood-tolerant and high-yield-variety seeds on these same outcomes by comparing the seed control group to the flood-tolerant and high-yield-variety seed treatment groups (across all insurance and pure control villages). Third, we can examine the interaction between insurance and specialty seeds by comparing pure control households (those in both the insurance control group and the seed control group) to households who receive either type of speciality seed, insurance, or both. Finally, if floods occur during the 2024 growing season, we can estimate the impact of flooding on households with and without insurance, and households with and without each type of specialty seed (and the combination) by combining plausibly-exogenous variation in flooding with our experimental treatment arms.[15]

We outline all of these comparisons in further detail, including estimating equations, in Section 6 below.

Timeline Figure 3 presents the planned timeline for the experiment. We will conduct baseline surveys in early June 2024, before planting has begun. During this survey, we collect information about households' demographics, past farming behavior, flood exposure, risk preferences, consumption, assets and loans, and off-farm work. As part of the baseline survey, we also provide households with their insurance offer (if applicable), and elicit willingness-to-pay (WTP) for flood-tolerant and high-yield-variety seeds. Farmers who purchase the seeds through the WTP elicitation game will receive their seed kits once the survey is completed.

[14]Because this experiment – in an effort to mimic real-world policy – provides heavily subsidized insurance, another possible channel through which our insurance product could impact WTP for seeds is through an income effect: fully-subsidized insurance provides farmers with additional income in expectation. Our main analysis will not be able to distinguish between this and the crowd-in effect, though we believe that the combination of these two impacts is the policy-relevant quantity. One way of testing this is to estimate heterogeneity along the margin of farmers' risk aversion: a fully risk-neutral farmer should experience no crowd-in effect, and should treat the subsidized insurance product the same way they would treat an unconditional cash transfer with a future payout equal to the expected value of the insurance product. A risk-averse farmer, in contrast, will respond to both the income effect and the crowd-in effect.

[15]While flooding itself may be correlated with village characteristics, once we control for historical flood probability, realized flooding can be expected to be as-good-as-randomly-assigned.

Figure 3: Experimental timeline

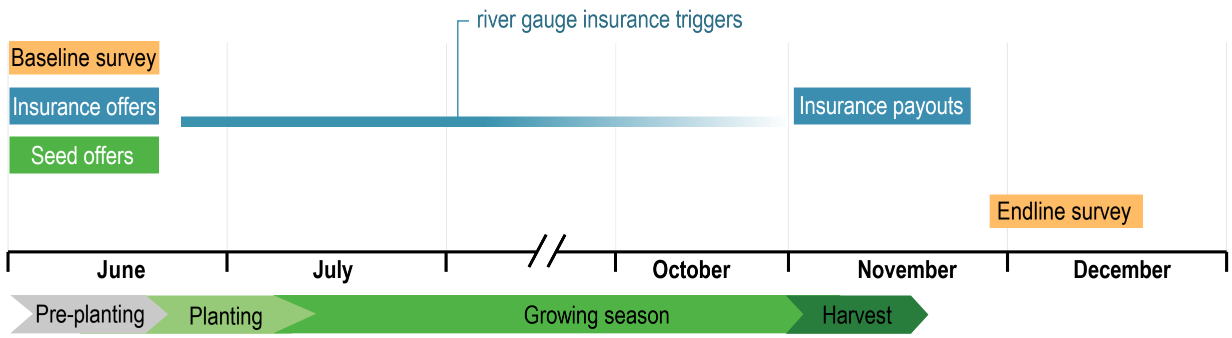

Notes: This figure presents the planned timeline of our experiment in relation to the agricultural cycle. We expect to run the experiment during the 2024 Kharif season. We will implement the baseline survey, provide insurance offers, and provide seed offers in early June. Insurance payouts will be triggered by river gauge measurements between June 20 and October 31, and insurance payouts will occur in November. Finally, we will conduct our endline survey in November/December.

Applicable insurance payouts will occur in November, and our endline survey will take place after harvest, in November - December 2024.

Statistical power In order to ensure that our design is powered to detect reasonably-sized treatment effects, we conduct a series of power calculations (shown in Figure 4. All power calculations use a standardized outcome variable, with control mean 0 and SD 1; are based on a two-sided hypothesis test with a 5% significance level; and use a sample of 300 villages containing a total of 1,800 households. We present two sets of power calculations: a "two-group" calculation, which compares villages without insurance to villages with insurance (Panel A), and a "three-group" calculation, comparing no insurance to low-payout and high-payout insurance (Panel B). In the two-group calculation, we compare a control group of 125 villages to a treatment group of 175 villages. In the three-group calculation, we compare a control group of 125 villages to a treatment group of 125 villages (representing the control vs. high-payout insurance comparison, our main effect of interest). In both cases, we present ICCs of 0, 0.05, 0.1, and 0.15. In the two-group calculation, we are powered to 80% for effects of approximately 0.15–0.2 SD. In the three-group calculation, we are powered to 80% for effects of approximately 0.2–0.25 SD. Since the insurance treatment is free to farmers, we expect take-up to be close to 100%. In practice, these calculations are likely somewhat conservative, as we will use specifications which control for baseline data, removing residual variation in the outcome.

These calculations give us confidence that the experiment is powered to detect treat-

Figure 4: Power calculations

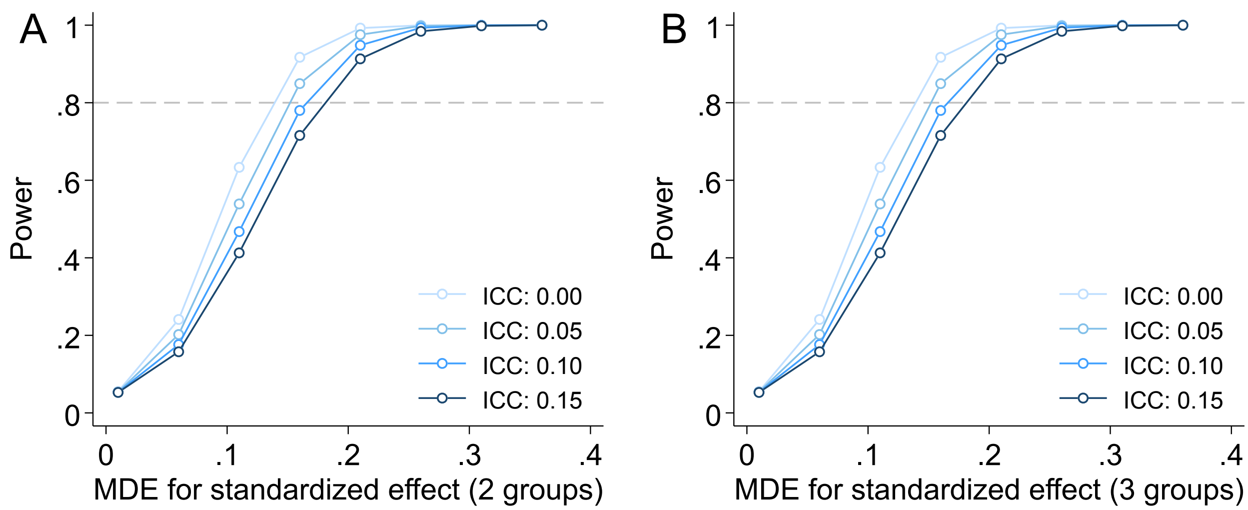

Notes: This figure plots power calculations for treatment effects on standardized outcomes (control mean = 0, standard deviation = 1), using a two-sided hypothesis test at 5% significance, and a sample of 1,800 households within 300 villages. Panel A plots the results of these calculations for two groups, comparing a control group (125 villages) to a pooled insurance treatment group (175 villages). Panel B plots the results of these calculations for three groups, comparing a control group (125 villages) to the high-payout insurance treatment group (125 villages). In both cases, we present ICCs of 0 (lightest blue) to 0.15 (darkest blue).

ment effects within the literature of impacts across prior agricultural studies in low-income countries (e.g., Mobarak and Rosenzweig (2014); Karlan et al. (2014); Emerick et al. (2016); Carter et al. (2017); Cole and Xiong (2017)). As perhaps a particularly helpful benchmark, Burlig et al. (2024) estimated that an index insurance product which had a maximum payout two-thirds the size of that described in our proposed experiment, which was provided 50 villages (with a control group of 100 villages), yielded impacts on agricultural investment of 0.12 SD, statistically significant at the 10% level. Because we are providing insurance to 175 villages (125 at the high payout level and 50 at the low payout level), we expect to be able to estimate precise treatment effects on agricultural outcomes. We ultimately conclude that these power calculations support our study design.

5 Data

5.1 Main outcomes

5.1.1 Baseline main outcomes: Willingness-to-pay

BDM As described above, we will use the BDM mechanism to elicit farmers' willingness to pay for both the high-yield-variety seeds and the flood-tolerant seeds at baseline. These WTP measures (and their difference, as described in Section 6 below) are our primary outcomes of interest for the study.

5.1.2 Endline main outcomes: *Ex ante* outcomes

Agricultural activity We hypothesize that farmers who receive insurance and/or specialty seeds could change which plots they cultivate; which crops they choose to plant; and/or the inputs they apply to these crops. We measure these outcomes as follows.

We ask farmers how many plots they own and cultivate (including which plots they rent out and rent in). For each plot we ask the farmer to describe the plot; list the size of the plot; describe how they obtained the plot; whether it was cultivated or left fallow last year; list which crops were cultivated on each plot (and which rice varieties were selected), and what share of each plot was used for each crop during the Kharif season.

Next, we ask farmers about inputs: seeds, fertilizer, labor and irrigation. For seeds and fertilizer, we ask how much they apply across all plots, which plots receive that input, and the share they apply to each plot. This allows us to compute inputs per land area. This does not provide a measure of inputs by crop, which we determined was too complicated for farmers to report. For paddy, we also capture which seed variety each farmer planted. To capture labor investments, we ask how many person-days the household hired external labor, and how many person-days family members worked on the plots. We also ask what it cost to hire 1 person to work the farm. Finally, we ask the total cost of irrigation per plot.

We both use these *ex ante* outcomes directly, and construct an *ex ante* input index, consisting of land under cultivation, crop/variety choice, and total input expenditure.

5.1.3 Endline main outcomes: *Ex post* outcomes

Agricultural activity In addition to the *ex ante* agricultural outcomes described above, we also measure *ex post* outcomes. In particular, for each plot, we ask how much of each crop they harvested. We then ask the amount of the harvested crop that is sold, consumed, spoiled, and slated for future sale. We use these data to construct crop revenue (including crops that were actually sold, as well as the value of crop production regardless of sale) and profits.

Consumption We ask respondents to report total consumption expenditure across three categories:

1. "Frequent" (within the last 7 days): cereals (rice, suji/rawa), milk, and tobacco
2. "Infrequent" (within the last 30 days): meat (eggs, chicken, goat, fish), and mobile phone charges
3. "Rare" (within the last 5 months): clothing, medicine/doctor, and celebrations

We can also scale and aggregate these measures to produce a single measure of average consumption per day. We complement this measure with the World Food Programme's Food Consumption Score (FCS). This indicator is a composite score derived from households' dietary diversity, food consumption frequency, and the nutritional value of different food groups.

Assets We ask respondents whether or not they own individual items across a list of assets, and the value of these assets. We ask about a set of representative assets which includes: TV, Bicycle, Motorbike, Telephone/Mobile phone, Computer, Sewing Machine, Thresher, Bullock Cart.

Income generating activities In addition to the agricultural income described above, farmers may receive non-agricultural income. To capture this, we ask respondents about their non-farm business, wage employment and livestock. For non-farm business, we ask the type of business, the revenue and profits they earned. For wage employment, we ask whether anyone in the household was engaged in any paid employment on/off the farm not including self-employment. We also ask respondents about the current stock of livestock, and profits from livestock over the previous 12 months.

Household finance Similarly, we expect that forecasts and/or insurance may impact overall household finances. We ask farmers about how much money they have in savings, if any. We then ask if they have taken any loans, the amount, and whether they have any outstanding balance. Finally we ask about whether they received any pension/government transfers/insurance money or lottery; the amount they received from these sources; and whether they received any money from a relative/friend (including remittances) and the total value of these transfers.

Migration We ask whether anyone in the household has left the village for work during the last agricultural season.

5.2 Heterogeneity

As we describe in Section 6 below, we test for heterogeneous treatment effects in farmer WTP for seeds along a series of dimensions.

5.2.1 Risk preferences

As described in Section 3 above, risk preferences play an important role in the take-up and impact of forecasts and insurance. We measure risk preferences (risk aversion) with multiple price list decision tasks (adapted from Holt and Laury (2002)).

5.2.2 Beliefs about seed effectiveness

To measure this, we ask farmers their general belief about the effectiveness of three types of seeds (traditional, high-yielding, and flood-tolerant) under two types of conditions (normal year, flood year). For each scenario we ask the production they would expect to have if they planted one hectare of that seed and the season were either a normal one or a flood affected one. The comparison of the high-yielding seed to traditional seed in a normal year will capture farmer perceptions of the general effectiveness of high-yield varieties. The comparison of traditional seed to flood-tolerant seed in a flood-affected year will capture farmer perceptions of the flood-tolerant trait.

5.2.3 Past experience with each seed type

We ask each farmer whether they themselves have ever used flood-tolerant seeds and high-yield-variety seeds. In each category, we also ask which specific seed variety they have used. If they have used Pratiksha or Swarna-Sub1, we ask what price they paid for the seed and what quantity they purchased.

5.2.4 Past flood exposure

We ask each farmer whether, in the past five years, they have experienced a flood that damaged their crops. In addition, for each plot the farmer reports farming, we ask how many times in the last five years the plot has been flood-affected.

5.2.5 River gauge trigger

We collect data on river water levels and each river gauge's "danger level" (at which flood damages occur) from the Goverment of West Bengal's Irrigation and Waterways Department from 2014 to 2023 for each river gauge used in our sample, where available. We compute the share of years that the river level reaches or exceeds the danger level.

5.2.6 Non-farm business and wage work

We ask each farmer whether the household also operates a non-farm business or has a household member who has (non-farm) wage work.

5.2.7 Savings and liquidity

We ask each farmer about their current total savings and if the household could find funds to cover an emergency costing amounts ranging from INR 100 to INR 5,000 without borrowing.

5.3 Secondary survey outcomes

In addition to the main outcomes and dimensions of heterogeneity described above, we are interested in conducting secondary analysis on several additional measures. We are not directly pre-specifying the analyses on these additional outcomes.

5.3.1 Summary statistics

Insurance exposure We want to measure farmers' exposure to the types of products that we are working with in this experiment. We measure specialty seed exposure as described above. We also ask about whether farmers have previously purchased crop insurance, and if so, when, what crops it covered (if crop-specific), the cost of this insurance, whether they received a payout, and if so, how much. We ask whether farmers intend to purchase (or have already purchased) crop insurance for the 2024 Kharif season. We also ascertain whether farmers have previously purchased insurance for their home, livestock, or other assets.

5.3.2 Additional outcomes

Migration Forecasts may impact whether household members decide to work on the farm or migrate. We ask households at baseline whether anyone in the household intends to migrate, when and for what reasons. At endline, we ask about whether anyone did migrate in the past year, why, and whether they intend to return.

Mental health Farmer suicides have been linked to agricultural outcomes in India (Carleton (2017)). To establish whether forecasts improve mental health, we measure farmers' depression levels using the standard PHQ-9 questionnaire.

6 Hypotheses and analysis

Throughout this section, we present three sets of specifications: those which pool both insurance treatment arms, those which separately estimate impacts of high-payout insurance vs. low-payout insurance, and those which estimate linear effects by insurance coverage level. Our main test of interest is the difference between the high-payout insurance and the control group, estimated using the specifications which estimate separate effects for the two insurance arms. We are also interested, in the pooled estimates, linearized estimates, and testing the difference between low-payout and control, and low-payout and high payout, but treat these as less important.

6.1 Willingness-to-pay

Impact of insurance on WTP We estimate the effect of insurance coverage on farmers' WTP for each type of seed. We estimate this using:

$$\begin{aligned} WTP\_FT_{ivs} &= \alpha + \beta Insurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \\ WTP\_HYV_{ivs} &= \alpha + \beta Insurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned} \tag{5}$$

where WTP_FT_{ivs} (WTP_HYV_{ivs}) is WTP for the flood-tolerant (high-yield variety) seed for farmer i in village v in strata s, and $Insurance_{vs}$ is an indicator for our (randomly-assigned) flood insurance product.[16] For this and all other regressions in this section, unless otherwise specified, X_{ivs} is a set of controls, chosen from available baseline variables using double-selection LASSO, as well as enumerator fixed effects. ψ_s is a strata fixed effect. ε_{ivs} is an error term which we cluster at the village level.

Here we are interested in whether β is different from zero for each type of seed. We are also interested in the *relative* effect of insurance coverage on WTP for the flood-tolerant seed vs. the high-yield variety seed. To estimate this, we estimate:

$$(WTP\_FT - WTP\_HYV)_{ivs} = \alpha + \beta Insurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \tag{6}$$

[16] For this specification, we pool low-payout and high-payout insurance into a single indicator. In the next set of treatment effects, described below, we estimate impacts separtely by insurance coverage level.

where we are primarily interested in whether β is not equal to zero. If $\beta > 0$, this indicates that insurance increases demand for flood-tolerant seeds relatively more than high-yield variety seeds: the "crowd-in" effect dominates. If $\beta < 0$, insurance increases demand for the high-yield variety seed relatively more than for the flood-tolerant seeds: the "crowd-out" effect dominates. If $\beta = 0$, either insurance does not impact demand for either seed, or insurance impacts demand for both seeds equally (depending on the results from Equations (5) above.)

Effects of insurance coverage level on WTP Our main object of interest is the comparison between WTP in the high-payout insurance arm vs. the control group. To estimate this, and the elasticity of seed WTP with respect to insurance coverage levels (i.e., payout generosity) across the seed types, we will estimate a version of Equations (5) and (6) where we allow for differential effects for low- vs. high-payout insurance offers:

$$
\begin{aligned}
WTP\_FT_{ivs} &= \alpha + \beta LowInsurance_{vs} + \beta_2 HighInsurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \\
WTP\_HYV_{ivs} &= \alpha + \beta LowInsurance_{vs} + \beta_2 HighInsurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \\
(WTP\_FT - WTP\_HYV)_{ivs} &= \alpha + \beta_1 LowInsurance_{vs} + \beta_2 HighInsurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

where $LowInsurance_{vs}$ and $HighInsurance_{vs}$ are indicators for receiving a low-payout insurance product or a high-payout insurance product, and all other regressors are as described above.

We also estimate a version where we linearize insurance coverage level:

$$
\begin{aligned}
WTP\_FT_{ivs} &= \alpha + \beta InsurancePayout_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \\
WTP\_HYV_{ivs} &= \alpha + \beta InsurancePayout_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \\
(WTP\_FT - WTP\_HYV)_{ivs} &= \alpha + \beta_1 InsurancePayout_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

where $InsurancePayout_{vs}$ is a equal to the payout of the insurance product in the event of a flood (i.e., INR 10,000 for high-payout group, 5,000 for low-payout group, and 0 for control), and all other regressors are as described above.

In addition to these specifications, we will plot simple demand curves for each seed, both pooled and separately by insurance coverage level.

Heteroegeneity in WTP The extent to which insurance affects demand for seeds varies with a number of factors. In particular, we are interested in whether 1) risk preferences, 2) beliefs about the effectiveness of each seed type, 3) past experience with each seed type, 4) past flood exposure, 5) historical flooding at the village's river gauge, 6) existence of non-farm business (or wage work), and 7) savings and liquidity change the relationship between insurance and WTP for each type of seed. Our base specification to test this will be:

$$\begin{aligned} Y_{ivs} = \alpha + &\beta_1 Insurance_{vs} + \beta_2 Insurance_{vs} \times HETERO_{vs} \\ &+ \beta_3 HETERO_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

We will also run analogous heterogeneity tests disaggregating each level of insurance:

$$\begin{aligned} Y_{ivs} = \alpha + &\beta_1 LowInsurance_{vs} + \beta_2 LowInsurance_{vs} \times HETERO_{vs} \\ &+ \beta_3 HighInsurance_{vs} + \beta_4 HighInsurance_{vs} \times HETERO_{vs} \\ &+ \beta_5 HETERO_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

And linearized:

$$\begin{aligned} Y_{ivs} = \alpha + &\beta_1 InsurancePayout_{vs} + \beta_2 Insurance_{vs} \times HETERO_{vs} \\ &+ \beta_3 HETERO_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

Seed take-up The last primary outcome of interest is how take-up of the seeds which each farmer was offered (either high-yield or flood-tolerant) at baseline (through the BDM exercise) varies with BDM price draw. We estimate baseline seed take-up via the following specification:

$$Y_{ivs} = \alpha + \beta_1 Free_{ivs} + \beta_2 Free_{ivs} * FT_{ivs} + \beta_3 FT_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

where the outcome is a binary indicator for receiving the seed kit during the baseline survey, FT_{ivs} is an indicator for receiving an offer to buy flood-tolerant seeds, and $Free_{ivs}$ is an indicator for receiving a zero price offer for the mini-kit.

6.2 *Ex ante* agricultural outcomes

Our next set of main hypotheses concerns farmers' *ex ante* (i.e., before the growing season realization) agricultural outcomes.

Interaction of insurance and specialty seeds We begin by estimating the effect of both insurance and specialty seed (offers) on outcomes. We first run a pooled specification which groups together the two types of seed:

$$\begin{aligned} Y_{ivs} = \alpha + & \beta_1 Insurance_{vs} + \beta_2 Insurance_{vs} \times FreeSeed_{ivs} \\ & + \beta_3 FreeSeed_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

where Y_{ivs} includes: land under cultivation, crop choice, rice variety choice, input expenditure (including on fertilizer, labor, seeds, irrigation), and an index of the above outcomes.

$Insurance_{vs}$ is defined as in Equation (5) above; $FreeSeed_{ivs}$, and is an indicator for having been offered either the high-yield-variety seed or the flood-tolerant seed for free.

We will then also disaggregate to test for differential responses by seed type:

$$\begin{aligned} Y_{ivs} = \alpha + & \beta_1 Insurance_{vs} + \beta_2 Insurance_{vs} \times FreeFT_{ivs} + \beta_3 Insurance_{vs} \times FreeHYV_{ivs} \\ & + \beta_4 FreeFT_{ivs} + \beta_5 FreeHYV_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

where $FreeFT_{ivs}$ and $FreeHYV_{ivs}$ are indicators for receiving offers of free flood-tolerant and free high-yield-variety seeds, respectively.

We also estimate the extent to which these impacts vary by insurance coverage level:

$$\begin{aligned} Y_{ivs} = \alpha + & \beta_1 LowInsurance_{vs} + \beta_2 LowInsurance_{vs} \times FreeSeed_{ivs} \\ & + \beta_3 HighInsurance_{vs} + \beta_4 HighInsurance_{vs} \times FreeSeed_{ivs} \\ & + \beta_5 FreeSeed_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

and

$$\begin{aligned} Y_{ivs} = \alpha + & \beta_1 LowInsurance_{vs} + \beta_2 HighInsurance_{vs} \\ & + \beta_3 LowInsurance_{vs} \times FreeFT_{ivs} + \beta_4 LowInsurance_{vs} \times FreeHYV_{ivs} \\ & + \beta_5 HighInsurance_{vs} \times FreeFT_{ivs} + \beta_6 HighInsurance_{vs} \times FreeHYV_{ivs} \\ & + \beta_7 FreeFT_{ivs} + \beta_8 FreeHYV_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

Linearized as:

$$\begin{aligned} Y_{ivs} = \alpha + & \beta_1 InsurancePayout_{vs} + \beta_2 InsurancePayout_{vs} \times FreeSeed_{ivs} \\ & + \beta_3 FreeSeed_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

and

$$Y_{ivs} = \alpha + \beta_1 InsurancePayout_{vs} + \beta_2 InsurancePayout_{vs} \times FreeFT_{ivs} + \beta_3 Insurance_{vs} \times FreeHYV_{ivs}$$
$$+ \beta_4 FreeFT_{ivs} + \beta_5 FreeHYV_{ivs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

Pooled effects of insurance As a secondary test, we estimate the effect of insurance (and insurance coverage level) on *ex ante* inputs, pooling across seed treatment groups. where all terms in the specification are defined as above.

$$Y_{ivs} = \alpha + \beta_1 Insurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

and

$$Y_{ivs} = \alpha + \beta_1 LowInsurance_{vs} + \beta_2 HighInsurance_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

and

$$Y_{ivs} = \alpha + \beta_1 InsurancePayout_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

Pooled effects of specialty seeds As a secondary test, we estimate the effect of specialty seeds, and then flood-tolerant and high-yield-variety seeds separately, on *ex ante* inputs, pooling across insurance treatment groups.

$$Y_{ivs} = \alpha + \beta_1 FreeSeed_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

and

$$Y_{ivs} = \alpha + \beta_1 FreeFT_{vs} + \beta_2 FreeHYV_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

Finally, for all regressions including seeds, we will estimate LATEs. To do this, rather than using an indicator for free seeds, free high-yield-variety seeds, or free flood-tolerant seeds as the regressors of interest, we replace these with indicators for either take-up (at baseline) or usage (at endline) of these seeds, instrumenting with the free offers.

6.3 *Ex post* outcomes

Our final set of main hypotheses concerns farmer's *ex post* (i.e., after the growing season realization) outcomes. We will repeat all of the tests from the *ex ante* outcome section, with the following outcomes: crop production, value of production, crop sales, yield, agricultural profits, savings, debt, off-farm labor, non-agricultural business, non-agricultural investment, business profits, assets, consumption per capita, food security, and mental health.

Impacts of flooding In addition to our main *ex post* tests, we are also interested in the extent to which flooding impacts outcomes. We first estimate the overall impact of floods on the same set of *ex post* outcomes described above:

$$Y_{ivs} = \alpha + \beta_1 Flood_{vs} + \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}$$

where $Flood_{vs}$ is an indicator for flooding in village v, and $FloodHazard_{vs}$ is a flood-hazard-category fixed effect, based on the West Bengal Flood Hazard Atlas (Department of Space (2021)). Though whether a given village floods in a given year is plausibly exogenous, villages in our sample differ in their flood probabilities. We therefore control for historical flooding to purge this source of potential bias, and argue that conditional on historical flooding, flooding in 2024 is plausibly random.

We also test for interactions between flooding and our experiment through three sets of specifications. First, we measure the interaction between flooding and insurance:

$$\begin{aligned} Y_{ivs} = {} & \alpha + \beta_1 Insurance_{vs} \times Flood_{vs} + \beta_2 Insurance_{vs} + \beta_3 Flood_{vs} \\ & + \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

and

$$\begin{aligned} Y_{ivs} = {} & \alpha + \beta_1 LowInsurance_{vs} \times Flood_{vs} + \beta_2 HighInsurance_{vs} \times Flood_{vs} \\ & + \beta_3 LowInsurance_{vs} + \beta_4 HighInsurance_{vs} + \beta_5 Flood_{vs} \\ & + \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

and

$$\begin{aligned} Y_{ivs} = {} & \alpha + \beta_1 InsurancePayout_{vs} \times Flood_{vs} + \beta_2 InsurancePayout_{vs} + \beta_3 Flood_{vs} \\ & + \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs} \end{aligned}$$

Next, we measure the interaction between each seed type and flooding:[17]

$$
\begin{aligned}
Y_{ivs} = \alpha &+ \beta_1 FreeHYV_{ivs} \times Flood_{vs} + \beta_2 FreeFT_{ivs} \times Flood_{vs} \\
&+ \beta_3 FreeHYV_{ivs} + \beta_4 FreeFT_{ivs} + \beta_5 Flood_{vs} \\
&+ \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

Finally, we estimate the full interaction of insurance, seeds, and flooding:

$$
\begin{aligned}
Y_{ivs} = \alpha &+ \beta_1 Insurance_{vs} + \beta_2 Insurance_{vs} \times FreeFT_{ivs} + \beta_3 Insurance_{vs} \times FreeHYV_{ivs} \\
&+ \beta_4 FreeFT_{ivs} + \beta_5 FreeHYV_{ivs} \\
&+ \beta_6 Insurance_{vs} \times Flood_{vs} \\
&+ \beta_7 Insurance_{vs} \times FreeFT_{ivs} \times Flood_{vs} + \beta_8 Insurance_{vs} \times FreeHYV_{ivs} \times Flood_{vs} \\
&+ \beta_9 FreeFT_{ivs} \times Flood_{vs} + \beta_{10} FreeHYV_{ivs} \times Flood_{vs} + \beta_{11} Flood_{vs} \\
&+ \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

and

$$
\begin{aligned}
Y_{ivs} = \alpha &+ \beta_1 LowInsurance_{vs} + \beta_2 HighInsurance_{vs} \\
&+ \beta_3 LowInsurance_{vs} \times FreeFT_{ivs} + \beta_4 LowInsurance_{vs} \times FreeHYV_{ivs} \\
&+ \beta_5 HighInsurance_{vs} \times FreeFT_{ivs} + \beta_6 HighInsurance_{vs} \times FreeHYV_{ivs} \\
&+ \beta_7 FreeFT_{ivs} + \beta_8 FreeHYV_{ivs} \\
&+ \beta_9 LowInsurance_{vs} \times Flood_{vs} + \beta_{10} HighInsurance_{vs} \times Flood_{vs} \\
&+ \beta_{11} LowInsurance_{vs} \times FreeFT_{ivs} \times Flood_{vs} + \beta_{12} LowInsurance_{vs} \times FreeHYV_{ivs} \times Flood_{vs} \\
&+ \beta_{13} HighInsurance_{vs} \times FreeFT_{ivs} \times Flood_{vs} + \beta_{14} HighInsurance_{vs} \times FreeHYV_{ivs} \times Flood_{vs} \\
&+ \beta_{15} FreeFT_{ivs} \times Flood_{vs} + \beta_{16} FreeHYV_{ivs} \times Flood_{vs} \\
&+ \beta_{17} Flood_{vs} + \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

[17] Note that we do not pre-specify estimates pooling both seed offer types, as given the agronomy underlying the two seeds, we expect that they will perform differently under flood vs. normal conditions.

and

$$
\begin{aligned}
Y_{ivs} = \alpha &+ \beta_1 InsurancePayout_{vs} + \beta_2 InsurancePayout_{vs} \times FreeFT_{ivs} \\
&+ \beta_3 InsurancePayout_{vs} \times FreeHYV_{ivs} \\
&+ \beta_4 FreeFT_{ivs} + \beta_5 FreeHYV_{ivs} \\
&+ \beta_6 InsurancePayout_{vs} \times Flood_{vs} \\
&+ \beta_7 InsurancePayout_{vs} \times FreeFT_{ivs} \times Flood_{vs} \\
&+ \beta_8 InsurancePayout_{vs} \times FreeHYV_{ivs} \times Flood_{vs} \\
&+ \beta_9 FreeFT_{ivs} \times Flood_{vs} + \beta_{10} FreeHYV_{ivs} \times Flood_{vs} + \beta_{11} Flood_{vs} \\
&+ \eta FloodHazard_{vs} + \gamma X_{ivs} + \psi_s + \varepsilon_{ivs}
\end{aligned}
$$

a As with the *ex ante* outcomes, we will also estimate analogous IV specifications, instrumenting for specialty seed usage with free offers.

6.4 Multiple hypothesis testing

Because we are testing multiple endline outcomes in this experiment, in addition to reporting standard p-values, we will also present sharpened False Discovery Rate (FDR) q-values, which control for the expected proportion of rejections that are Type I errors, following Anderson (2008). We will apply these q-values to outcomes that we measure using multiple questions. This includes the full set of *ex ante* agricultural activities; the full set of *ex post* agricultural activity measures; *ex post* consumption measures; *ex post* asset measures; and *ex post* income-generating opportunity measures. Multiple hypothesis testing corrections are not relevant for our WTP outcomes, since we only have one measure of WTP for each seed type per farmer.

7 Conclusion

In this project, we study the extent to which an increasingly-popular policy – large subsidies for crop insurance – impact farmers' demand for other risk-coping strategies. In theory, subsidized insurance could either increase demand for flood-tolerant seeds, as insurance protects farmers against the downside risk of experimenting with a new technology, or reduce demand for flood-tolerant seeds, as insurance protects the household against the same risk as the seeds.

We propose a cluster-randomized trial that randomizes villages into one of three groups:

a control group, a group that is offered fully-subsidized insurance that pays out a low amount in the event of a flood, and a group that is offered fully-subsidized insurance that pays out a high amount in the event of a flood. We use this variation to estimate the impact of insurance on farmers' demands for specialty seeds. We then induce random variation in the take-up of these seeds at the farmer level. We use both sources of randomization to estimate the effects of insurance, high-yield-variety and flood-tolerant seeds, and their interaction on *ex ante* agricultural investments and *ex post* outcomes, including welfare metrics.

Despite the increasing prevalence of crop insurance subsidies in low-income countries, there is limited evidence on the impacts of these subsidies on private investment in risk mitigation. As a result, we expect this project to make a meaningful contribution to the literature on the full consequences of these technologies and agricultural policy in the developing world.

References

Aker, Jenny and B. Kelsey Jack. 2023. “Harvesting the rain: The adoption of environmental technologies in the Sahel.” *Review of Economics and Statistics* Accepted.

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

Annan, Francis and Wolfram Schlenker. 2015. “Federal crop insurance and the disincentive to adapt to extreme heat.” *American Economic Review* 105 (5):262–266.

Auffhammer, Maximilian. 2022. “Climate Adaptive Response Estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption.” *Journal of Environmental Economics and Management* 114:102669.

Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro. 2016. “Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century.” *Journal of Political Economy* 124 (1):105–159.

Baylis, Patrick and Judson Boomhower. 2021. “Mandated vs. voluntary adaptation to natural disasters: The case of U.S. wildfires.” NBER working paper no. 29621.

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

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

Boomhower, Judson, Meredith Fowlie, Jacob Gellman, and Andrew Plantinga. 2023. “How are insurance markets adapting to climate change? Risk selection and regulation in the market for homeowners insurance.” Working paper.

Boucher, Stephen R., Michael R. Carter, Jon Einer Flatnes, Travis J. Lybbert, Jonathan G. Malacarne, Paswel Marenya, and Laura A. Paul. 2022. “Bundling genetic and financial technologies for more resilient and productive small-scale agriculture.” NBER working paper no. 29234.

Burchardi, Konrad B., Jonathan de Quidt, Selim Gulesci, Benedetta Lerva, and Stefano Tripodi. 2021. "Testing willingness to pay elicitation mechanisms in the field: Evidence from Uganda." *Journal of Development Economics* 152:102701.

Burgess, Robin, Olivier Deschênes, Dave Donaldson, and Michael Greenstone. 2017. "Weather, climate change, and death in India." Working paper.

Burlig, Fiona, Amir Jina, Erin M. Kelley, Gregory Lane, and Harshil Sahai. 2024. "Long-range forecasts as climate adaptation: Experimental evidence from developing-country agriculture." NBER Working Paper No. 32173.

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

Carleton, Tamma A. 2017. "Crop-damaging temperatures increase suicide rates in India." *Proceedings of the National Academy of Sciences* 114 (33):8746–8751.

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

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

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

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

Dar, Manzoor H., Alain de Janvry, Kyle Emerick, David Raitzer, and Elisabeth Sadoulet. 2013. "Flood-tolerant rice reduces yield variability and raises expected yield, differentially benefitting socially disadvantaged groups." *Scientific Reports* 3:3315.

Department of Space. 2021. "Flood hazard atlas: West Bengal – A geospatial approach." Tech. rep. Government of India.

Deryugina, Tatyana. 2017. "The fiscal cost of hurricanes: Disaster aid versus social insurance." *American Economic Journal: Economic Policy* 9 (3):168–198.

Ehrlich, Isaac and Gary S. Becker. 1972. "Market insurance, self-insurance, and self-protection." *Journal of Political Economy* 80 (4):623–648.

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

Government of West Bengal. 2024. "West Bengal state portal: Agriculture." URL https://wb.gov.in/departments-details.aspx?id=D170907140022669&page=Agriculture.

Hallegatte, Stephane, Adrien Vogt-Schilb, Mook Bangalore, and Julie Rozenberg. 2017. *Unbreakable: Building the resilience of the poor in the face of natural disasters*. World Bank Group.

Hirabayashi, Yukiko, Roobavannan Mahendran, Sujan Koirala, Lisako Konoshima, Dai Yamazaki, Satoshi Watanabe, Hyungjun Kim, and Shinjiro Kanae. 2013. "Global flood risk under climate change." *Nature Climate Change* 3 (9):816–821.

Holt, Charles A. and Susan K. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review* 92 (5):1644—1655.

Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D. J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345):1362–1369. URL https://www.science.org/doi/abs/10.1126/science.aal4369.

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

India Meterological Department. 2023. "Drought condition assessment based on drought indices generated at India Meteorological Department: Period - September 2023 and June-September 2023." Tech. rep.

Indian Council of Agricultural Research. 2024. "West Bengal." URL https://icar.org.in/node/17296. Ministry of Agriculture and Farmers' Welfare.

Irrigation and Waterways Department. 2024. "Flood scenario of West Bengal." URL https://wbiwd.gov.in/index.php/applications/flood_mgmt. Government of West Bengal.

Ismail, Abdelbagi M., Uma S. Singh, Sudhanshu Singh, Manzoor H. Dar, and David J. Mackill. 2013. "The contribution of submergence-tolerant (Sub1) rice varieties to food security in flood-prone rainfed lowland areas in Asia." *Field Crops Research* 152:83–93.

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

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

Kousky, Carolyn, Erzo F. P. Luttmer, and Richard J. Zeckhauser. 2006. "Private investment and government protection." *Journal of Risk and Uncertainty* 33 (1):73–100.

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

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

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

Ostriker, Abigail and Anna Russo. 2024. "The effects of floodplain regulation on housing markets." Working paper.

Peoples' Republic of China. 2024. "Wider cover promises food security." *The State Council Information Office* URL http://english.scio.gov.cn/m/pressroom/2024-06/03/content_117230980.htm#:~:text=The%20Finance%20Ministry%20launched%20an,growth%20rate%20of%2022%20percent.

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

Raghu, Prabhakaran T., Prakashan Chellattan Veettil, and Sukanya Das. 2022. “Smallholder adaptation to flood risks: Adoptino and impact of Swarna-Sub1 in Eastern India.” *Environmental Challenges* 7:100480.

Rode, Ashwin, Tamma Carleton, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E. Kopp, Kelly E. McCusker, Ishan Nath, James Rising, and Jiacan Yuan. 2021. “Estimating a social cost of carbon for global energy consumption.” *Nature* 598:308–314.

Singh, Sudhanshu, David J. Mackill, and Abdelbagi M. Ismail. 2009. “Responses of SUB1 rice introgression lines to submergence in the field: Yield and grain quality.” *Field Crops Research* 113 (1):12–23.

———. 2011. “Tolerance of longer-term partial stagnant flooding is independent of the SUB1 locus in rice.” *Field Crops Research* 121 (3):311–323.