Environmental Beliefs and Adaptation to Climate Change

Dev Patel

December 13, 2024*

Abstract

I examine the causes and consequences of environmental beliefs. Rice farmers in Bangladesh must learn about their soil's salinity to choose the appropriate seed. Comparing expectations to agronomic readings, I document both significant over- and underestimation of salt levels. A simple identification problem explains this pattern: farmers learn about multiple unobserved threats from the same set of ambiguous signals. Bayesian farmers endogenously process data in support of their priors, e.g., someone worried about high salinity will interpret low yield as a sign of too much salt. Climate change amplifies this process by systematically altering the risks considered most threatening. I confirm the framework's predictions using a lab-in-the-field exercise and natural experiments that isolate salient shocks that capture attention (e.g., tidal flooding) and subtle shifts that go unnoticed (e.g., irrigation water contamination through rising sea-levels). Despite equal effects on true salt levels, salient saltwater floods increase salinity beliefs substantially more than does subtle irrigation intrusion. These experiences shape how farmers interpret new data: past exposure to salient shocks increases the mental link between low yield and salinity while subtle shocks reduce the perceived diagnosticity of salinity clues. In large-scale field experiments, correcting misperceptions significantly alters farmers' demand for salinity-tolerant seeds with substantial consequences for profits.

*This is a heavily revised version of chapter one of my dissertation, entitled "Learning About a Warming World: Attention and Adaptation in Agriculture." I am especially grateful to Emily Breza, Claudia Goldin, Asim Khwaja, and Gautam Rao for their unwavering kindness and brilliance. I thank Shawn Cole, John Conlon, Melissa Dell, Natalia Emanuel, Siddharth George, Ed Glaeser, Rema Hanna, Zoë Hitzig, Larry Katz, Gabriel Kreindler, Eliana La Ferrara, Steven Lee, Ross Mattheis, Matthew Rabin, G. M. Mustafizur Rahman, Ben Roth, Justin Sandefur, Josh Schwartzstein, Brit Sharoni, Andrei Shleifer, Clara Sievert, Evan Soltas, Stefanie Stantcheva, Arvind Subramanian, S. M. Shahriar Zaman, and numerous seminar participants for invaluable feedback. I am indebted to all of the farmers who participated in this research for their time and insights and my field team for their exceptional efforts. Roshni Islam and Noshin Othoi provided superb research assistance. This research underwent human subjects review at Harvard University. I gratefully acknowledge financial support from the Chae Family Economics Research Fund, the Center for International

1 Introduction

Environmental change stemming from global warming poses an existential threat, especially to the 600 million smallholder farmers whose livelihoods directly depend on the climate. To mitigate these dangers, individuals must adapt to the evolving world around them. Optimal decision-making requires the formation of accurate expectations about environmental conditions. Yet we know little about what these beliefs are, where they come from, and how they matter.

I study environmental beliefs and their role in climate change adaptation. I ask three questions: First, how do environmental beliefs compare to the ground truth? Second, how do people learn about their environment? Third, how do these expectations shape economic decision-making? I study the belief formation process and its impacts among rice farmers living in Bangladesh, ground zero for the harmful consequences of climate change. I investigate these issues in the context of soil salinity, a threat exacerbated by global warming that impacts agricultural productivity on 30% of irrigated land worldwide (Hopmans et al., 2021). Rice farmers can mitigate salt's damage by planting salinity tolerant seeds, which perform better than alternative seeds in high salt conditions but relatively worse in low ones.

I first assess the accuracy of environmental beliefs by collecting information on farmers' perceptions and the ground truth. Soil salinity has several appealing features for answering this research question. First, agronomic sensors can directly measure soil conditions on farmers' own plots, allowing me to accurately measure the decision-relevant environmental state. Second, farmers lack access to this soil measurement technology, creating scope for inaccurate beliefs. Third, perceptions about soil salinity directly enter into the high-stakes decision of which seed to plant, and therefore farmers have strong incentive to learn the correct environmental conditions. I collect data on true soil salinity and farmers beliefs in a large-scale panel survey with nearly 2,300 rice farmers across 250 villages. To overcome the scarcity of existing salinity data, I use agronomic sensors to test soil conditions on farmers' own plots during repeated readings over the agricultural season. I develop a visual method to precisely elicit probabilistic beliefs in this low-numeracy population on the same scale as the agronomic sensors.

I find both significant over- and under-estimation of salinity levels, with the vast majority of variation in farmers' perceptions left unexplained by the ground truth. As evidence that

Development/Malcolm Wiener Center for Social Policy, the Harvard Climate Change Solutions Fund, the Harvard Mellon Urban Initiative, the Institute for Humane Studies under grants IHS016778 and IHS017455, the Mind, Brain, and Behavior Initiative, a National Science Foundation Graduate Research Fellowship under grant DGE1745303, the National Science Foundation Doctoral Dissertation Research Improvement Grant under grant 2242263, the Warburg Fund, and The Weiss Fund for Research in Development Economics at the University of Chicago.

this dispersion reflects more than noise, I show that gaps between soil salinity beliefs and the agronomic readings strongly predict farmers' decisions to plant salinity-tolerant seeds. Farmers with more land, more experience, and viewed as more skilled by their neighbors hold more accurate beliefs, suggesting that learning frictions might play a key role.

To explain this equilibrium distribution of environmental misperceptions, I use a simple conceptual framework of belief formation amid climate change. Guided by the results of narrative qualitative interviews and open-ended questions with my main sample, the framework features farmers who infer environmental conditions by observing the output of their harvest and physical characteristics of their crops. Different potential factors impacting yield manifest themselves in the same way, creating an identification problem (Acemoglu et al., 2016). A farmer observing low yield, for example, cannot distinguish between high soil salinity and insufficient fertilizer as the root cause. As a result, I show that Bayesian farmers with arbitrarily small differences in initial beliefs will nevertheless fail to converge to agreement about their local environment if they initially disagree on which environmental threat is more likely. This gap occurs because farmers' (potentially incorrect) ranking of priors exhibit persistence in this underidentified setting. Initial expectations about which environmental features matter most shape the interpretation of signals, so new data reinforces prior beliefs, similarly to the literature on misspecified learning (Heidhues et al., 2021). For example, a farmer who initially expects soil salinity to be high will (rationally) interpret a bad harvest as evidence of too much salt, even if another factor might actually be the culprit.

How does global warming shape this belief formation process? I focus on how farmers learn from two classes of signals characteristic of climate change: salient shocks and subtle shifts in the environment.[1] In the case of soil salinity, experiencing a flood with saline water (perhaps due to a tidal surge, for instance) constitutes a salient shock by drawing farmers' attention to the salt that has just been deposited onto the plot's surface. Indeed, farmers recall floods more than any other factor when recounting determinants of their soil's salt levels. By contrast, rising sea levels that cause the intrusion of ocean water into a farmer's irrigation source subtly shifts soil salinity, often occurring underground.

Experiencing a salient event instead of a subtle one can generate persistent gaps in learning. I show how this same prediction can emerge under three scenarios: a purely Bayesian model with no behavioral friction, a Bayesian learner with imperfect information, and a

[1] Although I use the terms salient shocks and subtle shifts, conceptually, both are "shocks" in the sense of being unexpected deviations to the climate. The distinction between these two aspects of climate change has been discussed in a psychology literature discussing learning and the adaptation process (Gilbert, 2006; Johnson and Levin, 2009; Gifford, 2011). More broadly, economists have documented that behavior responds more to drastic shifts than to gradual change and the difficulties of learning from infrequent hazards (Davis, 2004; Greenstone and Gallagher, 2008; Da et al., 2014; Wagner, 2022; Heft-Neal et al., 2023).

farmer facing attention costs. The most appropriate assumption will depend on the specific context. In the pure Bayesian benchmark, if the informational content of the two types of events differentially changes farmers' beliefs going into the next period, then farmers will endogenously interpret new data in different ways. This can sustain the initial disagreement perpetually because farmers have shifted their expectations about the relative danger of different environmental risks, which shapes how they process new information in this under-identified setting. Second, even when the underlying informational content of salient shocks and subtle shifts are the same, if farmers fail to notice the subtle changes in their environment (because, for instance, the shifts occur underground), their beliefs will differentially change, again leading to a gap in priors entering the next harvest. Third, I show that even in the case where the core informational content is the same and farmers notice both the salient shock and the subtle shift such that their posterior beliefs are identical, learning differences can still emerge in equilibrium if the salient shock changes which environmental threats come to mind upon seeing new data. The relevant mechanism depends on the environmental context, e.g., farmers might always notice any time it rains when forming monsoon beliefs but not an incremental increase in groundwater salinity when learning about their soil. Regardless of the channel, however, I show that in all cases, experiencing a salient shock instead of a subtle shift can have permanent impacts on beliefs by altering how farmers interpret new data, endogenously changing the diagnosticity of harvest signals with respect to different environmental threats.

To empirically test these predictions, I first use a lab-in-the-field exercise to document the presence of this identification problem. I show farmers pictures of rice plants, each of which suffers from a difficult-to-observe threat, e.g., pests, disease, fungus, soil salinity, etc. I elicit farmers' beliefs about the potential underlying causes that could cause the plant to look as it does in the image, for example by creating white spots on the leaves. I first find substantial disagreement among farmers about how to interpret the physical characteristics of a rice plant, even among images that an agronomist would characterize as solely diagnostic of one particular issue. I next show that the errors are systematic: the farmers most worried about salinity also tend to *correctly* identify salinity as the cause of salinity images and *incorrectly* mention salinity as an explanation for the non-salinity images. These patterns match the predictions of the conceptual framework, highlighting that farmers face uncertainty about the meaning of the signals they observe and systematically interpret these data in favor of the threats about which they are most concerned.

To provide causal evidence on the predictions of the conceptual framework, I use two quasi-random natural experiments that separately identify the impact of salient shocks and subtle shifts in soil salinity. The large geographic scope of my surveys—spanning approxi-

mately 15 percent of the land area of Bangladesh—provides substantial variation in underlying climate to provide the statistical power for these analyses. First, I examine flooding with salty water as an example of a salient shock. These floods deposit salt directly onto plots. Underscoring the attention-grabbing nature of these environmental events, farmers mention them more than any other factor when recalling the determinants of their soil's salinity. Second, I use salinity intrusion into irrigation water as an illustration of a subtle shift. As sea levels rise, salty water from the ocean can contaminate the water used by farmers to irrigate their rice plants, increasing soil salinity. This process—much of which occurs underground—can easily go overlooked.

I use two distinct identification strategies to estimate the consequences of these environmental events. First, I use a difference-in-differences design to isolate the causal impact of experiencing a flood with more instead of less saline water—the salient environmental shock. I hold constant any other general impacts of floods (unrelated to the saltiness of the water) and the propensity to experience a salty flood by conditioning on a machine-learning generated measure of flood risk and the salt content of local river systems I calculate from data collected at river stations scattered throughout the country. Second, I use a triple difference-in-differences framework to capture idiosyncratic variation in salinity intrusion into irrigation water—the subtle environmental shift. I compare villages with differentially higher exposure to quasi-random deviations in sea-level rise and ocean salinity modulated by the village's distance to the coast, using satellite data to measure ocean height and salt levels.

Consistent with the conceptual framework's predictions, I find asymmetric impacts of these two types of climate signals on beliefs. Although both saltier floods and contaminated irrigation water cause equivalent increases in true soil salinity as measured by the direct agronomic sensors, only saltier floods—the salient shocks—move farmers' perceptions of their soil's salt content, whereas salinity intrusion in irrigation—the subtle shift—has no detectable impact on beliefs. The increase in salinity beliefs following a salty flood far exceeds the true impact on soil conditions, ultimately generating overestimation of salt levels. In line with the model's key mechanisms, farmers also alter their interpretation of data on crops and yield depending on the nature of their past environmental experiences. Salient salty floods increase the likelihood that farmers use generic plant characteristics (e.g., plant death) to learn about salinity. Groundwater intrusion, by contrast, increases the chance that farmers attribute salinity-specific plant features to non-salinity explanations, reducing the diagnosticity of salinity clues.

These environmental misperceptions have economic costs if they distort important decisions, such as adaptation to climate change. To estimate the causal impact of environmental beliefs on technology adoption, I conduct large-scale field experiments among the same

sample of farmers in Bangladesh and find that beliefs have large impacts on agricultural production decisions and profits. I first measure demand for the soil salinity information I gathered from agronomic readings on farmers' plots, using the Becker-DeGroot-Marschak (BDM) method to preserve incentive compatibility. More than 80% of farmers have a positive willingness-to-pay for these data. Embedding an experiment into this elicitation, I randomly provide some farmers with data on soil readings. This information has large impacts on farmers predictions of future soil salinity conditions, especially among those who held more inaccurate beliefs initially. The treatment effects persist in a follow-up phone survey six months later. To evaluate the consequences for technology adoption, I elicit willingness-to-pay for a salinity-tolerant seed variety, again using BDM. Using treatment assignment as an instrument, I find that 1 s.d. higher soil salinity beliefs causes a 41% increase in demand for salinity-tolerant seeds.

In a separate experiment, I show that the decision to adapt this climate adaptation technology on the margin has high economic stakes. Farmers must match their seed choice to the appropriate environmental conditions: salinity-tolerant varieties grow better than alternatives in high salt conditions but relatively worse in low ones. Prior to planting, I randomly offer some farmers a small amount of these salinity-tolerant seeds for free. This offer causes a significant increase in the share of land planted with a salinity-tolerant variety, especially among farmers who initially overestimated salinity conditions. I return to farmers after harvest and find that farmers who overestimated salt levels and therefore planted more salinity-tolerant seeds in inappropriate (low salinity) soil experience a large reduction in agricultural profits, underscoring the risk of mismatching adaptation and environmental conditions.

Finally, I quantify the role of misperceptions about the local environment in climate change adaptation by using these randomized experiments to discipline a structural model of seed choice. Given the non-linear relationship between technology adoption and beliefs, the reduced form results alone do not reflect the overall effect of expectations on adaptation. I estimate a mixed multinomial logit model with random coefficients to allow for correlation of choices across different seeds (McFadden, 1974; Train, 2009). I estimate the parameters among the treatment group from the seed experiment and then validate the predictions in the control group.[2] The model fit performs extremely well in this out-of-sample test: regressing true seed choices on the predicted choice probabilities yields a coefficient indistinguishable from 1. Using this structure, I simulate the adoption of salinity tolerant seeds under al-

[2] My approach follows an impressive set of papers combining experiments with structural estimation including Todd and Wolpin (2006); Kremer et al. (2011); Duflo et al. (2012), and Attanasio et al. (2012). See Todd and Wolpin (2023) for a review.

ternative environmental beliefs. Conditional on the truth, moving from the 10^{th} percentile of salinity beliefs (significant underestimation of salt levels) to the 90^{th} percentile (overestimation of salinity) increases take-up of salinity-tolerant seeds by 81%. I then examine a scenario in which farmers can perfectly forecast their soil salinity conditions, and applying the reduced form results from my seed experiment, I estimate that agricultural profits would have been 16% higher during the 2022-23 season with correct beliefs.

Overall, these results suggest that policymakers can increase appropriate adaptation to climate change by targeting information to places based off their previous exposure to salient shocks versus subtle shifts. The cost of collecting high-quality local data presents one potential practical challenge. For instance, gathering salinity readings from farmers' own plots may be infeasible at scale. Motivated by this concern, I embed an additional arm into my information experiment that provided farmers solely with geographically aggregated data on salinity without any information specific to their own plots. Nevertheless, I find that farmers have a high demand for this aggregate data and that the treatment effects on behavior are even larger than those from providing the information specifically about farmers' own land. Collecting and distributing environmental data at scale presents a promising policy for improving adaptation to climate change.

Related literature This study builds on three strands of literature. First, I contribute to the rich set of papers examining adaptation to climate change and environmental threats.[3] A growing body of work has considered these issues specifically in low- and middle-income countries, where both the climate dangers and relevant margins of adaptation differ considerably from richer settings.[4] Much of the research on developing countries has focused on rainfall

[3] See, among others, Bardhan (1983), Nordhaus (1991), Fafchamps (1993), Rosenzweig and Binswanger (1993), Mendelsohn et al. (1994), Nordhaus (1994), Kahn (2005), Agrawal (2008), Deschênes and Greenstone (2011), Acemoglu and et al. (2012), Boustan et al. (2012), Gallagher (2014), Hornbeck and Naidu (2014), Hsiang and Jina (2014), Annan and Schlenker (2015), Currie et al. (2015), Desmet and Rossi-Hansberg (2015), Bakkensen and Mendelsohn (2016), Barreca et al. (2016), Burke et al. (2016), Burke and Emerick (2016), Auffhammer (2018), Kahn and Zhao (2018) Kocornik-Mina et al. (2020), Conte et al. (2021), Biardeau et al. (2020), Giuliano and Nunn (2021), Kahn (2021), Peri and Robert-Nicoud (2021), Shapiro (2021), Baylis and Boomhower (2022), Carleton et al. (2022), Cicala et al. (2022), Gandhi et al. (2022), Greenstone et al. (2022), Moscona (2022), Nath (2022), Ostriker and Russo (2022), Acemoglu et al. (2023), Harstad (2023), Bilal and Rossi-Hansberg (2023), Jedwab et al. (2023), Hornbeck (2023), Moscona and Sastry (2023), and Grosset et al. (2023).

[4] For examples of this work, see Jayachandran (2009), Burgess and Donaldson (2010), Burgess et al. (2012), Tazhibayeva and Townsend (2012), Dar et al. (2013), Duflo et al. (2013), Mattoo and Subramanian (2013), Mobarak and Rosenzweig (2013), Arceo-Gomez and Oliva (2014), Greenstone and Hanna (2014), Elliott et al. (2015), Hanna and Oliva (2015), Pande et al. (2015), Cattaneo and Peri (2016), Costinot et al. (2016), Taraz (2017), Auffhammer and Carleton (2018), Barwick et al. (2018), Casaburi and Willis (2018), Mullins et al. (2018), Adhvaryu et al. (2019), Balboni (2019), Chang et al. (2019), Rangel and Vogl (2019), Blakeslee et al. (2020), Ito and Zhang (2020), McGuirk and Nunn (2020), Khanna et al. (2021), Pople et al. (2021), Chatterjee et al. (2022), Jack and et al. (2022), Macours et al. (2022), Ryan and Sudarshan (2022), Aker

and temperature, perhaps due to the relative accessibility of data.[5] I focus on soil salinity, an environmental threat endangering agricultural households across the globe.[6] Existing work has examined a host of important frictions to individual investment in adaptation, including credit constraints (Lane, 2022), information asymmetries (Beaman et al., 2014; Mahadevan et al., 2023), financial market imperfections (Karlan et al., 2014), supply-side frictions (Emerick et al., 2016), and market structures (Bhandari et al., 2022). I build on this research by studying the role of environmental beliefs, a distinct potential barrier to coping with climate change.[7] I provide some of the first direct measurements of decision-relevant environmental beliefs using frontier methods in the elicitation of subjective expectations and directly link them to farmers' choices.[8] I show that these environmental beliefs have large consequences for climate change adaptation even amid the host of other frictions documented in the literature, with important implications for policy. For example, my results suggest increasing the measurement of soil salinity and disseminating those results could generate substantial economic gains.

Second, I add to the literate on the economics of beliefs. Despite the important role environmental perceptions in particular play in the economic decisions of smallholder farmers, we have extremely little evidence on how individuals form these expectations and learn from the signals around them (Kala, 2019). I am able to make progress on this question by combining precise belief measurement across a large geographic area with quasi-random natural experiments in past climate shocks. In doing so, I build on a growing set of papers studying the role of past experiences in shaping beliefs.[9] In my conceptual framework, past experiences

and Jack (2023), Balboni et al. (2023), Barwick et al. (2023), Berkouwer and Dean (2023), Glennerster and Jayachandran (2023), Greenstone et al. (2023), Hsiao (2023), Kochhar and Song (2023), Pelli et al. (2023), and Kala et al. (2023) for a review.

[5] See Jayachandran (2006), Deschênes and Greenstone (2007), Guiteras (2009), Dell et al. (2012), Mueller et al. (2014), Burke et al. (2015), Burgess et al. (2017), Shah and Steinberg (2017), Kaur (2019), Garg et al. (2020), Bharadwaj and Mullins (2021), and Somanathan et al. (2021), among others.

[6] The papers considering the economics of salinity intrusion (all of which happen to focus in Bangladesh) include Chen and Mueller (2018) on the migration consequences, Chen et al. (2022) on the effects on nighttime lights, and Guimbeau et al. (2023) on the health impacts.

[7] Most direct measurement of climate beliefs has taken place in rich countries, and often on more general aspects of the environment as opposed to those expectations that directly impact decision-making of the surveyed individual (Baldauf et al., 2020; Decehzleprêtre et al., 2022). The other papers examining the consequences of local environmental beliefs include Kala (2019), which inferred farmers beliefs about the monsoon onset using a revealed preference argument about their planting behavior, Bakkensen and Barrage (2022), which elicited flood risk beliefs in Rhode Island, and Zappalá (2023), which studied farmers' perceptions about past droughts.

[8] To my knowledge, the only other papers eliciting local, decision-relevant environmental beliefs in a manner that can be precisely compared to true conditions are Gine et al. (2015) and Burlig et al. (2024), in which the authors document predictors of farmers' beliefs about monsoon onset in India, measuring expectations using similar techniques to the ones I use here. See Delavande (2023) for a broader discussion of eliciting beliefs in low-numeracy populations.

[9] See, for instance, Mullainathan (2002), Malmendier and Nagel (2016), Enke et al. (2020), Malmendier

shape how Bayesian farmers interpret signals because they operate in a data-scarce environment, giving priors outsized influence on the interpretation of new data (Acemoglu et al., 2016; Gentzkow et al., 2023; Kapons and Kelly, 2023). The core features of my conceptual framework—many hypotheses, limited information, and heterogeneous experiences—apply to many economic settings beyond the agricultural decision-making process I examine. For example, stock brokers learning about the fundamentals of a firm by observing the stock price face a very similar problem. By focusing on the salient shocks and subtle shifts characteristic of climate change, however, I speak directly to how individuals learn about the warming world, which has implications for mitigation and adaptation policies more broadly.

Finally, I build on a long literature on the adoption of new techniques and products in developing countries. This important body of work has documented many significant determinants of the decision to invest in new technology—particularly in agricultural settings—such as social learning (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Beaman et al., 2021), comparative advantage (Suri, 2011), factor market failures (Jones et al., 2022), limited attention (Hanna et al., 2014), and present bias (Duflo et al., 2011). I test whether beliefs about the local environment play a role in shaping farmers' choices above and beyond these other factors. The results of my experiments and structural estimation show these beliefs can explain a large share of variation in technology adoption.[10]

2 How Accurate Are Environmental Beliefs?

This section provides background on soil salinity, presents the data on the environment and farmers' perceptions that I use to answer these questions, and compares expectations to the truth to test for evidence of learning frictions.

2.1 Soil Salinity

I focus on the amount of salt in the soil, a growing problem in both senses of the phrase. First, the salinity content of the soil can substantially impact plant health, particularly for rice—the most common and important crop in Bangladesh.[11] Farmers widely recognize this

(2021), Graeber et al. (2022), Fudenberg et al. (2022), and Bordalo et al. (2023).

[10] A growing set of papers has explored the value of information in environmental contexts, including Fishback et al. (2011), Rosenzweig and Udry (2013), Rosenzweig and Udry (2014), Weinberger et al. (2018), Barwick et al. (2019), Rosenzweig and Udry (2019), Fabregas et al. (2019), Molina and Rudik (2022), Leeffers (2023), Fairweather et al. (2023), Kruttli et al. (2023), Rudder and Viviano (2023), Shrader et al. (2023), and Burlig et al. (2024).

[11] Extending far beyond rice production in Bangladesh, soil salinity poses one of the greatest threats to food security worldwide. Across most crops, soil featuring a share of soluble salts in the root zone exceeding a plant's tolerance can significantly inhibit growth. Salt-affected soils occur globally—with an estimated

critical component of the agricultural production function, with 98.24% of respondents in my sample reporting that a lot of salt in the soil harms crops. Second, climate scientists forecast significant increases in soil salinity under global warming through a combination of environmental forces including rising sea levels, increased evaporation, droughts, and floods (Mukhopadhyay et al., 2021).[12] Projections in Bangladesh suggest soil salinity alone could reduce output by 15 percent in some areas by 2050 (Clarke et al., 2015; Dasgupta et al., 2015, 2018).

Farmers can adapt to this threat by altering their fertilizer, irrigation method, seed variety, or crop choice, or by shifting away from agriculture altogether. Salinity-tolerant seeds in particular offer a promising avenue for adaptation and are widely encouraged by the government in saline-prone areas. To understand the trade-offs associated with this seed choice, I consider two complementary pieces of evidence: farmers' own perceptions of the returns to these seeds, and objective agronomic assessments of their performance. I first elicit farmers' beliefs about the returns to switching their seed to a salinity-tolerant one (or vice versa if they are already planting a salt-resistant seed), and scale this by their expectation about their total harvest.[13] Appendix Figure A.1 plots the relationship between the perceived marginal return to salinity-tolerant seeds against farmers' beliefs about soil salinity on their plots (described further below). The results show a strong positive relationship that emerges above the government recommended threshold for adopting these resistant varieties, consistent with farmers understanding the value of this technology. To supplement these subjective assessments, I next build a new database of rice varieties by collating information on its resistance and tolerance to various environmental threats, typical plant height, yield, and year of release from various official sources.[14] Appendix Table D.1 presents simple bi-variate regressions illustrating how other seed dimensions vary with salinity-tolerance. Consistent with common understanding, the data show salinity-tolerant seeds perform worse than other varieties in ideal, non-salty conditions. The adoption decision therefore requires matching the

impacted area of over a billion hectares (Hopmans et al., 2021). From Sudan to Peru, Vietnam to Pakistan, salinity's negative impacts especially jeopardize the livelihoods of agricultural households in low- and middle-income countries (Singh, 2022).

[12] According to the Soil Resources Development Institute, the amount of salinity-affected land already increased more than 25 percent between 1973 and 2009 to reach 1.05 million hectares.

[13] I use a visual belief elicitation to elicit these expectations in a quantitative manner (Delavande et al., 2011). See Appendix Section C.2 for details.

[14] First, I use records from the website of the Digital Herbarium of Crop Plants from the Department of Crop Botany at Bangabandhu Sheikh Mujibur Rahman Agricultural University, which itself aggregates information from the seed developers. I supplement this data with additional information on the growth duration of each seed from the Food and Agriculture Organization on Bangladesh. For the remaining seeds without information from either of these sources, I execute online searches for each variety, conduct phone surveys with four seed dealers across the Khulna division, and use information directly elicited from farmers in my survey. For further details, see Appendix Section D.1.

technology to the appropriate conditions: in soil with high salinity, farmers achieve higher yields with salinity-tolerant rice but would be better off with a different seed choice amid low salt levels.

Identifying salinity levels on their own plots presents a substantial challenge for this population who lack access to accurate measurement technology. Table A.2 presents the indicators farmers report looking for to assess salt levels in their own soil, based off of responses to an open-ended question in the survey. Almost no farmer has ever used an electrical conductivity sensor or relied on a government agricultural officer who might have access to more accurate salinity information. Farmers most heavily rely on literally seeing the salt on the surface of the soil—which often represents an extreme form of salt inundation and cannot be seen once plots have been flooded—or on salient visual characteristics of the plant—such as ultimate height or characteristics of the leaves. Using plant characteristics can present a challenging inference problem when a host of other factors can also impact how plants look. To examine the extent of this challenge, I measure farmers ability to diagnose high levels of salt in the soil by showing a random selection of five images of rice plants damaged by salinity, disease, or pests, and asking them what may have caused the issues. When shown photos of rice grown in too much salt, 59.93 percent of farmers fail to report salinity as a potential explanation for any of the images, and 32.34 percent misidentify salinity as the cause of photos of plants suffering from pests or disease.

Consistent with the difficulty of directly measuring soil salinity, many farmers hold incorrect beliefs about the underlying environmental process behind soil salinity and effective adaptation. Although 83.63 percent correctly reported that soil salinity changes over the course of the calendar year, just 31.20 percent correctly name the months when it is highest.[15] Along the same vein, spraying sugar water across plots has emerged as a popular adaptation method, with 24.62 percent of farmers reporting that spraying sugar water helps to combat salinity and 10.93 percent having taken that measure in the past. Yet from an agronomic perspective, this method does nothing to improve the soil while costing farmers both time and money from paying for the sugar.

2.2 Measuring True Salinity Conditions

The scarcity of existing information on the local environment presents a significant challenge to farmers, government officials, and researchers. As a result, I primarily rely on salinity measurements I collected directly from farmers' plots. I equipped enumerators with handheld

[15] Relatedly, 34.40 percent correctly identified a month when salinity is typically lowest, and only 22.07 correctly answered both questions. See Appendix Section C.2 for further details on the construction of these variables.

electrical conductivity (EC) meters which I used to measure salinity twice during the 2022-2023 season on a randomly selected plot for each farmer in my main sample.[16] I take the first reading around the time of the baseline survey, prior to the planting of Boro season rice, and the second reading during the spring, prior to harvest. I convert these measurements taken at specific points in time during the calendar year into an average seasonal exposure to salinity based on a simple linear model I estimate from historical government soil samples collected monthly from nearby soil sites. The 2022-23 soil salinity season ultimately exhibited quite unusual patterns relative to historical norms, as shown in Appendix Figure A.5. The fall salinity readings indicate salt levels fairly typical of this area, with 82.58% of farmers' plots in my sample expected to experience salinity above the government-recommended threshold for adopting salinity-tolerant seeds. The salinity measured in February and March 2023 was significantly lower than previous springs, and after incorporating these additional measurements into the model, the resulting overall salinity was abnormally low for the season, with just 4.61% of farmers' plots ultimately above that same threshold. I verified these trends in the soil sample data I obtained from the government and confirmed the overall accuracy of the handheld sensors I used through laboratory tests.[17]

2.3 Measuring Farmers' Beliefs about Salinity Conditions

To capture farmers' perceptions of the local climate, I conducted surveys with 2,279 farmers who serve as the primary decision-makers on their plots across 250 villages spanning the Khulna division of Bangladesh.[18] The widespread geographic coverage of sampled villages throughout this area—which stretches more than 8,600 miles—allows me to capture significant spatial heterogeneity in climate conditions. I first interviewed households at the beginning of the Boro rice growing season and conducted an endline survey after harvest. My primary measures consist of farmers' past, present, and future perceptions of soil salinity. To elicit quantitative beliefs in this low numeracy population, I use a visual elicitation method

[16] I randomly selected plots, weighting by the area of the plot. The median number of plots is 2. When eliciting size, I asked about the five largest plots (which only binds for a very small minority of farmers): the median plot size is 24 decimals, where one decimal equals $1/100^{th}$ of an acre.

[17] For details on these tests and further details on this data collection, see Appendix Section B.2.

[18] For further details on sampling, see Appendix Section C.1. Appendix Table A.1 presents basic demographic information about this sample. Appendix Figure C.2 presents a timeline of data collection. The sample restrictions pose a potential threat to the interpretation of results. By limiting respondents to Boro season rice farmers still living in each village, I could be selecting for those farmers who—due to their beliefs—have not adapted by switching to a different economic activity or permanently migrating away. Two factors help to address this concern. First, those who have yet to adapt along these dimensions constitute the policy-relevant population. Second, an extremely low share of farmers in the sample report any likelihood of these extensive margin adaptation measures. For instance, 96.05 percent of respondents say no when asked, "If the amount of salt in your soil became very high, would you consider stopping planting rice in the Boro season?"

in which enumerators instruct farmers to place buttons on images in accordance with their perceived likelihood of each pictures scenario, a method which the vast majority of farmers understood well.[19] I use this approach to measure perceptions of soil salinity on farmers' own plots during the baseline survey for the 2022-23 season and during the endline survey for the 2023-24 season.

Eliciting comparable soil salinity beliefs to the agronomic measures presents a challenge because no one naturally thinks in deciSiemens per meter (dS/m), the units of the electrical conductivity sensors. I overcome this using an image of rice plants taken from the experiment in Grattan et al. (2002) in which researchers grew a standard rice seed in different growing conditions, randomly changing the salt content of the soil across treatment arm. The image shows seven plants from different arms at the end of the growing season, where the healthiest, tallest plants correspond to the lowest salt levels, and the least healthy plants grew in the highest soil salinity. I explain this image in detail to farmers, ensure their understanding that the photos capture different salinity levels, and then ask them to predict through the allocation of 10 buttons across the images which of the pictures they expect to look most like a plant grown by these researchers who used a not saline-tolerant seed on their plot, replicating all aspects of the farmer's soil such as the water, fertilizers, and weather over the course of the season.[20] Because I know the salinity conditions used in the original experiment to grow the rice plants in each picture, I can back out from farmers' answers their expected salinity in dS/m.

I designed this question to address several potential concerns in the elicitation of soil salinity beliefs. First, salinity evolves over the course of the season due to many factors, including the types of fertilizer used by farmers, the amount of rain (that can potentially wash away salts from the surface), and the irrigation water. As a result, I made sure farmers understood that these factors should be taken into consideration when making their prediction. Of course, one potential risk with this approach is that other factors can also impact the height of the rice plant separately from salinity, e.g., a certain pest or disease. Since I do not want to capture predictions about these other risks, I explicitly say to farmers,

[19] See Appendix Section C for details on this survey including links to the full questionnaires for the baseline and endline surveys, both in English and Bangla. Farmers answered a series of comprehension questions to gauge understanding of the belief elicitation method: across both the baseline and the endline, 63.67 percent of respondents answered these questions perfectly, and just 3.82 percent of respondents answered all three comprehension questions incorrectly. Just 11.63 percent reported that they found the example belief elicitation to be either "a lot" or "very" confusing. I pre-specified conducting cuts of the data by how well respondents answer these practice questions, their education, and their socio-economic status. This belief elicitation approach has been widely used to measure beliefs in similar settings (Delavande et al., 2011; Gine et al., 2015; Delavande, 2023).

[20] Appendix Section C.2 provides further details on the belief elicitation for soil salinity, rainfall, and for flooding and on the construction of these key variables.

"We are asking this question because we are trying to understand how much salt you think is in your soil." Piloting indicated that farmers clearly understood this question to be specifically about salinity, and answered accordingly. Second, I highlight that farmers should assume they are growing a seed that is not salinity tolerant to match the experiment of Grattan et al. (2002), a key step given that nearly half of farmers at baseline expect to plant a salinity-tolerant seed. Interactions across inputs pose a threat to the elicitation if, for instance, farmers irrigate their land with different water sources depending on their seed choice because I am asking them to hold those other factors constant. Reassuringly, however, I ask farmers directly in the endline, and the vast majority report that they do not change any of their other agricultural input decisions based on their seed choice. Third, one might worry that farmers simply predict what they expect their own plants to look like or what they recall from last year instead of thinking through the salinity prediction in particular. To address this, I explicitly ask farmers prior to the main salinity question about the picture that best describes their last harvest and their expectation about their harvest this coming season. Appendix Figure A.2 presents the full distribution of the mean salinity belief that farmers report in dS/m units. As a benchmark, the government recommends switching to a salinity-tolerant seed for values above 4.0 ds/m. In my sample, 43.86 percent believe the salt content on their soil will exceed this threshold over the course of the 2022-23 season.

2.4 Assessing Belief Accuracy and Testing for Learning Frictions

Linking the beliefs and environmental data together, I can compare how farmers' expectations compare to the ground truth and examine the extent to which these patterns suggest frictions in the learning process. On average, farmers exhibit strikingly accurate beliefs about the salt content of their soil. Figure 1 plots a binned scatter plot in blue comparing farmers' beliefs to the truth. The dashed gray line denotes the 45-degree line. The dots generally fall along the gray line. Consistent with this graphical evidence, Table 1 presents bi-variate regressions of beliefs against true salinity levels as measured via the agronomic sensors. At the 10 percent level, I fail to reject that the slope of this relationship equals 1 and the intercept equals 0, again consistent with accuracy on average. This holds both for the pre-specified sample of those who passed the belief comprehension checks in column (1) and including those who did not in column (2).

This average accuracy masks significant heterogeneity, however. As one way to see this, consider the explanatory power of the regression model. With simply a linear term of the truth and a constant, the agronomic truth can explain less than four percent of the total variance in beliefs. Even using a flexible function of the truth interacting the linear slope

with 20 ventiles, the explanatory power of the model still only ever reaches just over six percent.

Some of this unexplained gap between soil salinity beliefs and the truth likely stems from measurement error—but not all. I conduct several empirical tests illustrating that the residual belief error captures a meaningful economic object. Appendix Table A.4 shows results from specifications altering the scope of classical measurement error are inconsistent with noise alone explaining the results.[21] Table 2 and the corresponding plots in Figure 3 show that errors in farmers' beliefs explains their actual planting behavior with respect to salinity-tolerant seed choice above and beyond the true salinity content of their land. Farmers who overestimate the salt content of their soil by a larger amount also plant salinity-tolerant seeds on a larger share of their land ($p-$value $< .01$). This provides evidence that the gaps between beliefs and the true soil conditions capture more than just noise. Next, I examine the individual-level predictors of belief accuracy in Table 3. I focus on four dimensions that we expect might capture skill or farmer's ability to learn: the amount of land they farm, their age, whether their neighbors' view them as knowledgeable farmers, and whether their peers follow their agricultural advice. Across all of these dimensions, I find that accuracy is higher among farmers with more data (i.e., more land or experience) and viewed as better farmers. These facts are inconsistent with measurement error explaining the gap between beliefs and the truth and additionally suggest an important scope for learning frictions. Second, I examine the spatial component of errors. Figure 4 shows strong spatial covariance of errors, again inconsistent with classical measurement error.

These results raise an important puzzle: given the strong incentives farmers have to learn the truth, why do mistaken environmental beliefs persist in equilibrium? To shed light on that question, I next turn to a simple conceptual framework of farmer learning.

3 Learning About the Environment

This section presents a simple conceptual framework to illustrate how farmers learn about the environmental inputs into the agricultural production function amid the signals

[21] First, column (1) of Appendix Table A.4 shows that classical measurement error alone cannot explain the gap between measured salinity beliefs and the truth by regressing farmers' answer to a separate, binary question about whether they think the plot's soil is salty on continuous belief captured via the visual elicitation method. The results show a strong, positive relationship (p-value $< .01$). Measurement error would have to be correlated across these two very different questions in order to explain this relationship. Second, restricting the sample to those who understood the belief elicitation method better (as measured by their self-reported understanding) in column (2), the explanatory power of the model in fact decreases by more than 50 percent, inconsistent with a story in which measurement error in beliefs drives the errors. Third, error in the agronomic measure of salinity is unlikely to explain the results: regressing individuals' beliefs on the village leave-out mean, I find essentially no change in the explanatory power of the model.

characteristic of global warming.[22] The key assumptions I make stem directly from the quantitative findings from my surveys, the farmers' responses to the open-ended survey questions, and the narrative qualitative interviews I conducted on a separate sample. I include quotations and references to these statistics in footnotes throughout this section.[23]

3.1 Set-Up

A farmer grows rice in two periods $t \in \{1, 2\}$. Output in period t is given by the binary indicator $y_t \in \{0, 1\}$, where $y_t = 0$ denotes low harvest, and $y_t = 1$ denotes high harvest. Harvest is subject to a random productivity shock ξ that can be either negative ($\xi = -1$), positive ($\xi = 1$), or neutral ($\xi = 0$). I assume productivity shocks are distributed symmetrically with mean zero such that the positive and negative shocks occur with equal, positive probability denoted by $\rho > 0$ and that neutral shocks occur with positive probability such that $\rho < .5$. In the first period, farmers make no decisions about inputs and plant the standard seed. In the second period, salinity tolerant seeds are introduced, and farmers decide whether to plant salinity tolerant seed or plant the standard seed. This decision is given by the binary indicator $d_t \in \{0, 1\}$. Planting a standard seed is given by $d_t = 0$, where $d_1 = 0$ by default because salinity tolerant seeds are not available in the first period. In the second period, farmers may plant a salinity tolerant seed, denoted by $d_2 = 1$. Seed choice costs $c(d_t)$, where I normalize such that planting a non-salinity tolerant seed is free $c(d_t = 0) = 0$. I assume planting a salinity tolerant seed costs $c(d_t = 1) > 0$, where $c(d_t = 1)$ is positive yet small to capture the notion that salinity tolerant seeds perform relatively better in high salt environments yet relatively worse than standard seeds amid low salinity.

Two independent and unchanging environmental conditions denoted by the set $\{S, B\}$ can impact harvest, where S is the soil salinity and B is blast, an important fungus threatening rice.[24] I use lower case letters to denote the true, binary environmental states in these respective domains, given by $s \in \{0, 1\}$, where $s = 0$ denotes low salt levels and $s = 1$ denotes high salt levels, and by $b \in \{0, 1\}$, where $b = 0$ denotes no blast and $b = 1$ denotes the presence of blast.

I assume the agricultural production function follows a particular, simple functional form given by Equation 1. The maximization and minimization expressions ensure the binary

[22] For a formal treatment of all results in this section, please see Appendix Section E.

[23] See Appendix Section F.1 for details on the narrative interviewing approach I used and coding procedure for these qualitative data. For additional quotations supporting these assumptions, see Appendix Section F.2.

[24] The same intuition applies in the case of a large number of inputs. I focus on the case of unchanging environmental states: introducing the potential for unobservable shifts in the climate makes the farmers' problem even more challenging.

support of output $y_t \in \{0, 1\}$. Planting salinity tolerant seeds $d_t = 1$ mitigates the damage from soil with high salt content ($s = 1$).

$$y_t(s, b, d_t) = \max\left(\min\left(1 - (s - d_t)^2 - b + \xi, 1\right), 0\right) \tag{1}$$

I assume farmers cannot directly observe the environmental states.[25] As a result, farmers are uncertain about how their decision d_2 impacts their output y_2 entering the second period. This uncertainty is captured by their prior over the states of soil salinity and blast. I use $\widehat{\cdot}$ and lower case letters to denote beliefs in the environmental domains S and B, such that $\widehat{s}_t$ denotes a farmer's belief entering period t about the probability that true salinity levels are high $\widehat{s}_t = P(s = 1)$, and $\widehat{b}_t$ denotes a farmer's belief in period t about the probability that blast is present $\widehat{b}_t = P(b = 1)$. Other than potential misperceptions about the environmental conditions, I assume farmers are correctly specified in their understanding of the production function in Equation 1.

I assume that farmers use Bayes' rule to learn about these unobserved environmental conditions by updating using the harvest in period 1.[26] Note that a key feature of the production function is that farmers face an *identification problem*: low yield is consistent with both high soil salinity and blast, and high yield is consistent with low salinity and the absence of blast. In other words, I have assumed that poor harvest performance could reflect multiple potential issues in the agricultural production function, and these underlying threats manifest themselves in identical ways.[27] Given initial non-degenerate priors about the likelihood of high salinity $\widehat{s}_1$, Equation 2 gives the posterior at time t about the likelihood of high salinity levels, with the symmetric expression for blast beliefs $\widehat{b}_2$.

[25] As one farmer said during the endline survey, "We do not have the equipment to test the salinity of the land. Understanding the amount of soil salinity depends on our assumptions." Appendix Table A.2 illustrates this difficulty by showing the rarity of farmers using salinity sensors or consulting with government officials to learn about salt on their soil.

[26] Appendix Table A.2 shows this is the most commonly used way farmers learn about soil salinity, for instance. As one farmer explained, "We understand the amount of salt in the land by calculating the crop. If the crop is good, it means that the salt content in the land is low." Additionally, some farmers also use physical characteristics of the plants or the land, though this is less common and less universally applicable. For instance, one farmer pointed out that, "Salinity is easily recognized when the soil is dry and whitish in color. But I can not understand the salinity level if there is water in the land." Appendix Figure A.9 shows a photograph I took of this type of visible surface salt. This type of signal only appears at extremely high salt levels and is only visible—as this farmer points out—when the land is dry. This makes it impossible, for instance, to assess salinity using this indicator during the growing season because farmers keep their plots perpetually submerged with irrigation water, a particular issue given that salinity increases dramatically over the course of the season, and so is in fact lowest prior to irrigation when surface salts might be possibly viewed.

[27] As one farmer said when discussing crops being destroyed, "I don't know if it's because of the salinity or something else."

$$\widehat{s}_2 = P(s = 1|\widehat{s}_1, y_1) \tag{2}$$

Risk-neutral farmers choose seeds to maximize output in period 2 given their beliefs, as shown in Equation 3. To isolate the key mechanism of this framework, I have assumed that farmers stop growing after the second period, implying that farmers are myopic, maximizing only the current period's output without regard towards the potential learning benefits of experimentation.

$$U = \max_{d_2} \mathbb{E}\left[y_2(\widehat{s}_2, \widehat{b}_2, d_2) - c(d_2)\right] \tag{3}$$

Finally, I define the default hypothesis as the environmental threat that a farmer perceives to be the most likely to occur, as shown in Definition 1.[28]

Definition 1 An environmental factor $E \in \{S, B\}$ is the *default* hypothesis in period t when the corresponding belief $\widehat{e}_t = \max(\widehat{s}_t, \widehat{b}_t)$.

3.2 Implications for Learning

In this conceptual framework, beliefs exhibit path dependence: environmental threats that are perceived to be more likely remain relatively more likely after new data is observed. Remark 1 formalizes this statement, and Appendix Section E provides a formal treatment of this result.

Remark 1 *If two farmers i and j have different default hypotheses* $(\max(\widehat{s_1^i}, \widehat{b_1^i}) \neq \max(\widehat{s_1^j}, \widehat{b_1^j}))$*, then even if their priors are arbitrarily close* $(|\widehat{s_1^i} - \widehat{s_1^j}| < \varepsilon)$ *and* $(|\widehat{b_1^i} - \widehat{b_1^j}| < \varepsilon)$*, after observing identical data* $(y_1^i = y_1^j)$*, their posterior beliefs will exhibit the same difference in default hypotheses* $(\max(\widehat{s_2^i}, \widehat{b_2^i}) \neq \max(\widehat{s_2^j}, \widehat{b_2^j}))$*.*

Because low (high) yield is consistent with both high (low) salinity and (no) blast, the same signal can be interpreted in different ways while still being consistent with the data and correctly specified model.[29] Farmers endogenously process the same raw data differently in accordance with their default hypothesis: low yield is seen as evidence that the environmental threat they initially deemed more likely is indeed true, while high yield is interpreted just the opposite. Note that this divergence does not come from a misspecified mental model or a

[28] I assume no ties ($\widehat{s}_t \neq \widehat{b}_t$): that is, farmers never hold exactly the same belief about the likelihood of both environmental threats.

[29] This finding echoes a related result in Chen et al. (2023) in which by virtue of overestimating one technology, farmers endogenously adjust their beliefs about another to justify their observed outcome, leading to underestimation of the second.

behavioral friction. Instead, this reflects the Bayesian response and the fact that a farmer's priors affect how they interpret the ambiguous signal of harvest.[30] In fact, Appendix Section E shows that this effect scales with the degree of "defaultness": that is, when the gap in priors ($|\widehat{s_1^j} - \widehat{b_1^j}|$) is larger, this skew in interpretation of new data is exacerbated.

Figure 5 provides a graphical illustration of this path dependence in beliefs separately for the cases of high and low yields. Figures 5b and 5a plot the changes in beliefs for a given prior for high and low yield, respectively, and Figures 5c and 5d scale these changes to account for the mechanical limitations to magnitudes given the support of possible beliefs. All graphs exhibit a striking divergence around the 45°-line—the boundary along which default hypotheses switch. One way of interpreting Remark 1 is that two farmers on either side of that line will remain on separate sides, even after observing the exact same data.

3.3 Extensions

I now discuss various extensions of this conceptual framework that relax some of the main assumptions.

Uncertainty About the Agricultural Production Function Thus far, I have assumed that farmers perfectly observe how features of the environment impact yield. In practice, farmers learn about both the levels and the marginal costs of climate conditions. One immediate consequence of relaxing perfect information about the agricultural production function is the exacerbation of the identification problem. Farmers who observe a bad harvest now cannot distinguish between (a) low salinity and high salt damages, (b) high salinity and low salt damages, (c) low blast and high blast damage, (d) high salinity and low salt damages.[31] Recent work on convergence in misspecified learning by Heidhues et al. (2021) and Chen et al. (2023) illustrates that even after abstracting away from this identification issue, an incorrect mental model of the agricultural production function can persist in equilibrium as long as farmers can find beliefs that justify the observed output.

Experimentation A key aspect of this set-up has been the lack of experimentation. In practice, farmers in fact trade off between benefits today and learning for tomorrow (Foster and Rosenzweig, 1995). Indeed, the empirical results that farmers with more land and more experience hold more accurate beliefs suggests an important role for experimentation

[30] More general results on the failure of asymptotic agreement in under-identified settings can be found in Acemoglu et al. (2016).

[31] This occurs even abstracting away from bottom-coding issues given the impossibility of negative harvest. Non-linearity of this type in the production function (as featured in Equation 1) further exacerbates this issue as because farmers then cannot distinguish between high salinity and high or low salt damages.

to overcome this identification problem. Challenges can nevertheless occur that reduce the efficacy of experimentation in leading to correct beliefs. First, I have assumed that farmers are correctly specified about every other aspect of the production function, including notably the ways in which environmental conditions impact yield and the distribution of idiosyncratic production shocks. In reality, these primitives are also unknown to farmers, and uncertainty over those dimensions will exacerbate the identification challenges. Second, when experimentation is costly (due to a subsistence constraint or fixed cost in technology investment, for instance), then farmers may choose not to experiment *despite* knowing that their beliefs may be wrong because given their priors, the expected benefit of learning may be small.

Social Learning The conceptual framework also omits another important feature of this setting: social learning. In many ways, one's neighbors act as a form of experimentation as discussed above, particularly given the high spatial covariance of environmental threats. One factor that may inhibit the role of social learning to correct beliefs is the nature of errors in this setting. I have thus far been agnostic as to the origins of priors, but as I show below, past environmental experiences play a key role in shaping farmers' initial beliefs. Because environmental experiences also exhibit high spatial covariance, the neighbors one might turn to for advice likely hold a similar set of biased beliefs.

Higher Dimensional Signals I have focused on the simple case in which farmers only observe yield which can either be high or low. In practice, farmers can observe a much richer set of data, such as the color of leaves or a continuous measure of yield. This will help farmers learn when they correctly know that certain physical characteristic is linked with a particular threat, and therefore the cumulative distribution functions of the two environmental beliefs are in fact disjoint, thereby alleviating the identification problem. I find evidence that farmers often do not have knowledge of, for example, the particular leaf patterns emblematic of blast. Without this type of external information, higher dimensional signals do not alleviate the learning challenges in this conceptual framework.

Learning in the Long-Run This set-up has focused on the two-period case. What happens to beliefs in the long run? Under the assumption that farmers have full information about the distribution of idiosyncratic production shocks, farmers in this set-up will be able to distinguish between a world in which both environmental threats are present and only one, yet in the one threat case, they are not guaranteed to ever be able to learn *which* threat is prevalent. In the case where farmers are also uncertain about the distribution of shocks,

then even the case of both threats will be fully consistent with the data when in reality only one threat is relevant, or vice versa.

3.4 The Formation of Priors

This conceptual framework has illustrated the power of priors in this setting. How do farmers form their initial beliefs about the environment? Many potential sources of information could influence which factors farmers deem important *ex ante*, ranging from wisdom passed along from their peers to guidance provided by agricultural extension workers.

I study how the nature of past environmental shocks shapes the priors held by farmers and thereby the ways in which they interpret signals about their yield. Under the upheaval of changes to the climate under global warming, understanding how farmers learn from shifts in the environment holds particular policy relevance. I therefore focus on two types of signals characteristic of climate change: salient shocks and subtle shifts. Almost all environmental threats impacted by global warming feature these two classes of changes. For dimensions like rainfall, temperature, and pollution, I consider a daily or seasonal extreme as a salient shock, while secular trends in the distribution—while nevertheless important and observable—I classify as subtle shifts. For natural disaster risk, marginal increases in determinants of the hazard (e.g., riverbank erosion or sea-level rise in the case of flooding) constitute subtle shifts, while an occurrence of the natural disaster (e.g., a flood itself) counts as a salient shock. Appendix Table A.5 provides examples of salient shocks and subtle shifts across a host of different environmental dimensions.

How might salient shocks and subtle shifts vary? I discuss three channels through which experiencing one event rather than the other can lead to persistent learning differences. The relevant mechanism will depend on the context and assumptions made about the types of data to which farmers attend, as I discuss further below. In all cases, however, the main empirical predictions remain the same.

First, remaining in the purely Bayesian world with no restrictions on farmers' attention, a salient shock may result in different equilibrium beliefs than a subtle one if the statistical content of the signal varies. For example, imagine that experiencing a flood (a salient shock) indicates a significantly higher future flood risk than an incremental increase in river height (a subtle shift) because of what the occurrence of a flood indicates about the data generating process. In this case, a Bayesian farmer who happens to experience a flood as a salient shock will rationally enter the next period with a higher prior belief about future flood risk than one who experienced a subtle shift. This gap in initial beliefs can be amplified by precisely the same mechanism as discussed above in which farmers endogenously alter how they interpret

new data in line with their priors. The model predicts that whichever signal suggests higher posterior environmental risk will cause farmers to overestimate that risk by more.[32] Whether or not the true statistical content of the signal varies across these two events is a difficult question to answer and of course highly context specific. It is not obvious in this example, for instance, that experiencing a flood (a salient shock) should indeed increase beliefs about future flood risk more than an incremental increase in river height (a subtle shift).

I now focus on the case where both the salient shock and subtle shifts carry the same statistical content but nevertheless lead to learning gaps. To capture this, I introduce limited attention in two ways. First, I consider a case where due to the subtlety of incremental change, that class of shifts can be simply difficult for farmers to notice. In that world, because farmers fail to attend to this information, they do not update their priors about environmental risk in the same way as they do after experiencing a salient shock, even though both events inherently have the same information about the underlying data generating process. This causes farmers to enter the next period with different initial beliefs, a gap which can be exacerbated by the same mechanisms in this underidentified setting as before. This assumption about failing to notice subtle shifts may be more plausible in some settings than others. On the one hand, incremental increases in the salinity of irrigation water may be too small to be detected by tasting the water, which means that even if farmers use all the tools available to them in an attempt to notice such a change, they may fail to detect such subtle shifts. On the other, farmers might always notice if its raining, and so even if the number of days of rain does change from year to year, they will always have attended to this data.

Even in cases when the subtle shifts and salient shocks feature the same statistical information and farmers attend to both, differences in the nature of the stimulus can still generate learning gaps amid a different form of limited attention. Specifically, instead of only attending to some environmental data, I instead assume that limited attention restricts the set of potential hypotheses that farmers consider when interpreting a new signal. As discussed in (Hanna et al., 2014), farmers face a high-dimensional space of factors potentially impacting agricultural production, and keeping track of all of these inputs can be challenging in of itself, let alone conducting Bayesian inference across so many dimensions. Salient shocks attract attention automatically and involuntarily, or "bottom up" (Bordalo et al., 2022a).[33]

[32] Consistent with the empirical evidence I find in the setting of soil salinity, I focus on the case where salient shocks lead to overestimation, and thus I assume it is the salient shock that suggests higher environmental risk.

[33] Psychologists have noted the evolutionary origins of this cognitive mechanism, which features clear survival benefits. The psychology literature on climate change has particularly noted the challenges this instinct poses when applied to the mitigation of and adaptation to environmental harm (Gilbert, 2006; Johnson and Levin, 2009; Gifford, 2011).

In doing so, they change which the set of hypotheses that come to mind by making some environmental culprits more salient.[34] This system 1 thinking occurs near-instantaneously: upon seeing new data, farmers apply Bayes' rule to update their beliefs only among the factors that come to mind. By neglecting those potential explanations that are less available (Tversky and Kahneman, 1973), farmers effectively put more weight on those factors for which they experienced a salient shock.[35]

The key empirical prediction distinguishing these two limited attention frameworks and the Bayesian benchmark stems from how salient shocks differentially impact posterior beliefs as compared to subtle shifts, *conditional* on their true statistical content being equivalent. Without attention costs, this ancillary feature of the salient shock is irrelevant for farmers beliefs and ultimate decisions. By contrast, amid limited attention that causes farmers to either fail to notice the subtle shift entirely or make it less likely that environmental explanation comes to mind, learning gaps can persist.

4 Learning from Salient Shocks and Subtle Shifts

To test the key predictions of this conceptual framework in the domain of soil salinity, I study two quasi-random natural experiments to capture salient shocks and subtle shifts in the environment.

4.1 Empirical Strategies

Salient Shocks: Saline Floods In my framework, a salient shock alters the environment in a manner that captures farmers' attention. Under this definition, I can take a data-driven approach to identifying the ideal treatment to answer this question. I ask all of the farmers in my sample to recount any event that has affected the salinity of the soil on the plot about which I measure beliefs and from which I take agronomic readings. At the time of the baseline survey, 17.60% of framers reported ever experiencing at least one event. Flooding is the modal answer, constituting 38.53% of 436 total events. Farmers report salinity increasing in 91.67% of these cases, and specifically mention flooding with salty water 73.21% of the

[34] Differently from other models of limited attention and belief formation, in my set-up, farmers always attend to all data from which they might learn. Limited attention acts on how farmers interpret these signals and which features are associated with a given hypotheses by limiting the number of hypotheses farmers entertain. In the language of Tversky's (1977) similarity framework, my model makes the similarity function linking signals and hypotheses endogenous to attention. Previous work, by contrast, treats the link between data and hypotheses as exogenously determined, and instead operationalizes limited attention through agents attending to some signals more than others (Schwartzstein, 2014; Hanna et al., 2014; Gagnon-Bartsch et al., 2021; Bordalo et al., 2022b).

[35] This result echoes the empirical findings in Enke (2020).

time.

Based on this overwhelming dominance of saline floods among recollected salinity events, I focus on identifying the causal impact of these climate occurrences to test for the role of salient shocks in the model. I adopt a difference-in-differences specification to capture quasi-random variation in the incidence of saline floods, as shown in Equation 4, where $v(i)$ denotes that the variable is defined at the village level. This main coefficient of interest β_1 captures the differential impact of experiencing a flood with more salty water, holding constant underlying flood risk, local water saltiness, and any other broader impacts of experiencing a flood. I cluster standard errors at the village level, the level of treatment. The quasi-random determinants of flooding and water saltiness allow me to interpret β_1 as the causal impact of happening to actually experience this salient shock. Conditional on the same underlying flood risk, otherwise identical places may or may not experience a flood due to idiosyncratic variation in factors like rainfall and tidal surges. Similarly, the salt content of water fluctuates randomly due to factors like cloud cover changing the evaporation rate. Water salinity also evolves over the calendar year, and as such the precise timing of a flood can dramatically shape the salt content of the floodwaters. Appendix Figures B.8 and B.11 visualize this variation in salinity in Bangladesh's rivers and in the ocean off the coast.

$$Y_i = \alpha + \beta_1 Floods_{v(i)} \times Water\ Saltiness_{v(i)} + \beta_2 Flood_{v(i)} + \beta_3 Water\ Saltiness_{v(i)} + \psi Flood\ Risk_{v(i)} + \varepsilon_i \tag{4}$$

To measure past exposure to flooding, I use the same data as before in which I combine methods from machine learning and geophysics to estimate local floods, described further in Patel (2023). I use a supervised machine learning model trained on underlying geographic factors to derive my measure of flood risk, as discussed further in Appendix Section B.1, and control for fixed effects of quintiles of this variable. To capture the saltiness of water, I combine two sources. First, I use remote sensing data from satellites that detects fluctuations in ocean salinity from space. Second, I use data I obtain from the Bangladesh Water Development Board from river stations scattered across the country at which the government directly collects water salinity measurements. Together, these data allow me to measure the saltiness of local water.[36] For ease of interpretation, I convert water salinity into standard deviation units.

Subtle Shifts: Rising Sea Levels I cannot use farmers' own assessments to measure subtle shifts as almost by definition, they will not recall or easily identify these events.

[36] See Appendix Section B.3 for details on these data.

Instead, I rely on the scientific literature on soil salinity and climate change, which specifically points to rising sea levels and the associated impact on irrigation water as a driving force behind increasing salinity (Mukhopadhyay et al., 2021). As global warming increases sea levels along the Bangladeshi coast, the water farmers use to grow rice can become saltier, either through underground contamination of groundwater (as illustrated in Appendix Figure A.8), or through river systems themselves becoming saltier upstream.

To isolate the causal impact of subtle shifts, I estimate a triple-differences design. For the first difference, I calculate deviations in ocean water salinity from the village-by-calendar-month mean, thereby comparing villages with relatively higher or lower exposure to local water saltiness compared to their historical average. Second, I calculate deviations in ocean sea level elevation from the village-by-calendar-month mean, thereby comparing villages exposed to higher vs. lower sea level rise relative to their historical average. Finally, I compare villages closer vs. farther from the coast, under the intuition that rising sea levels will infiltrate irrigation water more the closer the village lies to the ocean. This yields the estimating specification shown in Equation 5. I cluster standard errors at the village level, the level of treatment.

$$Y_i = \alpha + \beta_1 Closer \times Saltier \times Higher + \beta_2 Closer \times Saltier + \beta_3 Closer \times Higher + \beta_4 Saltier \times Higher + \beta_5 Closer + \beta_6 Saltier + \beta_7 Higher + \varepsilon_i \quad (5)$$

The coefficient of interest β_1 captures the degree to which villages closer to the coast happens to be exposed to higher sea levels when the ocean is saltier have differential outcomes. The quasi-random nature of ocean salinity (which can be influenced by factors like cloud cover and evaporation) and sea level rise give a causal interpretation to this regression. I use data from satellites to capture ocean salinity and sea level elevation; see Appendix Section 2.2 for details on these data and variable constructions. For ease of interpretation, I convert all units into standard deviations.

4.2 Differential Impacts of Salient Shocks and Subtle Shifts

As a first step in the analysis of these natural experiments, I examine the impact of experiencing a salient shock and subtle shift on the true soil salinity conditions as measured by the agronomic sensors. Columns (1) and (2) of Table 4 present the results of this analysis, reporting the main interaction terms from each difference-in-differences specification (the

primary coefficients of interest).[37] Coincidentally, it happens to be the case in this context that experiencing a 1 s.d. saltier flood has equivalent impacts on true soil salinity as does experiencing a 1 s.d. increase in ocean elevation interacted with a 1 s.d. saltier ocean among villages 1 s.d. closer to the coast. The salient shock increases measured salinity on farmers' plots by .153 dS/m while the subtle shift increases it by .113 dS/m, and these effects are not statistically different from one another.

In contrast to the comparable impacts on true salinity, I find markedly asymmetric impacts on farmers' beliefs. Columns (3) and (4) of Table 4 present the two main interaction terms from the difference-in-differences specifications and shows that only salient shocks increase beliefs, while subtle shocks have no detectable effects. Furthermore, salient shocks increase beliefs by more than the corresponding impact on true salinity, leading to overestimation.

The model predicts that this asymmetry occurs through the differential interpretation of signals in this underidentified environment. I directly test for this mechanism by combining other survey questions and a lab-in-the-field elicitation with this natural experiment. First, I simply ask farmers in the endline survey about the signs that they use to learn about salinity on their soil. Table 5 shows that consistent with the theory, farmers experiencing a salient shock disproportionately interpret more "generic" signals as indications of high salinity, particularly as compared to those having experienced a subtle shift. This is consistent with the identification problem: after a salty flood, farmers expect salinity to be a likely environmental threat, and as a result, they interpret even generic indicators of crop health like plant death and stunted height as indications that consistent with their prior belief, salinity is high. Farmers experiencing a subtle shift show no similar pattern.

By the same underlying channel, farmers experiencing salinity intrusion into their irrigation water via a rising sea level shock are more likely to attribute physical indicators of salinity to other causes. After the question referenced above in which I ask farmers to list all features of the crop that might indicate salinity, I ask farmers to list any other environmental causes that could generate the same symptoms. Despite reporting fewer symptoms overall, especially among generic indicators of plant health (as indicated by the results in Table 5), I find that farmers who experienced a subtle shift list significantly more other environmental explanations that could disguise themselves as soil salinity.

[37] Appendix Tables A.6 and A.7 present coefficients for the full difference-in-differences specifications for the outcome of true soil salinity. Appendix Tables A.8 and A.9 does the same for the outcome of farmers' beliefs.

5 Consequences of Beliefs for Climate Adaptation

How do these patterns of beliefs shape decision-making? This section presents the results from large-scale field experiments to measure the economic consequences of environmental beliefs.[38]

5.1 The Causal Impact of Beliefs on Technology Adoption

I estimate the causal effect of environmental beliefs on farmers' adoption of climate technology using an information experiment centered around the salinity data I collect from farmers' plots. The script introduced this information to farmers by explaining that the salt levels would be provided using both numbers and pictures and also in relative terms to the government's official recommendation regarding the soil salinity level above which farmers should adopt saline-tolerant seeds. I embed this randomization into a Becker-DeGroot-Marschak elicitation of farmers' willingness-to-pay (WTP) for information about the amount of salinity on their soil.[39] To measure demand using this approach, beginning with a low price, I ask farmers if they would be willing to purchase the item for that price.[40] If they say yes, I increase the price and ask again, repeating until the respondent says no. I then randomly select a price, and if the respondent was willing to purchase the good for that amount, the transaction takes place. Following Dizon-Ross and Jayachandran (2022), I elicit WTP for a plate as a benchmark good to include as a control.

In the case of soil salinity information, Farmers faced a price randomly drawn price from a bi-modal distribution with nearly all of the mass split evenly between 0 and 500 BDT, while prices in between were drawn with probability .0001 each to preserve incentive compatibility. Prior to the price list elicitation, most farmers viewed this information as useful: 69.42 percent of farmers said it would be extremely or very helpful, while just 1.98 percent said it would not be helpful at all. Farmers' demand reflects these attitudes: the median farmer offered up to 30 BDT for the information, with an average value of 41.74 BDT, and 84% of farmers had a positive willingness-to-pay. Within this information experiment, I also randomize the type of information: some farmers were only offered average salinity levels for their upazila, while others were in addition offered data on their own plots. I delay

[38] For a discussion of the ethical considerations around all of the experiments in this project, see Appendix Section G.

[39] I pre-registered this experiment here. All randomization occurred within strata, which I define based on the intersection of village with direction of soil salinity belief error, determined based off of farmers' baseline beliefs.

[40] During piloting, it became clear that norms against accepting gifts for free distorted the results of the BDM when beginning at 0 BDT. To circumvent this issue, I begin the price list at a small positive value, and only ask about 0 BDT if the respondent says no to that initial price.

discussion of this aspect of the experiment until the conclusion, but note the distinction for now to explain the specification for estimating treatment effects.

As a first-stage, this information does shift farmers' perceptions about the amount of salt in their soil next year. I test for this effect by estimating equation 6, where $Prior$ denotes farmer i's prediction about the 2022-23 soil salinity level on their plot, $OwnTruth$ signifies the true salinity as measured by the EC-meters, $UpaTruth$ denotes the average of these salinity readings across all participating farmers in i's upazila, $OwnTreatment_i$ indicates whether farmer i received information about $OwnTruth$, and $UpaTreatment_i$ does the same for $UpaTruth$. Table 9 presents the results from these regressions, showing large F-statistics regardless of whether I restrict the sample to those who pass the belief elicitation comprehension checks or not. Consistent with soil salinity ultimately being quite low during the 2022-23 season, most farmers in the treatment group updated downwards, expecting lower salt levels in the future.

$$Y_i = \alpha + Prior_i + \beta_1 OwnTreatment_i + \beta_2 OwnTreatment_i \times OwnTruth_i + \beta_3 OwnTreatment_i \times UpaTruth_i + \beta_4 UpaTreatment_i + \beta_5 UpaTreatment_i \times UpaTruth_i + \varepsilon \tag{6}$$

Using treatment as an instrument, I can measure how demand for salinity tolerant seed varieties varies with beliefs. I show that randomly providing farmers with information about salinity content on their soils as measured by the EC-meters (using the same visual as the belief elicitation to convey the information) shifts both their incentivized willingness-to-pay for a salinity-tolerant seed variety and their expectations about agricultural behavior moving forward. In this context, the information typically shifted farmers' beliefs downward, with a resulting decrease in the seed willingness-to-pay and a decline perceived likelihood of planting a saline tolerant seed next season, as shown in Table 10.

5.2 Salinity-Tolerant Seed Experiment

How much does adaptation to salinity through seed choice impact farmers' ultimate economic well-being? To answer that question, I conduct a separate experiment, randomly varying the marginal cost of planting saline-tolerant seeds during the baseline survey.[41] Sim-

[41] The endline survey exhibited an extremely low attrition rate of 1.14 percent. Farmers receiving a saline-tolerant seed BDM price of 0 during the baseline were slightly more likely to not be found in the endline (0.77 percentage points), an effect which is marginally statistically significant (p-value=0.08) but amounts to a difference of just 9 farmers, too small to meaningfully impact any results.

ilar to the soil salinity information above, I embed this randomization into the price drawn during the Becker-DeGroot-Marschak elicitation of farmers' WTP for a saline-tolerant variety.[42] I draw the seed price from a distribution of prices that equals 0 BDT with probability .49905, 200 BDT with probability .49905, and each price in between with probability .0001 to preserve incentive compatibility. Let $Free_i$ denote an indicator for whether farmer i was randomly assigned a price of 0. To estimate the impact of adoption of the saline-tolerant seed variety $Tolerant_i$ on agricultural output Y_i, I estimate the regression specification in equation 7, instrumenting for adoption of seed 67 using $Free_i$. Following the pre-specified heterogeneity, my primary specification excludes the 31.07 percent of households who had already purchased their seeds for the season at the time of the interview.

$$Y_i = \alpha + \beta Tolerant_i + \varepsilon_i \tag{7}$$

I "reversed" the experimental design in this format (randomizing seeds in the baseline and information in the endline) in order to minimize the rounds of data collection necessary. In order to interpret the impact on harvest as a potential consequence of information, I must assume that the set of compliers is comparable across the two experiments: that is, that the marginal seed user induced via the information treatment is not differentially selected from the marginal seed user induced via the lower BDM price. In fact, an even weaker assumption still allows me to take an intent-on-the-treated interpretation: if shifting beliefs about salinity is equivalent to changing the marginal cost of purchasing saline-tolerant seeds, then I can link these two experiments together to draw conclusions about the impact of information on profits.

I first find that 85.9 percent of farmers have a positive WTP for these seeds, and 38.7 percent have a WTP greater than the market price.[43] Randomly assigning half of farmers a price of zero, I find that being assigned a price of 0 during the BDM procedure increases the likelihood of planting a salinity-tolerant seed by 13.87 percentage points, or 32 percent of the control group mean of 42 percentage points. In both the reduced form and instrumental variables framework, I find evidence that the treatment (and subsequent increase in salinity tolerant planting) reduced agricultural profits, consistent with soil salinity being especially low during the 2022-23 season. Together, these results show that environmental beliefs on the margin have large economic consequences for farmers' bottom-line.

[42] I pre-registered this experiment here.

[43] Because I was delivering the seed directly to the farmer and the BDM procedure is itself quite different than the experience of purchasing in the market, the relative levels to the market price should be interpreted cautiously.

6 Learning Counterfactuals and Policy Simulations

How large of a role do beliefs play in adaptation to climate change? The experimental results show that environmental expectations have large causal impacts on choice, but characterizing the overall magnitude is challenging from those reduced form estimates alone for several reasons. First, my primary outcome is willingness-to-pay as captured via the BDM mechanism. The nature of this elicitation (for instance, giving farmers 500 BDT immediately prior to the measurement) likely introduces level effects, as suggested by the significant share of people willing to pay above market price. Second, the impact of environmental beliefs on seed choice is non-linear. Third, the experiment captures how demand for salinity-tolerant seeds varies with beliefs about next season's salinity, measured over the summer. Farmer's beliefs likely evolve between the endline and the actual planting decision in the winter as they observe more information about their land, and the ideal elasticity relates to beliefs at the time of planting. Fourth, I can only identify the local average treatment effect of those farmers whose beliefs were influenced by the experiment.

Because of those limitations, in this section I turn to a structural model of demand for climate adaptation. Specifically, I estimate a random utility model of seed choice (McFadden, 1974; Train, 2009). In doing so, I can explicitly account for the non-linear role of beliefs to answer questions like: how would seed choice change if farmers had completely accurate beliefs about soil salinity? What would happen if farmers learned from subtle shifts in the same way as salient shocks? How effective would a government information intervention policy be at increasing climate adaptation?

6.1 Model of Seed Choice

I expand the set-up from Section 3 to allow farmers indexed by i living in village v to choose among $\mathcal{J}_v$ dry-season rice varieties. To simplify the estimation, I focus on the seed choice that each farmer plants on the largest share of their land, such that each farmer chooses one seed among the set $\mathcal{J}_v$. Because I cannot observe the true set of potential seed choices faced by each farmer, I define $\mathcal{J}_v$ to be the superset of all seeds planted on any amount of land by any farmer $i \in v$.

Equation 8 captures the indirect utility of farmer i choosing seed j, where $\mathbf{x}_{ij}$ is a vector of seed-specific characteristics, β_i are random coefficients that vary over farmers in the population, $\mathbf{w}_{ij}$ is a vector of seed-specific characteristics, $\mathbf{z}_i$ is a vector of farmer-specific characteristics, c_j are intercepts, and ε_{ij} are unobserved random taste shocks, which I model

as independent type I (Gumbel-type) extreme-value random variables.

$$U_{ij} = V_{ij} + \varepsilon_{ij} = \mathbf{x}_{ij}\beta_i + \mathbf{w}_{ij}\alpha + \mathbf{z_i}\delta_j + c_j + \varepsilon_{ij} \tag{8}$$

By including random coefficients, I allow for correlation of choices across alternatives and therefore do not have to assume independence of irrelevant alternatives. I do not directly estimate β_i but rather assume they have a multivariate normal distribution $\beta_i \sim N(\mu, \Sigma)$ and estimate μ and Σ. Given this utility function, the probability of the ith farmer choosing seed j is given by equations 9 and 10. Because the integral in Equation 10 has no closed-form solution, I compute it using maximum simulated likelihood.

$$P_{ij}(\beta) = \frac{e^{V_{ij}}}{\sum_{k \in \mathcal{J}_v} e^{V_{ik}}} \tag{9}$$

$$P_{ij} = \int P_{ij}(\beta) f(\beta) d\beta \tag{10}$$

6.2 Estimating the Demand Model

To estimate this model, I begin by building a database of seed characteristics. For the vector $\mathbf{w}_{ij}$, I include seed price, indicators for whether the modal farmer I survey reported that a feature of that seed was processing ease, market demand, good taste, and nice color, along with whether the modal farmer reported that seed being resistant to pests and resistant to salinity.[44] For the vector $\mathbf{x}_{ij}$, I include the interaction of salinity-tolerant status with the error in farmers' beliefs about the amount of salt in their soil. In my default specification, I do not include any further farmer-specific controls $\mathbf{z}_i$. I cluster all standard errors at the farmer level. I estimate the model exclusively among the treatment group from the seed experiment, and always exclude those who fail the beliefs elicitation comprehension checks at baseline. This provides a natural out-of-sample goodness-of-fit test by examining how the predictions of this model perform among the control group.

The results show that the model performs well and that beliefs about salinity levels strongly impact adoption of salinity tolerant seeds. First, coefficients in general tend to be consistent with our priors about how adoption should respond to seed attributes. Market demand for the harvested rice and the associated selling price are both strongly positively related with seed adoption, for example. Second, the standard deviation of the random coefficients for the interaction of salinity tolerance and belief error is positive and statistically different from zero, allowing me to reject that independence from irrelevant alternatives

[44] Note that I use farmers' beliefs about seeds because that is the characteristic of the variety upon which they make their decision.

would apply in this setting. Third, the model's predictions out of sample in the control group strongly predict farmers' actual behavior in that population. Table 8 presents simple regressions of actual seed choice on the model's predicted probabilities, separately for the treatment group in the seed experiment (upon which the demand model was estimated) and the control group. In both cases, I cannot reject a coefficient of 1: in other words, as the model's prediction about whether seed is selected increase by 1, so to does the actual behavior of farmers.

How does seed adoption change under alternative beliefs? I estimate these counterfactuals by calculating the marginal predicted choice probability for each seed under alternative variables using the structure of the demand model. Figure 9 plots these predicted adoption probabilities by error separately for farmers whose soil salinity initially fell below the government recommended threshold of 4.00 dS/m for switching to salinity tolerant seeds and those for whom their measurement fell above. Those facing high salinity levels always exhibit higher adoption rates regardless of their beliefs, reflecting differences in the seed market across places. Yet for both groups, beliefs play a large role in shaping adoption. On average, moving from the 10^{th} percentile to the 90^{th} percentile in belief error corresponds to an 80.87 percent increase in the probability of adoption a salinity-tolerant seeds.

7 Conclusion

This paper has documented the importance of environmental beliefs in shaping adaptation to climate change.

One clear policy recommendation emerging from these results is the value of collecting and disseminating information about local climate conditions. From a policy perspective, providing farmers with soil salinity information collected from their own plots presents a significant cost—requiring individual-level analysis of farmers' own plots. Instead, providing farmers with aggregate averages at a higher-geographic level presents an appealing possible alternative. In this context, upazila represents the relative geographic level for this policy question as the administrative unit at which the Bangladesh government currently collects and reports soil salinity data. The success of this type of policy hinges on the answer to two questions. First, do true salinity conditions covary enough within these areas to meaningfully speak to farmers' own agricultural decisions? I examine this question using the new soil salinity data I collected, and find that for the most part, the aggregate value accurately captures farmers' own experiences. The correlation between the salinity measured from farmers' own plots and the mean value across all plots from which I collect data in their upazila is 0.81. To put this magnitude in perspective, I provide two benchmarks. First, the

correlation between a farmers' own plot and their village average is 0.92.[45] Second, I take advantage of the three measurements taken by enumerators on each plot during each round of soil salinity data collection. This gives me six pairwise correlations of salinity on the *same* land on the *same date*, helping to provide a sense of the maximum correlation given the sensor's measurement error. On average, the average pairwise correlation from these measurements is 0.90. To view this another way, the government threshold for switching to saline tolerant seeds is 4.5; the recommendation based on data from a farmer's own plot disagrees with that from the upazila average just 2.57 percent of the time.

Do farmers themselves recognize this strong agronomic link between their own plots and that of their peers in the upazila? To speak to that second key question for this overall policy, I consider two types of evidence. First, I directly measure farmers' subjective views of the similarity of their plots to that of their neighbors, both within their village and within their larger upazila. Consistent with the soil salinity measurements, farmers perceive a lower spatial covariance at higher geographic levels: 25.48 percent say that the salt on their plot is extremely or very similar to their village average, as compared to 9.23 percent for the upazila average. In fact, farmers for whom the difference in absolute value between their plot and their upazila mean minus the absolute value of their plot and village mean are also significantly more likely to qualitatively answer that their salt is more similar to their village neighbors (p-value = 0.02). Second, I take a revealed preference approach and embed an additional arm within the information experiment, offering a random 75 percent of farmers information both about their own soil salinity and the average salt levels for all farmers in their upazila, while the remaining farmers could only purchase information about the latter upazila mean figure. I find no evidence that farmers value these aggregate figures less than data about their own plots. Appendix Figure A.6 plots the distribution of farmers' demand for information separately by those offered the upazila-average only compared to those offered their own plot's data as well, and the distributions almost entirely overlap. In OLS regressions predicting demand for information conditional on WTP for the plate, I find no strong evidence of a systematic impact of the format of information provided, with a 95 percent confidence interval of -6.53 to 2.61. Overall, these results suggest that although farmers indeed recognize that information at higher geographic aggregates applies less to their own agricultural conditions, providing this (much less costly) information can nevertheless still shift beliefs.

[45] Note that both the upazila and village correlations are upwardly biased towards the truth because of imputations used replace missing salinity measurements (due to flooding or farmer unavailability), though this procedure reflects a reasonable implementation of this policy in practice.

References

Acemoglu, Daron and et al., "The Environment and Directed Technical Change," *American Economic Review*, 2012, *102* (1), 131–166.

__ , Philippe Aghion, Lint Barrage, and David Hémous, "Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution," Technical Report, National Bureau of Economic Research 2023.

__ , Victor Chernozhukov, and Muhamet Yildiz, "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 2016, *11* (1), 187–225.

Adhvaryu, Achyuta, Prashant Bharadwaj, James Fenske, Anant Nyshadham, and Richard Stanley, "Dust and death: evidence from the West African Harmattan," 2019.

Agrawal, Arun, "The role of local institutions in adaptation to climate change," 2008.

Aker, Jenny C. and Kelsey Jack, "Harvesting the Rain: The Adoption of Environmental Technologies in the Sahel," *Review of Economics and Statistics*, 2023.

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

Arceo-Gomez, Hanna and Paulina Oliva, "Does the Effect of Pollution on Infant Mortality Differ Between Developing and Developed Countries? Evidence from Mexico City," *Economic Journal*, 2014, *126* (591), 257–280.

Asiedu, Edward, Dean Karlan, Monica Lambon-Quayefio, and Christopher Udry, "A Call for Strucutred Ethics Appendices in Social Sciences Papers," *Proceedings of the National Academy of Sciences*, 2021, *118* (29).

Attanasio, Orazio P, Costas Meghir, and Ana Santiago, "Education choices in Mexico: using a structural model and a randomized experiment to evaluate Progresa," *The Review of Economic Studies*, 2012, *79* (1), 37–66.

Auffhammer, Maximilian, "Quantifying Economic Damages From Climate Change," *Journal of Economic Perspectives*, 2018, *32* (4), 33–52.

__ and Tamma A. Carleton, "Regional Crop Diversity and Weather Shocks in India," *Asian Development Review*, 2018, *35* (2), 113–130.

Bakkensen, Laura A and Lint Barrage, "Going underwater? Flood risk belief heterogeneity and coastal home price dynamics," *The Review of Financial Studies*, 2022, *35* (8), 3666–3709.

__ and Robert O Mendelsohn, "Risk and adaptation: Evidence from global hurricane damages and fatalities," *Journal of the Association of Environmental and Resource Economists*, 2016, *3* (3), 555–587.

Balboni, Clare Alexandra, "In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities," *Working Paper*, 2019.

Balboni, Clare, Johannes Boehm, and Mazhar Waseem, "Firm Adaptation in Production Networks: Evidence from Extreme Weather Events in Pakistan," *Working Paper*, 2023.

Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis, "Does climate change affect real estate prices? Only if you believe in it," *The Review of Financial Studies*, 2020, *33* (3), 1256–1295.

Bardhan, Pranab K, "Labor-tying in a poor agrarian economy: A theoretical and empirical analysis," *The Quarterly Journal of Economics*, 1983, *98* (3), 501–514.

Barreca, Alan, Karen Clay, Olivier Dêschenes, Michael Greenstone, and Joseph S. Shapiro, "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," *Journal of Political Economy*, February 2016, *124* (1).

Barwick, Panle Jia, Dave Donaldson, Shanjun Li, Yatang Lin, and Deyu Rao, "Improved Transportation Networks Facilitate Adaptation to Pollution and Temperature Extremes," *Working Paper*, 2023.

__ , Shanjun Li, Deyu Rao, and Nahim Bin Zahur, "The healthcare cost of air pollution: Evidence from the world's largest payment network," 2018.

__ , __ , Liguo Lin, and Eric Zou, "From fog to smog: The value of pollution information," 2019.

Baylis, Patrick W and Judson Boomhower, "Mandated vs. voluntary adaptation to natural disasters: The case of us wildfires," Technical Report, National Bureau of Economic Research 2022.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak, "Can network theory-based targeting increase technology adoption?," *American Economic Review*, 2021, *111* (6), 1918–1943.

__ , Dean Karlan, and Bram Thuysbaert, "Saving for a (not so) rainy day: A randomized evaluation of savings groups in Mali," Technical Report, National Bureau of Economic Research 2014.

Bergman, Peter, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *American Economic Review*, Forthcoming.

Berkouwer, Susanna B. and Joshua T. Dean, "Private Actions in the Presence of Externalities: The Health Impacts of Reducing Air Pollution Peaks but not Ambient Exposure," Technical Report w31614, National Bureau of Economic Research 2023.

Bhandari, Humnath, Ujjayant Chakravorty, Muhammad Ashraful Habib, and Kyle Emerick, "Targeted Subsidies for Water Conservation in Smallholder Agriculture," 2022.

Bharadwaj, Prashant and Jamie T. Mullins, "Weather, Climate, and Migration in the United States," *Working Paper*, 2021.

Biardeau, Léopold T., Lucas W. Davis, Paul Gertler, and Catherine Wolfram, "Heat Exposure and Global Air Conditioning," *Nature Sustainability*, 2020, *3* (1), 25–28.

Bilal, Adrien and Esteban Rossi-Hansberg, "Anticipating Climate Change Across the United States," Technical Report, National Bureau of Economic Research 2023.

Blakeslee, David, Ram Fishman, and Veena Srinivasan, "Way Down in the Hole: Adaptation to Long-Term Water Loss in Rural India," *American Economic Review*, 2020, *110* (1), 200–224.

Bordalo, P., N. Gennaioli, and A. Shleifer, "Salience," *Annual Review of Economics*, 2022, *14*, 521–544.

Bordalo, Pedro, Giovanni Burro, Katherine B Coffman, Nicola Gennaioli, and Andrei Shleifer, "Imagining the future: memory, simulation and beliefs about COVID," Technical Report, National Bureau of Economic Research 2022.

__ , John J. Conlon, Nicola Gennaioli, Spencer Y. Kwon, and Andrei Shleifer, "Memory and Probability," *Quarterly Journal of Economics*, 2023, *138* (1).

Boustan, Leah Platt, Matthew E. Kahn, and Paul W. Rhode, "Moving to Higher Ground: Migration Response to Natural Disasters in the Early Twentieth Century," *American Economic Review*, 2012, *102* (3), 238–244.

Burgess, Robin and Dave Donaldson, "Can Openness Mitigate the Effects of Weather Shocks? Evidence from India's Famine Era," *American Economic Review*, 2010, *100* (2), 449–453.

__ , Matthew Hansen, Benjamin Olken, Peter Potapov, and Stefanie Sieber, "The Political Economy of Deforestation in the Tropics," *Quarterly Journal of Economics*, November 2012, *127* (4), 1707–1754.

__ , Olivier Dêschenes, Dave Donaldson, and Michael Greenstone, "Weather, Climate Change and Death in India," 2017. Working Paper.

Burke, M., M. Craxton, C.D. Kolstad, C. Onda, H. Allcott, E. Baker, L. Barrage, R. Carson, K. Gillingham, J. Graff-Zivin, and M. Greenstone, "Opportunities for Advances in Climate Change Economics," *Science*, 2016, *352* (6283), 292–293.

Burke, Marshall and Kyle Emerick, "Adaptation to Climate Change: Evidence from US Agriculture," *American Economic Journal: Economic Policy*, 2016, *8* (3).

__ , Solomon Hsiang, and Edward Miguel, "Global Non-Linear Effect of Temperature on Economic Production," *Nature*, 2015, *527*, 235–239.

Burlig, F., A. Jina, E. M. Kelley, G. Lane, and H. Sahai, "Long-Range Forecasts as Climate Adaptation: Experimental Evidence from Developing-Country Agriculture," 2024.

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, "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*," *The Quarterly Journal of Economics*, 2022, *137* (4), 2037–2105.

Casaburi, Lorenzo and Jack Willis, "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya," *American Economic Review*, 2018, *108* (12), 3778–3813.

Cattaneo, Cristina and Giovanni Peri, "The Migration Response to Increasing Temperatures," *Journal of Development Economics*, 2016, *122*, 127–146.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell, "The Effect of Pollution on Worker Productivity: Evidence from Call Center Workers in China," *American Economic Journal: Applied Economics*, 2019, *11* (1), 151–172.

Chatterjee, Shoumitro, Rohit Lamba, and Esha Zaveri, "The role of farm subsidies in changing India's water footprint," 2022.

Chen, Binkai, Wei Lin, and Ao Wang, "Misspecified Learning in Technology Adoption: Evidence from Fertilizer Use in China," *Working Paper*, 2023.

Chen, J. and V. Mueller, "Coastal climate change, soil salinity and human migration in Bangladesh," *Nature Clim Change*, 2018, *8*, 981–985.

__ , __ , F. Durand, and et al., "Salinization of the Bangladesh Delta worsens economic precarity," *Population and Environment*, 2022, *44*, 226–247.

Cicala, Steve, David Hémous, and Morten G Olsen, "Adverse selection as a policy instrument: unraveling climate change," Technical Report, National Bureau of Economic Research 2022.

Clarke, D., S. Williams, M. Jahiruddin, K. Parks, and M. Salehin, "Projections of on-farm salinity in coastal Bangladesh," *Environ. Sci.: Processes Impacts*, 2015, *17*, 1127–1136.

Conley, Timothy G and Christopher R Udry, "Learning about a new technology: Pineapple in Ghana," *American economic review*, 2010, *100* (1), 35–69.

Conte, Bruno, Klaus Desmet, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg, "Local sectoral specialization in a warming world," *Journal of Economic Geography*, 2021, *21* (4), 493–530.

Costinot, Arnaud, Dave Donaldson, and Cory Smith, "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world," *Journal of Political Economy*, 2016, *124* (1), 205–248.

Cummings, J. A. and O. M. Smedstad, "Variational Data Assimilation for the Global Ocean," *Data Assimilation for Atmospheric, Oceanic, and Hydrologic Applications*, 2013, *2*, 303–343.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker, "Environmental health risks and housing values: evidence from 1,600 toxic plant openings and closings," *American Economic Review*, 2015, *105* (2), 678–709.

Da, Zhi, Umit G Gurun, and Mitch Warachka, "Frog in the pan: Continuous information and momentum," *The review of financial studies*, 2014, *27* (7), 2171–2218.

Dar, M., A. de Janvry, K. Emerick, and et al., "Flood-tolerant rice reduces yield variability and raises expected yield, differentially benefitting socially disadvantaged groups," *Sci Rep*, 2013, *3*, 3315.

Dasgupta, Susmita, Farhana Akhter Kamal, Zahirul Huque Khan, Sharifuzzaman Choudhury, and Ainun Nishat, "River salinity and climate change: evidence from coastal Bangladesh," in "World Scientific Reference on Asia and the World Economy," World Scientific, 2015, pp. 205–242.

__ , Md Moqbul Hossain, Mainul Huq, and David Wheeler, "Climate change, salinization and high-yield rice production in coastal Bangladesh," *Agricultural and Resource Economics Review*, 2018, *47* (1), 66–89.

Davis, Lucas W, "The effect of health risk on housing values: Evidence from a cancer cluster," *American Economic Review*, 2004, *94* (5), 1693–1704.

Decehzleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva, "Fighting Climate Change: International Attitudes Toward Climate Policies," *Working Paper*, 2022.

Delavande, Adeline, "Expectations in development economics," in "Handbook of Economic Expectations," Elsevier, 2023, pp. 261–291.
__ , **Xavier Giné, and David McKenzie**, "Measuring subjective expectations in developing countries: A critical review and new evidence," *Journal of development economics*, 2011, *94* (2), 151–163.
Dell, Melissa, Benjamin Jones, and Benjamin Olken, "Temperature Shocks and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics*, 2012, *4* (3).
DeLuca, Stefanie, Susan Clampet-Lundquist, and Kathryn Edin, *Coming of Age in the Other America*, Russel Sage Foundation, 2016.
Deschênes, Olivier and Michael Greenstone, "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather," *American economic review*, 2007, *97* (1), 354–385.
__ **and** __ , "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US," *American Economic Journal: Applied Economics*, 2011, *3* (4), 152–185.
Desmet, Klaus and Esteban Rossi-Hansberg, "On the Spatial Economic Impact of Global Warming," *Journal of Urban Economics*, 2015, *88*, 16–37.
Dizon-Ross, Rebecca and Seema Jayachandran, "Improving Willingness-to-Pay Elicitation by Including a Benchmark Good," *American Economic Association Papers and Proceedings*, 2022, *112*.
Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India," *The Quarterly Journal of Economics*, 2013, *128* (4), 1499–1545.
__ , **Michael Kremer, and Jonathan Robinson**, "Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya," *American economic review*, 2011, *101* (6), 2350–2390.
__ , **Rema Hanna, and Stephen P Ryan**, "Incentives work: Getting teachers to come to school," *American economic review*, 2012, *102* (4), 1241–1278.
Elliott, Robert J., Eric Strobl, and Puyang Sun, "The local impact of typhoons on economic activity in China: A view from outer space," *Journal of Urban Economics*, 2015, *88*, 50–66.
Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H. Dar, "Technological Innovations, Downside Risk, and the Modernization of Agriculture," *American Economic Review*, 2016, *106* (6).
Enke, Benjamin, "What you see is all there is," *The Quarterly Journal of Economics*, 2020, *135* (3), 1363–1398.
__ , **Frederik Schwerter, and Florian Zimmermann**, "Associative Memory and Belief Formation," Technical Report w26664, National Bureau of Economic Research 2020.
Fabregas, Raissa, Michael Kremer, Matthew Lowes, Robert On, and Giulia Zane, "SMS-extension and farmer behavior: lessons from six RCTs in East Africa," *Online at: https://www. atai-research. org/wp-content/uploads/2020/05/textfarmers1. pdf*, 2019.
Fafchamps, Marcel, "Sequential labor decisions under uncertainty: An estimable household model of West-African farmers," *Econometrica: Journal of the Econometric Society*, 1993, pp. 1173–1197.
Fairweather, Daryl, Matthew E. Kahn, Robert D. Metcalfe, and Sebastian Sandoval-Olascoaga, "The Impact of Climate Risk Disclosure on Housing Search and Buying Dynamics: Evidence from a Nationwide Field Experiment with Redfin," 2023.
Fishback, Price V, Werner Troesken, Trevor Kollmann, Michael Haines, Paul W Rhode, and Melissa Thomasson, "Information and the impact of climate and weather on mortality rates during the Great Depression," in "The Economics of Climate Change: Adaptations Past and Present," University of Chicago Press, 2011, pp. 131–167.
Foster, Andrew D. and Mark R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 1995, *103* (6), 1176–1209.
Fudenberg, Drew, Giacomo Lanzani, and Philipp Strack, "Selective memory equilibrium," *Available at SSRN 4015313*, 2022.
Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein, "Channeled Attention and Stable Errors," *Working Paper*, 2021.
Gallagher, Justin, "Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States," *American Economic Journal: Applied Economics*, 2014, *6* (3), 206–233.
Gandhi, Sahil, Matthew E Kahn, Rajat Kochhar, Somik Lall, and Vaidehi Tandel, "Adapting to Flood Risk: Evidence from a Panel of Global Cities," 2022, (30137).
Garg, Teevrat, Maulik Jagnani, and Vis Taraz, "Temperature and human capital in India," *Journal of the Association of Environmental and Resource Economists*, 2020, *7* (6), 1113–1150.
Gentzkow, Mathew, Michael B. Wong, and Allen T. Zhang, "Ideological Bias and Trust in Information Sources," *Working Paper*, 2023.
Gifford, Robert, "The Dragons of Inaction: Psychological Barriers That Limit Climate Change Mitigation and Adaptation," *American Psychologist*, 2011, *66* (4), 290.
Gilbert, Daniel, "If only gay sex caused global warming," *Los Angeles Times*, 2006, *2*, 02–07.
Gine, Xavier, Robert M. Townsend, and James Vickery, "Forecasting When it Matters: Evidence from Semi-Arid India," *Working Paper*, 2015.
Giuliano, Paola and Nathan Nunn, "Understanding Cultural Persistence and Change," *The Review of Economic Studies*, 2021, *88* (4), 1541–1581.
Glennerster, Rachel and Seema Jayachandran, "Think Globally, Act Globally: Opportunities to Mitigate Greenhouse Gas Emissions in Low- and Middle-Income Countries," *Journal of Economic Perspectives*, 2023, *37* (3), 111–136.
Global Administrative Areas, "GADM database of Global Administrative Areas," *Version 3.4*, 2018.
Graeber, Thomas, Florian Zimmermann, and Christopher Roth, "Stories, Statistics, and Memory," 2022.
Grattan, Stephen R., Linghe Zeng, Michael C. Shannon, and Stacy R. Roberts, "Rice Is More Sensitive to Salinity than Previously Thought," *California Agriculture*, 2002, *56* (6), 189–198.

Greenstone, Michael and Justin Gallagher, "Does hazardous waste matter? Evidence from the housing market and the superfund program," *The Quarterly Journal of Economics*, 2008, *123* (3), 951–1003.
__ **and Rema Hanna**, "Environmental regulations, air and water pollution, and infant mortality in India," *American Economic Review*, 2014, *104* (10), 3038–3072.
__ **, Andrew Hultgren, Tamma Carleton, Michael Delgado, Diana R. Gergel, 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**, "Estimating Global Impacts to Agriculture from Climate Change Accounting for Adaptation," November 2022.
__ **, Rohini Pande, Anant Sudarshan, and Nicholas Ryan**, "Can pollution markets work in developing countries? Experimental evidence from India," 2023.
Grosset, F., A. Papp, and C. Taylor, "Rain Follows the Forest: Land Use Policy, Climate Change, and Adaptation," *Climate Change and Adaptation*, 2023.
Guimbeau, Amanda, James Ji Xinde, Zi Long, and Nidhiya Menon, "Ocean Salinity, Early-Life Health, and Adaptation," *Working Paper*, 2023.
Guiteras, Raymond, "The Impact of Climate Change on Indian Agriculture," *Working Paper*, 2009.
Hanna, Rema and Paulina Oliva, "The Effect of Pollution on Labor Supply: Evidence from a Natural Experiment in Mexico City," *Journal of Public Economics*, 2015, *122*, 68–79.
__ **, Sendhil Mullainathan, and Joshua Schwartzstein**, "Learning Through Noticing: Theory and Evidence from a Field Experiment," *Quarterly Journal of Economics*, 2014, *129* (3), 1311–1353.
Harstad, Bård, "The conservation multiplier," *Journal of Political Economy*, 2023, *131* (7), 1731–1771.
Heft-Neal, Sam, Carlos F. Gould, Matthew V. Kiang, Kari Nadeau, Mark Duggan, Eran Bendavid, and Marshall Burke, "Behavior Mediates the Health Effects of Extreme Wildfire Smoke Events," *Working Paper*, 2023.
Heidhues, Paul, Botond Kőszegi, and Philipp Strack, "Convergence in models of misspecified learning," *Theoretical Economics*, 2021, *16* (1), 73–99.
Hopmans, Jan W., A.S. Qureshi, I. Kisekka, R. Munns, S.R. Grattan, P. Rengasamy, A. Ben-Gal, S. Assouline, M. Javaux, P.S. Minhas, P.A.C. Raats, T.H. Skaggs, G. Wang, Q. De Jong van Lier, H. Jiao, R.S. Lavado, N. Lazarovitch, B. Li, and E. Taleisnik, "Chapter One - Critical knowledge gaps and research priorities in global soil salinity," in Donald L. Sparks, ed., *Donald L. Sparks, ed.*, Vol. 169 of *Advances in Agronomy*, Academic Press, 2021, pp. 1–191.
Hornbeck, R., "Dust Bowl Migrants: Environmental Refugees and Economic Adaptation," *The Journal of Economic History*, 2023, *83* (3), 645–675.
Hornbeck, Richard and Suresh Naidu, "When the Levee Breaks: Black Migration and Economic Development in the American South," *American Economic Review*, 2014, *104* (3), 963–990.
Hsiang, Solomon M. and Amir S. Jina, "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence from 6,700 Cyclones," *Working Paper*, 2014.
Hsiao, Allan, "Sea Level Rise and Urban Adaptation in Jakarta," *Working Paper*, 2023.
Ito, K. and S. Zhang, "Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China," *Journal of Political Economy*, 2020, *128* (5), 1627–1672.
Jack, B. Kelsey and et al., "Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning," Technical Report w30690, National Bureau of Economic Research 2022.
Jayachandran, Seema, "Selling labor low: Wage responses to productivity shocks in developing countries," *Journal of political Economy*, 2006, *114* (3), 538–575.
__ , "Air Quality and Early-Life Mortality: Evidence from Indonesia's Wildfires," *Journal of Human Resources*, 2009, *44* (4), 916–954.
Jedwab, Remi, Federico Haslop, Roman Zarate, and Carlos Rodríguez-Castelán, "The Effects of Climate Change in the Poorest Countries: Evidence from the Permanent Shrinking of Lake Chad," 2023.
Johnson, D. and S. Levin, "The Tragedy of Cognition: Psychological Biases and Environmental Inaction," *Current Science*, 2009, pp. 1593–1603.
Jones, Maria, Florence Kondylis, John Loeser, and Jeremy Magruder, "Factor Market Failures and the Adoption of Irrigation in Rwanda," *American Economic Review*, 2022, *112* (7).
Kahn, Matthew E., "The Death Toll from Natural Disasters: The Role of Income, Geography, and Institutions," *Review of Economics and Statistics*, 2005, *87* (2), 271–284.
Kahn, Matthew E, *Adapting to climate change*, Yale University Press, 2021.
__ **and Daxuan Zhao**, "The impact of climate change skepticism on adaptation in a market economy," *Research in Economics*, 2018, *72* (2), 251–262.
Kala, N., C. Balboni, and S. Bhogale, "Climate Adaptation," *VoxDevLit*, 2023, *7* (3).
Kala, Namrata, "Learning, Adaptation, and Climate Uncertainty: Evidence from Indian Agriculture," *Working Paper*, 2019.
Kapons, Martin and Peter Kelly, "Biased Inferece due to Prior Beliefs: Evidence from the Field," *Working Paper*, 2023.
Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry, "Agricultural decisions after relaxing credit and risk constraints," *The Quarterly Journal of Economics*, 2014, *129* (2), 597–652.
Kaur, Supreet, "Nominal wage rigidity in village labor markets," *American Economic Review*, 2019, *109* (10), 3585–3616.
Khanna, Gaurav, Wenquan Liang, Ahmed Mushfiq Mobarak, and Ran Song, "The productivity consequences of pollution-induced migration in China," *Working Paper*, 2021.
Kochhar, Rajat and Ruozi Song, "Does Market Power in Local Agricultural Markets Hinder Farmer Climate Change Adaptation?," *Working Paper*, 2023.
Kocornik-Mina, Adriana, Thomas K. J. McDermott, Guy Michaels, and Ferdinand Rauch, "Flooded Cities," *American Economic Journal: Applied Economics*, April 2020, *12* (2), 35–66.

Kremer, Michael, Jessica Leino, Edward Miguel, and Alix Peterson Zwane, "Spring cleaning: Rural water impacts, valuation, and property rights institutions," *The Quarterly Journal of Economics*, 2011, *126* (1), 145–205.
Kruttli, Mathias S, Brigitte Roth Tran, and Sumudu W Watugala, "Pricing Poseidon: Extreme weather uncertainty and firm return dynamics," 2023.
Lane, Gregory, "Adapting to Floods with Guaranteed Credit: Evidence from Bangladesh," *Working Paper*, 2022.
Leeffers, Stefan, "It Will Rain: The Effect of Information on Flood Preparedness in Urban Mozambique," 2023.
Macours, Karen, Patrick Premand, and Renos Vakis, "Transfers, Diversification and Household Risk Strategies: Can productive safety nets help households manage climatic variability?," *The Economic Journal*, 2022, *132* (647), 2438–2470.
Mahadevan, Meera, Robyn Meeks, and Takashi Yamano, "Reducing information barriers to solar adoption: Experimental evidence from India," *Energy Economics*, 2023, *120*, 106600.
Malmendier, Ulrike, "FBBVA Lecture 2020 Exposure, Experience, and Expertise: Why Personal Histories Matter in Economics," *Journal of the European Economic Association*, 2021, *19* (6), 2857–2894.
__ **and Stefan Nagel**, "Learning from Inflation Experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.
Mattoo, Aaditya and Arvind Subramanian, *Greenprint: a new approach to cooperation on climate change*, Brookings Institution Press, 2013.
McFadden, Daniel, "The measurement of urban travel demand," *Journal of public economics*, 1974, *3* (4), 303–328.
McGuirk, Eoin F and Nathan Nunn, "Transhumant pastoralism, climate change, and conflict in africa," Technical Report, National Bureau of Economic Research 2020.
Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw, "The impact of global warming on agriculture: a Ricardian analysis," *The American economic review*, 1994, pp. 753–771.
Mobarak, Ahmed Mushfiq and Mark R Rosenzweig, "Informal risk sharing, index insurance, and risk taking in developing countries," *American Economic Review*, 2013, *103* (3), 375–380.
Molina, Renato and Ivan Rudik, "The social value of predicting hurricanes," 2022.
Moscona, Jacob, "Environmental Catastrophe and the Direction of Invention: Evidence from the American Dust Bowl," *Working Paper*, 2022.
__ **and Karthik Sastry**, "Does Directed Innovation Mitigate Climate Damage? Evidence from US Agriculture," *Quarterly Journal of Economics*, 2023.
Mueller, Valerie, Clark Gray, and Katrina Kosec, "Heat Stress Increases Long-Term Human Migration in Rural Pakistan," *Nature Climate Change*, 2014, *4*, 182–185.
Mukhopadhyay, Raj, Binoy Sarkar, Hanuman Sahay Jat, Parbodh Chander Sharma, and Nanthi S. Bolan, "Soil Salinity Under Climate Change: Challenges for Sustainable Agriculture and Food Security," *Journal of Environmental Management*, 2021, *280*.
Mullainathan, Sendhil, "A memory-based model of bounded rationality," *The Quarterly Journal of Economics*, 2002, *117* (3), 735–774.
Mullins, J., J. G. Zivin, A. Cattaneo, A. Paolantonio, and R. Cavatassi, "The adoption of climate smart agriculture: the role of information and insurance under climate change," in "Climate Smart Agriculture: Building Resilience to Climate Change" 2018, pp. 353–383.
Nath, Ishan, "Climate Change, The Food Problem, and the Challenge of Adaptation through Sectoral Reallocation," *Working Paper*, 2022.
Nordhaus, William D, "To slow or not to slow: the economics of the greenhouse effect," *The economic journal*, 1991, *101* (407), 920–937.
__ , *Managing the global commons: the economics of climate change*, Vol. 31, MIT press Cambridge, MA, 1994.
Ostriker, Abigail and Anna Russo, "The effects of floodplain regulation on housing markets," Technical Report 2022.
Pande, R., M. Greenstone, J. Nilekani, A. Sudarshan, A. Suganathan, and N. Ryan, "Lower Pollution, Longer Lives: Life Expectancy Gains if India Reduced Particulate Matter Pollution," *Economic and Political Weekly*, 2015, *L* (8).
Patel, Dev, "Floods," *Working Paper*, 2023.
Pelli, Martino, Jeanne Tschopp, Natalia Bezmaternykh, and Kodjovi M. Eklou, "In the eye of the storm: Firms and capital destruction in India," *Journal of Urban Economics*, 2023.
Peri, Giovanni and Frédéric Robert-Nicoud, "On the Economic Geography of Climate Change," *Journal of Economic Geography*, 2021, *21* (4), 487–491.
Pople, A., R. Hill, S. Dercon, and B. Brunckhorst, "Anticipatory Cash Transfers in Climate Disaster Response," 2021.
Rangel, Marcos A and Tom S Vogl, "Agricultural fires and health at birth," *Review of Economics and Statistics*, 2019, *101* (4), 616–630.
Rosenzweig, Mark and Christopher R Udry, "Forecasting profitability," Technical Report, National Bureau of Economic Research 2013.
Rosenzweig, Mark R and Christopher R Udry, "Assessing the benefits of long-run weather forecasting for the rural poor: Farmer investments and worker migration in a dynamic equilibrium model," Technical Report, National Bureau of Economic Research 2019.
__ **and Christopher Udry**, "Rainfall forecasts, weather, and wages over the agricultural production cycle," *American Economic Review*, 2014, *104* (5), 278–283.
__ **and Hans P Binswanger**, "Wealth, weather risk and the composition and profitability of agricultural investments," *The economic journal*, 1993, *103* (416), 56–78.
Rudder, Jess and Davide Viviano, "Learning from Weather Forecasts and Adaptation among Cotton Farmers in Pakistan," *Working Paper*, 2023.
Ryan, Nicholas and Anant Sudarshan, "Rationing the Commons," *Journal of Political Economy*, 2022, *130* (1), 210–257.
Schwartzstein, Joshua, "Selective attention and learning," *Journal of the European Economic Association*, 2014, *12* (6), 1423–1452.

Shah, Manisha and Bryce Millett Steinberg, "Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital," *Journal of Political Economy*, 2017, *125* (2), 527–561.
Shapiro, Joseph S, "The environmental bias of trade policy," *The Quarterly Journal of Economics*, 2021, *136* (2), 831–886.
Shrader, Jeffrey G, Laura Bakkensen, and Derek Lemoine, "Fatal Errors: The Mortality Value of Accurate Weather Forecasts," Technical Report, National Bureau of Economic Research 2023.
Singh, Ajay, "Soil Salinity: A Global Threat to Sustainable Development," *Soil Use and Management*, 2022, *38*, 39–67.
Somanathan, E., Rohini Somanathan, Anant Sudarshan, and Meenu Tewari, "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing," *Journal of Political Economy*, June 2021, *129* (6).
Suri, Tavneet, "Selection and comparative advantage in technology adoption," *Econometrica*, 2011, *79* (1), 159–209.
Taraz, Vis, "Adaptation to Climate Change: Historical Evidence from the Indian Monsoon," *Environment and Development Economics*, 2017, *22* (5), 517–545.
Tazhibayeva, Kamilya and Robert Townsend, "The Impact of Climate Change on Rice Yields: Heterogeneity and Uncertainty," in "IIES conference on Climate Change in" Citeseer 2012.
Todd, Petra E and Kenneth I Wolpin, "Assessing the impact of a school subsidy program in Mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility," *American economic review*, 2006, *96* (5), 1384–1417.
__ and __ , "The best of both worlds: combining randomized controlled trials with structural modeling," *Journal of Economic Literature*, 2023, *61* (1), 41–85.
Train, Kenneth E, *Discrete choice methods with simulation*, Cambridge university press, 2009.
Tversky, Amos, "Features of Similarity," *Psychological Review*, 1977, *84* (4), 327–352.
__ and Daniel Kahneman, "Availability: A heuristic for judging frequency and probability," *Cognitive psychology*, 1973, *5* (2), 207–232.
Wagner, Katherine RH, "Adaptation and adverse selection in markets for natural disaster insurance," *American Economic Journal: Economic Policy*, 2022, *14* (3), 380–421.
Weinberger, Kate R, Antonella Zanobetti, Joel Schwartz, and Gregory A Wellenius, "Effectiveness of National Weather Service heat alerts in preventing mortality in 20 US cities," *Environment international*, 2018, *116*, 30–38.
Zappalá, Guglielmo, "Adapting to Climate Change Accounting for Individual Beliefs," *Working Paper*, 2023.

8 Figures and Tables

Figure 1: Soil Salinity Beliefs vs. Agronomic Truth

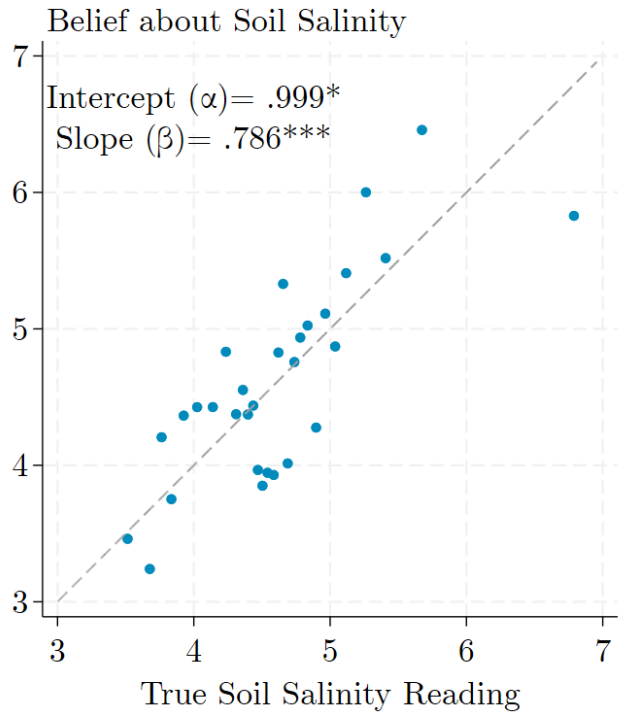
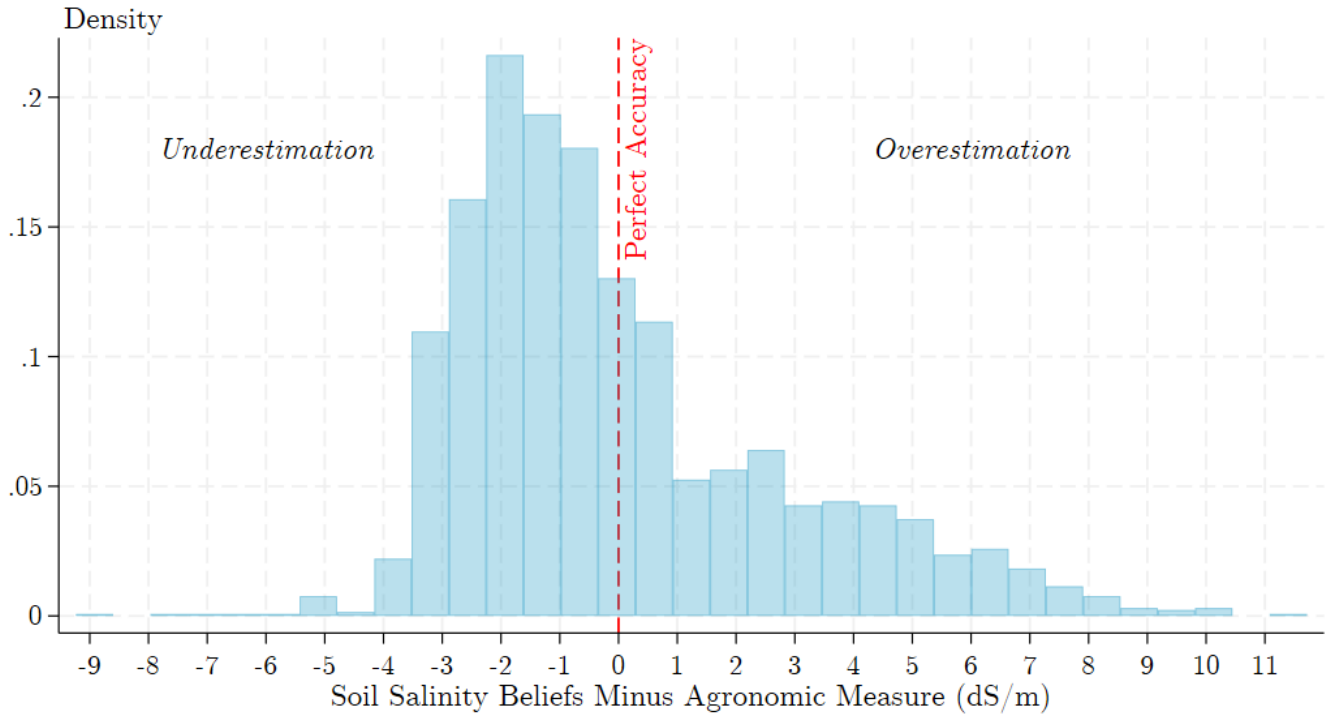


Figure 2: Histogram of Belief Errors



Note: Figure 1 plots a binned scatter plot of farmers' beliefs about the 2022-23 season's average soil salinity against the ground truth based on direct agronomic readings of their soil. The dashed gray line denotes the 45-degree line. Figure 2 plots a histogram of the errors in beliefs. The dashed vertical red line denotes perfect accuracy. Mass to the right of that line captures overestimation, and mass to the left captures underestimation.

Table 1: Soil Salinity Beliefs vs. Agronomic Truth

	(1) Salinity Belief	(2) Salinity Belief
Salinity Truth	0.786***	0.805***
	(0.121)	(0.111)
Constant	0.999*	0.913*
	(0.547)	(0.504)
Include Comp. Check Failures	No	Yes
Outcome Mean	4.617	4.633
Observations	2,068	2,271
R^2	0.034	0.038
p-value: β True Salinity = 1	0.078	0.080
p-value: β Constant = 0	0.068	0.070

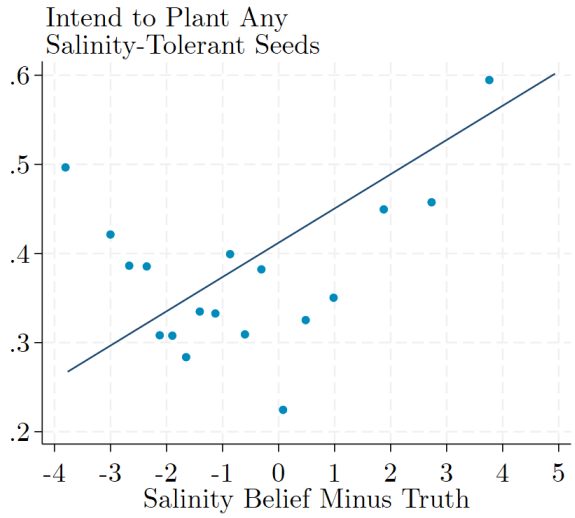
Note: Table 1 presents bi-variate regressions of farmers' beliefs about salinity on the true soil salinity measurements. The agronomic truth is based off of the best prediction of seasonal exposure given the fall 2022 measurements. The outcome in each regression is farmers' beliefs, which captures farmers' expectations about soil salinity over the course of the 2022-23 season, measured during the fall. Column (1) includes the pre-specified cut of excluding those farmers who fail the comprehension checks during the belief elicitation practice. Column (2) includes the full sample. All specifications report heterogeneity-robust standard errors.

Table 2: Planting Behavior and Salinity Belief Errors

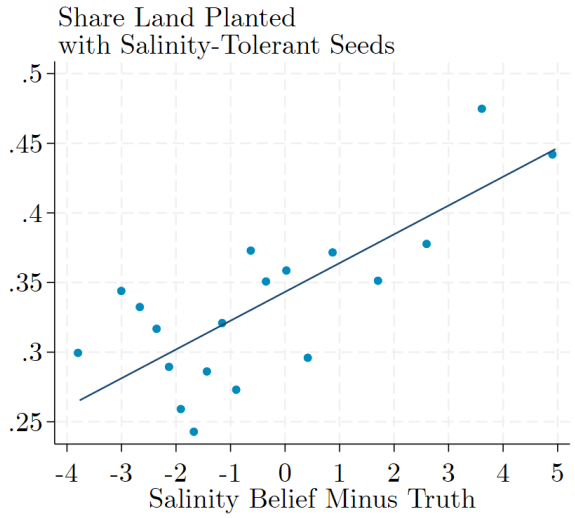
	(1) Intend To Plant Salinity-Tolerant Seeds	(2) Actual Share Land Planted With Salinity-Tolerant Seeds (Perceived Status)	(3) Actual Share Land Planted With Salinity-Tolerant Seeds (Actual Status)
Salinity Belief - Truth	0.0384***	0.0207***	0.0181***
	(0.00355)	(0.00365)	(0.00294)
True Salinity	0.180***	0.123***	0.135***
	(0.0164)	(0.0167)	(0.0146)
Observations	2,068	2,008	2,008
Outcome Mean	0.412	0.168	0.168

Figure 3: Planting Behavior and Salinity Belief Errors

(a) Salinity-Tolerant Seed Intention

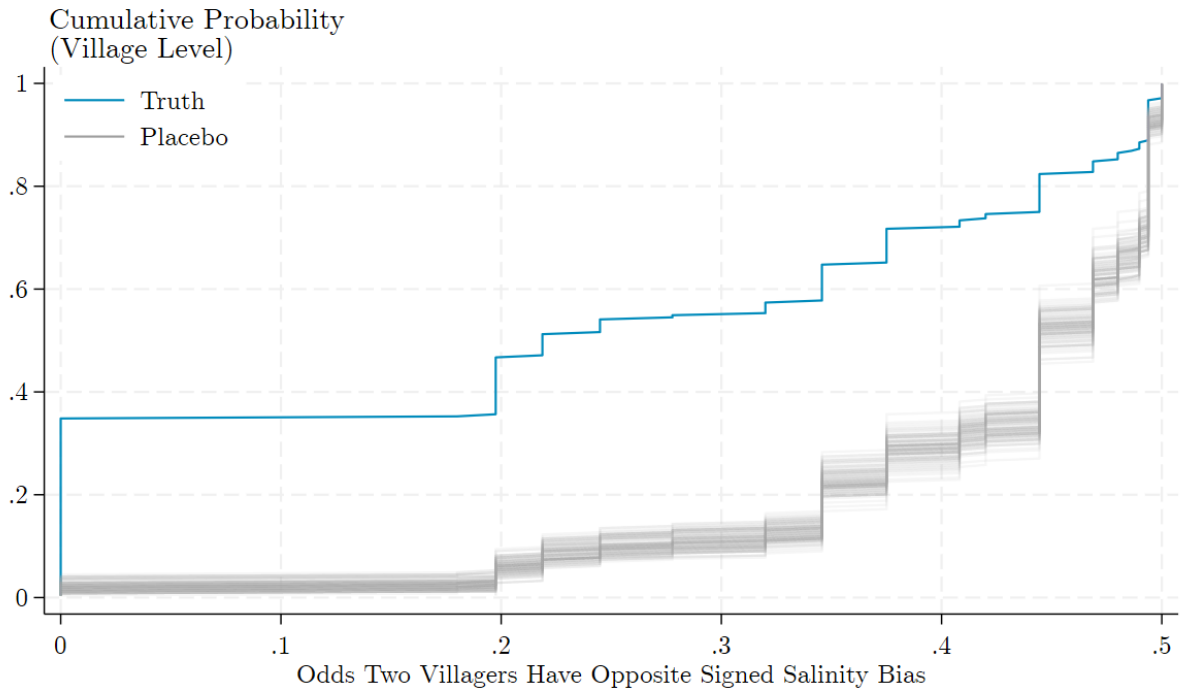


(b) Share Land with Salinity-Tolerant Seeds



Note: Table 2 presents regressions of the farmers' planting behavior on the error in their salinity beliefs (defined as beliefs minus the truth) and the true agronomic salinity on their plots. Column (1) uses as an outcome whether farmers report intending to plant any salinity-tolerant variety on their plot. Column (2) reports the total share of land across all plots on which farmers planted a seed that they consider to be salinity-tolerant. Column (3) shows the same, this time using my assessment of whether the seed is salinity-tolerant based on official agronomic records as opposed to the farmers' own reports. All regressions exclude those who fail the beliefs comprehension checks during the practice elicitation and report heteroskedasticity-robust standard errors. Figures 3a and 3b visualize the relationships in Columns (1) and (2) of Table 2 using binned scatter plots.

Figure 4: The Spatial Concentration of Salinity Belief Errors



Note: Figure 4 plots cumulative distribution functions at the village level of the probability that two randomly drawn people from that village have opposite signed salinity bias, calculated as 1 minus the Hirschman-Herfindahl Index where the group shares are the proportion of people over- and underestimating the salinity on their plot. The blue line shows the true data. The gray lines plot the cumulative distribution functions for 100 placebo distributions after randomly re-assigning villages.

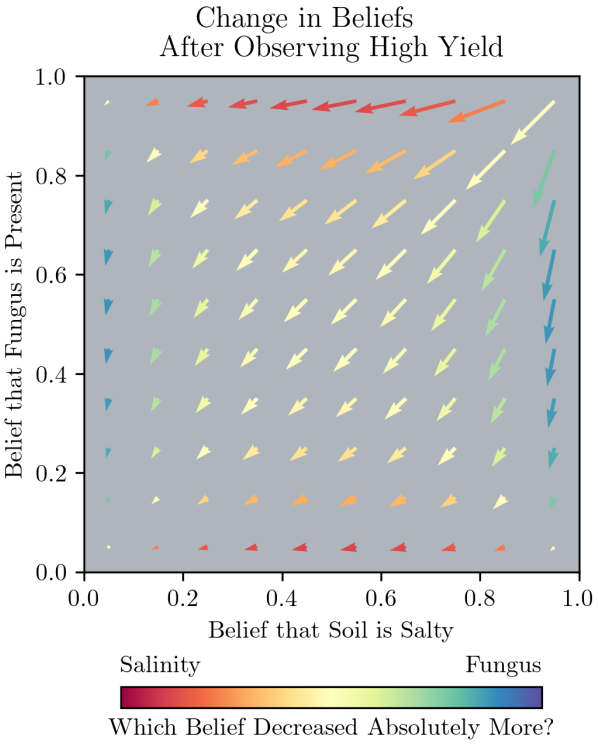
Table 3: Predictors of Soil Salinity Belief Inaccuracy

	|Salinity Beliefs - Truth| (Standardized)			
	(1)	(2)	(3)	(4)
Log Plot Area	-0.119***			
	(0.0146)			
Age (10 years)		-0.0881***		
		(0.0126)		
Peers View Extremely Knowledgeable			-0.251***	
			(0.0914)	
Peers Follow Seed Advice				-0.674***
				(0.0706)
Observations	2,008	2,068	2,035	2,035
Outcome Mean	-0.020	-0.002	0.001	0.001

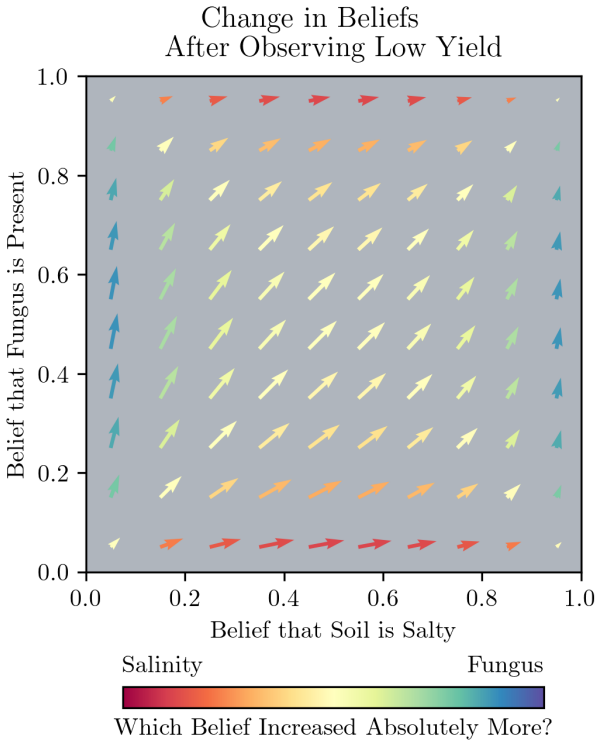
Note: Table 3 presents regressions prediction the absolute value of salinity beliefs minus the agronomic truth, normalized to have mean equal to 0 and standard deviation equal to 1. Column (1) regresses this outcome on the log of total harvested land as measured in the endline survey. Column (2) regresses this outcome on the farmer's age, scaled by 10 years for readability. Column (3) regresses this outcome on the share of other farmers in the village who say that the respondent is extremely knowledgeable about farming. Column (4) regresses this outcome on the share of other farmers in the village who say that they would consider following the respondent's advice about what seed to plant. All specifications include the true salinity level interacted with fixed effects for each ventile. All regressions exclude those who fail the beliefs comprehension checks during the practice elicitation and report heteroskedasticity-robust standard errors.

Figure 5: Simulated Belief Movement in Conceptual Framework

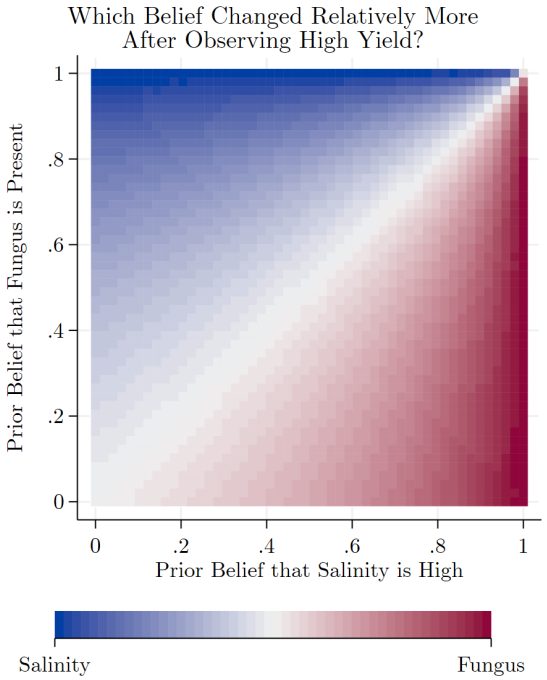
(a) Absolute Changes in Beliefs: High Yield



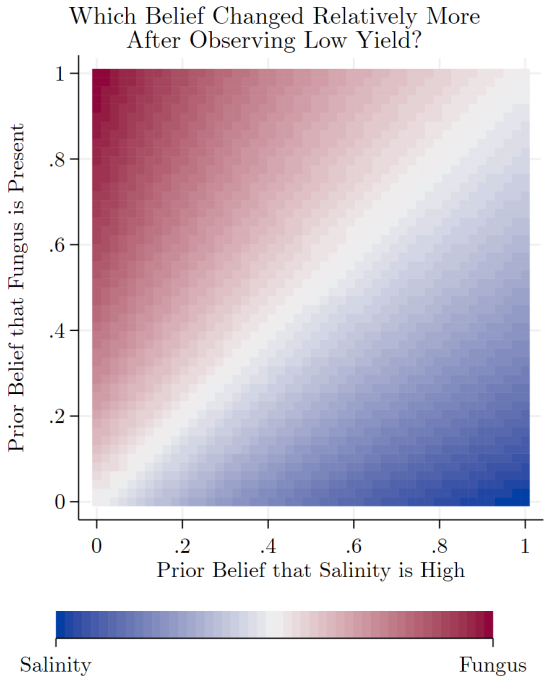
(b) Absolute Changes in Beliefs: Low Yield



(c) Relative Changes in Beliefs: High Yield



(d) Relative Changes in Beliefs: Low Yield



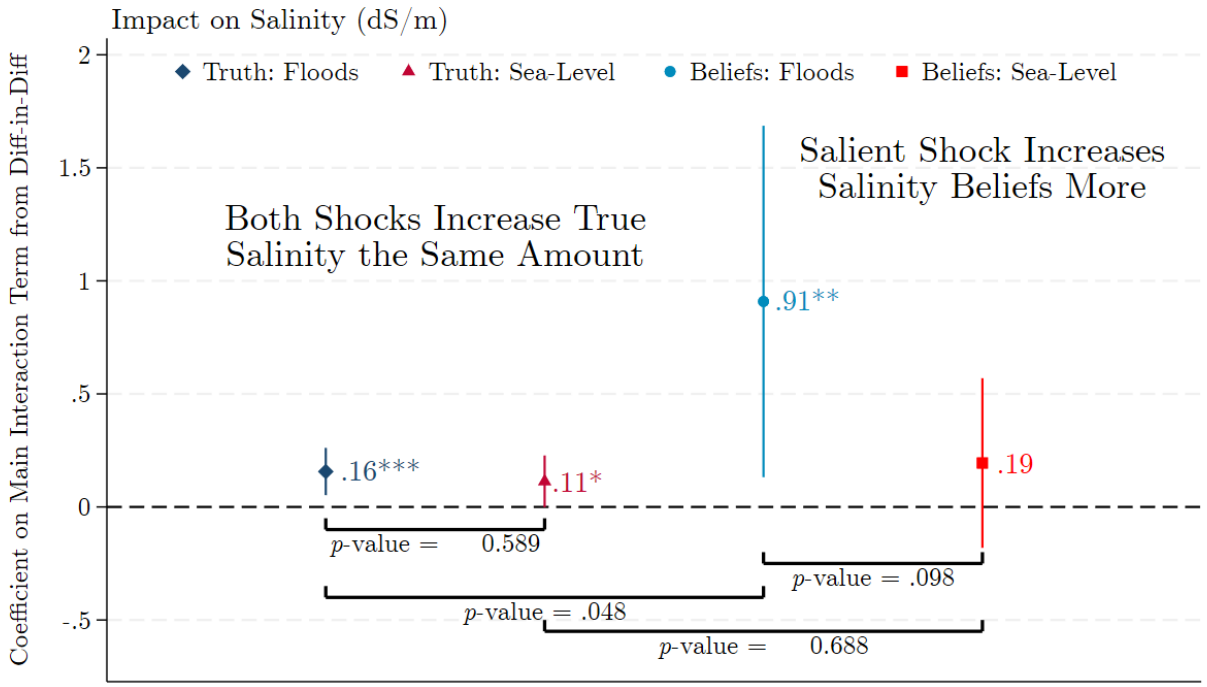
Note: Figure 5 shows changes in beliefs for simulated data based on the model. In all simulations, I assume $\rho = .33$, that is their is an equal chance of a high, low, and neutral productivity shock. The left column corresponds to observing high yield in the first period, and the right column corresponds to observing low yield. Figures 5a and 5b illustrate how beliefs change after observing the new data, where the arrows begin at the prior. Figures 5c and 5d scale this change to account for the mechanical limitations of belief updating given the bounds, shading based on the *relative* change in beliefs at each prior.

Table 4: Impact of Climate Shocks on True Salinity and Salinity Beliefs

	True Salinity		Salinity Beliefs	
	(1)	(2)	(3)	(4)
Flood × Saltier	0.157***		0.909**	
	(0.0531)		(0.395)	
Closer to Ocean × Sea Level Rise × Ocean Salinity		0.113*		0.194
		(0.0581)		(0.190)
Observations	2075	2075	2068	2068
Clusters	250	250	250	250
Control Mean	4.605	4.605	4.617	4.617

Note: Table 4 displays only the main interaction term of the difference-in-differences. Standard errors are clustered at the village level. The outcome in columns (1) and (2) is the true soil salinity as measured by agronomic sensors. The outcome in columns (3) and (4) is farmers' mean belief about soil salinity. Independent variables Saltier, Closer to Ocean, Sea Level Rise, and Ocean Salinity are all measured in standard deviations. Sample restricted to farmers who passed belief comprehension checks.

Figure 6: Impact of Climate Shocks on True Salinity and Salinity Beliefs



Note: Figure 6 plots the coefficients from Table 4 and reports $p-$values testing for equality between coefficients.

Table 5: Impact of Climate Shocks on Interpretation of Environmental Signals

	Use of Generic Signals to Learn About Salinity		Attribution of Salinity Symptoms To Other Causes	
	(1)	(2)	(3)	(4)
Closer to Ocean × Sea Level Rise × Ocean Salinity	-0.0474		0.457**	
	(0.0646)		(0.187)	
Flood × Saltier		0.187***		-0.0413
		(0.0529)		(0.198)
Observations	2056	2056	2056	2056
Clusters	250	250	250	250
Control Mean	1.074	1.074	2.451	2.451

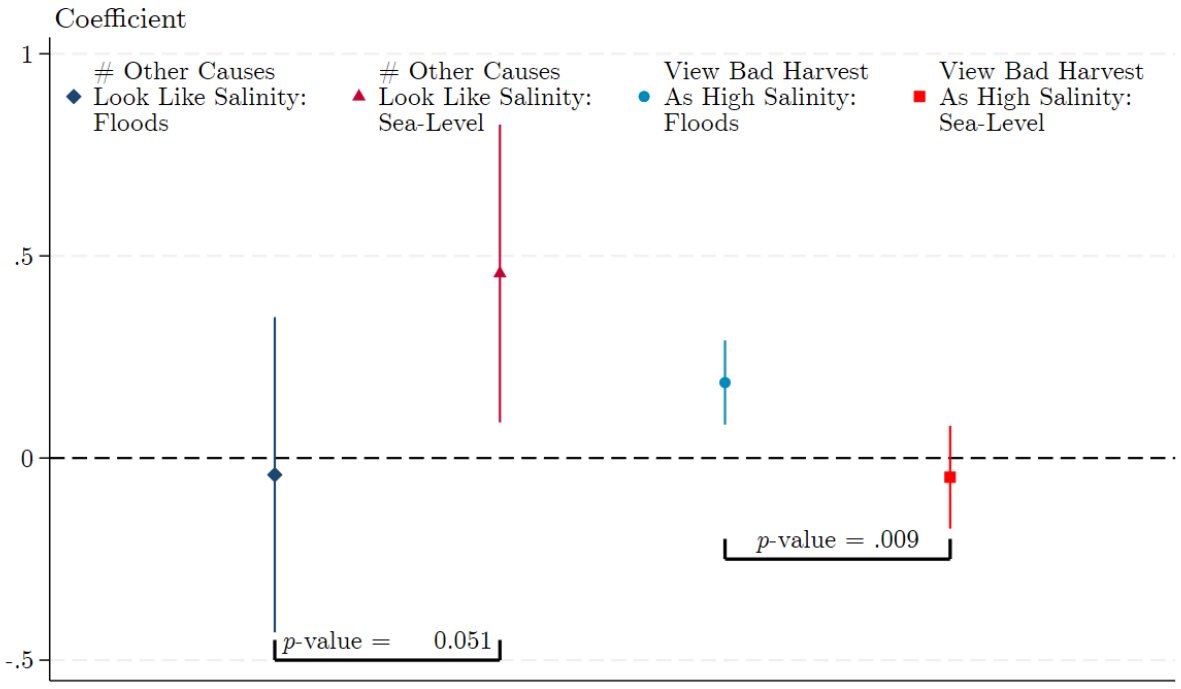
Note: Table 5 displays only the main interaction term of the difference-in-differences. Standard errors are clustered at the village level. The outcome in columns (1) and (2) is the sum of indicators for whether the farmer reports using the height of plants or plant death as an indicator from which they can learn about soil salinity. The outcome in columns (3) and (4) is the total number of other potential causes that farmers report saying could exhibit the same symptoms as high soil salinity. Independent variables Saltier, Closer to Ocean, Sea Level Rise, and Ocean Salinity are all measured in standard deviations. Sample restricted to farmers who passed belief comprehension checks.

Table 6: Impact of Salinity-Tolerant Seeds

	(1) Tolerant Seed Share (FS)	(2) Profits (OLS)	(3) Profits (RF)	(4) Profits (IV)
Offered Free Seeds	0.0898***		-0.0918**	
	(0.0143)		(0.0428)	
Salinity Tolerant Seed Share		-0.388***		-1.022**
		(0.0605)		(0.483)
Observations	2015	2015	2015	2015
Control Mean	0.124	0.021	0.021	0.021
First Stage F-Stat				39.938

Note: Table 6 presents results on the salinity-tolerant seed experiment. All specifications exclude those farmers who fail the baseline belief elicitation checks and report heteroskedasticity robust standard errors. Agricultural profits in Columns (2) through (4) have been standardized.

Figure 7: Impact of Climate Shocks on Interpretation of Environmental Signals



Note: Figure 7 plots the coefficients from Table 5 and reports $p-$values testing for equality between coefficients.

Table 7: Random Coefficients Demand Estimates

	(1)	(2)
Seed		
Seed Price (100s BDT)	0.962	0.963
	(0.0391)	(0.0395)
Salinity Tolerant	1.040	1.026
	(0.145)	(0.147)
Salinity Tolerant × Belief Error (dS/m)	1.112**	1.115**
	(0.0385)	(0.0418)
Selling Price (100s BDT)	1.030	1.030
	(0.0191)	(0.0195)
Growth Duration	0.999	0.999
	(0.00477)	(0.00480)
Processing Ease	0.686	0.683
	(0.539)	(0.541)
Market Demand	2.611***	2.645***
	(0.365)	(0.381)
Good Taste	0.643***	0.642***
	(0.0771)	(0.0774)
Nice Color	0.399*	0.387*
	(0.179)	(0.180)
Pest Tolerant	1.192	1.189
	(0.160)	(0.162)
/Normal		
sd(seed_saltolbelief)		1.201
		(0.221)
Observations	4469	4469
Farmers	724	724

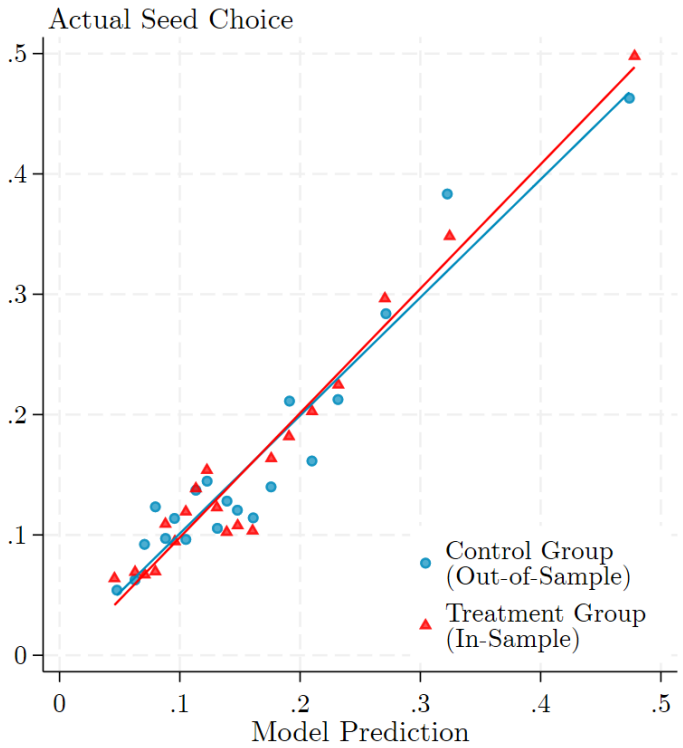
Note: Table 7 presents the estimates from the random coefficients demand estimation of seed choice. The sample only includes farmers in the treatment group from the seed experiment and excludes all those who fail the baseline belief elicitation comprehension checks. I report odds-ratios, but the stars denote statistical significance of the original coefficients (unexponentiated). I allow random coefficients for the interaction between salinity tolerant seeds and belief errors, where belief errors are defined as the difference between farmers' predictions of soil salinity on their plots for 2022-23 and the best agronomic prediction for the same period based on the fall EC-meter readings.

Table 8: Demand Model Out-of-Sample Performance

	(1) Control (Out-of-Sample)	(2) Treatment (In-Sample)
Model Prediction	0.981***	1.034***
	(0.0359)	(0.0364)
Constant	0.00310	-0.00547
	(0.00579)	(0.00592)
Observations	4678	4455
Farmers	754	725

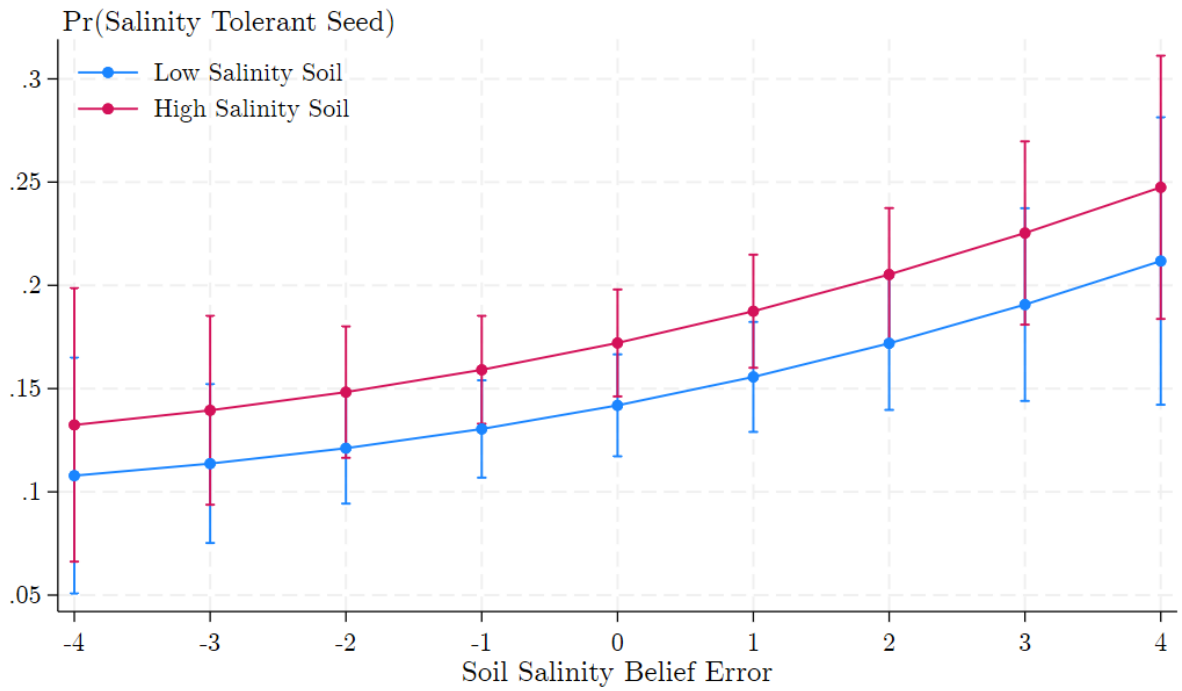
Note: Table 8 provides out-of-sample tests for the performance of the demand model. The outcome is a binary indicator for whether that seed is the one that farmer's planted the largest share of their land in the 2022-23 season. The predictor is the model's predicted probability of that seed being chosen.

Figure 8: Demand Model Out-of-Sample Performance



Note: Figure 8 plots a binned scatter plot comparing actual seed choice to the predicted probability of selecting a seed from the structural demand model, separately by the seed experiment treatment group (on which the model was estimated) and the control group. Those who failed the baseline belief comprehension checks are excluded.

Figure 9: Predicted Choice Probability Under Alternative Salinity Beliefs



Note: Figure 9 uses the marginal predicted choice probability from the structure of the model to evaluate the predicted choice probabilities for the primary salinity tolerant seed used in my sample under alternative belief errors. Here, belief error is defined as the gap between farmers' prediction for soil salinity over the 2022-23 season and the best agronomic prediction based on the fall reading.

Table 9: Soil Salinity Information Experiment: First-Stage

	(1)	(2)	(3)
Prior	1.361***	1.403***	1.377***
	(0.0829)	(0.0863)	(0.0894)
Own Truth - Prior	1.108***	1.153***	1.134***
	(0.0881)	(0.0924)	(0.0942)
Own Treatment	0.268	0.229	0.329
	(0.190)	(0.203)	(0.211)
Own Treatment x (Own Truth - Prior)	-0.832***	-0.865***	-0.842***
	(0.208)	(0.218)	(0.228)
Own Treatment x (Upa Truth - Prior)	0.980***	0.997***	0.992***
	(0.210)	(0.220)	(0.230)
Upa Treatment	0.244	0.180	0.305
	(0.283)	(0.296)	(0.313)
Upa Treatment x (Upa Truth - Prior)	0.243***	0.255***	0.276***
	(0.0813)	(0.0845)	(0.0872)
Constant	1.162***	1.069***	1.119***
	(0.228)	(0.238)	(0.247)
Sample	Full	Pass EL	Pass BL+EL
Observations	2245	2035	1864
$(\beta 3 + \beta 4) = \beta 5 : p - val$	0.262	0.165	0.162
R^2	0.138	0.148	0.140
F-statistic	44.390	43.905	38.274

Note: Table 9 presents the first-stage regressions of the impact of receiving information about soil salinity over the 2022-23 season on farmers' predictions about soil salinity over the 2023-24 season. Column (1) includes all respondents; column (2) includes only those farmers who passed the endline belief elicitation comprehension checks; and column (3) includes only those who passed both the baseline and the endline comprehension checks. All specifications report heteroskedasticity robust standard errors.

Table 10: Soil Salinity Information 2SLS: Demand for Salinity-Tolerant Seeds

	(1)	(2)	(3)
Salinity Belief 2023-24	3.043**	3.012**	3.136*
	(1.487)	(1.513)	(1.602)
Observations	2245	2035	1864
First Stage F-statistic	50.62	49.91	43.07

Note: Table 10 regresses demand for salinity-tolerant seeds as measured via the BDM against farmers beliefs about soil salinity in the upcoming season, instrumenting for beliefs using the experimental variation from the information treatment shown in Table 9. Regressions control for their willingness-to-pay for salinity-tolerant seeds at baseline and treatment status in the seed experiment. Column (1) includes all respondents; column (2) includes only those farmers who passed the endline belief elicitation comprehension checks; and column (3) includes only those who passed both the baseline and the endline comprehension checks. All specifications report heteroskedasticity robust standard errors.

Appendices

Table of Contents

A Additional Figures and Tables **53**

B Environmental Data **67**
B.1 Estimating Flood Risk 67
B.2 Soil Salinity 68
B.3 Water Salinity 70
B.4 Sea-Level Rise 73

C Farmer Perceptions Survey **76**
C.1 Sampling 76
C.2 Variable Construction 77

D Additional Data Sources **82**
D.1 New Seed Database 82

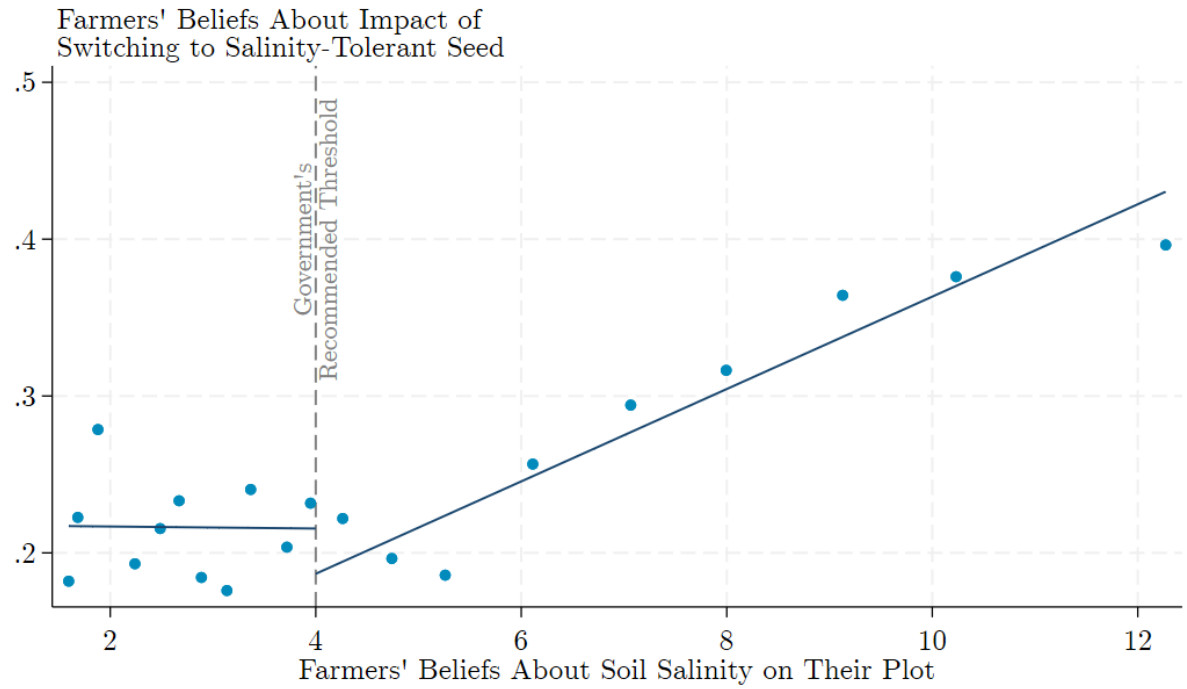
E Conceptual Framework Appendix **83**
E.1 Proofs of Theoretical Results 83

F Qualitative Evidence **86**
F.1 Methodology 86
F.2 Qualitative Evidence Motivating Modeling Assumptions 88

G Structured Ethics Appendix **90**

A Additional Figures and Tables

Figure A.1: Farmers' Beliefs About the Returns to Switching to Salinity-Tolerant Seeds



Note: Figure A.1 plots how farmers beliefs about the returns to planting salinity-tolerant seeds as measured as a share of their total expected yield relates to their belief about the salt content of their soil. The vertical dashed gray line denotes the threshold above which the government recommends adopting salinity-tolerant varieties.

Figure A.2: Distribution of Soil Salinity Beliefs

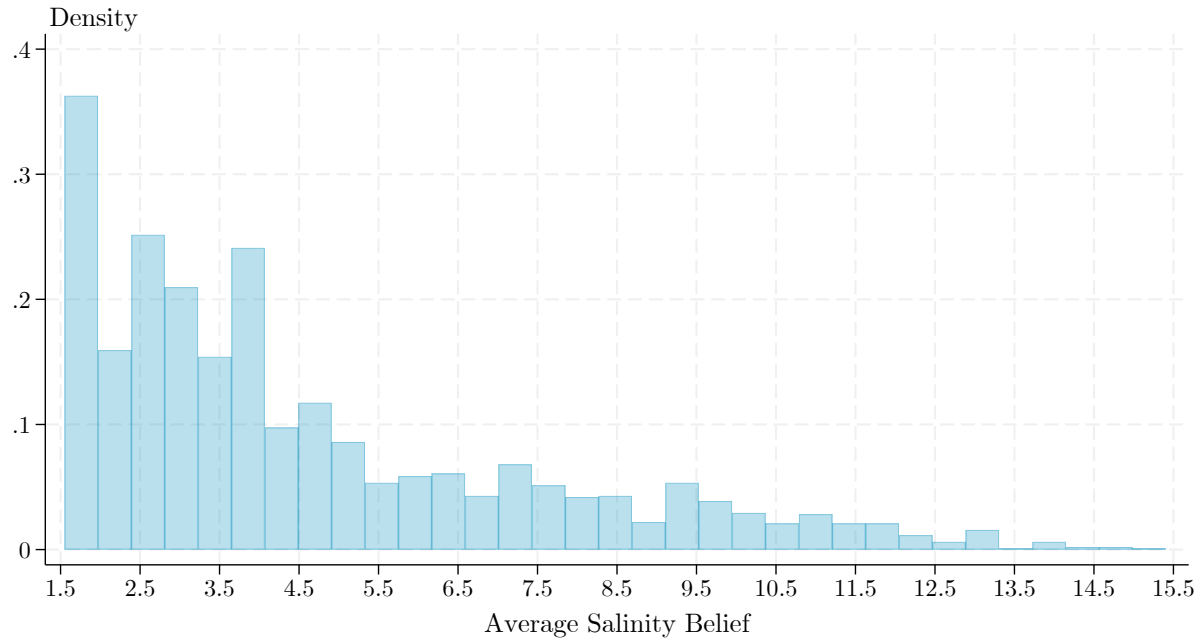
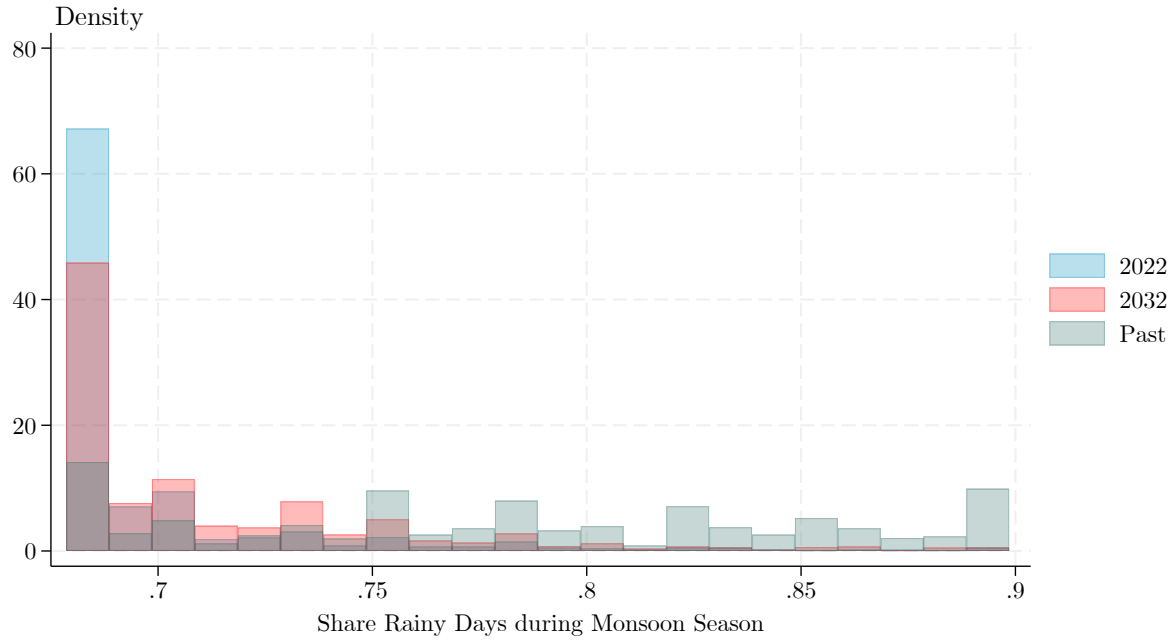
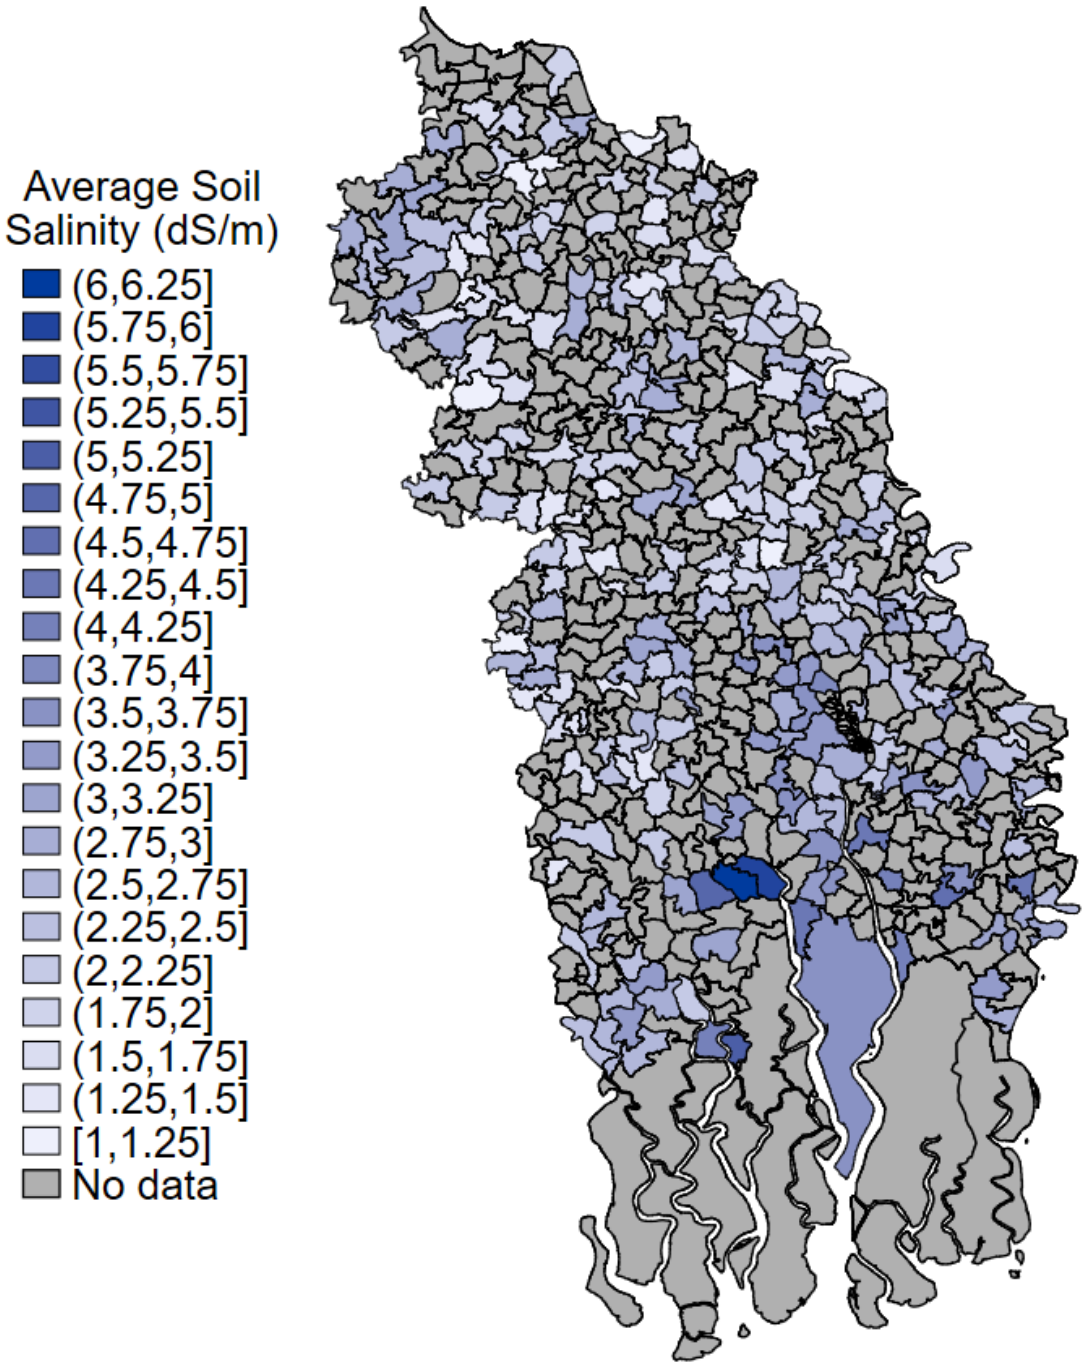


Figure A.3: Distribution of Rainfall Intensity Beliefs



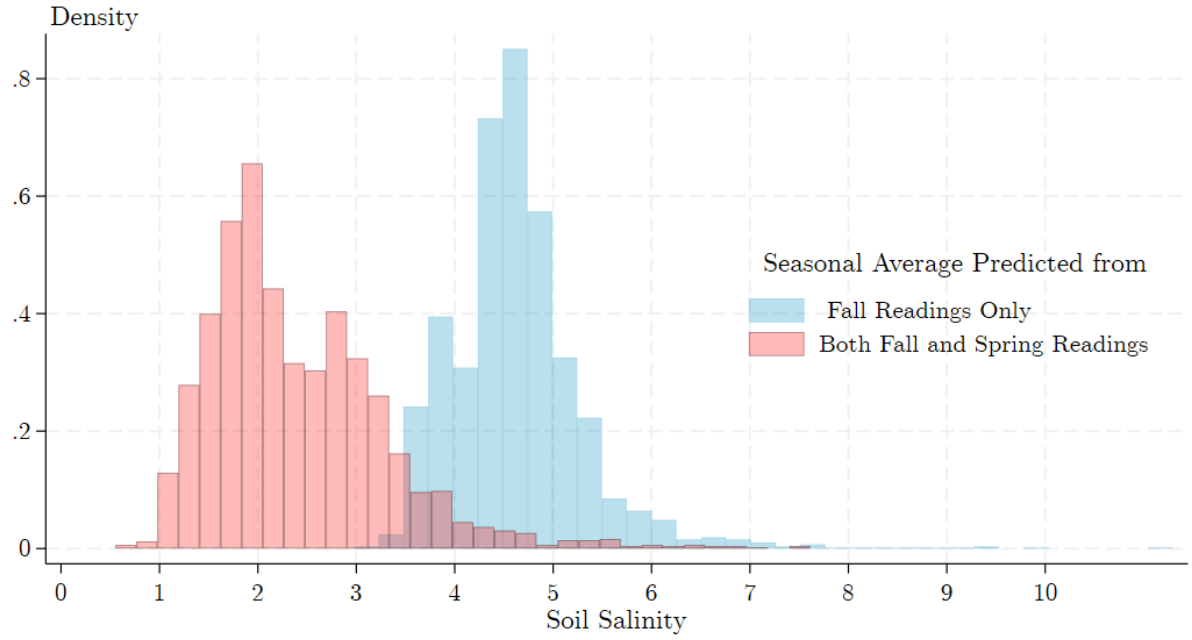
Note: Figure A.2 plots the distribution of farmers' mean salinity belief as elicited during the baseline survey. Figure A.3 plots the distribution of farmers' beliefs about the share of rainy days during the monsoon season for three different periods: the past (five or ten years earlier than the elicitation), the present (the most recent season), and the future (10 years from the elicitation year).

Figure A.4: Map of Average Soil Salinity from EC-Meter Measurements



Note: Figure A.4 maps the average soil salinity level within each union based on the measurements collected directly from farmers' plots over the course of the 2022-2023 season.

Figure A.5: Change in Soil Salinity over 2022-23 Season



Note: Figure A.5 graphs the distributions of the best prediction of the average soil salinity on farmers plots. The blue bars denote the histogram of seasonal salinity predicted based on the fall soil readings only. The red bars denote the same, now including the spring readings into the prediction as well.

Figure A.6: Distribution of Farmers' Willingness-to-Pay for Soil Salinity Information

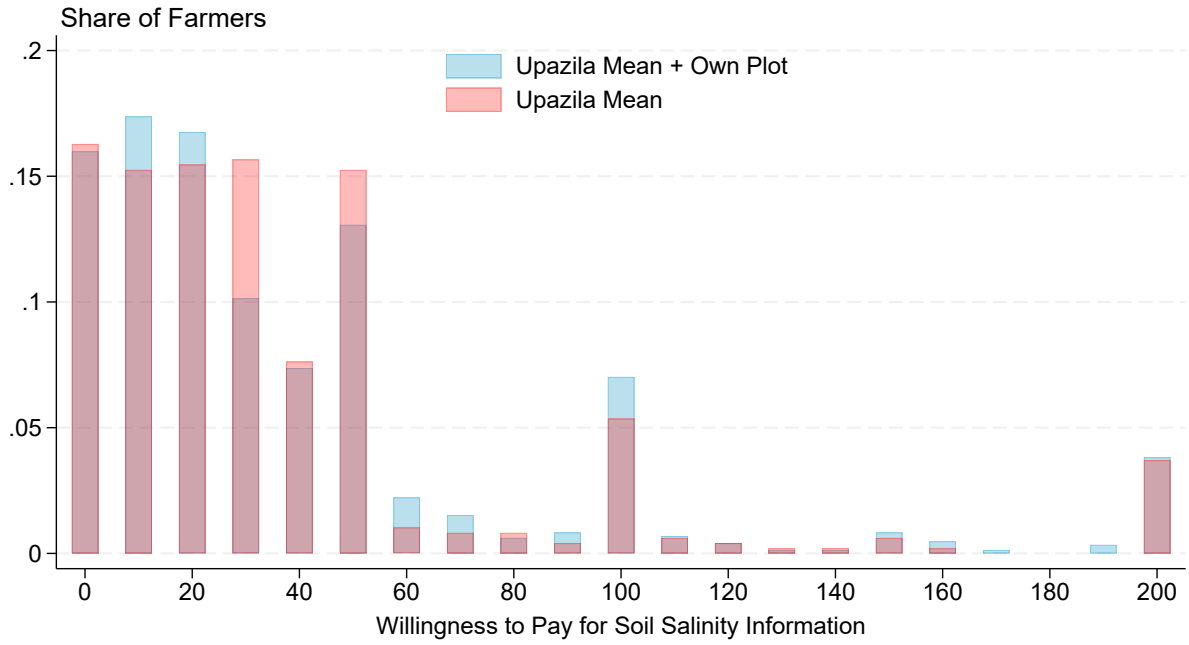
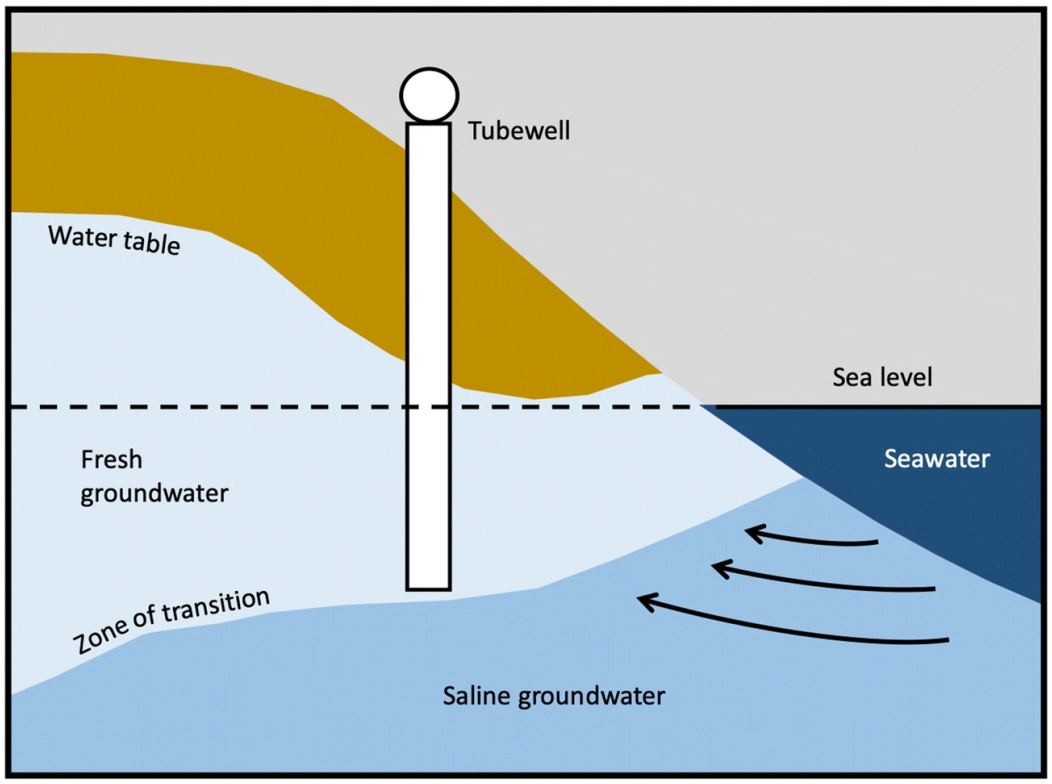


Figure A.7: Figure A.6 plots histograms showing the distribution of farmers' WTP for the soil salinity information as elicited via the Becker-DeGroot-Marschak price list. The blue bars denote the distribution for farmers offered information about both the average soil salinity level for all farmers in their distribution and their own plots, while the red bars shows the equivalent for farmers only offered information about the upazila average. For the purposes of this graph, I winsorized WTP at 200 BDT.

Figure A.8: Illustration of Rising Sea Levels' Impact on Soil Salinity



Note: Figure A.8 presents a diagram from the United States Environmental Protection Agency illustrating the dynamics between rising sea levels and the contamination of irrigation sources.

Table A.1: Farmer Survey Demographics

	Mean	Median
Household Size	4.7	4
Female	.033	0
Age	46	45
Years of Schooling	6.4	5
No Schooling	.21	0
Fewer than 5 Years of Schooling	.52	1
Completed Primary Schooling	.31	0
High School or Above	.13	0
Annual Earnings (USD)	1,656	1,365
Number of Plots	2.9	2
Survey Plot Size	46	33
Years Farming	22	20
Years Farming on Survey Plot	13	10
Lived Elsewhere in Past	.049	0
N	2,261	

Note: Table A.1 presents summary statistics for the main sample of farmers. The survey plot is the plot on which most of the survey questions focused, which was randomly selected with weights proportional to size.

Figure A.9: Photograph of Salt Visible on Soil Surface



Note: Figure A.9 shows a photograph I took of soil with salt visible on the surface from a village in Khulna.

Table A.2: Signs Used by Farmers to Assess Soil Salinity

	Share Using Sign
Look for white powder on the ground	0.432
Taste the soil	0.027
Taste the water	0.086
Red patches on the leaves	0.354
White patches on the leaves	0.072
Brown patches on the leaves	0.105
Brown leaves	0.108
Small plants/stunted height	0.549
Plant death	0.232
Use sensor	0.002
Ask Agricultural Officer/Dealer	0.004
Ask friends/family members/neighbors	0.026
Other	0.057
Don't use any signs	0.070
N	2,279

Note: Table A.2 presents the share of farmers using each indicator as a sign to assess soil salinity on their plots, as measured during the baseline survey. These answers come from the question, "What signs do you use to figure out the amount of salt in the soil?" Enumerators were instructed not to read out possible answers to the respondents.

Table A.3: Farmers' Perceptions of Symptoms of High Soil Salinity

	Share Using Sign
Red Patches on the Leaves	0.641
White Patches on the Leaves	0.084
Brown Patches on the Leaves	0.128
Brown Leaves	0.134
Small Plants/Stunted Height	0.784
Plant Death	0.503
N	2,253

Note: Table A.3 presents the share of farmers mentioning each response to the question, "If you had too much salt in your soil, how do you think that would impact what your rice plants look like" in the endline survey.

Table A.4: Role of Measurement Error in Salinity Belief Inaccuracy

	(1) Binary Belief	(2) Continuous Belief	(3) Continuous Belief
Beliefs - Truth	0.0286***		
	(0.00396)		
True Salinity		0.454***	
		(0.141)	
Leave-Out Mean Salinity			1.180***
			(0.126)
Constant	0.381***	2.262***	-0.821
	(0.0140)	(0.644)	(0.567)
Observations	2,068	1,151	2,068
Outcome Mean	0.444	4.364	4.617
R^2	0.025	0.012	0.053

Note: Table A.4 presents regressions testing for the role of measurement error in explaining the lack of explanatory power of true soil salinity on farmers' beliefs. Column (1) regresses whether farmers answered "Yes" in response to the question, "Do you think the soil on plot [X] is salty?" on the error between farmers' continuous belief and the agronomic measure. Columns (2) and (3) both use this continuous belief as the outcome. In column (2), the specification restricts to only those farmers who say they were not at all confused during the belief elicitation. In column (3), the predictor is the village leave-out mean of the agronomic soil salinity. All regressions exclude those who fail the beliefs comprehension checks during the practice elicitation and report heteroskedasticity-robust standard errors.

Table A.5: Examples of Salient Shocks and Subtle Shifts for Different Environmental Threats

Environmental Risk	Salient Shock	Subtle Shift
Flood risk	Flood	Marginal riverbank erosion
Soil salinity	Saline flood	Salinity intrusion into groundwater
Rainfall	Drought	Fewer days of rain
Temperature	Heat wave	Slightly higher temperatures
Cyclone risk	Cyclone	More intense storms in ocean
Air Pollution	Vision-impairing pollution	Marginally higher pollution
Wildfire risk	Wildfire	Longer periods without rain

Note: Table A.5 provides examples of salient shocks and subtle shifts across a host of different environmental threats.

Table A.6: Salient Shock: Salty Flood Impacts on True Soil Salinity

	(1)	(2)	(3)
Flood	-0.314*	-0.309*	0.0377
	(0.181)	(0.167)	(0.0842)
Saltier		0.406***	0.517***
		(0.0862)	(0.0905)
Flood × Saltier		0.157***	0.117**
		(0.0531)	(0.0557)
Constant	4.644***	4.650***	4.613***
	(0.0405)	(0.0386)	(0.0349)
Observations	2075	2075	2075
Clusters	250	250	250
Control Mean	4.605	4.605	4.605
Flood Risk Controls	Yes	Yes	No

Note: Table A.6 presents results from the difference-in-differences specification comparing the differential impact of salty floods on the agronomic measure of soil salinity. Column (1) presents regression results of the impact of floods alone, controlling for the machine learning generated flood risk measure. Columns (2) and (3) estimate the full difference-in-differences specification, where saltier is measured in terms of standard deviations. Column (2) is my preferred specification; column (3) does not control for baseline flood risk. All standard errors are clustered at the village level. Regressions exclude farmers who failed the baseline beliefs elicitation comprehension checks.

Table A.7: Subtle Shift: Salinity Intrusion's Impact on True Soil Salinity

	(1)	(2)
Sea Level Rise	0.286***	0.284***
	(0.0913)	(0.0991)
Ocean Salinity	-0.441***	-0.372***
	(0.0927)	(0.103)
Sea Level Rise × Ocean Salinity	0.192***	0.0432
	(0.0541)	(0.0638)
Closer to Ocean		0.179**
		(0.0711)
Closer to Ocean × Sea Level Rise		0.181**
		(0.0911)
Closer to Ocean × Ocean Salinity		-0.297***
		(0.0976)
Closer to Ocean × Sea Level Rise × Ocean Salinity		0.113*
		(0.0581)
Constant	4.430***	4.530***
	(0.0582)	(0.0757)
Observations	2075	2075
Clusters	250	250
Control Mean	4.605	4.605

Note: Table A.7 presents results from the triple difference-in-differences specification comparing the differential impact of being closer to the coast while being exposed to relatively higher sea levels and ocean salinity on the agronomic measure of soil salinity. Column (1) presents regression results of a difference-in-difference specification ignoring distance to the coast and just isolating the shocks to sea level elevation and ocean salinity. Column (2)—my preferred specification—additionally interacts with distance to the coast. All standard errors are clustered at the village level. Regressions exclude farmers who failed the baseline beliefs elicitation comprehension checks. All variables are measured in terms of standard deviations.

Table A.8: Salient Shock: Salty Flood Impacts on Soil Salinity Beliefs

	(1)	(2)	(3)
Flood	0.433	0.143	0.661
	(0.776)	(0.763)	(0.503)
Saltier		0.466	0.737**
		(0.346)	(0.352)
Flood × Saltier		0.909**	0.811**
		(0.395)	(0.396)
Constant	4.563***	4.569***	4.520***
	(0.166)	(0.167)	(0.151)
Observations	2068	2068	2068
Clusters	250	250	250
Control Mean	4.617	4.617	4.617
Flood Risk Controls	Yes	Yes	No

Note: Table A.8 presents results from the difference-in-differences specification comparing the differential impact of salty floods on farmers' beliefs about soil salinity on their plots. Column (1) presents regression results of the impact of floods alone, controlling for the machine learning generated flood risk measure. Columns (2) and (3) estimate the full difference-in-differences specification, where saltier is measured in terms of standard deviations. Column (2) is my preferred specification; column (3) does not control for baseline flood risk. All standard errors are clustered at the village level. Regressions exclude farmers who failed the baseline beliefs elicitation comprehension checks.

Table A.9: Subtle Shift: Salinity Intrusion's Impact on Soil Salinity Beliefs

	(1)	(2)
Sea Level Rise	0.274	0.383
	(0.344)	(0.470)
Ocean Salinity	-1.169***	-1.100**
	(0.348)	(0.433)
Sea Level Rise × Ocean Salinity	0.00356	-0.300
	(0.163)	(0.217)
Closer to Ocean		0.413*
		(0.244)
Closer to Ocean × Sea Level Rise		0.226
		(0.439)
Closer to Ocean × Ocean Salinity		-0.529
		(0.400)
Closer to Ocean × Sea Level Rise × Ocean Salinity		0.194
		(0.190)
Constant	4.591***	4.791***
	(0.204)	(0.270)
Observations	2068	2068
Clusters	250	250
Control Mean	4.617	4.617

Note: Table A.9 presents results from the triple difference-in-differences specification comparing the differential impact of being closer to the coast while being exposed to relatively higher sea levels and ocean salinity on farmers' beliefs about the amount of salt in their soil. Column (1) presents regression results of a difference-in-difference specification ignoring distance to the coast and just isolating the shocks to sea level elevation and ocean salinity. Column (2)—my preferred specification—additionally interacts with distance to the coast. All standard errors are clustered at the village level. Regressions exclude farmers who failed the baseline beliefs elicitation comprehension checks. All variables are measured in terms of standard deviations.

B Environmental Data

B.1 Estimating Flood Risk

This section describes the approach I use to calculate true flood risk for each union. Perhaps the simplest measure is simply the historical incidence of flooding in each polygon, which I calculate using the new measure of flooding that I derived from remote sensing data. The primary challenge with relying on that measure is that idiosyncratic factors may have caused a flood in one area and not the other in recent years, yet both places could still have the same fundamental flood risk in the future. The goal therefore is to use existing data to identify sets of locations where floods *might* happen using an objective scale. In addition to the simple historical exposure rates, I consider three separate types of indicators. First, I calculate "traditional" measures of exposure based on standard hydrological models. Second, I use my new measure of flood exposure derived from satellites to build a measure of flood risk. Third, I take a model-free approach by linking other geographic data with that historical measure and using supervised machine-learning to generate a flood risk measure.

For the first set, I defer to the government's estimates and calculate flood hazard and proneness from shapefiles produced by the Bangladesh Agricultural Research Council (BARC). When controlling for this set of flood risk, I include fixed effects for each of the government's flood proneness categories and quadratic polynomials of the flood hazard measure.

For the second, I calculate—for each union—the average flood exposure of its neighbors. The intuition behind this method is that neighboring areas likely experience very similar flood risk, and thus by calculating the average incidence among them, I can proxy for a given union's own flooding likelihood without giving undue influence to one place's own historical experience. I calculate two measures: one that includes the primary union's own past incidence in addition to its neighbors, and one that does not.

Finally, I calculate a predicted measure of flood risk using machine learning to dictate how underlying geographic features impact flood risk. This approach has the advantage of not relying on hydrological models, which can often be quite sensitive in their predictions, while also allowing me to let the data tell me which unions are comparable with one another. To flexibly estimate predicted flood, I train an algorithm to predict true flood experience based on geographic characteristics. First, I calculate the daily flood hazard rate based on the full panel of flooding experience. To account for the long tail of the distribution, I assign each union its percentile rank in this distribution of hazard rates, though results are broadly similar using the raw value. This variable constitutes the main flood risk outcome that I predict in the algorithm. As inputs, I calculate the mean and standard deviation of elevation in each union, binned latitude and longitude at the tenth of a degree level, the length of major rivers through that union, average drainage characteristics and flooding depth from BARC shapefiles, and each of the flood measures from the first set described in the previous paragraph. I then randomly split the sample of 5,158 unions into a training dataset and a testing dataset, with 20 percent reserved for the latter. Finally, I train a random forest algorithm to predict flood rank using the full set of inputs. Applying this model out of sample to the hold-out unions, I can explain 0.71 percent of the variation in true historical incidence ranking using my predicted measure. The fact that this R^2 is less than one can be

viewed as an advantage in this case because I can always control for the true past experience. This measure, by contrast, captures a data-driven notion of similarity.

B.2 Soil Salinity

I focus on two complementary data sources to measure soil salinity.

Direct Agronomic Data Collection First, enumerators took soil salinity measurements on the randomly selected plot using handheld soil sensors that measure electrical conductivity, with a range from 0 to 10 dS/m. For each plot, measurements were taken in three places, at least six inches from the edge of the soil, three feet apart from one another. I calculate the average reading from these three samples to measure each plots salinity on a given day. For repeat measurements, enumerators use photographs, GPS readings, and the farmer's location to return to the same point in each plot and collect follow-up data. Figure B.1 shows an image from the baseline survey of enumerators using the soil sensors to measure salinity. Enumerators used Yieryi EC-98361 Digital Soil EC Meter Testers with a resolution of 0.01 ms/cm.

Figure B.1: Salinity Measurement



Note: Figure B.1 shows an image of soil salinity being collected by enumerators.

Salinity evolves over the course of the dry season, but due to budget restrictions, high-frequency salinity measurement was not possible. Instead, enumerators took salinity measurements twice: first during the baseline survey, and second during a midline visit during the season.[46] Based on multiple photos and detailed GPS coordinates taken during each measurement, I verify that the enumerators indeed returned not only to the same plot but to the same point in each plot. Ideally, measurements would be taken continually during the growing season to understand average exposure. Given budget restrictions, the midline was timed to maximize predictive power of the average soil salinity from November through April. Most of the baseline survey measurements took place in November. Based on data from the SRDI's soil monitoring sites described further below, the additional month that maximizes predictive power in a simple linear regression is March, with an R^2 of .901. The second round of measurements was timed so as to collect as much data as possible in March, subject to constraints around the timing of Ramadan and the geographic scope required of enumerators, with some data collected in the second-half of February.

To convert these snapshots of salinity into a measure of seasonal exposure, I estimate

[46] Additionally, enumerators measured salinity from a random sample of 30 villages during a third visit at the end of May and beginning of June. These data confirm the trends shown from the second-round of measurements.

the link between seasonal average and the relevant calendar months using the 142 full seasons of salinity from the SRDI data and apply this relationship out of sample to the soil measurements I collect. For a small number of households, enumerators could not collect salinity measurements during one of the two visits: 164 during the baseline (largely due to flooding making the plot inaccessible) and 31 during the midline (typically because of migration or other household unavailability). For these households, I replace the missing measurement with the mean village value prior to predicting seasonal salinity exposure; if the entire village lacks the measurement, then I use the mean for the upazila.

The gold-standard method for measuring soil salinity requires taking soil samples to a lab for analysis—prohibitively expensive at the scale of this study. This raises an important concern with the handheld sensors used to measure salinity in this sample of farmers: any systematic bias in these sensors would distort the results. In particular—given that the salinity levels I measure tend to be systematically lower than historical salinity levels measured by the SRDI—one might worry that my sensors underestimate the truth, especially at higher salt levels. To test the scope of this potential issue, I conduct laboratory tests with the handheld sensors used in the study and compare the results to other more expensive sensors available in the market. This approach features the advantage of being able to know the ground truth: by carefully measuring out grams of NaCl and mixing it with a solution of distilled water, I can calculate the exact electrical conductivity of the solution and compare the measured values to this ground truth.[47] Of course, these sensors are designed for direct soil measurement as opposed to this controlled solution, and thus discrepancies could indicate differences in the type of sample instead of indicating errors with the sensor.

Figure B.2: Salinity Sensor Lab Experiments



Note: Figure B.2 shows an image of one of the soil salinity sensors used in this study being tested in a solution of distilled water and NaCl.

I test for differences in measurement across sensors and salinity levels in these lab experiments. I conduct the evaluations in 250 milliliter solutions featuring grams of NaCl ranging from 0 to 1.5 in increments of .1 grams, equivalent to a dS/M range of 0 to 10.27. I compare three sensors used in the field, two sensors of the same model but brand new (in case the sensors change over time or by use), and two brand new high-end models. I randomized the order in which each sensor was placed in the solution, and include fixed effects for this order. I first document that among the high-end models—which are much too expensive to feasibly be used at scale—the salinity sensors nevertheless tend to overestimate the objective salinity level by an average value of (0.73 dS/m).[48] The sensors I use during the midline overestimate by a smaller amount, with a mean of 0.69 dS/m. To help put this magnitude in perspective, the within

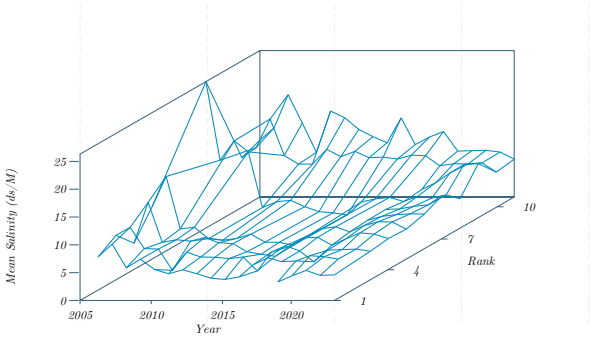
[47] Specifically, the gram molecular weight (GMW) of sodium (Na) is 22.99, the GMW of chloride is 33.45, so I calculate the true conductivity as grams of NaCL divided by 58.44 times 1000 divided by millileters of distilled water.

[48] One of the high-end models top-codes at 2.00 dS/m, and thus I exclude all measures above that point for that model.

plot, within day standard deviation in measured salinity in my sample is 0.15 ds/M on average. In a regression specification allowing for a differential linear bias in true salinity, I find no statistically significant evidence that the cheaper sensors I use perform differently than the more expensive ones. The fact that the overall discrepancy suggests that if anything, the sensors I used overestimate salinity, contrary to what would be expected given the low levels I observe in the field compared to historical trends.

Government Salinity Panel Second, SRDI provided measurements from monthly soil samples from 11 monitoring sites in and around Khulna.[49] Most of these samples date back to 2004. Figure B.3 visualizes the average salinity value during the Boro season for these sites over time. The graph illustrates the considerable variation both over time and across places in soil salinity. Figure B.4 presents the simple average time series in annual salinity among a constant sample of stations. Salt levels also exhibit significant variation *within* the year in predictable seasonal patterns. Figure B.5 plots this seasonal variation using data from the government's soil stations.

Figure B.3: Historical Boro Season Salinity from Soil Measurement Sites



Note: Figure B.3 plots salinity data from soil monitoring sites that I acquired from the Bangladesh Soil Research Development Institute. To calculate annual soil salinity, I take the average value between November of the previous calendar year and April, with each month given equal weight. I arrange monitoring sites by their average salinity value across all available years.

The government's data provide reassurance that the patterns in seasonal salinity I document are true. Figures B.6 and B.7 show that consistent with my readings, February and March of the 2022-23 season were abnormally low. This is particularly surprising given the low amounts of rainfall, but the results are striking even in this relatively underpowered analysis.

B.3 Water Salinity

To measure the local water salinity, I combine two data sources.

First, I obtain monthly water salinity levels from 133 river station monitoring sites maintained by the Bangladesh Water Development Board (BWDB). Figure B.10a maps these stations, which I assign to unions based on the closest centroid. I acquire all available data since 2011, yielding 5,955 observations total. For each station and month, I observe the maximum, minimum, and average salinity at high tide. Figure B.8 plots the average values across month, illustrating substantial seasonal variation.

Second, I estimate exposure to ocean salinity levels every day since January 1^{st}, 2002 by first calculating surface sea water salinity from the Hybrid Coordinate Ocean Model (HY-

[49] A settlement was built on one of these sites during the 2022-23 season, so I only have data on 10 for the most recent months.

Figure B.4: Trends in Annual Soil Salinity—SRDI Soil Stations (Constant Sample)

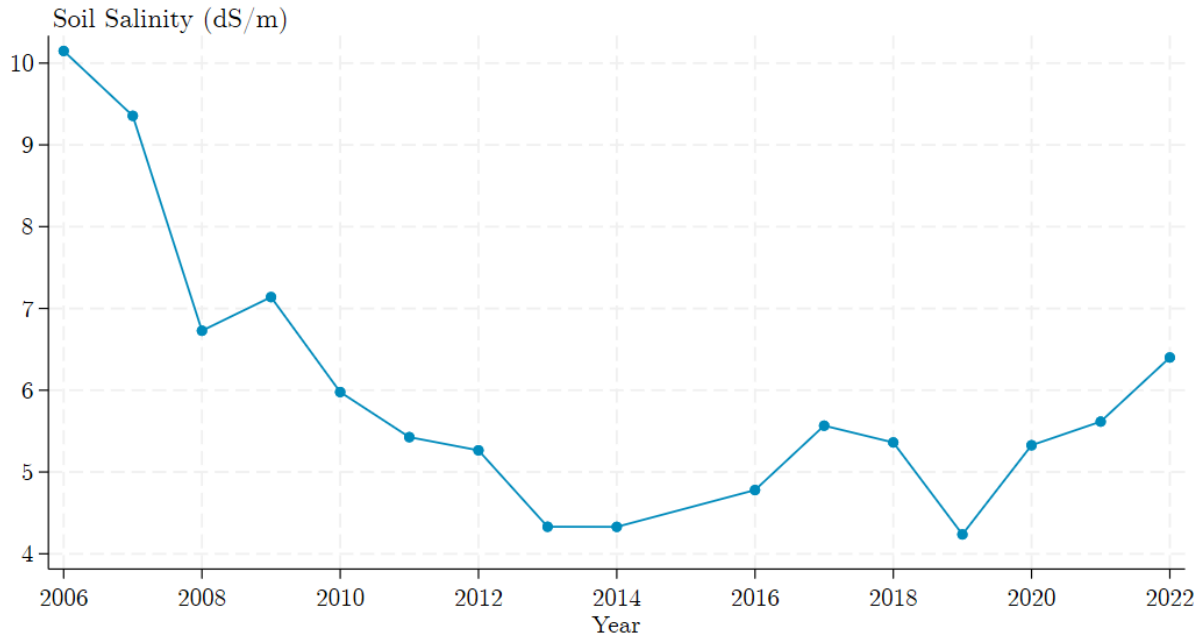
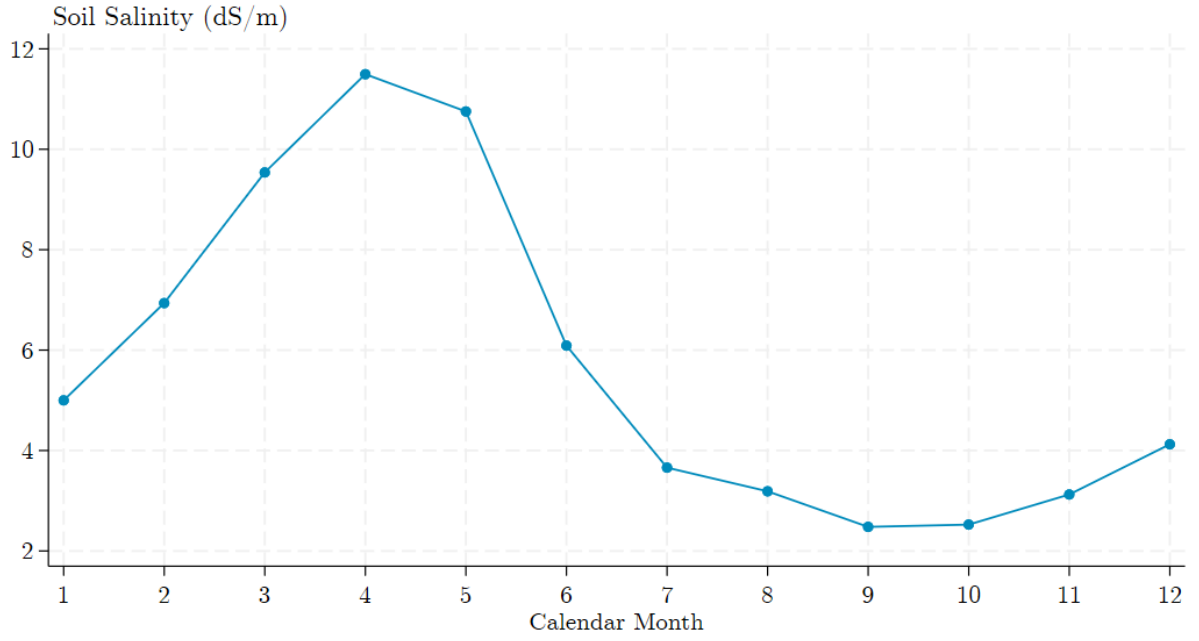


Figure B.5: Seasonal Patterns in Soil Salinity—SRDI Soil Stations



Note: Figure B.5 plots the average soil level across the 11 locations monitored monthly by the Bangladesh government's Soil Resource Development Institute. Seven of the stations have data for most months since 2004; four new ones were added in 2018. Figure B.4 takes the constant sample of these stations with data available for all twelve months of each year, and plots the average salinity for all available years.

Figure B.6: 2022-23 Soil Salinity Anomaly: SRDI Soil Sites Raw Data

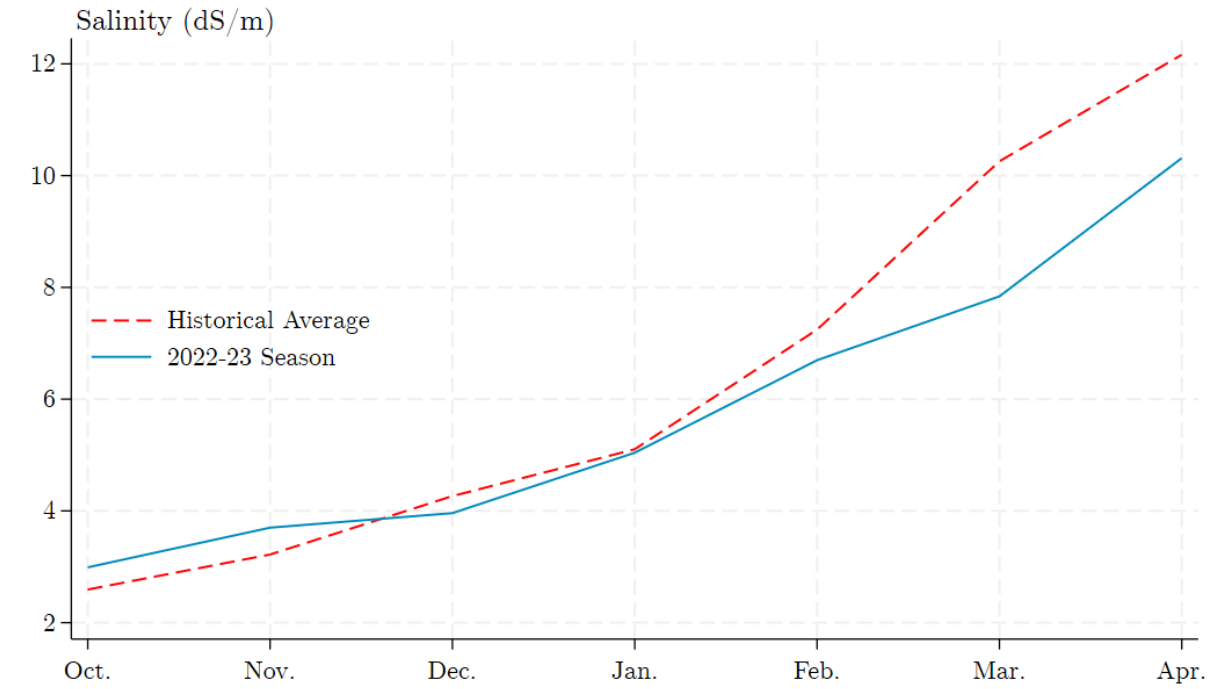
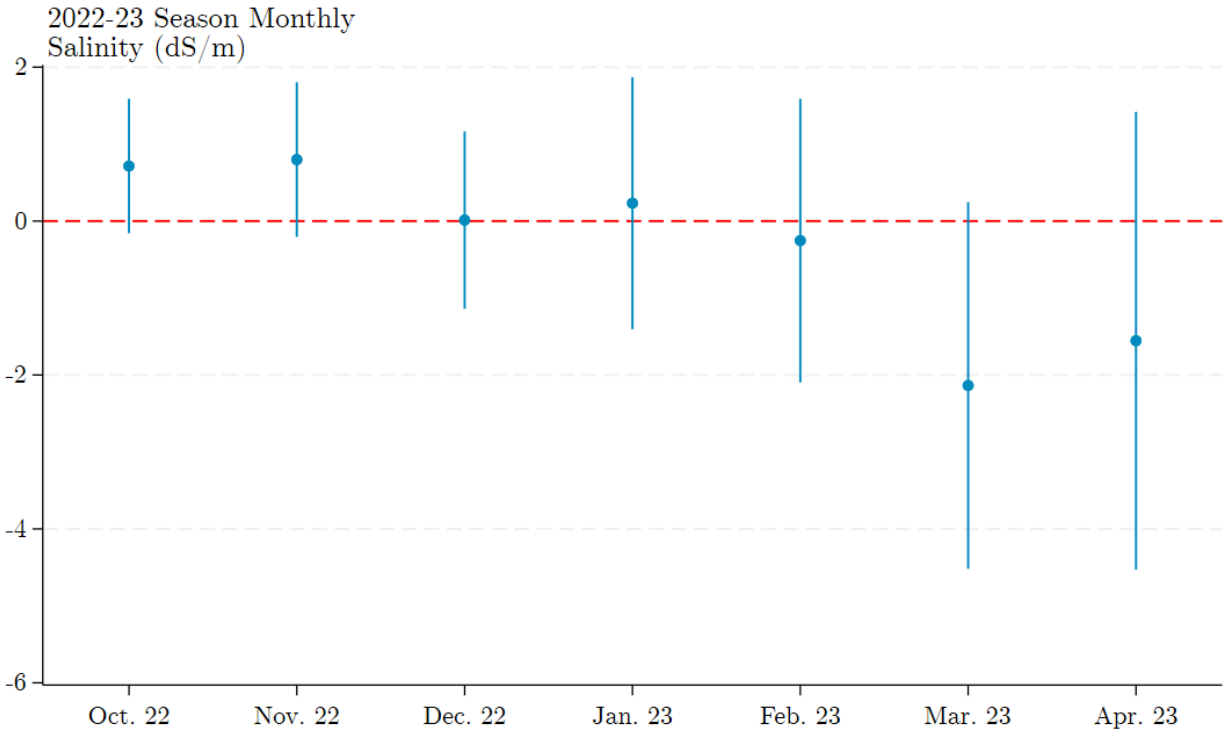


Figure B.7: 2022-23 Soil Salinity Anomaly: SRDI Soil Sites Raw Data



Note: Figures B.6 and B.7 use data for the 10 sites for which I have historical data from the SRDI to assess the extent to which 2022-23 exhibit abnormal soil salinity patterns. Figure B.6 plots the raw means of the 2022-23 season against the historical average by calendar month across all sites. Figure B.7 plots coefficients from a regression of calendar month interacted with a dummy for the 2022-23 season controlling for site and calendar month fixed effects, clustering standard errors by site.

COM) along 10 kilometer-wide grids off the Bangladeshi coast (Cummings and Smedstad, 2013). Figure B.10b illustrates the coastal grids over which I calculate daily ocean salinity from this remote sensing data in red, and the nearest straight line distance to unions in blue.

The BWDB does not collect data during all months of the year—and in particular, not during the critical flooding times of the monsoon season. To overcome this missing data challenge, I train a supervised machine learning algorithm to predict river salinity using predictors constructed from ocean salinity measurements, which are available year-round. I first construct a vector of predictors from the ocean salinity data. Following Guimbeau et al. (2023), I identify the five closest coastal polygons to each BWDB station, and then construct five month leads and lags of mean and maximum ocean salinity from those five polygons. I include features for both the unweighted salinity, weighted according to the inverse squared distance of the five closest locations, and a final set multiplying this weighted version by the minimum distance from the salinity station to the coast. Using a random forest to predict the truth, I can explain 82.41 percent of the variation in a 20 percent hold-out testing sample.

B.4 Sea-Level Rise

To measure village-level exposure to sea-level rise, I again use data from the Hybrid Coordinate Ocean Model (HYCOM) along 10 kilometer-wide grids off the Bangladeshi coast (Cummings and Smedstad, 2013). I calculate the sea surface elevation anomaly relative to the modeled elevation mean. Figure B.11 plots the evolution of average elevation across these polygons over time, with darker colors denoting earlier years, to illustrate the variation I use in the natural experiments.

Figure B.8: River Station Salinity Data

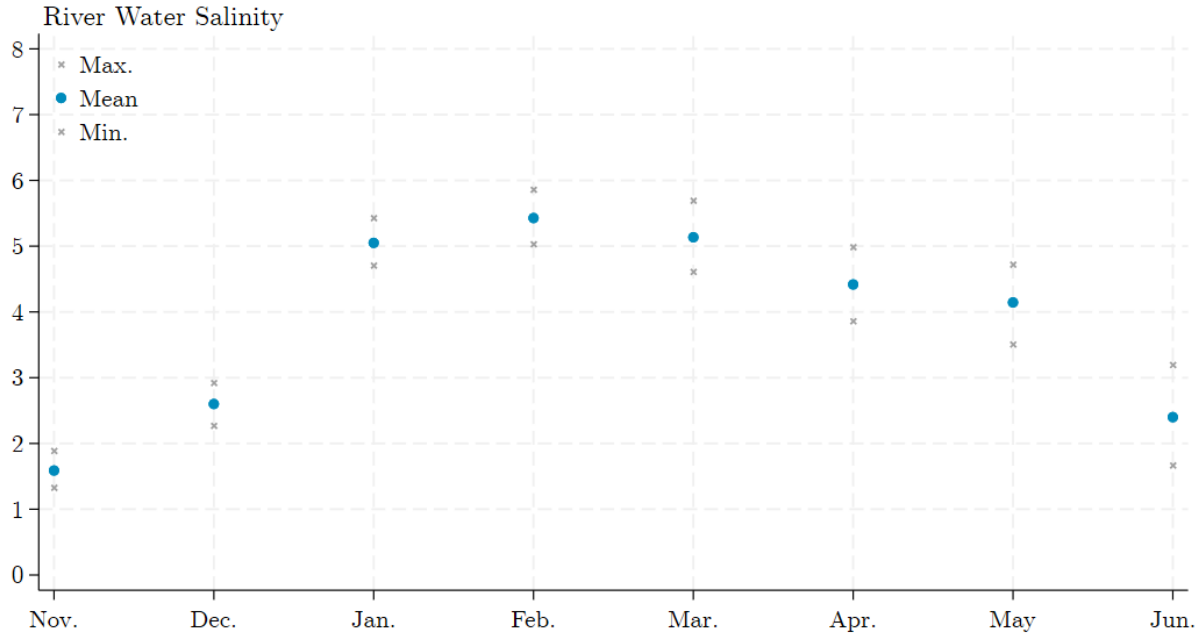
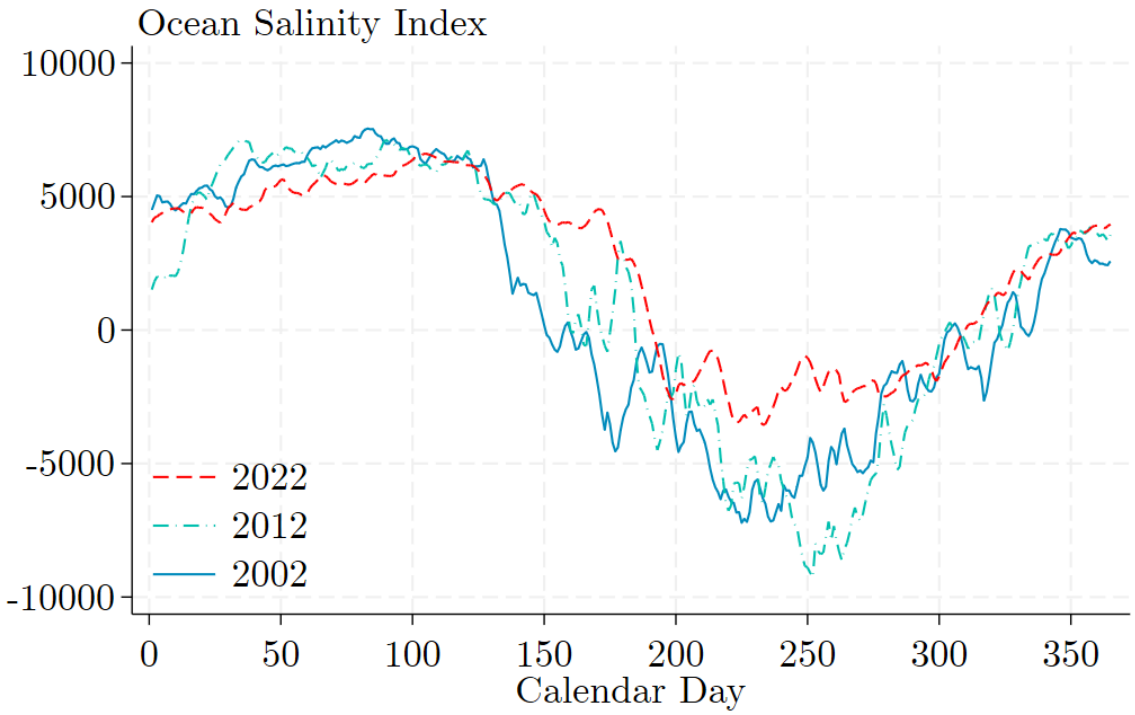


Figure B.9: Ocean Salinity Over Time



Note: Figure B.8 plots the average maximum, minimum, and mean salinity at high tide across river stations by calendar month using the data I obtained from the Bangladesh Water Development Board. Figure B.9 plots average ocean salinity across polygons for three sample years over the course of the calendar year.

Figure B.10: Mapping Water Salinity Sources

(a) BWDB River Salinity Stations

(b) Linking Unions to Coastal Grids

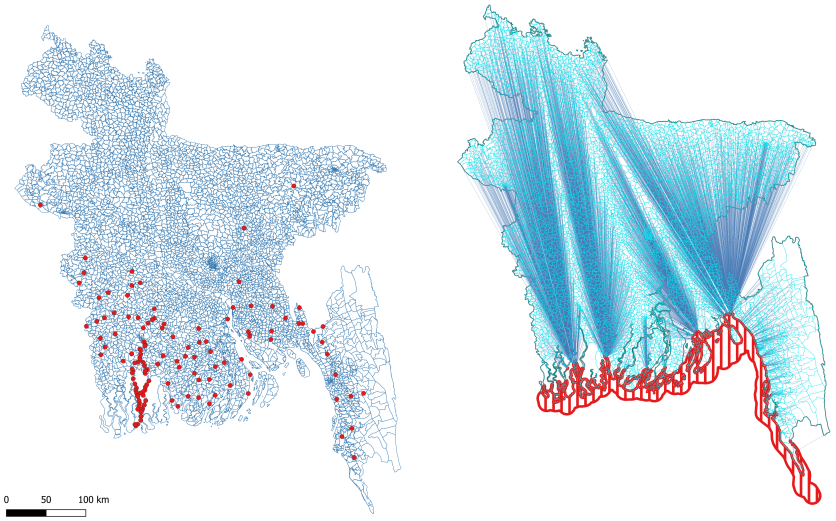
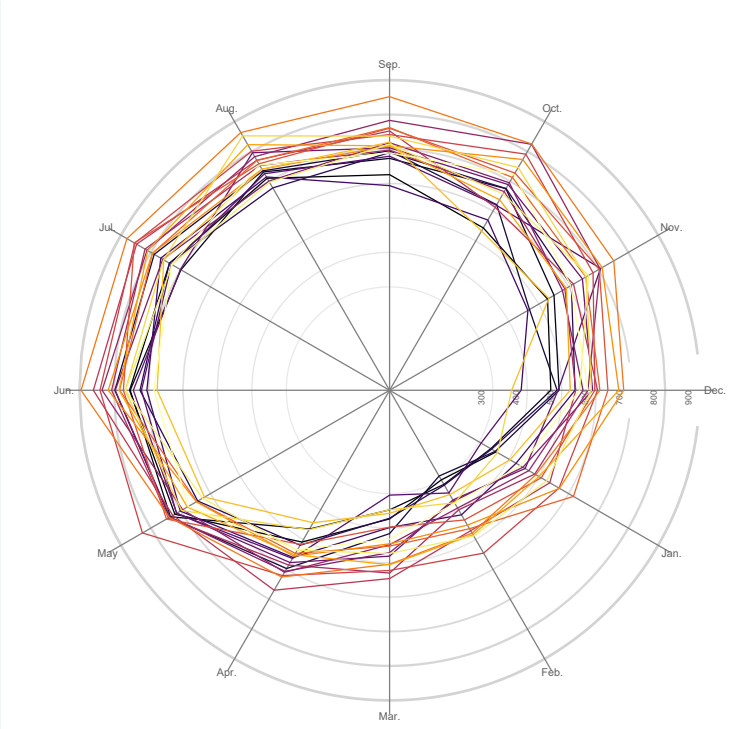


Figure B.11: Rising Sea Levels in Bangladesh



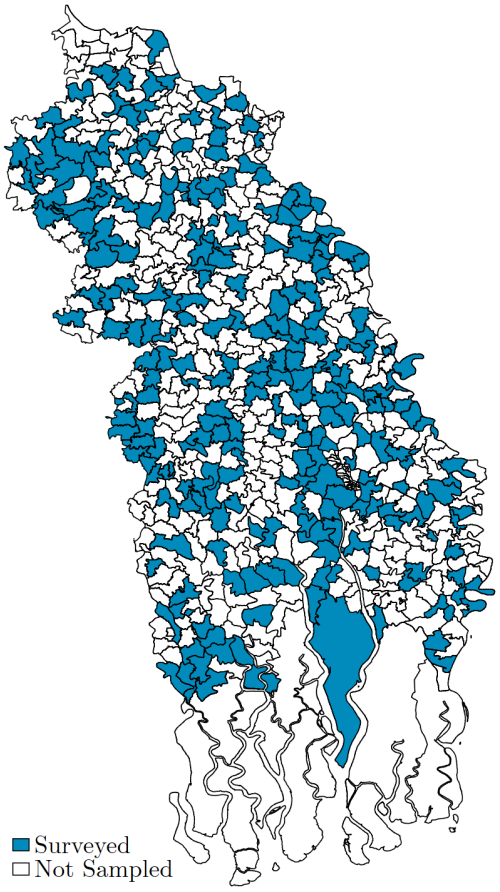
Note: Figure B.10a maps the BWDB river salinity stations from which I obtain data. Figure B.10b illustrates the ways I link unions to coastal grids via the closest cell. Figure B.11 plots how sea levels have risen along the coast of Bangladesh based on my calculations from satellite data of the HYCOM (Cummings and Smedstad, 2013), averaging across the 10 kilometer-wide grids I have created off the Bangladeshi coast.

C Farmer Perceptions Survey

C.1 Sampling

Union Sampling The survey was conducted across 250 unions in the Khulna division of Bangladesh. From the 642 unions in the Global Administrative Areas (2018) data, I exclude 32 urban areas and then select nine unions from which I have government salinity station data, 37 unions with water stations from the Bangladesh Water Development Board, 48 unions that are also included in the Bangladesh Integrated Household Survey sampling frame, and 121 unions that are also included in the 2016-2017 Bangladesh Labor Force Participation Survey sampling frame. This yields 185 unique unions. Both the Bangladesh Integrated Household Survey and the Bangladesh Labor Force Participation Survey were designed to be representative, and just 29 of the 185 initially selected unions fall outside of both of those survey's sampling frames as exclusively part of the government salinity or water stations. Then, I randomly sort the remaining unions, and choose the next 65.[50]

Figure C.1: Map of Surveyed Unions

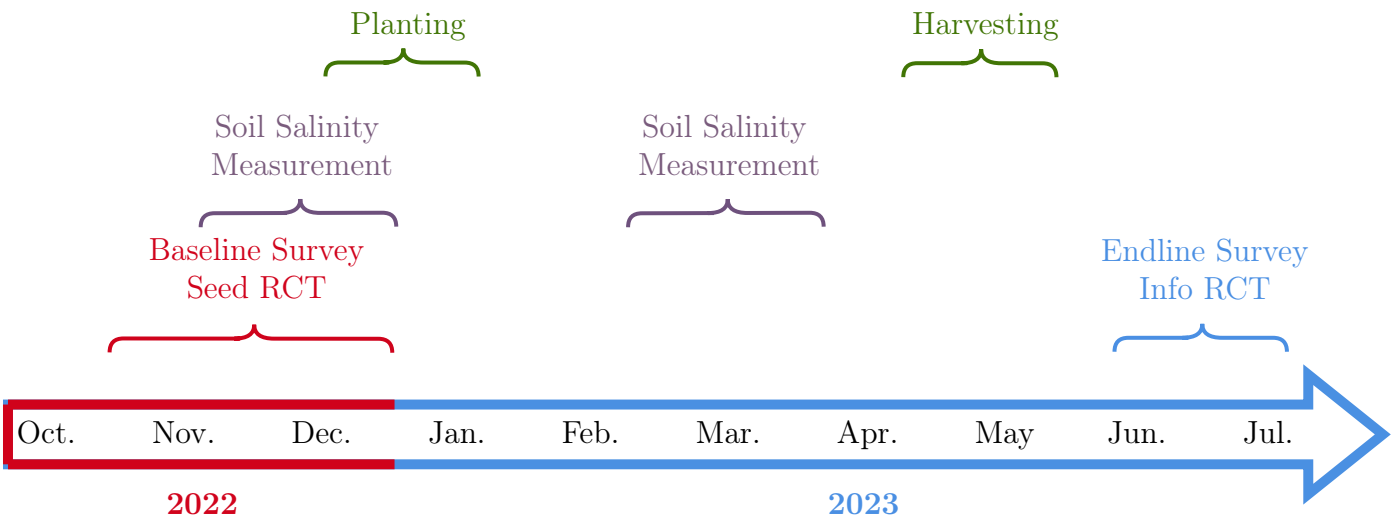


Note: Figure C.1 presents a map of Khulna division with blue denoting unions included in the survey.

Farmer Sampling Enumerators visited each union and did an initial listing of 50 households who were planning on harvesting rice during the upcoming Boro season and made the primary agricultural decisions on their land. In almost all unions, this goal of 50 households was achieved and typically within a single village. From this initial list, farmers were randomly ordered to be selected for an interview. Initially, 10 households were selected per union, though this number was revised down to nine given survey length concerns after the first week. On average, 9.1 farmers were surveyed in each union. Before a household was deemed unavailable and a replacement household was selected from the randomized listing order, enumerators attempted to contact them multiple times over multiple days via their phone number collected during the

[50] The original sampled list included 250 unions. After enumerators attempted to conduct the listing exercise, they could not find a sufficient number of Boro rice farmers in either Nalian Range or Satkhira Range, reducing the total sample to 248, which was the pre-registered sample size. During the course of the baseline survey, it was discovered that in two additional unions, farmers had a different interpretation of the term "Boro" and did not harvest rice during the relevant season, reducing the sample size to 246. After additional funding was received partway through data collection, however, four new replacement unions were added to bring the total back to 250, following the initial randomization order.

Figure C.2: Data Collection Timeline



Note: Figure C.2 shows the timeline of data collection activities along with farmers' typical planting and harvest times.

listing. Of the endline respondents, 97.07 percent were also interviewed during the baseline; in the small number of cases when the baseline respondent was unavailable, another person from the household who makes decisions about farming was interviewed. Among the 66 households successfully interviewed in the endline but for which the primary respondent was unavailable, 28 of the baseline respondents had migrated, six had passed away, and 32 had another conflict.

Timeline Data collection spanned from Fall 2022 through Summer 2023, as pictured in Appendix Figure C.2. Villages were assigned to enumerators based on location to minimize staff travel time. Within enumerator, the order of villages was randomized for both the baseline and endline surveys. Enumerators typically completed three surveys a day, spending three days in each village.

C.2 Variable Construction

Please visit the following links to view the complete survey instruments: Baseline (English translation), Baseline (Bangla translation), Endline (English translation), and Endline (Bangla). The plot-specific questions and the salinity measurements took place on a plot selected at random with a probability proportional to the plot size.

Demographics I construct a continuous years of education measure based on the categorical answers given to the education question as follows: no schooling or informal education only (0), class 5 or below (5), class 8 or below (8), SSC/Dakhil or below (10), HSC/Alim

or below (12), Degree (12), Graduate (16), Technical/Vocational education (16), and Post-graduate (18).

Environmental Knowledge *Salinity:* Respondents were first asked, "Does the amount of salt in the soil change throughout the year?" Those who answered yes were then asked in which month the salt in the soil was highest and lowest. Based off of monthly data from the Bangladesh Soil Resource Development Institute's soil stations in Khulna and nearby areas, I classify correct answers generously: any month between mid-March to mid-June (in the Bangla calendar) for the peak and any month between mid-July and mid-November for the trough.

Environmental Beliefs *Salinity:* To elicit beliefs about soil salinity, I ask farmers to distribute buttons across the picture shown in Figure C.3. I first introduce this image by telling farmers, "This image shows pictures of rice plants at the end of the season, once they are fully grown. They are arranged from the best growing to the worst growing. The smallest ones grew the worst. The biggest plants grew the best. Plant number 1 is the least healthy, and plant number 7 is the most healthy." To ensure that farmers do not simply misinterpret the question as asking about last year's or next year's crops, I next ask, "First, think about last year. Which of these pictures best matches the plants that grew on your plot last year? I would now like to know how you think your own crops will fare this year. Think about the end of the season. What are your guesses about what your grown plant will look like on your plot? Place the highest number of buttons on the image that best matches your guess. Remember, plant 7 is the healthiest and plant 1 is the least healthy." Then, I introduce the salinity study specifically by instructing farmers, "This is not a picture of your own plant, it is taken from a previous study. Researchers have grown rice seedlings under different conditions. This rice variety is not specially adapted for saline soils. *Instructions: Point to picture that has the biggest plant.* This picture shows the seed grown in soil with the least amount of salt. *Instructions: Point to picture that has the smallest plant.* This picture shows the seed grown in soil with the most amount of salt. *Instructions: Point to the pictures in the middle.* These pictures show seeds grown in increasing amounts of salt, from largest to smallest. Do you have any questions about these plants?" After answering any lingering questions, the main belief elicitation question asks farmers, "This photo comes from researchers who planted rice that is not saline tolerant in different soils with different amounts of salt. If they used your soil from your plotplot, which of these pictures do you think would look most like the plant at the end of the season? We are asking this question because we are trying to understand how much salt you think is in your soil. You should assume that the researchers copy all aspects of your soil, such as the water and fertilizers you use over the season and the weather on your plot. Please place more buttons on the pictures that you think are more likely." In follow up questions, I elicit beliefs about soil conditions five years in the past and in the future.

The image from Figure C.3 comes from the field experiments testing salinity's impact on rice yield from Grattan et al. (2002). Based on the growing conditions from that paper, I convert beliefs into average salinity perceptions based on equation 11, where s_i^k denoting the number of buttons places on the k^{th} rice plant, where the healthiest plant is 7 and the

Figure C.3: Soil Salinity Belief Elicitation Visual Aid



Note: Figure C.3 shows the image for measuring expectations about soil salinity among farmers, which comes from field experiments run to test the impact of salinity on rice yield from Grattan et al. (2002).

least healthy is 1. This equation incorporates both the conversion from the plant health to the seasonal average of field water salinity and the translation of field water salinity to salinity from the top three inches of soil. In both steps, Grattan et al. (2002) find strong linear relationships with high R^2.

$$\text{Expected Salinity}_i = \widehat{s}_i = 1.17 * \left(\frac{\sum_{i=1}^{7} (.3 + 1.95(7-k)) s_i^k}{10} \right) + 1.2 \tag{11}$$

Due to the discrete nature of the button elicitation method, farmers cannot express beliefs that span the complete support of possible salinity expectations. To account for this measurement error, I adjust beliefs under the following structural assumption. This ensures that any bunching I observe in the data is not an artifact of the elicitation method. Let s^r denote a possible salinity expectation that can be expressed using buttons, where r ranks values from smallest to largest. The key assumption I make is that the true expected belief s_i^* equals their reported value $\widehat{s}_i = s^{r(i)}$ plus an adjustment drawn from a uniform distribution

spanning between ranks, following equation 12, yielding the final belief in equation 13.

$$\mu_i \sim U\left(\frac{s^{r(i)-1} - s^{r(i)}}{2}, \frac{s^{r(i)+1} - s^{r(i)}}{2}\right) \tag{12}$$

$$s_i^* = \widehat{s}_i + \mu_i \tag{13}$$

One limitation of this measurement approach is bottom-coding from the Grattan et al. (2002) experiments. Indeed, I find soil salinity levels below the lowest amount considered in that study on several hundred plots. To account for this, when comparing beliefs to the truth, I also bottom-code the true salinity level at the same level as these beliefs.

Perceived Returns to Salinity-Tolerant Seeds To measure beliefs about the returns to salinity-tolerant seeds, I first ask farmers to estimate the amount of rice they expect to harvest from the target survey plot. Then, I ask them by how much they expect this harvest to change if they switched to a salinity-tolerant seed. If they already intend to plant a salinity-tolerant seed, then I instead ask about switching to a non-salinity tolerant seed. In all questions, I use the visual belief elicitation method described above that uses buttons to elicit probabilistic beliefs. I simply use the mean return scaled by the total mean expected harvest in the analysis of the returns to salinity-tolerant seeds.

Agricultural Profits Capturing agricultural profits presents a variety of challenges in this setting. I therefore calculate several complementary measures, each with its own advantages and drawbacks.

Separately from farmers' self-reported profits, I calculate a measure of greenness based on satellite data. First, I take the average latitude and longitude of each farmers' plots from the salinity sensor readings. Second, I construct a circle around this point of .0005 degrees, or approximately 55 meters. Note that because plots vary in size and the soil salt measurements did not occur in the middle of the plot, this step certainly generates measurement error, though it should be independent of treatment status. Third, I calculate a normalized difference vegetation index (NDVI) using remote sensing data from the Sentinel-2 satellite. I use the Harmonized MultiSpectral Instrument Level 2-A series made available on Google Earth Engine. From the endline survey, I take the harvest week of the Boro crop and calculate average greenness in the three weeks leading up to that date.[51]

I compile every satellite image between February 15^{th} and March 15^{th} 2023. I calculate the average value of the red and near infra-red (NIR) wavelengths for each image, excluding any observations with any cloud coverage. Finally, I calculate the NDVI on each day according to equation 14 and take the average value across this time period for each plot.

$$NDVI_i = \frac{NIR - RED}{NIR + RED} \tag{14}$$

[51] I use responses from the 98.33 percent of farmers who report a date between the second half of February and the first half of June. For the others (including those who did not harvest in Boro), I assign them the modal date from their treatment status by village response. There is no difference between treatment and control group in reporting an answer within the relevant date (p-value=0.922). In one village, all harvest dates are missing, and I exclude respondents from this village from this analysis.

$$EVI = \frac{2.5 * (NIR - RED)}{NIR + 6 * RED - 7.5 * BLUE + 1} \tag{15}$$

Agricultural Issues: What Comes to Mind? The first question asked to respondents after the consent form was: "What are the biggest issues that you face in cultivating Boro rice?" I recorded the audio of farmers' responses to this question, and then coded up their answers according to the issues mentioned by each farmer. I construct 13 categories, in addition to an other category: Pests/Disease, Salinity, Input prices, Non-salinity water problems, Input shortages, Seed issues, Rice sale price, Non-salinity soil problems, Other rice plant health problems, Farmers' own health, Floods, Rain/droughts, and Temperature.

D Additional Data Sources

D.1 New Seed Database

Rice farmers in Bangladesh plant a large number of different seeds, many of which are local varieties. Unfortunately, no systematic data includes information on these different varieties. To overcome this issue, I construct a new database of seed characteristics by compiling data from several different sources. I take records from the Digital Herbarium of Crop Plants from the Department of Crop Botany at Bangabandhu Sheikh Mujibur Rahman Agricultural University, which itself aggregates information from the seed developers. I supplement these data with information from three government sources: the Bangladesh Ministry of Agriculture Bangladesh Agricultural Research Council's Agri-Advisory Portal, the Seed Certification Agency, and the Bangladesh Rice Research Institute. I add additional information on the growth duration of each seed from the Food and Agriculture Organization on Bangladesh. I execute online searches for each variety and, in some cases, call the seed importers directly to collect additional information.[52] Finally, local varieties constitute an important share of planted seeds in this setting, yet official records detailing their characteristics do not exist. I use phone surveys with four seed dealers from across the Khulna division in an attempt to fill in this data gap. I focus on salinity tolerance status, the most important seed characteristic for my analysis.[53]

The complete set of seeds for which I attempt to find data spans 330 different varieties, of which I classify 17 as salinity tolerant. As a benchmark, I classify 42 as resistant to at least one pest. Table D.1 presents information on how other seed characteristics vary by whether the variety is tolerant to saline soil. Salinity-tolerant seeds tend to produce smaller yields and enter the market more recently.

[52] The Hera-2 seed is somewhat salinity-tolerant, but the original Hera seed is not. However, because farmers sometimes use the term "Hera" to refer to either, I defer to their judgment on a case-by-case basis as to whether their variety of "Hera" is indeed resistant to high salinity levels.

[53] The information from seed dealers sometimes contradicts the official government records or one another for this feature, and thus I only defer to the seed dealers' classification in the case of local varieties for which no official source exists, and when at least two dealers report the same seed being saline tolerant.

Table D.1: Seed Characteristics by Salinity Tolerance Status

	(1) Boro Yield (Tons/Hectare)	(2) Plant Height (cm.)	(3) Release Year	(4) Growth Duration (days)
Salinity Resistant	-1.024*	-1.841	14.36***	4.508
	(0.536)	(3.070)	(2.969)	(3.302)
Non-Resistant Mean	6.689	109.832	1997.891	135.824
Observations	64	119	147	139

Note: Table D.1 presents ordinary least squares regressions of key seed characteristics against each variety's salinity tolerance, based on the new database I assemble. Observations vary depending on what information I am able to track down for each seed.

E Conceptual Framework Appendix

E.1 Proofs of Theoretical Results

This section provides a formal treatment of the results from Section 3. To make this section self-contained, I begin by repeating the key notation from the set-up as discussed in the main text and then derive the main predictions.

Set-Up Farmer grows rice in two periods, $t \in \{1, 2\}$. Output in period t is given by the binary indicator $y_t \in \{0, 1\}$, where $y_t = 0$ denotes low harvest, and $y_t = 1$ denotes high harvest. Harvest is subject to a random productivity shock ξ that can be either negative ($\xi = -1$), positive ($\xi = 1$), or neutral ($\xi = 0$). I assume productivity shocks are distributed symmetrically with mean zero such that the positive and negative shocks occur with equal, positive probability denoted by $\rho > 0$ and that neutral shocks occur with positive probability such that $\rho < .5$. In the first period, farmers make no decisions about inputs and plant the standard seed. In the second period, salinity tolerant seeds are introduced, and farmers decide whether to plant salinity tolerant seed or plant the standard seed. This decision is given by the binary indicator $d_t \in \{0, 1\}$. Planting a standard seed is given by $d_t = 0$, where $d_1 = 0$ by default because salinity tolerant seeds are not available in the first period. In the second period, farmers may plant a salinity tolerant seed, denoted by $d_2 = 1$. Seed choice costs $c(d_t)$, where I normalize such that planting a non-salinity tolerant seed is free $c(d_t = 0) = 0$. I assume planting a salinity tolerant seed costs $c(d_t = 1) > 0$, where $c(d_t = 1)$ is positive yet small to capture the notion that salinity tolerant seeds perform relatively better in high salt environments yet relatively worse than standard seeds amid low salinity. Two independent and unchanging environmental conditions denoted by the set $\{S, B\}$ can impact harvest, where S is the soil salinity and B is blast, an important fungus threatening rice. I use lower case letters to denote the true, binary environmental states in these respective domains, given by $s \in \{0, 1\}$, where $s = 0$ denotes low salt levels and $s = 1$ denotes high salt levels, and by $b \in \{0, 1\}$, where $b = 0$ denotes no blast and $b = 1$ denotes the presence of blast. I assume the agricultural production function follows a particular, simple functional form given by Equation 16. The maximization and minimization expressions ensure the binary support of output $y_t \in \{0, 1\}$. Planting salinity tolerant seeds $d_t = 1$ mitigates the damage from soil with high salt content ($s = 1$).

$$y_t(s, b, d_t) = \max\left(\min\left(1 - (s - d_t)^2 - b + \xi, 1\right), 0\right) \tag{16}$$

I assume farmers cannot directly observe the environmental states. As a result, farmers are uncertain about how their decision d_2 impacts their output y_2 entering the second period. This uncertainty is captured by their prior over the states of soil salinity and blast. I use $\widehat{\cdot}$ and lower case letters to denote beliefs, such that $\widehat{s}_t$ denotes a farmer's belief entering period t about the probability that true salinity levels are high $\widehat{s}_t = P(s = 1)$, and $\widehat{b}_t$ denotes a farmer's belief in period t about the probability that blast is present $\widehat{b}_t = P(b = 1)$. I assume that farmers use Bayes' rule to learn about these unobserved environmental conditions by updating using the harvest in period 1. Risk-neutral farmers choose seeds to maximize output in period 2 given their beliefs, as shown in Equation 17.

$$U = \max_{d_2} \mathbb{E}\left[y_2(\widehat{s}_2, \widehat{b}_2, d_2) - c(d_2)\right] \tag{17}$$

Finally, I define the *default* domain to be the environmental domain $E \in \{S, B\}$ corresponding to the most likely threat $\widehat{e} = \max(\widehat{s}, \widehat{b})$.

Learning About the Environment To characterize farmers' posterior beliefs entering the second period, I simply apply Bayes' rule. First, consider the case shown in Equation 18 of updating about the likelihood of high salinity after observing a bad harvest.

$$\widehat{s}_2 = P(s = 1|y_1 = 0) = \frac{P(s_1)P(y_1 = 0|s = 1)}{P(y_1 = 0)} \tag{18}$$

I now derive each of these components. The first term is simply farmers' priors about salinity $P(s_1) = \widehat{s}_1$. The second and third terms can be calculated using the agricultural production function, as shown in Equations 19 and 20, respectively.

$$P(y_1 = 0|s = 1) = \widehat{b}_1 + (1 - \widehat{b}_1)(1 - \rho) \tag{19}$$

$$P(y_1 = 0) = \widehat{b}_1\widehat{s}_1 + \widehat{s}_1(1 - \widehat{b}_1)(1 - \rho) + \widehat{b}_1(1 - \widehat{s}_1)(1 - \rho) + (1 - \widehat{b}_1)(1 - \widehat{s}_1)\rho \tag{20}$$

Substituting these expressions into Equation 18 gives Equation 21.

$$\widehat{s}_2 = \frac{\widehat{s}_1\widehat{b}_1 + \widehat{s}_1(1 - \widehat{b}_1)(1 - \rho)}{\widehat{b}_1\widehat{s}_1 + \widehat{s}_1(1 - \widehat{b}_1)(1 - \rho) + \widehat{b}_1(1 - \widehat{s}_1)(1 - \rho) + (1 - \widehat{b}_1)(1 - \widehat{s}_1)\rho} \tag{21}$$

This expression can be simplified to Equation 22.

$$\widehat{s}_2 = \frac{\widehat{s}_1 - \rho\widehat{s}_1 + \rho\widehat{b}_1\widehat{s}_1}{\widehat{b}_1 + \widehat{s}_1 + \rho + 3\rho\widehat{b}_1\widehat{s}_1 - \widehat{b}_1\widehat{s}_1 - 2\rho\widehat{b}_1 - 2\rho\widehat{s}_1} \tag{22}$$

The difference between posterior and prior beliefs about salinity is then given by Equation 23.

$$\widehat{s}_2 - \widehat{s}_1 = \frac{\widehat{s}_1 - 2\rho\widehat{s}_1 - \widehat{b}_1\widehat{s}_1 + 3\rho\widehat{b}_1\widehat{s}_1 - (\widehat{s}_1)^2 + 2\rho(\widehat{s}_1)^2 + \widehat{b}_1(\widehat{s}_1)^2 - 3\rho\widehat{b}_1(\widehat{s}_1)^2}{\widehat{b}_1 + \widehat{s}_1 + \rho + 3\rho\widehat{b}_1\widehat{s}_1 - \widehat{b}_1\widehat{s}_1 - 2\rho\widehat{b}_1 - 2\rho\widehat{s}_1} \tag{23}$$

The expression for the difference between posterior and prior beliefs about blast is given by the symmetric expression, substituting $\widehat{b}_1$ for $\widehat{s}_1$ and vice versa. Note that the denominators are identical and always positive since it is simply $P(y_1 = 0)$, so I focus exclusively on the numerator.

Due to the bounds of beliefs $0 \leq \widehat{s}_t, \widehat{b}_t \leq 1$, the amount that a farmer can update depends on their prior. To account for this mechanical limitation, the numerator from Equation 23 can be scaled to become a share of the potential scope for belief updating; in other words, by dividing by $|1 - y_t - \widehat{s}_1|$.

$$\frac{\widehat{s}_2 - \widehat{s}_1}{|1 - y_1 - \widehat{s}_1|} \propto \frac{\widehat{s}_1 - 2\rho\widehat{s}_1 - \widehat{b}_1\widehat{s}_1 + 3\rho\widehat{b}_1\widehat{s}_1 - (\widehat{s}_1)^2 + 2\rho(\widehat{s}_1)^2 + \widehat{b}_1(\widehat{s}_1)^2 - 3\rho\widehat{b}_1(\widehat{s}_1)^2}{1 - \widehat{s}_1} \tag{24}$$

I factor out $(1-\widehat{s}_1)$ from the numerator in Equation 25 to arrive at the simplified Equation 26.

$$\frac{\widehat{s}_2-\widehat{s}_1}{|1-y_1-\widehat{s}_1|} \propto \frac{(1-\widehat{s}_1)\widehat{s}_1+(1-\widehat{s}_1)(-2\rho\widehat{s}_1)+(1-\widehat{s}_1)(-\widehat{b}_1\widehat{s}_1)+(1-\widehat{s}_1)(-3\rho\widehat{s}_1\widehat{b}_1)}{1-\widehat{s}_1} \tag{25}$$

$$\frac{\widehat{s}_2-\widehat{s}_1}{|1-y_1-\widehat{s}_1|} \propto \widehat{s}_1-2\rho\widehat{s}_1-\widehat{b}_1\widehat{s}_1-3\rho\widehat{s}_1\widehat{b}_1 \tag{26}$$

Since the expression for blast is symmetric, the expression for the relative difference between posterior and prior beliefs along the two unobserved environmental dimensions is therefore given by Equation 27.

$$\frac{\left(\widehat{b}_2-\widehat{b}_1\right)}{|1-y_1-\widehat{b}_1|}-\frac{(\widehat{s}_2-\widehat{s}_1)}{|1-y_1-\widehat{s}_1|} \propto (\widehat{b}_1-\widehat{s}_1)(1-2\rho) \tag{27}$$

Since $\rho < .5$ by assumption, the sign depends on the term $(\widehat{b}_1-\widehat{s}_1)$. When prior beliefs about the likelihood of high salinity exceed initial beliefs about the chance of blast, then observing low yield leads the farmer to disproportionately increase their beliefs about salinity relative to blast.

Now, consider updating about the likelihood of high salinity after observing a good harvest.

$$\widehat{s}_2 = P(s=1|y_1=1) = \frac{P(s_1)P(y_1=1|s=1)}{P(y_1=1)} \tag{28}$$

Following the same procedure as in the case of a bad harvest, the posterior likelihood of salt levels being high is given by Equation 29.

$$\widehat{s}_2 = \frac{\widehat{s}_1(1-\widehat{b}_1)\rho}{\widehat{s}_1(1-\widehat{b}_1)\rho+\widehat{b}_1(1-\widehat{s}_1)\rho+(1-\widehat{s}_1)(1-\widehat{b}_1)(1-\rho)} \tag{29}$$

$$\widehat{s}_2 = \frac{\widehat{s}_1\rho-\widehat{b}_1\rho\widehat{s}_1}{1-\widehat{b}_1-\widehat{s}_1-\rho+2\widehat{b}_1\rho+2\widehat{b}_1\widehat{s}_1+\widehat{s}_1\rho+\widehat{b}_1\rho-3\widehat{b}_1\rho\widehat{s}_1} \tag{30}$$

Again noting that the denominator is always positive and identical in the case of salinity and blast since it is simply $P(y_1=1)$, I focus on the numerator. The change in salinity beliefs is given by Equation 31.

$$\widehat{s}_2-\widehat{s}_1 \propto (\widehat{s}_1)^2+2\widehat{s}_1\rho-2\widehat{b}_1\rho\widehat{s}_1-\widehat{s}_1+\widehat{b}_1\widehat{s}_1-(\widehat{s}_1)^2\rho-2\widehat{b}_1\rho\widehat{s}_1-2\widehat{b}_1(\widehat{s}_1)^2+3\widehat{b}_1\rho(\widehat{s}_1)^2 \tag{31}$$

Scaling to account for the mechanical restrictions due to the bounds on beliefs gives Equation 33.

$$\frac{\widehat{s}_2-\widehat{s}_1}{|1-y_1-\widehat{s}_1|} \propto \widehat{s}_1+2\rho-2\widehat{b}_1\rho-1+\widehat{b}_1-\widehat{s}_1\rho-2\widehat{b}_1\rho-2\widehat{b}_1\widehat{s}_1+3\widehat{b}_1\rho\widehat{s}_1 \tag{32}$$

As before, the change between posterior and prior beliefs about the likelihood of blast is symmetric, substituting $\widehat{s}_1$ for $\widehat{b}_1$ and vice versa. The relative change is given by Equation

33.

$$\frac{(\widehat{b}_2 - \widehat{b}_1)}{|1 - y_1 - \widehat{b}_1|} - \frac{(\widehat{s}_2 - \widehat{s}_1)}{|1 - y_1 - \widehat{s}_1|} \propto (\widehat{b}_1 - \widehat{s}_1)3\rho \tag{33}$$

As before, the sign of this expression hinges on $\widehat{b}_1 - \widehat{s}_1$. Because this is the case of experiencing a low yield (and thus farmers update their beliefs downward), the relative reduction in likelihood for the default domain is smaller than that in the non-default domain. In other words, a farmer who thinks salinity is more likely than blast going into period 1 will still think that salinity is more likely than blast going into period 2 after observing a good yield, even though both posteriors will be lower than the priors.

Combining the results from Equations 27 and 33 illustrates that after observing either low yield or high yield, the rank ordering of beliefs is preserved. In other words, the default domain exhibits path dependence and will always be the default domain. This proves Remark 1.

F Qualitative Evidence

F.1 Methodology

Main Sample Open-Ended Response Audio Recordings

I focus on audio recordings of four open-ended questions asked to the main sample during the endline survey:

1. What are the biggest issues that you face in cultivating Boro rice?
2. How do you learn about the amount of salt in the soil of your plots?
3. What do you find hard about figuring out how much salt is in your soil? What do you find easy about figuring out how much salt is in your soil?
4. Are you worried or not worried about salinity on your plots? Why or why not?

I transcribe and translate the responses to these questions.

Qualitative Interview Data Collection

I conducted qualitative interviews in six villages in the same upazilas of the main survey, but no farmer in the main sample lived in any of these villages. The sampling frame mirrored the main survey to include all Boro season rice farmers, though within that set, the 40 farmers constituting the qualitative interview sample were recruited by convenience instead of randomization from a listing. Interviews were conducted in June 2023, coinciding with the launch of the endline survey. Most interviews lasted between 20 and 30 minutes and were held outside, typically in common areas of the village. Shahriaz Ahmmed—a lecturer in Development Studies at Khulna University—conducted consecutive translation to facilitate the interviews between myself and the farmers. All interviews were recorded, and the transcriptions of these interviews form the basis of the data used in the coding procedure.

Interviews followed an in-depth narrative approach, which has a long and rich history in social science research (see DeLuca et al. (2016) for more on this method and Bergman et al. (Forthcoming) for a recent example from the economics literature). This method—which

emphasizes a natural, in-depth conversation over a series of questions and short answers—has been shown to allow a wide range of responses to emerge through initial, open-ended questions that lead to targeted follow-up. In the case of this study, the conversations primarily focused on how farmers made agricultural decisions and their related economic concerns. All interviews started and ended with the same pair of questions. I began by asking each farmer, "How did your harvest go?" and concluded by asking "Do you have anything else to share?" In between, the conversation ranged but largely focused on what farmers identified as their biggest agricultural challenges, the choices they make to adapt to these obstacles, and how they gather information to arrive at these decisions.

Qualitative Interview Data Coding Protocols

I took an inductive approach to coding the qualitative data, by first reviewing the corpus of all interviews and noting the aspects of decision-making and learning that emerged as most salient from farmers' own accounts. I then coded all transcripts according to these themes, recording the frequency with which they were mentioned across all farmers in the sample. The following paragraphs provide further detail on the coding protocols for each theme.

Inference by Harvest Outcome: This code describes farmers who use the outcome of their harvest to assess the degree of environmental threats on their land. I code interviews under this theme even if farmers also use other signals to formulate their beliefs, as long as they mention this particular data point explicitly as well.

Low Belief Uncertainty: This code covers farmers who reported who—when describing environmental conditions on their land or the effectiveness of technology—emphasized the certainty of their expectations. For instance, when describing the amount of salt on their soil, a farmer in this classification would be very confident in their assessment.

Unwillingness to Switch from Boro Cultivation: This code captures farmers who—when discussing potential margins of adaptation to environmental threats on their land—stressed their unwillingness to ever stop cultivating rice during the Boro season, even if conditions become particularly dire.

Confidence in Success of Salinity Adaptation Methods: This code applies to farmers who report taking an adaptation step that they say fully corrects the salinity problem after pointing out salinity issues on their land. Note that by construction, this code can only possibly apply to farmers who recount a salt threat to their agricultural productivity. Farmers falling under this category explicitly acknowledge that salinity would harm their output had it not been for some measure they take, which they say has fully eradicated the potential harm.

Reliance on Social Learning: This code includes farmers explicitly pointing to input from others as the major factor in determining their agricultural decisions. I further classify

interviews by the source of this information: neighbors or family members in the same village, or relevant local agricultural figures (agricultural extension officers or seed/fertilizer dealers). Farmers typically mention this in the context of seed, fertilizer, and pesticide decisions.

Demand for Government Assistance: This code covers any farmer who brought up a request for government support. This typically—though not always—was directly related to agricultural production, and in particular often in the form of financial assistance, either directly or through price controls for inputs. Sometimes, this also included demand for information provided by the government. I also classify farmers under this theme if they hold a pessimistic attitude towards government assistance; that is, they say bring up the topic in the context of decrying the government's current activities (for instance, in the case of providing some farmers with free seeds but not others through a opaque or unfair procedure).

F.2 Qualitative Evidence Motivating Modeling Assumptions

This section presents additional quotations motivating the modeling assumptions used in my theoretical framework. I organize the quotes based on the the same points made in the set-up of Section 3.

Farmers cannot directly observe salinity

"How do we understand the amount of salt in the soil? We are not engineers. No, we have no way of knowing whether the soil has salinity or not."

"I find it difficult to estimate the level of salinity in the soil. I don't understand easily because I don't have a machine."

"I don't have a soil test or salinity tester, and I can't tell the salinity of the land by looking at the soil."

"If there is a machine then the matter of measuring salinity is easy. And if there is no salt measuring machine then the matter seems difficult to me."

"It is impossible to understand without testing the salinity. And we can't understand by looking."

" We can't understand if there is salinity in the land where there is grass without testing the machine."

"We do not have the equipment to test. So we find it difficult to estimate the salinity level."

" It is difficult to understand salinity by looking at bare land."

Farmers rely on harvest output and physical plant characteristics to infer salinity

"It seems difficult to estimate the salinity level if the crop is good in the field. I can easily estimate the level of salinity by looking at the appearance of the crop and the color change of the soil."

"When the paddy of my land is red in color, I can easily understand that my land is definitely affected by salt. I don't find anything difficult here. I can tell by looking at the

rice plant that it is affected by salt."

"If the rice plant does not grow when planted in the land, the rice plant turns red and dies, it is easy to understand that there is salinity in the land. It is difficult to understand salinity by looking at bare land."

"When the bunch of paddy plant turns red, we know that there is salinity in the land. Besides, paddy plants become smaller and white salt floats on the ground. By looking at these I basically estimate the salinity of the land."

"We can understand the salinity level of the land by looking at the appearance of the crops on the land. When the leaves of the rice plant turn red, white spots appear on the soil and the rice plants wither. Then it is easy to understand that there is salinity in the land."

"Rice plants do not grow due to salinity, rice plants turn red and rice plants die. Then I realized that the amount of salt in the land is high."

"The amount of salt in the soil can be understood. White salt appears on the soil after planting rice plants when the rice plants die and the soil dries up. Then it is understood that the salt content in the land is high. The amount of salt can be estimated by looking at these two things: soil and crops."

""After planting the rice plant, when the rice plant becomes small, we understand that the salt content in the land is high. And in saline land, the soil becomes excessively muddy and soft. Besides, earthworm infestation is seen in saline land. By looking at these two signs, we basically estimate the amount of salinity."

"I find it difficult to understand salinity by looking at bare land. And when crops are planted in the land and the appearance of the crops, the salinity of the land can be easily understood. By seeing the crop after planting, we can tell whether the land is salinity or not."

The link between these signals and salinity is indirect

"When the leaves of paddy plants turn red in the land, it is easy to understand that there is salinity in the land. And when rice plants are not growing well, it becomes difficult to know whether the problem is due to salinity, soil or fertilizer."

"After planting rice plants in the land, when the rice plants turn brown, it is easy to understand that the salinity of the land has increased. And when the paddy plant gets a little bigger and then if the paddy plant does not grow, it seems difficult to know whether the problem is due to salinity or some other reason."

"If white substance like salt is seen in the soil then it is easy to understand that salinity level is high. And when the paddy plant dies, it becomes difficult."

"After planting rice in the land, when the rice plants do not grow and are small in size, it is easy to understand that there is salinity in the land. But when the paddy leaves turn red, I find it difficult whether the problem is due to salinity or some other reason."

"When the rice turns red it is easy to understand that it is due to salinity. When after planting paddy in the land and applying proper fertilizers it is found that the crop is not growing well and the crop is not being nourished then it is difficult to understand whether it is actually due to salinity or some other reason."

"If a white substance like salt appears in the soil when the soil is dry, it is easy to understand that the salinity level is high. And if a disease occurs after planting paddy in the

land, it is difficult to understand whether the problem is due to salinity or for some other reason."

"If after planting paddy in the land, if the paddy does not grow and the paddy plants are stunted, then it is easy to understand that the land has salinity. But when rice plants turn yellow due to insect attack, it is not easy to understand whether the problem is due to salinity or insect attack."

G Structured Ethics Appendix

Following the guidance of Asiedu et al. (2021), this section presents a discussion of the ethical considerations of the human subjects research conducted as part of this paper. I discuss the two experiments involved in this project separately: RCT #1, the randomization of saline tolerant seed prices as part of the BDM procedure of the baseline, and RCT #2, the random provision of information about soil salinity during the endline survey.

Policy Equipoise *Is there policy equipoise? That is, is there uncertainty regarding participants' net benefits from each arm of the study relative to the other arms and to the best possible policy to which participants could have access? If not, ethical randomization requires two conditions related to scarcity: (1) Was there scarcity, i.e., did the inclusion of multiple arms change the expected aggregate value of the programs delivered? (2) Do all ex-ante identifiable participants have equal moral or legal claims to the scarce programs?*

For RCT #1, policymakers hold genuine uncertainty as to how effective these saline-tolerant seeds can be in real-world settings, and in particular, how the productivity may vary with saline level. The scale of this experiment provides substantial new evidence outside of controlled lab settings or example plots to assess the relative efficacy in a vulnerable population preserving their agricultural techniques and decision-making for other inputs.

For RCT #2, while there can be little doubt that providing more information to farmers is weakly beneficial, it remains an open policy question as to the magnitude of these benefits and the format of delivery. The cost-benefit calculation behind a policy to collect soil salinity data and provide that information to farmers hinges on the extent to which that information might lead to meaningful behavior change and whether farmer-specific soil measurement is necessary for achieving this impact or whether averages at higher geographic levels can also be effective. This experimental design targets both parameters.

Role of Researchers with Respect to Implementation *Are researchers "active" researchers, i.e. did the researchers have direct decision making power over whether and how to implement the program? If YES, what was the disclosure to participants and informed consent process for participation in the program? Providing IRB approval details may be sufficient but further clarification of any important issues should be discussed here. If NO, i.e., implementation was separate, explain the separation.*

The researcher is considered "active" and had direct decision-making power over all aspects of the experimental design for both RCT #1 and RCT #2. Informed consent was elicited from all participants, as approved by Harvard's IRB board.

Potential Harms to Participants or Non-Participants from the Interventions or Policies *Does the intervention, policy or product being studied pose potential harm to participants or non-participants? Related, are participants or likely affected non-participants particularly vulnerable? Also related, are participants' access to future services or policies changed because of participation in the study? If yes to any of the above, what is being done to mitigate such risks?*

For RCT #1, one potential concern is that the specific seed variety which I provide to farmers as part of the BDM procedure could cause harm to farmers' output. Several factors help alleviate this worry. First, the government is actively trying to distribute these same seed varieties in many of the areas I survey, often by giving out 10 kilogram packets for free. Due to scarce government resources however, the scope of this effort is significantly limited. In one sense, this experiment can be viewed as an extension of an existing government program, albeit at a smaller scale. Second, the enumerator instructions deliberately avoid making any normative statements about whether the farmer should use the seed. Instead, they simply report that "this seed has been designed to grow well even if the soil has a lot of salt in it", echoing the description used by the government's Agricultural Information Service when describing this variety. Third, the WTP elicitation for the seed comes after the BDM mechanisms for a pen, a plate, and a store voucher. To the extent that these standard goods do not have a clearly interpretable normative position, farmers were primed to treating the exercise as a true willingness-to-pay. Fourth, while BRRI 67 grows well in high salinity environments, it grows no worse in low salinity environments. Therefore, even if farmers are overestimating the amount of salt in their soil, their harvest will not be especially negatively impacted if they use the seed. However, it should be noted that other seeds that farmers may use instead in the case of a low saline soil could have other appealing attributes, such as taste or relatively higher yield, and thus harvest may still be lower compared to the counterfactual seed choice. Finally, by only providing one kilogram of seeds, I further help to mitigate this risk, as farmers would likely diversify to other seeds if they had not been planning on using BRRI 67 already.

For RCT #2, one potential concern is that I might accidentally provide erroneous information about soil salinity to respondents. This is particularly relevant given the surprising results of the handheld EC-meters showing low levels of salinity: if salinity is in fact higher than I measure, then I may inadvertently mislead farmers, which could result in worse decisions in the future. I take several steps in an attempt to mitigate this risk. First, I take considerable steps to ensure that the soil salinity measures accurately capture the amount of salt on each farmer's plot. These steps—detailed further in section B.2—include repeated measurements and validation using a controlled lab test. Second, when providing the information to farmers, enumerators emphasize the generalizability of the measurement, saying: "These numbers are based off of the measurements we took this past season. The amount of salt in the soil can change, so it may be different next year." Third, the records I obtain from the government's soil measurements (noting that they only collect data from 10 plots) match the same patterns I find, indicating that indeed, the 2022-23 season exhibited strikingly low salinity levels.

Potential Harms to Research Participants or Research Staff from Data Collection or Research Protocols *Are data collection and/or research procedures adherent to privacy, confidentiality, risk-management, and informed consent protocols with regard to human subjects? Are they respectful of community norms, e.g., community consent not merely individual consent, when appropriate? Are there potential harms to research staff from conducting the data collection that are beyond "normal" risks?*

Yes. This research received IRB approval from Harvard University. Data collection posed no unusual risks to research staff.

Financial and Reputational Conflicts of Interest *Do any of the researchers have financial conflicts of interest with regard to the results of the research? Do any of the researchers have potential reputational conflicts of interest?*

No.

Intellectual Freedom *Were there any contractual limitations on the ability of the researchers to report the results of the study? If so, what were those restrictions, and who were they from?*

No.

Feedback to Participants or Communities

Is there a plan for providing feedback on research results to participants or communities? If yes, what is the plan? If not, why not?

Yes. The results of both experiments and the trends in salinity over the 2022-23 season were shared with participants prior to the 2023-24 Boro planting. I consulted with G. M. Mustafizur Rahman of the Soil Research Development Institute to ensure that the information I provided farmers was consistent with the best available guidance from the government. The text of the script read by enumerators is as follows: "The amount of soil salinity can change from year to year. Last year, the amount of soil salinity in Khulna division was much lower than it had been in previous years. It is difficult to know whether it will be low again in the future. But it is clear that last year was low relative to what it had been in the past. Because last year's salinity levels were so low, the farmers who planted more of their land with salinity-tolerant seeds actually earned less money than the farmers who planted different seeds. If the amount of salinity on the soil had been higher, then it might have been better to plant the salinity tolerant seeds. But because it was low, that year, those seeds were not as profitable. It is hard to know what the soil salinity will be next year. In general, when soil salinity is very high, it is better to plant salinity tolerant seeds, and when soil salinity is very low, it is better to plant a different kind of seed. In coastal areas, the government recommends planting salinity tolerant seeds."

Foreseeable Misuse of Research Results *Is there a foreseeable and plausible risk that the results of the research will be misused and/or deliberately misinterpreted by interested parties to the detriment of other interested parties? If yes, please explain any efforts to mitigate such risk.*

No.

Other Ethics Issues to Discuss *Are there any other issues to discuss?*
No.