INTERNATIONAL
FOOD POLICY
RESEARCH
INSTITUTE

IFPRI

IFPRI Discussion Paper 01591

December 2016

The Role of Learning in Technology Adoption

Evidence on Hybrid Rice Adoption in Bihar, India

Jared Gars

Patrick S. Ward

Environment and Production Technology Division

INTERNATIONAL FOOD POLICY RESEARCH INSTITUTE

The International Food Policy Research Institute (IFPRI), established in 1975, provides evidence-based policy solutions to sustainably end hunger and malnutrition, and reduce poverty. The institute conducts research, communicates results, optimizes partnerships, and builds capacity to ensure sustainable food production, promote healthy food systems, improve markets and trade, transform agriculture, build resilience, and strengthen institutions and governance. Gender is considered in all of the institute's work. IFPRI collaborates with partners around the world, including development implementers, public institutions, the private sector, and farmers' organizations, to ensure that local, national, regional, and global food policies are based on evidence.

AUTHORS

Jared Gars (gars@wisc.edu) is a Ph.D. candidate in the Agricultural and Applied Economics Department at the University of Wisconsin, Madison, US.

Patrick S. Ward is a research fellow in the Environment and Production Technology Division of the International Food Policy Research Institute, Washington, DC.

Notices

[1.] IFPRI Discussion Papers contain preliminary material and research results and are circulated in order to stimulate discussion and critical comment. They have not been subject to a formal external review via IFPRI's Publications Review Committee. Any opinions stated herein are those of the author(s) and are not necessarily representative of or endorsed by the International Food Policy Research Institute.

[2.] The boundaries and names shown and the designations used on the map(s) herein do not imply official endorsement or acceptance by the International Food Policy Research Institute (IFPRI) or its partners and contributors.

Contents

Tables

Figures

ABSTRACT

Much empirical research has shown that individuals' decisions to adopt a new technology are the result of learning–both through personal experimentation through observing the experimentation of others. Yet even casual observation would suggest significant heterogeneity of learning processes, manifesting itself in widely varying patterns of adoption over space and time. This paper explores this heterogeneity in the context of early adoption of hybrid rice in rural India. Using specially designed experiments conducted as part of a primary survey in the field, we identify which of four broad learning heuristics most accurately reflects individuals' information processing strategies. Linking these learning heuristics with observed use of rice hybrids, we demonstrate that pure Bayesian learning is well suited for the tinkering and marginal adjustments that are required to learn about a technology like hybrid rice, but it is also more cognitively taxing than other learning styles requiring a longer memory and more complex updating processes. Consequently, only about 25 percent of the farmers in our sample can be characterized as pure Bayesian learners. Present-biased learning and relying on first impressions will likely hinder adoption of a technology like hybrid rice, even after controlling for access to credit and a rudimentary proxy for intelligence.

Keywords: learning heuristics, experimental economics, technology adoption, India

ACKNOWLEDGMENTS

This work was supported by the Cereal Systems Initiative for South Asia (CSISA), with funding provided by the Bill and Melinda Gates Foundation and the United States Agency for International Development. We have benefited from fruitful conversations with David Spielman and Laura Schechter. We also wish to acknowledge the assistance of Vartika Singh and Sourav Kumar, without whom data collection would not have been possible.

This paper was undertaken as a part of the CGIAR Research Program on Policies, Institutions, and Markets (PIM), which is led by IFPRI and funded by CGIAR Fund Donors. This paper has not gone through IFPRI's standard peer-review process. The opinions expressed here belong to the authors and do not necessarily reflect those of PIM, IFPRI, CGIAR.

1. INTRODUCTION

In many parts of the developing world, the transition from indigenous practices to modern technologies is often viewed as a critical step toward achieving broad agricultural development objectives such as food security or self-sufficiency. Experiences from the Asian Green Revolution from the mid-1960s through the end of the 1990s provide an almost textbook illustration of this process (for example Gollin et al. 2005). During its early years, most of the Green Revolution was concerned with the transition from traditional varieties and landraces to modern, high-yielding varieties, specifically the adoption of semidwarf varieties of rice and wheat. From 1965 to 1970, estimates suggest that the cultivation of modern varieties of rice and wheat in South Asia increased from an essentially negligible baseline to 10 percent and 39, respectively, percent of harvested area (Gollin et al. 2005). In India, as in much of South Asia, the adoption of these modern varieties provided the potential for increased yields, though arguably these new varieties did not reach their full yield potential until they were paired with complementary inputs, such as irrigation and fertilizers. Indeed, the diffusion of these modern varieties in many ways propelled the adoption of fertilizers and irrigation (Morris and Byerlee 1998; Gollin et al. 2005). But although the expansion of irrigation facilities and the increased use of chemical fertilizers and pesticides share some of the credit for the massive gains in food grain production during this period, such gains arguably might never have been possible without the widespread adoption of high-yielding rice and wheat varieties that were particularly responsive to these complementary inputs (for example Morris and Byerlee 1998).

This process of varietal modernization has been inconsistent over both space and time. Most of the benefits of these modern varieties were realized in the northwestern states of Punjab, Haryana, and western Uttar Pradesh, who had larger and more egalitarian farm structures and generally greater access to irrigation, in some cases due to a more favorable policy environment. In other parts of India–particularly eastern Indian states such as Bihar, Odisha, and West Bengal–varietal modernization has been considerably slower. Furthermore, while there was initially quite rapid adoption of these modern varieties, the cumulative level of adoption, especially for rice, remains incomplete to this day. And even where farmers have made the transition from traditional varieties or landraces to modern varieties (what Morris and Byerlee 1998 refer to as Type A varietal change), there has not necessarily been a subsequent transition from first generation modern varieties to newer modern varieties (what Morris and Byerlee 1998 refer to as Type B varietal change). This latter form of varietal modernization has been shown to be particularly important, as many of the genetic advantages conferred by the breeding efforts deteriorate over time, resulting in reduced productivity and increased susceptibility to various stresses (Krishna et al. 2016). Furthermore, more rapid varietal turnover among modern varieties allows farmers access to technological enhancements embodied in newer germplasm. As the private seed sector continues to develop within India, this varietal turnover will likely take the form of switching from inbred varieties to hybrids, which offer higher yields and lower seed rates, though at the expense of farmers' being able to save seed from year to year.

It has often been observed that the adoption of new technologies, including new cultivars, is a gradual process. When plotted against time, the cumulative proportion of a population adopting a given technology generally follows a sigmoid (S-shaped) trend, wherein the most rapid adoption (that is, at an exponential growth rate) occurs at some nontrivial time after the initial introduction of a technology (Griliches 1957; Rogers 1995). The shape of the diffusion curve implies a great deal of heterogeneity, which manifests itself through differentials in the timing of adoption among farmers in the population (for example, some farmers are innovators, some are laggards, and so on). If one considers differences in the shape and the location of such diffusion curves across subgroups in the population (for example, across different states), there emerges an even greater degree of cross-sectional heterogeneity, not just in terms of differences in the timing of adoption (that is, what Griliches 1957 refers to as the "date of origin") in these different subgroups, but also the pace of adoption (that is, the "rate of acceptance") and the timing and cumulative level of steady-state or equilibrium adoption rates.

What factors underlie the observed heterogeneity in the adoption of new technologies? Clearly, differences in resource constraints including access to credit, labor availability, and landownership affect individual adoption decisions. Additionally, various studies also attribute this heterogeneity to idiosyncratic differences in farmers' risk preferences (for example Feder 1980; Feder et al. 1985). When considering a new technology, farmers are confronted with considerable uncertainty, not only because there is risk associated with the yield or profitability of the technology, but also because the nature of this risk (that is, the underlying distribution of yields or potential farm profits) is unknown. The decision to transition from a traditional technology to a new technology requires some consideration of the relative benefits of the two.[1] Under this view, therefore, there must be some process or mechanism by which expectations of the long-term value of the technology are uncovered or inferred. This, in turn, implies that a precondition for technology adoption is that the potential adopter has, in some fashion or another, learned about the potential benefits of the new technology, compared these benefits with those of the traditional technology, and considered the cost of transitioning from the traditional to the modern technology.

In the context of modern seed varieties during the Green Revolution, Foster and Rosenzweig (1995) demonstrated the importance of learning, both from one's own experimentation and from in observing the experimentation of others. They suggested that imperfect or suboptimal input usage is the result of imperfect subjective beliefs, which are subsequently improved through Bayesian updating with increased experience with the new varieties. The profitability of modern varieties is increasing in farmers' own and their neighbors' experience with the modern varieties, but farmers with relatively richer neighbors are more likely to delay adoption and observe the costly experimentation of their neighbors. Foster and Rosenzweig (1995) suggested that these neighbor effects result in free riding on the experiences of more capitalized farmers with respect to modern seed varieties in India. In their study of Mozambican farmers, Bandiera and Rasul (2006) found an inverse U-shaped propensity to adopt relative to the numbers of family members and friends who adopt. There is an increasing network effect when there are relatively few adopters in the network, but as the number of adopters increases, the propensity to adopt declines, an effect attributed to strategic delay. In both cases, there is apparently a conflict between the uncovering of additional information on the benefits of the new technology and the temptation to delay one's own adoption until the full distribution of potential benefits is realized: more information is helpful in providing optimal input information, but having multiple observations potentially allows farmers to wait and observe heterogeneous outcomes before experimenting on their own.

Heterogeneity in individual learning may influence how farmers weigh information from others thus encouraging or inhibiting social learning. Conley and Udry (2010) considered the issue of information quality in social learning and allowed for a more flexible learning model that does not force farmers to learn an entire production function. Rather, according to the authors, famers engage in local learning, or learning just about the relevant outcomes of the production function for the level of inputs applied. Munshi (2004) exploited heterogeneity in growing conditions among rice farmers during the Green Revolution to show that rice farmers were more likely to experiment than wheat farmers because the quality of social information among wheat farmers was considerably lower. In contrast, wheat farmers responded strongly to neighbors' experiences as well as targeted extension efforts with groups of contact farmers. Social learning has the capacity to overcome the challenges farmers face in making adoption decisions, though there remain groups of farmers on both the intensive and extensive margin who rely on learning by doing and individual experimentation when making adoption and intensity decisions. Additionally, longer and more heterogeneous processes of diversification and adaptation to climate change may limit the scope for farmers to learn from their neighbors, forcing many farmers to rely on individual learning and experimentation.

There is a considerable literature addressing how individuals process information and update

[1] Traditionally, it was thought that such decisions were made on the basis of net present value. Carey and Zilberman (2002) have suggested that one source of heterogeneity in technology adoption may arise due to differences in the option value ascribed to postponing investment in the new technology (that is, differences in the critical value above which the expected long-term value of adoption must exceed the initial investment).

beliefs in repeated decision making and noncooperative games, as well as decades of research on the role of learning by doing and learning from others within one's social network in technology adoption. However, with the exception of Barham et al. (2015), there has been little work analyzing how individual learning heuristics vary across actual decision makers and how they affect adoption decisions. This study attempts to uncover some of the broad heuristics through which individuals form and update expectations about conditional probability distributions that facilitate or constrain the process of technology adoption in rural India. It utilizes field experiments to elicit and characterize individual learning styles, and combines these findings with risk and uncertainty preferences and observed adoption decisions to test whether heterogeneity in learning acts as a significant barrier to adoption. The study considers a set of simple individual learning models outlined in Gans et al. (2007), which are used to analyze the role of learning in adoption decisions as introduced by Barham et al. (2015). While there are potentially as many unique information processing strategies as there are heterogeneous agents being studied, these strategies may be classified according to a small group of key heuristics. These learning rules include strict Bayesian learning as well as more simplified processes that vary over the time horizon that people use to calculate expectations (first impressions versus short memory) and over classifications of outcomes into "good" and "bad" draws. Although Bayesian learning is widely assumed in economic research as a model for updating expectations, the other learning rules have legacies in the psychological literature. These non-Bayesian processes are less complicated heuristics of belief formation that may be more appropriate behavioral approximations of subjective belief updating in actual decision making. Previous research on belief-updating heuristics have found that subjects in the lab, including farmers in the United States, exhibit substantial heterogeneity in how they update beliefs, and that this heterogeneity is partially explained by observable levels of education and cognitive ability (Cheung and Friedman 1997; Camerer and Ho 1999; Gans et al. 2007; Barham et al. 2015). Additionally, Barham et al. (2015) found that Bayesian learning performed poorly among Wisconsin and Minnesota farmers, and that farmers who formed strong beliefs were slower adopters of new technologies, namely genetically modified maize and soybeans. One would expect Bayesian updating to be an even poorer model of learning among relatively less educated populations in rural areas in developing countries.

This present study builds upon this previous research using field experiments with farmers in rural Bihar, India. Ignorance about returns on inputs or optimal management of technologies has been argued to be a pertinent cause of underinvestment in inputs and new technologies in developing countries (Foster and Rosenzweig 2010). As previously mentioned, empirical evidence from India suggests that more highly educated farmers are "better" learners and may be more efficient at processing information due to their ability to do so across multiple dimensions of technology usage (Foster and Rosenzweig 1995; Hanna et al. 2014). The state of Bihar, India, has experienced some of the highest economic growth rates in India over the past decade but remains one of the poorest states in India. The lack of long-term growth has been attributed to the state's relatively low human capital stock alongside large variations in aggregate total factor productivity due to poor agricultural productivity. Chanda (2011) estimated average years of schooling among adults in Bihar to be 3.5 years, with nearly half of the adult population having no schooling.

In an effort to shed light on the impact of learning on technology adoption in rural Bihar, our empirical tests take advantage of a multiyear panel dataset of 576 farmers in rural India. During the initial round of data collection in 2013, we elicited farmers' risk and uncertainty preferences using lottery-based experiments, in which farmers were asked to make a series of choices between a lottery and a riskless payout. The results of this experiment were then used to estimate learning behavior based on a learning experiment conducted during the second round of data collection in 2014. In this learning experiment, farmers again made a series of decisions between a lottery and a riskless payout, but each subsequent round revealed more information about the underlying distribution of risk in the lottery. Farmers' actual payoffs in the lottery were based on the true distribution of beads in a bag, but their per-round expected utility payoffs were based on their updated beliefs as more information was revealed. We collected plot-level inputs and yields in both survey rounds, as well as individual characteristics, including age, caste, and literacy, that potentially affect the learning and technology adoption process. Results from the learning experiments reveal considerable heterogeneity in learning across farmers, assuming both risk neutrality and risk aversion. In the case of risk neutrality, we find that Bayesian learning is not an inappropriate model of belief updating for many farmers, but overweighting early events is a better approximation of their learning style, regardless of whether risk neutrality or risk aversion are assumed. We also find evidence that caste and cognitive measures explain some of the variation in learning rules, though much of the variation remains unexplained by observable characteristics. Finally, we link individual learning rules with adoption behavior and find that farmers who weight only recent information are less likely to be early adopters of new technology.

2. REVIEW OF PREVIOUS LITERATURE ON LEARNING STYLES

Farmers' adoption decisions rely on their processing information about the productivity or profitability of various technologies, updating their beliefs given this new information, and subsequently making choices based on their posterior, subjective distributions. Identifying sources of heterogeneity in the intensity and timing of the adoption of productive technologies has been the focus of empirical research for decades (Griliches 1957; Feder 1980; Foster and Rosenzweig 1995). The technology adoption literature has typically assumed Bayesian learning because it is empirically tractable and theoretically consistent (Foster and Rosenzweig 1995; Conley and Udry 2010). However, this is clearly an unrealistic assumption in almost any real-world context. Voluminous research suggests that individuals simplify otherwise complex cognitive tasks (for example Tversky and Kahneman 1974), and furthermore that there is substantial heterogeneity in the learning rules or heuristics that people employ, which may affect the adoption decision and interact with other characteristics (that is risk preferences) to encourage or prevent adoption (Gans et al. 2007; Barham et al. 2015). The literature on learning in noncooperative games has adapted a variety of models of learning and belief updating beliefs that may be more relevant in contexts with low levels of human capital (Cheung and Friedman 1997; Camerer and Ho 1999).

Differences in non-Bayesian updating or learning are the result of how people process or weight the information at their disposal. Belief-learning models characterize how players update beliefs and make decisions given their subjective expected distribution based on their history of observed outcomes. Two common belief-learning models are Cournot learning (short memory) and fictitious play (long memory). Cheung and Friedman (1997) developed a general one-parameter class of learning rules (for weighted fictitious play) that nests Cournot and fictitious play as special cases, with a range between them of adaptive learning rules whereby all observations may affect the expected state but the weight given to more recent information varies with the parameter. Importantly, Cheung and Friedman (1997) found that players exhibit a range of learning styles and are more likely to weight recent information in more informative environments. Camerer and Ho (1999) introduced a hybrid model of learning that includes aspects of both belief learning and reinforcement, or rote, learning. Their experience-weighted attraction model "wraps a parametric skin" around belief and reinforcement learning as boundaries of the parameter space. However, this model may have little relevance in purely individual games in which payoffs from other strategies are stochastic and not well known.

The learning models described above belong to the set of individual learning models, or what may be considered behavioral models of "learning by doing." Other prominent individual learning models include reinforcement learning and individual evolutionary learning, which have applications in a variety of settings but may be less relevant for technology adoption (Erev and Roth 1998; Arifovic and Ledyard 2011). The present study seeks to identify heterogeneous learning styles and determine their effect on technology adoption decisions in which farmers make choices over multiple technologies that have unknown yield or profit distributions. The farmers' problem, then, is similar to two-armed or multi-armed bandit problems, in which farmers are optimizing their decisions while simultaneously improving their information. The farmers face the inherent trade-off between experimentation (that is, on a temporary and reversible basis) and adoption (that is, on a more permanent and irreversible basis). Results from multi-armed bandit experiments provide evidence that people diverge from Bayesian behavior in many situations. Meyer and Shi (1995) found that players tend to be myopic in their updating and to make less-than-optimal decisions due to this type of updating. Anderson (2001) attributed less-than-optimal experimentation to risk aversion. Gans et al. (2007) evaluated multiple simple updating rules relative to the Gittins index, finding that people show substantial heterogeneity in their updating rules and that they outperform the optimal decisions given the environment.

Barham et al. (2015) developed a learning experiment in which subjects chose between a sure and a risky payoff as they obtained more information about the underlying probability distribution of "winning" and "losing" draws in a risky lottery, after a series of draws with replacement. After evaluating the learning rules used by Gans et al. (2007) and find evidence that farmers in their sample tended to be most influenced by recent draws. The authors concluded, however, that there is heterogeneity in learning rules, given the fact that some farmers acted in a way that was closer to the type of behavior that would be predicted by strict Bayesian updating. Farmers with higher levels of education were more likely to use more sophisticated learning rules, suggesting that increased human capital may diminish the impacts of suboptimal learning. Integrating their experimental results with Wisconsin farmers' recollection of the timing of their adoption of genetically modified maize and soybeans, Barham and colleagues (2015) found that farmers who develop strong beliefs (such as being sensitive to first impressions) are slower adopters because they neglect continual learning.

3. EMPIRICAL STRATEGY

Data and Sampling Methodology

The data used in this study come from laboratory-in-the-field experiments conducted as part of a larger, longitudinal data collection effort that gathered, among other things, data on household structure and household member characteristics; seasonal data on land utilization and season-, plot-, and variety-level data on rice production and input use. We employed a multistage sampling strategy. In the first stage, we selected 3 districts heavily dependent upon rice production but also displaying significant heterogeneity in terms of agroecological conditions. These districts were Bhojpur (in west-central Bihar), Nawada (in south Bihar, bordering Jharkhand) and Madhubani (in north Bihar, bordering Nepal). Figure 3.1 illustrates the geographic location of each of these districts. In the second stage, we selected 16 high-rice-producing blocks (sub-district administrative units) across the 3 districts, where the share of blocks drawn from each district was proportional to that district's share in overall rice production among the 3 selected districts. Within each of these 16 blocks, we randomly selected 2 villages. Finally, from each of these 36 villages, we randomly selected 18 rice-growing households.

Figure 3.1 Location of sample districts in Bihar, India

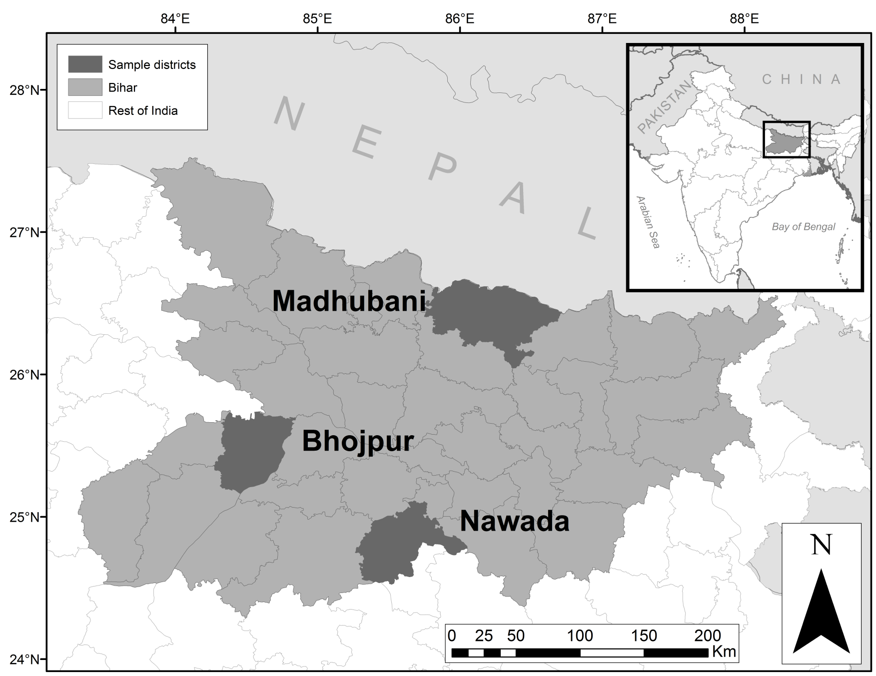

Source: Authors.

The first round of data collection commenced in April 2013, prior to the kharif (monsoon) season. This survey collected detailed information on agricultural production and input use for kharif 2012, along with risk and uncertainty preferences elicited through an experimental approach similar to the one used by Barham et al. (2014). In addition, the survey gathered information on characteristics of the household head, such as caste, level of education, age, and access to credit. During the follow-up survey in April 2014, we gathered information on input usage during kharif 2013 and planned input usage in kharif 2014, which, at the time, was roughly two months away. Input and production data were plot specific for up to five of a given farmer's plots on which rice was cultivated. The inputs covered in the survey included irrigation, fertilizers, and seed, with irrigation and fertilizer information collected for each application throughout the season. Plot-level characteristics included slope, soil type, erosion, and contract type. Seed varieties were delineated by name and whether they were local, hybrid, or high-yielding varieties. As a result of the two rounds of data collection, we have plot- and time-specific application amounts (planned for kharif 2014 and used during kharif 2013 and 2012). During the follow-up survey in 2014, we also conducted a novel experiment to identify the heuristic that would most closely characterize the way in which individuals in the sample processed information and refined beliefs.

Experimental Procedure

Before describing the experimental procedure for eliciting individual learning rules, we will first briefly introduce the experiments that were used to elicit risk and uncertainty preferences, because these preferences play an instrumental role in the identification of learning rules (Ward and Singh 2015). The identification of risk and uncertainty preferences involved two related experiments conducted as part of the first round of data collection in 2013. In each experiment, farmers were presented with a series of 11 choices between a riskless option and a risky option, which took the form of a simple lottery. The riskless option was always a guaranteed payment of 20 Indian rupees (Rs). The lotteries consisted of a "winning" draw paying Rs 40 and a "losing" draw paying some lesser amount. The amount of the "losing" draw monotonically decreased with each subsequent choice. Because this "losing" draw was monotonically decreasing, we might expect respondents' preferences between the risky and riskless options to change at most one time. Following Tanaka et al. (2010), we enforce monotonic preference switching by asking respondents only about the particular choice at which they would switch from preferring the risky option to preferring the riskless option.

Where the risk and the uncertainty experiments differed was in the information provided to the respondent about the probabilities associated with the "winning" and "losing" draws in the lottery. In the uncertainty experiment, farmers were asked to choose between the risky and riskless options without being provided any information on these probabilities, so that they had to rely instead on naïve subjective expectations. In the subsequent risk experiment, farmers were told that the odds of "winning" and "losing" the lottery were 50:50. The uncertainty experiment was completed first to ensure that the farmers did not base their beliefs about the distribution of "winning" and "losing" draws on the probabilities revealed in the risk experiment. To communicate this probability, respondents were shown a bag containing 10 chips, numbered 1 through 10. Chips 1 through 5 were considered to be "winning" draws, while chips 6 through 10 were considered "losing" draws.

Assuming that preferences exhibit constant relative risk aversion and constant relative uncertainty aversion (CRRA and CRUA, respectively) with isoelastic utility function $u(c) = \frac{c^{1-\rho}}{1-\rho}$ for $\rho \neq 1$ and $u(c) = \ln c$ for $\rho = 1$, in which ρ is the CRxA coefficient, then the payout at which the respondent switches from preferring the risky option to preferring the riskless option allows us to estimate an interval of possible aversion risk and uncertainty aversion coefficients. While we can identify CRxA coefficients within an interval, in the ensuing analysis we make the simplifying assumption that an individual's risk or uncertainty aversion coefficient takes the value of the upper limit of the interval within which the individual's choices would be consistent.

We attempted to make the choices in these experiments incentive compatible by paying respondents the specified amount for a random choice, depending upon their preference for the risky or riskless option in that particular situation, and if they preferred the risky option, allowing them to make a random selection of a chip from the bag containing 10 chips.[2] Table 3.1 illustrates the series of decisions that participants were asked to make.

---

[2]In reality, there were three other similar experiments that were conducted during the same interview, so the potential payout was not limited to these two experiments.

Table 3.1 General structure of risk and uncertainty experiments and corresponding coefficients of risk and uncertainty aversion for given switching decision

Decision	Riskless payout	Risky option "Winning" draw	Risky option "Losing" draw	Interval of CR*x*A coefficient for switching from risky to riskless
1	Rs 20	Rs 40	Rs 20	$(3.76, \infty)$
2	Rs 20	Rs 40	Rs 20	$(1.86, 3.76]$
3	Rs 20	Rs 40	Rs 20	$(1, 1.86]$
4	Rs 20	Rs 40	Rs 20	$(0.65, 1]$
5	Rs 20	Rs 40	Rs 20	$(0.52, 0.65]$
6	Rs 20	Rs 40	Rs 20	$(0.40, 0.52]$
7	Rs 20	Rs 40	Rs 20	$(0.31, 0.4]$
8	Rs 20	Rs 40	Rs 20	$(0.22, 0.31]$
9	Rs 20	Rs 40	Rs 20	$(0.09, 0.22]$
10	Rs 20	Rs 40	Rs 20	$(0.0, 0.09]$
11	Rs 20	Rs 40	Rs 20	$(-\infty, 0.0]$

Source: Authors.
Notes: The interval of CRxA coefficients in the last column was not shown to participants. CRxA = constant relative risk aversion or constant relative uncertainty aversion; Rs = Indian rupees.

The learning experiment was designed in such a way that we are able to identify the most likely heuristic by which individuals form and update beliefs. In this particular experiment, we learned about the rules farmers employed to update their beliefs about the distribution of green and blue beads in a bag containing a total of 100 beads. Before commencing with the actual experiment, farmers were given instructions on its overall structure, goals, and rules, as well as information on the real financial implications of their decisions. There was also a practice round in which the farmers had the opportunity to practice counting beads and making choices between risky and riskless options (described below). If they chose the risky option during the practice round, the enumerator asked them what their compensation would be if this had been a decision with real financial implications depending upon their draw of a bead from the bag. Even if they chose the risky option, they were subsequently asked to explain what their compensation would have been if they had chosen the riskless option instead, and whether this decision had real financial implications. The same procedure took place if they initially chose the riskless option. If they made any mistakes in understanding the process or compensation, the enumerator explained the rules again, and with the same beads they had already drawn, asked them to go through the process again. Finally, the enumerator recorded how well the respondents understood the experiment.[3]

The farmers were not aware of the actual distribution of green and blue beads prior to the commencement of the experiment. In each round, the farmers drew five beads at random from the bag. After each draw, the number of blue and green beads was recorded on a laminated experiment sheet. The farmers were allowed to contemplate the outcome of a draw before being asked to choose between a riskless option and a risky option for the next draw. As with the risk and uncertainty experiments, the riskless option in the learning experiment consisted of a guaranteed payment of Rs 20. The stakes of the risky option were specific to the individual: the high payout (with probability equal to the number of blue beads in the bag) was always Rs 40, but the low payout (with probability equal to the number of green beads in the bag) was determined by each farmer's switching value in the uncertainty experiment described above. Because they did not know the true number of blue and green beads in the bag, their expected payoffs in these situations were determined by their beliefs about the distribution of green and blue beads, based on their current and any previous observations of the draws. After farmers made their choice between the risky and riskless alternatives, the selection was recorded on the same laminated sheet as the number of blue and green beads from previous draws.

[3]Clearly, the practice round provided some information about the number of beads in the bag, and therefore the learning rules are calculated by including this information as the first round in the updating of subjective probabilities.

This process continued in the same fashion for 14 rounds. After completion of the 14th round, the farmers were asked to state their belief about the actual number of blue beads in the bag and were informed that they would be rewarded an extra Rs 5 if they were within 2 of the correct answer (72). After making this guess, the true number of beads was revealed, and farmers were asked one more time to choose between the risky and riskless options. To reduce hypothetical bias, farmers were informed prior to commencing with the experiment that they would be compensated based on one of their choices across the decisions. Their payment was based on a decision that was selected randomly after all decisions were made. If they chose the riskless option in that decision round, they would receive Rs 20 rupees. If they chose the risky option in that decision round, they were then asked to draw a single bead at random from the bag of blue and green beads to determine their payout from the lottery.

Decision Model

Assuming a random utility model, choices reflect utility-maximizing behavior with an additive random component. For each choice, agents assess the expected utility difference between choosing the risky option (lottery) and the sure payout using their subjective probabilities to infer the value of the risky option. Assuming that the random component for each choice is independently and identically distributed following an extreme-value type 1 (Gumbel) distribution, the choice model in each period reduces to a logit:

$$P(\text{Choose risky option}) = \frac{e^{(z)}}{1+e^{(z)}}, \tag{1}$$

where z is the difference in expected utility between choosing the risky option and choosing the riskless option.

Consider the case in which an individual is faced with a choice between a sure payment of Rs 20 and a risky prospect that pays Rs 40 if a blue bead is drawn and Rs 10 if a green bead is drawn. The expected utility difference between the risky prospect and the sure payment in round t is expressed as $z(t) = S(t)u(40) + (1 - S(t))u(10) - u(20)$, where $S(t)$ is the individual's subjective probability of drawing a blue bead in round t and $u(t)$ is the utility function, with isoelastic functional form exhibiting CRRA. We estimate the likelihood function using the risk aversion coefficient from the risk experiments defined previously.

Learning Rules

Individual learning processes are used to inform the farmers subjective value for $S(t)$, where each learning process corresponds to a weighting function that defines the weight of the ith draw of t total draws. Then the player's latent and intrinsic weighting will define the learning rule for each set of choices over t draws. While the nature of the experiment and environment are context specific, the literature on learning provides a variety of learning rules that individuals may employ (cf. Gans et al. 2007). To examine different learning processes, we follow the approach used by Barham et al. (2015) and specify four potential models for learning that have legacies in psychological studies of learning patterns. We refer to these four learning processes as Bayesian learning, impressionable learning, reactionary learning, and myopic updating. These learning rules are described in greater detail below.

Consider Bayesian learning, in which information from all rounds is weighted equally and farmers' subjective probability updates in each round given new information. The Bayesian subjective probability of drawing a blue bead in round t_k is therefore

$$S(t_k) = \frac{1}{t_k} \sum_{t=1}^{t_k} B_t.$$

In contrast, under impressionable learning, farmers consider only the ratio of blue beads to total beads from the first n rounds when making all of their subsequent choices. We refer to this model of learning as "impressionable" because it reflects the strength of first impressions, in which additional information revealed after the nth round is essentially ignored. In the extreme case, individuals' beliefs about the distribution are influenced by information revealed in the first–and only the first–round. Under this rule, the impressionable subjective probability in round t is given by

$$S(t) = B_1.$$

These are farmers who overweight their initial information and form strong beliefs based on their first impressions, forgoing any future learning.

Similarly, last-n learners neglect all of their previous information in favor of only the most recent n draws. These farmers may be characterized as having short memories or attention spans, or may simply choose to neglect information beyond the most recent n experiences because updating using this relatively distant information may be tedious and error prone. Again, in the extreme, individuals may attend only to the most recent information revealed, ignoring any other information that had previously been revealed. The subjective probability in round t based on this "reactionary" learning style is given by

$$S(t) = B_t.$$

Finally, the myopic updating learning rule is similar to last-n in that it considers only information from the most recent n rounds, but it differs in that it uses a simple classification of draws into "good" and "bad." Draws with four or five blue beads are considered "good" while draws of three or fewer blue beads are "bad." Farmers characterized by myopic updating both overweight only the most recent n rounds of information and do not fully consider the probabilities over which they are making their choices. The corresponding subjective probability weights are

$$S(t) = \begin{cases} 1 & \text{“good” draw} \\ 0.5 & \text{“bad” draw.} \end{cases}$$

Upon calculating these probabilities for each individual, we can fit and evaluate the models based on the Bayes information criterion (BIC), which equals $-2 * LL$ using the calculated logit log-likelihood. We rank each learning rule using the BIC, such that the first-best learning rule (that is, the learning rule that most accurately reflects the observed behavior) has the lowest BIC, the second-best learning rule has the second-lowest, and so on. Occasionally, learning rules may have equivalent BIC rankings if the sequence of draws decisions, or both is similar. When describing the distribution of these learning rules we include ties between rules in each of the rules' respective totals. After classifying individuals according to their first-best learning rules, we estimate a multinomial logit model (excluding ties) across the possible first best-rules to determine whether literacy, age, caste, and evaluated comprehension of the learning experiment are determinants of particular learning rules. Finally, we model the usage of hybrid rice in 2013 as a function of learning rules and individual characteristics that potentially impact adoption to investigate which learning rules encourage or prevent adoption.

4. RESULTS

Table 4.1 provides household summary statistics for the full sample that participated in the learning experiment and for whom we have complete information about their inputs, along with the subsamples of farmers for whom the elicited risk aversion coefficient was finite and those for whom it was infinite.[4] Column (4) includes t-tests of the difference between the finite and infinite risk aversion samples. Of the farmers who participated in the survey in 2013, nearly one-third (124) had an infinite risk aversion coefficient in the risk experiment. Farmers with a finite risk aversion coefficient had a mean risk aversion coefficient of 0.61, indicating a modest degree of risk aversion. This level of risk aversion is similar to the estimates reported by Binswanger (1981), also from India, though with a slightly different elicitation method. The Binswanger (1981) estimate of 0.71 suggests that the farmers in his sample from semiarid tracts of Maharashtra and Andhra Pradesh were slightly less risk averse than the farmers in our sample in Bihar. Furthermore, the estimated risk aversion coefficients reported here are roughly consistent with those reported by Cardenas and Carpenter (2008) from many other contexts around the globe. The mean risk aversion coefficient here is slightly lower than those reported by Barham et al. (2015) for Minnesota and Wisconsin farmers (0.77) using a virtually identical preference elicitation mechanism. When generating the likelihoods for risk-averse farmers in the following analysis, we exclude those with infinite risk aversion coefficients because we are unable to calculate their utility without using an arbitrary value. While this represents a large portion of the sample, as shown in Table 4.1 they differ primarily in caste makeup and evaluated comprehension of the learning experiment.

Table 4.1 Summary statistics for full sample and finite constant relative risk aversion sample

	(1)	(2)	(3)	(4)
Characteristic	Full sample	Finite risk aversion sample	Infinite risk aversion	Difference
Age	46.90	47.46	45.92	−1.535
	(13.0)	(12.5)	(13.79)	(1.271)
Gender (male=1)	0.96	0.98	0.94	−0.036*
	(0.20)	(0.15)	(0.24)	(0.018)
Can read and/or write	0.70	0.71	0.66	−0.055
	(0.46)	(0.45)	(0.48)	(0.045)
Comprehension: good	0.44	0.46	0.40	−0.060
	(0.50)	(0.50)	(0.49)	(0.048)
Comprehension: moderate	0.51	0.48	0.58	0.101**
	(0.50)	(0.50)	(0.50)	(0.048)
Comprehension: poor	0.04	0.06	0.02	−0.041**
	(0.21)	(0.24)	(0.13)	(0.020)
General caste	0.31	0.37	0.21	−0.159***
	(0.46)	(0.48)	(0.41)	(0.044)
Other backward caste	0.43	0.41	0.46	0.048
	(0.50)	(0.49)	(0.50)	(0.048)
Scheduled caste	0.22	0.20	0.27	0.073*
	(0.42)	(0.40)	(0.44)	(0.04)
Scheduled tribe	0.03	0.02	0.05	0.037**
	(0.17)	(0.13)	(0.23)	(0.017)

[4]As previously stated, the risk aversion coefficients elicited through the experimental procedure documented above fell within an interval. To operationalize analysis, we assumed that an individual's coefficient of risk aversion was the value at the upper limit of this interval. This implied that some individuals (that is, those individuals who preferred the riskless option in first round as report in Table 3.1) were assigned an infinite coefficient of risk aversion. Given that the expected value of the lottery is Rs 30 and, even with a "losing" draw, the lottery payout would be no less than the riskless option, individuals choosing the riskless option clearly displayed considerable aversion to risk. Whether it is fair to characterize these individuals as infinitely risk averse is debatable. Yet it is also not straightforward to study their information processing strategies under such extreme preferences.

Table 4.1 Continued

Characteristic	(1) Full sample	(2) Finite risk aversion sample	(3) Infinite risk aversion sample	(4) Difference
Access to credit (2013)	0.03	0.03	0.02	−0.016
	(0.17)	(0.18)	(0.13)	(0.016)
Blue ratio (actual)	65.06	65.21	64.80	−0.403
	(11.42)	(11.06)	(12.07)	(1.119)
Blue ratio (guess)	65.35	65.54	65.02	−0.512
	(7.70)	(7.65)	(7.78)	(0.753)
Cultivated hybrid rice (2013)	0.14	0.15	0.13	−0.016
	(0.35)	(0.36)	(0.34)	(0.034)
Risk aversion coefficient		0.61		
		(0.58)		
Observations	451	287	164	

Source: Authors.
Notes: * Significant at 10% level; ** significant at 5% level; *** significant at 1% level. Standard deviations in parentheses in columns 1–3, Standard errors in parentheses in column 4.

Rankings of Learning Rules

We begin by providing rankings of the first-, second-, and third-most-likely learning rules for the entire sample under the assumption of risk neutrality. These are reported in panel (a) of Table 4.2. Due to potential similarities between the weighting functions, we have included ties between rules in both of the corresponding rows. Ties occurred only for first-best rule and were rare, with only one tie in the risk-neutral case and three ties in the risk-averse case. These results suggest that the most common learning rules are the myopic updating and impressionable learning rules, so that farmers either are reactive to good draws or largely base their subsequent decisions on the first round, implying that first impressions are important to many farmers in the sample. Somewhat surprisingly, Bayesian learning is a more common first ranked learning rule than reactionary learning, and is the most common second-best rule. From this top panel of Table 4.2, it is clear that there is substantial heterogeneity in the learning rules used by farmers in Bihar, in both their complexity and their timing (initial versus recent information). Despite the low levels of education in the sample, the prevalence of Bayesian learning suggests it may not be a bad approximation for belief updating even in rural villages in developing countries.

Table 4.2 Ranking of learning rules

Sample and learning rule	(1) First	(2) Second	(3) Third
(a) Full sample, risk neutral			
Bayesian learning	18.4	41.7	24.6
Impressionable learning	33.5	21.1	14.2
Reactionary learning	14.9	18.8	40.4
Myopic updating	33.5	18.4	20.8
(b) Finite risk aversion sample, risk neutral			
Bayesian learning	17.8	39.4	24.4
Impressionable learning	28.9	21.6	17.4
Reactionary learning	16.0	20.2	37.3
Myopic updating	37.6	18.8	20.9
(c) Finite risk aversion sample, risk averse			
Bayesian learning	25.4	39.7	25.4
Impressionable learning	39.0	11.1	9.1
Reactionary learning	15.3	31.7	33.4
Myopic updating	21.3	17.4	32.1

Source: Authors.

The rankings of risk-neutral and risk-averse learning rules using the limited subsample of farmers with a finite risk aversion coefficient are provided in panels (b) and (c), respectively, of Table 4.2 Surprisingly, the exclusion of farmers with infinite risk aversion does not change the rankings of learning rules under the assumption of risk neutrality relative to the sample in panel (a). Specifically, we still estimate myopic updating and impressionable learning to be the most prevalent first-best learning rules. However, there are slight differences between the rankings when allowing for risk aversion in the utility functions used in the estimation of the learning rules. As before, fewer farmers exhibit reactionary learning as their first learning rule assuming risk aversion, though this rule may be of secondary or tertiary import, as shown by the increasing frequencies in columns (2) and (3). Notably, while impressionable learning is still prominent (used by nearly 40 percent of the sample), the number of Bayesian farmers (25 percent) is slightly higher than those relying upon myopic updating (21 percent), though if we consider reactionary and myopic-updating learners to be indicative of those with a present bias in their learning, then Bayesian learning is less common than relying upon first impressions or having present biases (36 percent). The fact that myopic updating is more prevalent than reactionary learning, despite the fact that both are biased toward recent observations, suggests that farmers in our sample are more likely to rely upon recent information if it provides a promising signal.

Determinants of Learning Rules

Next, we investigate whether observable farmer characteristics are correlated with particular learning rules. To do this, we estimate a series of multinomial logits under different sample specifications. Multinomial logit analysis generalizes logit regression to a multiclass setting, essentially allowing the analyst to estimate the marginal contribution of a series of covariates to the probability of a series of outcomes that are generally–though not necessarily–mutually exclusive. We include a set of covariates that may influence the individual's utilization of a particular learning rule.[5] Specifically, we include controls for age, gender, literacy, caste, and the enumerators' evaluations of how well respondents understood the rules and structure of the experiments. The age controls act as a proxy for experience and wisdom, which may affect the farmer's learning rule through patience or years of experience with learning processes. Previous studies have found gender differences in the choice of learning rules, which could be the result of innate or social factors that contribute to differences in how men and women either complete the learning experiment or otherwise process information (Barham et al. 2015; Gans et al. 2007). Formal education is relatively low in rural Bihar, particularly among the current generation of adults, though there is variation in the literacy level between farmers. Literacy may contribute to higher-level information processing that carries over into learning about distributions over time. Given the low level of education, enumerator evaluations of respondents' comprehension provide ancillary controls for intelligence and comprehension that are not accounted for with literacy.

The first set of results in Table 4.3 are for the full sample under the assumption of risk neutrality. The coefficients for reactionary learning and myopic updating suggest that evaluated comprehension is correlated with an individual's most likely learning rule, meaning that those with poor comprehension are much more likely to use either reactionary learning or myopic updating than Bayesian updating, with relative log odds ratios of 1.7 and 3.6, respectively. Similarly, farmers with only moderate comprehension are more likely to follow myopic updating than Bayesian updating (odds ratio of 1.05), but are less likely to adhere to either impressionable learning (odds ratio 0.71) or reactionary learning (odds ratio 0.50). When we use comprehension of the experiment as a proxy for intelligence, we find that those with lower comprehension exhibit more reactionary responses to favorable draws and are less likely to follow Bayesian updating or use the most recent draws to inform their belief updating. Generally, there is considerable unexplained variation in the choice of learning rules under risk neutrality after controlling for experience, gender, and potentially weak proxies for intelligence.

[5] We exclude the one farmer with a tie in his first-best rule so that the rules are mutually exclusive.

Table 4.3 Multinomial logit relative odds ratios: Determinants of first-best learning behavior

	Full sample, risk neutral			Finite risk aversion sample, risk neutral			Finite risk aversion sample, risk averse		
Constant	Impressionable	Reactionary	Myopic updating	Impressionable	Reactionary	Myopic updating	Impressionable	Reactionary	Myopic updating
	1.003	1.011	1.006	1.015	1.026	1.016	1.020	1.022	1.022
	(0.009)	(0.013)	(0.012)	(0.016)	(0.017)	(0.017)	(0.016)	(0.020)	(0.017)
	0.439	0.803	0.531	0.306	0.339	1.709	1.405	1.811	0.475***
	(0.424)	(0.815)	(0.425)	(0.450)	(0.466)	(2.746)	(1.266)	(2.637)	(0.456)
Can read and/or write	1.258	0.696	1.156	1.775	1.333	2.071*	1.032	0.750	0.771
	(0.364)	(0.249)	(0.374)	(0.721)	(0.630)	(0.843)	(0.422)	(0.384)	(0.360)
Other backward caste	1.016	0.571	0.711	0.859	0.560	0.518	0.468**	0.489	0.220***
	(0.282)	(0.202)	(0.256)	(0.353)	(0.274)	(0.234)	(0.171)	(0.187)	(0.086)
Scheduled caste/tribe	0.958	0.825	0.619	0.805	0.530	0.420*	0.813	0.796	0.686
	(0.447)	(0.439)	(0.268)	(0.431)	(0.328)	(0.220)	(0.315)	(0.403)	(0.387)
Comprehension: moderate	0.706	0.495*	1.051	0.680	0.562	1.179	0.985	0.965	1.645
	(0.176)	(0.191)	(0.277)	(0.257)	(0.277)	(0.387)	(0.329)	(0.460)	(0.584)
Comprehension: poor	0.248	1.689	3.637*	7.44×10^{-7}***	1.769	4.184*	0.779	1.84×10^{-7}***	0.300
	(0.310)	(1.694)	(2.708)	(6.56×10^{-7})	(1.828)	(3.270)	(0.436)	(1.22×10^{-7})	(0.281)
	3.660	1.374	2.759	2.467	1.186	0.477	0.663	0.263	1.36×10^{-7}***
	(3.033)	(1.335)	(2.252)	(3.421)	(1.828)	(0.732)	(0.636)	(0.423)	(1.54×10^{-7})
Observations	450			286			284		
Log pseudo likelihood	-581.853			-362.179			-361.580		

Source: Authors.

Notes: * Significant at 10% level; ** Significant at 5% level; *** Significant at 1% level. Standard errors adjusted for clustering at the village level. Bayesian learning is the reference category in all regressions. Caste effects are relative to general caste. Comprehension effects are relative to understanding well. Ties amongst most likely learning rules not included.

The second set of estimates in Table 4.3 are the multinomial logit estimates under risk neutrality for the subsample of farmers who have a finite risk aversion coefficient. Despite the similarities in the distribution of learning rules and household characteristics, there are evident differences in the determination of learning rules. Surprisingly, the relative log odds of myopic updating as opposed to Bayesian learning increase by 2 for farmers who can read, write, or both implying that literacy does not seem to increase the likelihood of using more computationally intensive learning rules. Being a member of a scheduled tribe or scheduled caste increases the likelihood of myopic updating, but caste does not seem to have robust effects across samples under risk neutrality. Farmers with poor comprehension are much more likely to learn according to myopic updating, with relative log odds of 4.18, and slightly more likely to be impressionable learners.[6] Taking comprehension as a current measure of intelligence, we find that farmers with lower evaluated comprehension are more likely to use computationally simple updating rules when making decisions, assuming risk neutrality.

Finally, the third set of estimates in Table 4.3 are the multinomial logit estimates for the subsample of farmers with finite risk aversion, where we estimate learning rules allowing for individuals to exhibit risk aversion. Across impressionable, reactionary, and myopic updating learning, we find that members of other backward castes are more likely to use Bayesian updating (odds ratios of 0.47, 0.49, and 0.22, respectively). Farmers with poor comprehension of the experiment are considerably less likely to be reactionary learners, suggesting that lower comprehension is not necessarily associated with reactive responses to draws.

Accuracy of Learning Rules

After identifying the rankings and determinants of learning rules, we investigate whether particular learning rules are better predictors of individuals' estimation of the number of blue beads in the bag. Recall that prior to learning the true number of blue beads in the bag, farmers were asked to guess how many they believed were in the bag and were awarded 5 extra rupees if their guess was within 2 of the actual number. The true number of blue beads in the bag was 72, but farmers observed only a sequence of subsamples over 15 draws (including the practice round). To analyze how accurate farmers were in their estimates, we calculate the absolute difference between their guesses and the average of the total draws that they observed. Then we multiplied this average by 100 to produce a value on the same scale as the guess. In other words, if an individual observed 57 blue beads during the course of the experiment, then considering there were 75 total beads drawn, the average number of blue beads observed was 0.76. Based on these observations, an individual might reasonably expect that there are 76 blue beads in total in the bag. These absolute differences are illustrated in Figure 4.1. We estimate a simple linear regression model with this absolute difference as the dependent variable, conditional on the learning rules and individual characteristics used previously. Village fixed effects are included to control for differences across villages at the time of the survey and potential sharing of the correct number of blue beads among village members.

[6]While the relative odds ratio is statistically significant with very low probability of type I error, the effect is so small as to likely not be a particularly relevant determinant of learning behavior.

Figure 4.1 Accuracy of guesses: Absolute difference between guess and revealed draws

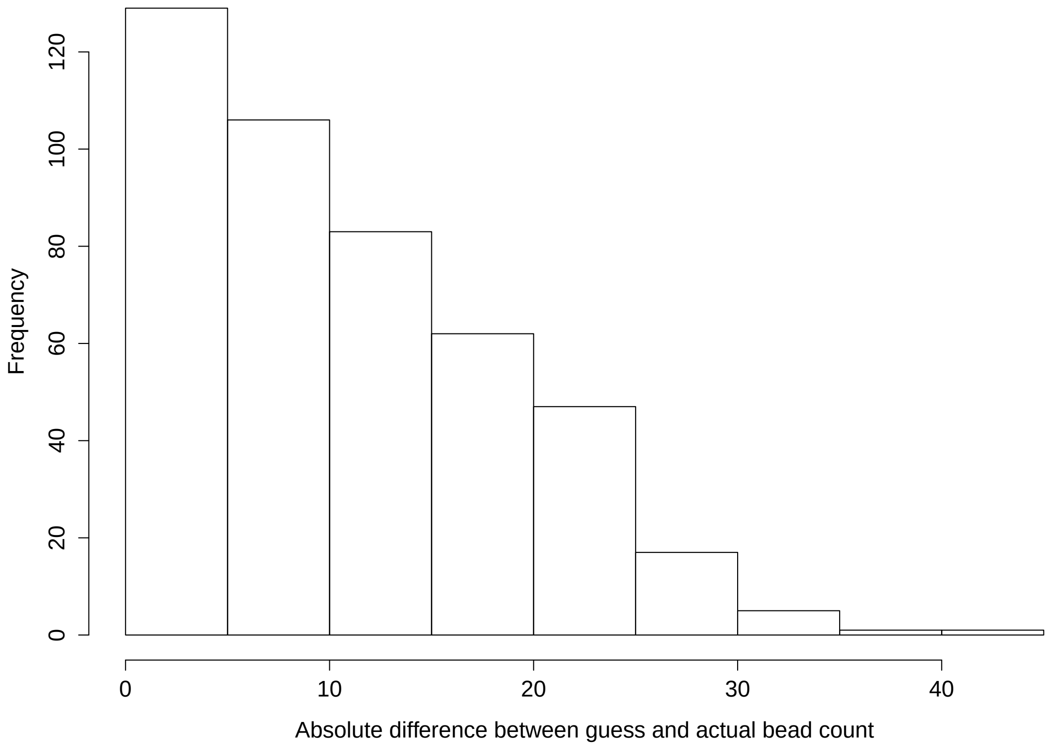

Source: Authors.
Notes: * Significant at 10% level; ** significant at 5% level; *** significant at 1% level. Robust standard errors in parentheses. Bayesian learning is the reference category in all regressions. Comprehension effects are relative to understanding well. Ties among most likely learning rules not included. All regressions contain intercepts; controls for caste, access to credit, age, gender, and literacy; and village fixed effects.

Results are provided in Table 4.4. The first column shows the risk-neutral learning rules for the full sample, the second column shows the risk-neutral learning rules for the subsample with finite risk aversion, and the third column shows the risk-averse learning rules for the subsample with finite risk aversion. Bayesian learning is the excluded rule in all three columns. Myopic-updating learners perform significantly worse in their guesses than Bayesian learners regardless of the subsample or risk aversion, with the average guess nearly 3 beads further away from the actual number of blue beads compared with Bayesian learners'. Given that myopic updating learners base their beliefs on promising signals (that is, they rely on information from a draw only if 4 or 5 blue beads drawn), this likely leads to overestimation of the number of blue beads, because their mean posterior belief is on the order of 80 blue beads or more. In contrast, reactionary learners perform significantly better than Bayesian learners. Reactionary learners' guesses, on average, are nearly 4 beads closer to the actual amount than Bayesian learners'. We note that there is little evidence to suggest that, after controlling for individual learning rules, perceived comprehension affects performance in guessing the actual number of blue beads. Across all three regressions, those individuals who did not appear to fully comprehend the game performed no worse than those who did.

Table 4.4 Difference in guess from revealed probability in learning game

Characteristic	Full sample, risk neutral	Finite risk aversion sample, risk neutral	Finite risk aversion sample, risk averse
Impressionable	−0.690	−0.479	−0.454
	(1.138)	(1.460)	(1.463)
Reactionary	−3.804***	−4.079***	−4.075***
	(1.310)	(1.560)	(1.562)
Myopic updating	3.000***	3.052**	3.094**
	(1.139)	(1.431)	(1.441)
Comprehension: moderate	1.115	1.301	1.339
	(0.767)	(1.004)	(1.010)
Comprehension: poor	−0.613	0.964	0.983
	(1.840)	(2.156)	(2.158)
Observations	450	286	284
R^2	0.152	0.222	0.221

Source: Authors.

Notes: * Significant at 10% level; ** significant at 5% level; *** significant at 1% level. Robust standard errors in parentheses. Bayesian learning is the reference category in all regressions. Comprehension effects are relative to understanding well. Ties among most likely learning rules not included. All regressions contain intercepts; controls for caste, access to credit, age, gender, and literacy; and village fixed effects.

Learning and Technology Adoption

Finally, we investigate the relationship between learning rules and the decision to adopt a new agricultural technology, specifically hybrid rice. In many ways, rice hybrids represent the next generation of the Green Revolution. Particularly in light of the world food price crisis in 2007-2008, many developing countries have shown increasing interest in finding solutions to increase productivity growth to ensure food security, and there are hopes that hybrid rice might be one such solution (Spielman et al. 2013, forthcoming). Previous research suggests that hybrids have had a significant effect on improving livelihoods and food security in several developing countries–most notably China–where rice is the principal food grain (for example Lin and Pingali 1994; Janaiah et al. 2002). Hybrids can contribute to increased food security both for producers (through higher productivity, resulting in increased own-consumption as well as larger marketable surpluses, which in turn results in higher farm incomes) and for consumers (because the increased quantity of rice on the market results in lower and more stable prices). Hybrids typically have considerably higher yields than conventional inbred varieties–even later generations of modern, high-yielding varieties arising from the Green Revolution. Much of these higher yields can be attributed to heterosis–the increase in the vigor of the rice crop resulting from the genetic contributions that result from crossing distinct parental lines. Not only does heterosis typically confer higher yields, but it also leads to a significant increase in genetic uniformity, which translates into an economic benefit through a lower seed rate needed to cultivate a given area (typically about one-third).

Despite these benefits, rice hybrids are not without significant downsides. Perhaps the most glaring disadvantage is the significantly higher seed price (on the order of about 10 times the price of modern varieties). Although partially offset by the lower seeding rate, this higher seed price–in conjunction with increased expenditures on complementary inputs like fertilizer and irrigation–still typically results in higher operational costs for hybrid rice production compared with cultivating modern varieties. Furthermore, yields and genetic uniformity decline dramatically after the first generation of seed (F1). There is therefore no benefit to farmers to save and store harvested grains to use as seed in subsequent seasons. Rather, farmers must typically purchase new F1 seed on a continual basis if they wish to avail themselves of the benefits of hybrids.

Nevertheless, the government of India has set ambitious targets to increase the area under hybrid rice. The National Food Security Mission set a goal to increase the area under hybrid rice to as much as 25 percent of all cultivated rice area by 2015, up from only about 6 percent in 2008–2009 (Spielman et al. 2012). As with any relatively new technology, however, farmers' ultimate decision to adopt hybrids depends crucially on subjective beliefs about the profitability of cultivation and on the ways in which farmers formulate their beliefs and update them with exposure to new evidence. From the summary statistics reported in Table 4.1 we see that some farmers in our sample have adopted rice hybrids, but the expansion is far from complete, with only 15 percent of farmers having adopted. We do not have a complete history of hybrid crop usage, so unfortunately we are unable to focus on the timing of the decision. Thus, our estimation of the effects of different learning rules on hybrid adoption includes only a snapshot of "earlier adopters" and potential determinants of adoption.

We estimated the decision to use hybrid rice on any of the farmers' rice plots during kharif 2013 using a simple probit model conditional on the different learning rules and other covariates. Access to credit is included as a covariate in this model because it can alleviate cash constraints and allow farmers to purchase the more expensive hybrid seeds. Results from estimating this probit model using maximum likelihood are reported in Table 4.5. As before, columns (1) and (2) show learning rule estimates under the assumption of risk neutrality, using the full sample and the subsample of farmers with finite risk aversion, respectively. Column (3) shows the learning rules estimated with risk aversion. Village fixed effects are included to control for variations within the village that may affect the decision to use hybrid rice during kharif 2013. The sample size decreases by nearly half when we include village fixed effects due to many of the villages' having no adoption.

Table 4.5 Probability of cultivating hybrid rice in kharif, 2013

Characteristic	Full sample, risk neutral	Finite risk aversion sample, risk neutral	Finite risk aversion sample, risk averse
Impressionable	−0.440	−0.675	−1.551***
	(0.286)	(0.430)	(0.417)
Reactionary	−1.442***	−2.079***	−1.982***
	(0.432)	(0.601)	(0.482)
Myopic updating	−0.0483	−0.189	−0.584
	(0.302)	(0.400)	(0.493)
Credit	0.706	1.425*	2.110***
	(0.619)	(0.774)	(0.621)
Understands ok	−0.423*	−0.166	−0.385
	(0.224)	(0.326)	(0.341)
Understands poorly	−0.527	−0.218	−1.308**
	(0.602)	(0.723)	(0.641)
Observations	233	136	136

Source: Authors.

Notes: * Significant at 10% level; ** significant at 5% level; *** significant at 1% level. Robust standard errors in parentheses. Bayesian learning is the reference category in all regressions. Comprehension effects are relative to understanding well. Ties among most likely learning rules not included. All regressions contain intercepts; controls for caste, access to credit, age, gender, and literacy; and village fixed effects.

Across all three columns in Table 4.5, we observe that farmers characterized by reactionary learning in the learning experiment are less likely to be early adopters relative to Bayesian learners. Thus, farmers who have more present-biased learning processes (reactionary learning) are less likely to be early adopters even after controlling for farmer comprehension, which was previously shown to be negatively correlated with reactionary learning. When learning rules are estimated assuming risk aversion–column (3)–we find that impressionable learners are similarly less likely to adopt hybrid rice than are Bayesian learners. These farmers rely upon first impressions, so they may be heavily influenced by the high up-front cost of hybrid seeds. The fact that Bayesian learners are generally more likely to adopt hybrids than those who rely upon other learning rules is not terribly surprising. With new technologies, including new seeds, there is typically a process of tinkering and making marginal adjustments to learn about the appropriate use of both the technology and it's complementary inputs. Bayesian learners are much more suited for this type of process than either those who are present biased or those who rely upon first impressions. But Bayesian learning is also considerably more cognitively taxing, requiring a much longer memory and a more complex updating process. It is not surprising, therefore, that most farmers within our sample rely on less complicated learning rules. But the fact that more than 25 percent of the farmers in our sample rely upon Bayesian learning as a first-best description of their learning process suggests that models assuming such learning may provide reasonable predictions and testable hypotheses about farmer behavior, even in rural settings in developing countries like India.

Unsurprisingly, credit access is positively correlated with early adoption of hybrid rice for the subsample of farmers with finite risk aversion coefficients. Farmers with lower evaluated comprehension in the experiment–which may be a reasonable proxy for intelligence–are less likely to be early adopters of rice hybrids, implying that current intelligence, as opposed to educational attainment, may be correlated with adoption, independent of learning processes.

5. CONCLUSION AND DISCUSSION

This paper has used experimental methods to identify various processes by which farmers in rural Bihar formulate and process information. The results suggest that there is a great deal of heterogeneity in farmers' learning heuristics, both in complexity and in the way they rely upon initial versus recent information. Generally speaking, the results suggest that farmers in the sample tend to either rely upon first impressions or react to recent information, but they are more likely to act upon recent information if it provides a promising signal as opposed to a simple, ambiguous message. Despite this tendency, however, roughly a quarter of the farmers in the sample can best be described as Bayesian learners, those who use each past observation to inform posterior beliefs. This is an interesting finding, given the cognitive tax that such processing imposes. But given that researchers often assume Bayesian learning processes in models of technology adoption, these results suggest that such models may provide reasonable predictions about farmer behavior.

The heterogeneity in learning patterns that we observe are rather difficult to ascribe to observable individual characteristics. In some ways, this suggests that learning rules are intrinsic to unique individuals and are not systematically determined by gender or age or, with some exceptions, caste. But we find convincing evidence that the nature of learning does impact technology adoption, with Bayesian learners typically much more likely to adopt hybrid rice than those characterized by impressionable or reactionary learning.

What do these results imply for Indian agricultural policies, particularly around hybrid rice, which the Indian government views as an important technology in its plans to increase food security? The results suggest a continued need for formal financial integration in rural India, because access to credit significantly increases adoption of hybrid rice and likely has similar effects with other agricultural technologies. Furthermore, given the heterogeneity in learning rules observed in our sample, efforts to promote hybrids (or likely any new technology) will probably require a more nuanced approach, rather than a one-size-fits-all extension message. For individuals who adhere to first impressions, messages and demonstrations will most likely have to exhibit an immediate impact. The same can largely be said for those with a present bias, but these learners would be more likely to forgive earlier failures if they are followed up by promising signals. Unfortunately, given the idiosyncratic nature of learning that we observe, it will be difficult to prescribe the appropriate message for an individual farmer based solely on observable characteristics. Identifying appropriate messaging strategies to satisfy the needs of different learners may be a fruitful avenue of future research.

REFERENCES

Anderson, C. 2001. "Behavioral Models of Strategies in Multi-Armed Bandit Problems." Doctoral thesis, California Institute of Technology, Pasadena, CA, US.

Arifovic, J., and J. Ledyard. 2011. "A Behavioral Model for Mechanism Design: Individual Evolutionary Learning." *Journal of Economic Behavior & Organization* 78 (3): 374–395.

Bandiera, O., and I. Rasul. 2006. "Social Networks and Technology Adoption in Northern Mozambique." *The Economic Journal* 116 (514): 869–902.

Barham, B., J.-P. Chavas, D. Fitz, V. {Rios Salas, and L. Schechter. 2014. "The Roles of Risk and Ambiguity in Technology Adoption." *Journal of Economic Behavior & Organization* 97: 204–218.

Barham, B. L., J.-P. Chavas, D. Fitz, V. Rios-Salas, and L. Schechter. 2015. "Risk, Learning, and Technology Adoption." *Agricultural Economics* 46: 11–24.

Binswanger, H. P. 1981. "Attitudes toward Risk: Theoretical Implications of an Experiment in Rural India." *The Economic Journal* 91 (364): 867–890.

Camerer, C., and T. Ho. 1999. "Experience-weighted Attraction Learning in Normal Form Games." *Econometrica* 67 (4): 827–874.

Cardenas, J. C., and J. Carpenter. 2008. "Behavioural Development Economics: Lessons from Field Labs in the Developing World." *Journal of Development Studies* 44 (3): 337–364.

Carey, J. M., and D. Zilberman. 2002. "A Model of Investment under Uncertainty: Modern Irrigation Technology and Emerging Markets in Water." *American Journal of Agricultural Economics* 84 (1): 171–183.

Chanda, A. 2011. *Accounting for Bihar's Productivity Relative to India's: What Can We Learn from Recent Developments in Growth Theory.* IGC Working Paper. London: International Growth Centre.

Cheung, Y.-W., and D. Friedman. 1997. "Individual Learning in Normal Form Games: Some Laboratory Results." *Games and Economic Behavior* 19 (1): 46–76.

Conley, T. G., and C. R. Udry. 2010. "Learning about a New Technology: Pineapple in Ghana." *American Economic Review* 100 (1): 35–69.

Erev, I., and A. E. Roth. 1998. "Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88 (4): 848–881.

Feder, G. 1980. "Farm Size, Risk Aversion and the Adoption of New Technology under Uncertainty." *Oxford Economic Papers* 32 (2): 263–283.

Feder, G., R. E. Just, and D. Zilberman. 1985. "Adoption of Agricultural Innovations in Developing Countries: A Survey." *Economic Development and Cultural Change* 33 (2): 255–298.

Foster, A. D. and M. R. Rosenzweig. 1995. "Learning by doing and learning from others: Human Capital and Technical Change in Agriculture." *Journal of Political Economy* 103 (6): 1176–1209.

Foster, A. D., and M. R. Rosenzweig. 2010. "Microeconomics of Technology Adoption." *Annual Review of Economics* 2 (1): 395–424.

Gans, N., G. Knox, and R. Croson. 2007. "Simple Models of Discrete Choice and their Performance in Bandit Experiments." *Manufacturing & Service Operations Management* 9 (4): 383–408.

Gollin, D., M. Morris, and D. Byerlee. 2005. "Technology Adoption in Intensive Post-Green Revolution Systems." *American Journal of Agricultural Economics* 87 (5): 1310–1316.

Griliches, Z. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica 25* (4): 501–522.

Hanna, R., S. Mullainathan, and J. Schwartzstein. 2014. "Learning through nTheory and Experimental Evidence in Farming." *Quarterly Journal of Economics* 129 (3): 1311–1353.

Janaiah, A., M. Hossain, and M. Husain. 2002. "Hybrid Rice for Tomorrow's Food Security: Can the Chinese Miracle Be Replicated in other Countries?" *Outlook on Agriculture* 31 (1): 23–33.

Krishna, V. V., D. J. Spielman, and P. C. Veettil. 2016. "Exploring the Supply and Demand Factors of Varietal Turnover in Indian Wheat." *Journal of Agricultural Science* 154 (2): 258–272.

Lin, J. Y., and P. L. Pingali. 1994. "An Economic Assessment of the Potential for Hybrid Rice in Tropical Asia: Lessons from the Chinese Experience." In *Hybrid Rice Technology: New Developments and Future Prospects*, edited by S. Virmani. Manila: International Rice Research Institute.

Meyer, R. J., and Y. Shi. 1995. "Sequential Choice under Ambiguity: Intuitive Solutions to the Armed-Bandit Problem. Management." *Science* 41 (5): 817–834.

Morris, M., and D. Byerlee (1998). Maintaining Productivity Gains in Post-Green Revolution Asian Agriculture. In *International Agricultural Development,* 3rd ed, edited by C. Eicher and J. Staatz, 458–473. Baltimore: Johns Hopkins University Press.

Munshi, K. 2004. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of Development Economics* 73 (1): 185–213.

Rogers, E. M. 1995. *Diffusion of Innovations, 4th ed.* New York: The Free Press.

Spielman, D., D. Kolady, and P. Ward. 2013. "The Prospects for Hybrid Rice in India." *Food Security* 5: 651–665.

Spielman, D., D. Kolady, P. Ward, H. Ar-Rashid, and K. Gulati. 2012. Public Expenditures, Private Incentives, and Technology Adoption: The Economics of Hybrid Rice in South Asia. IFPRI Discussion Paper 01233. Washington, DC: International Food Policy Research Institute.

Spielman, D., P. Ward, D. Kolady, and H. Ar-Rashid. Forthcoming. Public Incentives, Private Investment, and Outlooks for Hybrid Rice in Bangladesh and India." *Applied Economics and Policy Perspective*s.

Tanaka, T., C. F. Camerer, and Q. Nguyen. 2010. "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam." *American Economic Review* 100 (1): 557–571.

Tversky, A., and D. Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science185* (4157): 1124–1131.

Ward, P. S., and V. Singh. 2015. "Using Field Experiments to Elicit Risk and Ambiguity Preferences: Behavioural Factors and the Adoption of New Agricultural Technologies in Rural India." *The Journal of Development Studies* 51 (6): 707–724.

RECENT IFPRI DISCUSSION PAPERS

For earlier discussion papers, please go to www.ifpri.org/publications/discussion_papers.
All discussion papers can be downloaded free of charge.

1590. *Building effective clusters and industrial parks.* Xiaobo Zhang, 2016.

1589. *Nepal's 2072 federal constitution: Implications for the governance of the agricultural sector.* Jordan Kyle and Danielle Resnick, 2016.

1588. *Entitlement fetching or snatching?: Effects of arbitrage on India's public distribution system.* Suman Chakrabarti, Avinash Kishore, and Devesh Roy, 2016.

1587. *Changing preferences through experimental games: Evidence from sanitation and hygiene in Tamil Nadu.* Yaniv Stopnitzky, 2016.

1586. *Climate change, agriculture, and adaptation in the Republic of Korea to 2050: An integrated assessment.* Nicola Cenacchi, Youngah Lim, Timothy Sulser, Shahnila Islam, Daniel Mason-D'Croz, Richard Robertson, Chang-Gil Kim, and Keith Wiebe, 2016.

1585. *The dark side of competition: Gender differences.* Simon Chang, Kamhon Kan, and Xiaobo Zhang, 2016.

1584. *Harnessing net primary productivity data for monitoring sustainable development of agriculture.* Nathaniel P. Robinson, Cindy M. Cox, and Jawoo Koo, 2016.

1583. *Effects of agricultural mechanization on smallholders and their self-selection into farming: An insight from the Nepal Terai.* Hiroyuki Takeshima, Rudra Bahadur Shrestha, Basu Dev Kaphle, Madhab Karkee, Suroj Pokhrel, and Anjani Kumar, 2016.

1582. *The agricultural sector as an alternative to illegal mining in Peru: A case study of Madre de Dios.* Valeria Piñeiro, Jenny Thomas, and Pablo Elverdin, 2016.

1581. *What drives diversification of national food supplies?: A cross-country analysis.* Samira Choudhury and Derek Headey, 2016.

1580. *Identity, household work, and subjective well-being among rural women in Bangladesh.* Greg Seymour and Maria S. Floro, 2016.

1579. *Youth employment, agricultural transformation, and rural labor dynamics in Nigeria.* Margaret Adesugba and George Mavrotas, 2016.

1578. *Farm size and effects of chemical fertilizer price on farm households: Insights from Nepal terai.* Hiroyuki Takeshima, Sabnam Shivakoti, Binod Bhattarai, Madhab Karkee, Suroj Pokhrel, and Anjani Kumar, 2016.

1577. *Energy use and rural poverty: Empirical evidence from potato farmers in northern China*. Zihan Li, Yazhen Gong, and Kevin Z. Chen, 2016.

1576. *Does a "blue revolution" help the poor?: Evidence from Bangladesh.* Shahidur Rashid, Nicholas Minot, and Solomon Lemma, 2016.

1575. *Do development projects crowd out private-sector activities?: A survival analysis of contract farming participation in northern Ghana.* Isabel Lambrecht and Catherine Ragasa, 2016.

1574. *Strong democracy, weak state: The political economy of Ghana's stalled structural transformation.* Danielle Resnick, 2016.

1573. *Storage and handling among smallholder potato farmers in southwestern Uganda.* Bjorn Van Campenhout, Senne Vandevelde, Wiberforce Walukano, and Piet Van Asten, 2016.

1572. *What drives input subsidy policy reform?: The case of Zambia, 2002–2016.* Danielle Resnick and Nicole M. Mason, 2016.

1571. *Using household consumption and expenditure surveys to make inferences about food consumption, nutrient intakes and nutrition status: How important is it to adjust for meal partakers?* John L. Fiedler and Dena M. Mwangi, 2016.

1570. Improving *household consumption and expenditure surveys' food consumption metrics: Developing a strategic approach to the unfinished agenda.* John L. Fiedler and Dena M. Mwangi, 2016.

INTERNATIONAL FOOD POLICY RESEARCH INSTITUTE

www.ifpri.org

IFPRI HEADQUARTERS
2033 K Street, NW
Washington, DC 20006-1002 USA
Tel.: +1-202-862-5600
Fax: +1-202-467-4439
Email: ifpri@cgiar.org