

A Behavioral Social Learning Model for Studying the Dynamics of Forecast Adoption

Majid Shafiee-Jood¹, Tatyana Deryugina¹, and Ximing Cai¹

¹University of Illinois at Urbana-Champaign

November 24, 2022

Abstract

Drought forecasts, particularly at seasonal scales, offer great potential for managing climate risk in water resources and agricultural systems. In this context, the importance of assessing the economic value of such forecasts and determining whether a decision-maker should adopt them cannot be overstated. Value-assessment studies often, however, ignore the dynamic aspects of forecast adoption, despite evidence from field-based studies suggesting that farmers' forecast-adoption behavior fits the general framework of innovation diffusion, i.e. that forecast adoption is a dynamic learning process that takes place over time. In this study, we develop an agent-based model of drought forecast adoption to study the role played by heterogeneous economic and behavioral factors (i.e. risk aversion, wealth, learning rates), forecast characteristics (i.e. accuracy), and the social network structure (i.e. inter- and intra-county ties, change agents, self-reliance) in the process of forecast adoption and diffusion. We consider two learning mechanisms: learning by doing, represented by a reinforcement-learning mechanism, and learning from others, represented by a DeGroot-style opinion-aggregation model. Results show that, when social interactions between agents occur, forecast adoption follows a typical S-shaped diffusion curve. By contrast, when agents rely only on their own experience, the adoption pattern is close to linear. Our numerical experiment shows additionally that forecasts are never adopted if forecast accuracy drops below 65 percent. Finally, the proposed model also provides a flexible tool with which to test the effectiveness of extension targeting strategies in facilitating the diffusion of forecasts.

A Behavioral Social Learning Model for Studying the Dynamics of Forecast Adoption

Majid Shafiee-Jood¹, Tatyana Deryugina², and Ximing Cai^{1,3}

¹Ven Te Chow Hydrosystems Laboratory, Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA, ²Department of Finance, Gies College of Business, University of Illinois at Urbana-Champaign, Champaign, IL, USA, ³DOE Center for Advanced Bioenergy and Bioproducts Innovation, University of Illinois at Urbana-Champaign, Urbana, IL, USA.

Corresponding author: Ximing Cai (xmcai@illinois.edu)

Majid Shafiee-Jood (Orcid ID: 0000-0002-5808-3393)

Tatyana Deryugina (Orcid ID: 0000-0003-0870-8655)

Ximing Cai (Orcid ID: 0000-0002-7342-4512)

Key Points:

- Modeling the dynamic process of forecast adoption in a social setting to support the design of more effective targeting strategies
- Social interactions result in an S-shaped pattern of forecast-adoption diffusion
- The structure of a social network is of limited importance when agents' learning rates are high

Abstract

Drought forecasts, particularly at seasonal scales, offer great potential for managing climate risk in water resources and agricultural systems. In this context, the importance of assessing the economic value of such forecasts and determining whether a decision-maker should adopt them cannot be overstated. Value-assessment studies often, however, ignore the dynamic aspects of forecast adoption, despite evidence from field-based studies suggesting that farmers' forecast-adoption behavior fits the general framework of innovation diffusion, i.e. that forecast adoption is a dynamic learning process that takes place over time. In this study, we develop an agent-based model of drought forecast adoption to study the role played by heterogeneous economic and behavioral factors (i.e. risk aversion, wealth, learning rates), forecast characteristics (i.e. accuracy), and the social network structure (i.e. inter- and intra-county ties, change agents, self-reliance) in the process of forecast adoption and diffusion. We consider two learning mechanisms: learning by doing, represented by a reinforcement-learning mechanism, and learning from others, represented by a DeGroot-style opinion-aggregation model. Results show that, when social interactions between agents occur, forecast adoption follows a typical S-shaped diffusion curve. By contrast, when agents rely only on their own experience, the adoption pattern is close to linear. Our numerical experiment shows additionally that forecasts are never adopted if forecast accuracy drops below 65 percent. Finally, the proposed model also provides a flexible tool with which to test the effectiveness of extension targeting strategies in facilitating the diffusion of forecasts.

1 Introduction

Weather and climate forecasts, particularly at seasonal scales, can potentially play an important role in mitigating the negative impacts of climate variability in agriculture and water resources systems (Block, 2011; Hallstrom, 2004; Hansen, 2005). To realize this potential, forecasts should be used effectively and routinely by their recipients, which likely requires experimentation, practice, and reflection on experience regardless of how advanced or accurate the forecasts are (Hu et al., 2006; Whateley et al., 2015). This dynamic learning aspect of forecast adoption is often ignored in the literature, though, despite ample evidence from field-based studies suggesting that forecast-adoption behavior fits the general framework of *diffusion of innovation* (Luseno et al., 2003; Rubas et al., 2008; Tarnoczi & Berkes, 2010). That is, adoption of forecasts, like the adoption of any other *technology* or *innovation*, is a dynamic process that takes place over time and spreads across the social system (Rogers, 2003; Rubas et al., 2006; Ziervogel, 2004).

The central question that we address is: How do farmers (or water managers) make decisions about the use of forecast information, particularly when a forecast product is relatively unknown to them? Our approach deviates from the common modeling approach employed by forecast-valuation studies that assume that forecast users possess *perfect knowledge* of the characteristics of forecasts and can process forecast information in a statistically sophisticated manner (Millner, 2009). The perfect-knowledge assumption implies that forecast adoption is essentially a static individual decision-making problem that can be solved simply by computing the *ex-ante* value of forecasts (Millner, 2009; Rubas et al., 2008). Instead, we borrow from the literature on technology adoption and the diffusion of innovation theory (Rogers, 2003) and use a bottom-up approach to model farmers' forecast-adoption choices explicitly in an agent-based modeling (ABM) framework. By modeling and simulating individual farmers' heterogeneous

behavior as well as their interactions, ABM can capture macro-level emergent phenomena (Bonabeau, 2002), a capability that is particularly relevant in diffusion-of-innovation studies (Berger, 2001; Ng et al., 2011).

The decision-making context in our study is a stylized crop-allocation decision problem in which each farmer considers uncertainty about weather conditions during the crop season and chooses how to allocate land between two crops whose yields respond differently to drought conditions. We assume that farmers are rational decision-makers, but they cannot keep track of the history of their actions and experimental outcomes as well as those of their neighbors; they are in this sense *statistically unsophisticated* (Duffy, 2006; Millner, 2009). We also assume that farmers are not initially familiar with forecasts, but that they can learn about forecasts over time.

Learning has been recognized as a key driver in adopting a new technology. Using insights from field-based studies suggesting the importance of both individual experimentation (e.g. Hu et al., 2006; Ziervogel, 2004) and social influences (e.g. Crane et al., 2010; Hu et al., 2006; Ziervogel & Downing, 2004) in forecast-use decisions, we consider two learning mechanisms in the model: 1) learning by doing, in which users learn about an innovation largely through individual experimentation and observation (Arrow, 1962; Feder et al., 1985; Lindner et al., 1979); 2) learning from others, or social learning, in which users observe their neighbors' experiences and retain relevant information (Besley & Case, 1993; Foster & Rosenzweig, 1995; Manski, 1993; Munshi, 2004). To model how farmers learn from their own experience (i.e. learning by doing), we use a behavioral model motivated by the psychological theory of *reinforcement learning*. The cornerstone of reinforcement learning is the *law of effect* principle developed by Thorndike (Thorndike, 1911, 1932), which suggests that the tendency to repeat an action or a behavior that has succeeded will be reinforced whereas an action that has led to an unfavorable outcome will be incorporated less frequently (Roth & Erev, 1995; Tesfatsion, 2006).

In reinforcement learning, choice behavior is treated as a Markov stochastic process in which the tendency associated with each possible action (in this case, adoption or non-adoption of forecasts) is updated at every time step based on the consequences of a farmer's action in the previous time step (Brenner, 2006; Duffy, 2006). Furthermore, we assume that a farmer's adoption behavior is also influenced by the behavior of other farmers in his or her *social neighborhood* (i.e. a neighborhood defined by social interaction as opposed to geographic proximity). This form of social learning is also known as a *neighborhood effect* (Baerenklau, 2015; Manski, 1993). To express how farmers are influenced by their neighbors, we use a simple rule-of-thumb model based on the opinion-formation model of DeGroot (DeGroot, 1974; Jadbabaie et al., 2012).

This study makes several contributions to the literature. We develop a behavioral-learning model to represent farmers' forecast-adoption behavior by considering individual experimentation and the neighborhood effect. Because forecast performance, reflecting the probabilistic nature of forecasts, is more uncertain than other innovations (Agrawala & Broad, 2002), learning could play an even bigger role in the context of forecast adoption. The importance of learning has been documented by several field-based studies. For instance, a role-play exercise with smallholder farmers in Lesotho found that, as farmers became more familiar with the forecasts provided, "using a forecast no longer seemed foreign" and they were more willing to use them at the end of the experiment (Ziervogel, 2004). Millner (2009) used a behavioral-learning model based on reinforcement learning in the context of the cost-loss problem and showed that accounting for learning dynamics could significantly reduce the value

that the user obtains from forecasts. Our study extends and complements Millner's model by incorporating a social-learning mechanism, thereby accounting for the impact of social networks in forecast adoption and diffusion.

Our study also develops a flexible tool that makes it possible to better understand the temporal and spatial dynamics of forecast-use diffusion, which in turn can inform the design of economically efficient and effective strategies that facilitate forecast adoption. In the past two decades, there has been great interest in social learning as a key determinant of the diffusion process, especially in the context of agricultural technologies (Banerjee, 1992; Ellison & Fudenberg, 1993). Studies have found that social networks can play a major role in diffusion of innovation through both diffusion of knowledge (information) as well as diffusion of decisions (Cai et al., 2015; Holloway & Lapar, 2007; Sampson & Perry, 2019). Studies have found that stakeholder networks play a key role in the communication and dissemination of forecast information to farmers (Nidumolu et al., 2018; Ziervogel & Downing, 2004). Yet some critical elements in the diffusion process have not been carefully or rigorously studied in the context of forecast adoption: How do social interactions influence forecast-adoption behavior? To what extent is the structure of the social network important in diffusion of forecasts? What is the cumulative effect of social structure and individuals' characteristics in the forecast-use diffusion process? We identify these questions as gaps in the literature and address them in this study.

By explicitly modeling individual farmers' behavior as well as their learning from past forecast usage and from the experiences of others, the ABM we present derives the forecast-adoption path as an emergent property of collective behaviors. Therefore, our study fundamentally differs from top-down studies (e.g. Rubas et al., 2008) that impose adoption dynamics exogenously using widely accepted S-shaped adoption paths (Feder et al., 1985; Rogers, 2003). As such, our modeling paradigm is similar to that used in Ziervogel et al. (2005) and Bharwani et al. (2005), who investigated the impact of seasonal climate-forecast applications among smallholder farmers in Lesotho and South Africa, respectively, using agent-based social-simulation models. We depart from this approach, using reinforcement learning to model how farmers' tendency to adopt a forecast evolves over time as they experiment with forecasts, which also makes it possible to explore the impact of heterogeneous behavioral factors such as learning rates on the adoption and diffusion process.

The remainder of the article is organized as follows. In Section 2, we introduce the components of the agent-based model, including the crop-allocation decision problem and the learning process. We also explain the model dynamics. In Section 3, we design a numerical experiment that we use to demonstrate the adoption and diffusion of drought forecasts and present the critical assumptions of the model. We present the results in Section 4. First, we focus on the reinforcement-learning mechanism and investigate how risk aversion, wealth, and the learning rate influence an agent's tendency to adopt a forecast over time. Second, we use the agent-based model to investigate the temporal and spatial dynamics of forecast adoption and diffusion in a hypothetical agriculture-dominated case-study area. In Section 5, we use the model to demonstrate the effects of strategic targeting, asymmetrical learning, and forecast accuracy on the diffusion process. Finally, in Section 6, we conclude with a summary of the findings, the study limitations, and future work.

2 Model

In this section, we introduce a behavioral model of forecast adoption and diffusion in which decision-makers (DMs) learn about the usefulness of drought forecasts over time. DMs also decide whether to use forecasts when making planting choices, to which we refer as *adopting the forecast*. We focus our study on how the probability of forecast adoption evolves over time. Learning is stochastic and based on an agent's own experience (i.e. learning by doing) and on the experiences of *neighbors* in the agent's social network (i.e. social learning). To represent learning by doing, we use a behavioral model known as reinforcement learning (Bush & Mosteller, 1951, 1953; Roth & Erev, 1995). To account for learning from others, we use a DeGroot-style learning model of belief aggregation (DeGroot, 1974; Golub & Jackson, 2010; Jadbabaie et al., 2012). We now describe the decision-making problem, including these learning components, in greater detail (refer to Appendix C for a list of notations used in this study). Figure 1 presents a conceptual framework of the proposed model.

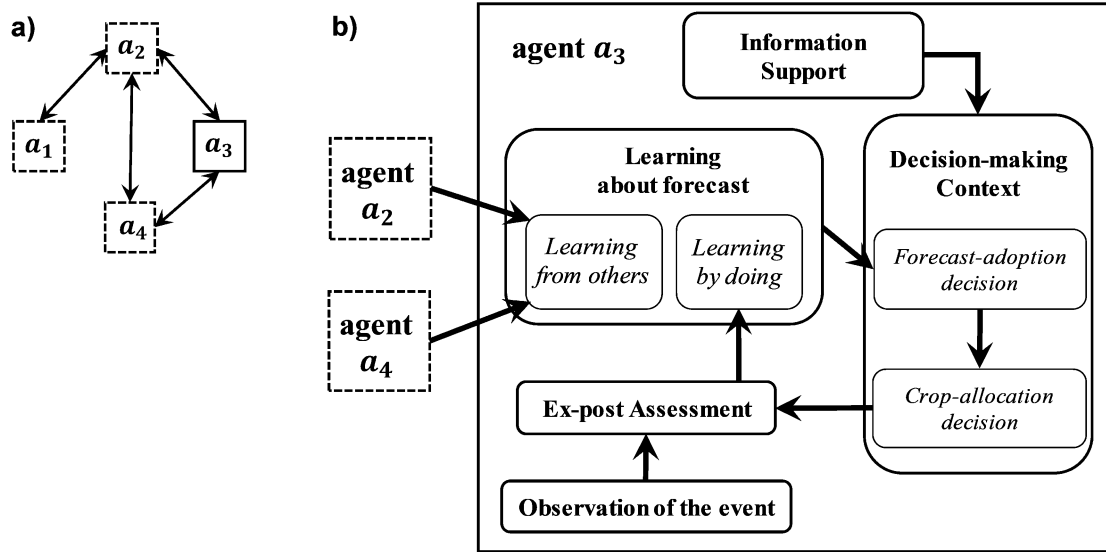


Figure 1. Conceptual framework of the model. (a) shows a simple representative social network structure (or topology) and (b) shows various components of the model for agent a_3 , as an example. Note that, in this network, agents a_2 and a_4 are neighbors of agent a_3 while agent a_1 is not.

2.1 Decision-making Context: Hedging against Drought

Consider a crop-allocation problem involving two crops, A and B , where a DM must determine what proportion of his or her land to allocate to each crop given the uncertainty associated with future weather. Without loss of generality, suppose that the weather event is a drought. Let $\theta \in \Theta = \{0,1\}$ be the random variable representing the state of the weather, where $\theta = 0$ and $\theta = 1$ correspond to no-drought (or normal) and drought conditions, respectively. Denote $p(\theta)$ as the DM's subjective belief regarding the probability that state θ occurs. As such, $p(\theta)$ embodies DM's knowledge about the uncertain event (Lawrence, 1999), which could be based on the event's historical probability (also called *climatological information*) and the DM's experience (Johnson & Holt, 1997; Sherrick et al., 2000). Suppose that the crop yield (per unit area of land) for crops A and B is a function of the weather alone, denoted by $y^A(\theta)$ and $y^B(\theta)$,

respectively. We assume that crop A is more drought-tolerant and has lower yield variability, while crop B is a high-yield variety whose yield falls off significantly in drought conditions, i.e. $y^B(1) < y^A(1) < y^A(0) < y^B(0)$.

Let $x \in [0,1]$ be the fraction of land that the DM allocates to crop A prior to the realization of θ (hence, $1 - x$ is the land fraction allocated to B). The DM's payoff function can be written as:

$$\pi(x, \theta) = \omega + x \cdot y^A(\theta) + (1 - x) \cdot y^B(\theta) - c(\theta), \quad (1)$$

where $\pi(x, \theta) = \pi$ is the normalized payoff given state θ and decision x , ω is the DM's normalized initial wealth, and $c(\theta)$ is the normalized total non-land cost of crop production (e.g. fertilizer and labor costs). We normalize ω and $c(\theta)$ by land area and crop price, respectively, and assume that the prices of the two crops are equal and do not depend on the occurrence of drought. Thus, π , ω , and c are all expressed in the same units as y (yield per unit area), which we denote by u . We assume that the DM is a utility maximizer whose risk preferences are characterized by an increasing von Neumann-Morgenstern utility function, $U = U[\pi(x, \theta)]$, as presented in Equation 2:

$$U = U[\pi(x, \theta)] = \begin{cases} \frac{\pi^{1-r}}{1-r} & r \neq 1 \\ \ln \pi & r = 1 \end{cases}, \quad (2)$$

where $r \geq 0$ is the Arrow-Pratt coefficient of relative risk aversion. The above-defined utility function belongs to a class of utility functions with constant relative risk aversion (CRRA) and is widely used in the economics literature (Mas-Colell et al., 2012). In CRRA utility functions, higher values of r correspond to more risk-averse behavior. One important feature that CRRA utility functions exhibit is that the risk premium for an absolute risk (a risk that is expressed in dollars as opposed to a share of the DM's wealth) is a decreasing function of wealth, i.e. wealthier individuals are more willing to take absolute risks. See Gollier (2001) for more information about risk characterization and utility functions; see Wakker (2008) for more details about the CRRA utility function. Therefore, the DM's optimization problem is given in Equation 3:

$$\max_x E_\theta[U] = \sum_{\theta=0}^1 p(\theta) \cdot U[\omega + x \cdot y^A(\theta) + (1 - x) \cdot y^B(\theta) - c(\theta)], \quad (3)$$

where E_θ is the expectation operator taken with respect to $p(\theta)$. Hence, the optimal allocation decision, x^* , must satisfy the following first-order condition:

$$\sum_{\theta=0}^1 p(\theta) \cdot \left(\frac{y^A(\theta) - y^B(\theta)}{(\omega + x^* \cdot y^A(\theta) + (1 - x^*) \cdot y^B(\theta) - c(\theta))^r} \right) = 0. \quad (4)$$

Because $\theta = \{0,1\}$, $p(\theta)$ can be characterized by a single parameter $p_1 := p(1)$, defined as the DM's *belief* that a drought will occur. Consequently, $p(0) = 1 - p_1$. See Appendix A for a sensitivity analysis demonstrating how optimal decision x^* changes with r , p_1 , or ω .

2.2 Learning-by-Doing

According to Brenner (2006), there are two fundamentally different ways of learning: reinforcement learning and cognitive learning. In reinforcement learning, the learning mechanism does not involve any conscious reflection on a problem, and therefore people are not always aware that they are learning. By contrast, cognitive learning is based on reflections about

actions and consequences, which requires active thinking and potentially involves processing statistical information (Brenner, 2006). Although people are able to reflect on their actions and consequences, in most cases they lack the cognitive capacity to reflect on *all* their actions. As a result, their reflections are likely to be distorted by cognitive biases (Brenner, 2006; Marx et al., 2007; Tversky & Kahneman, 1974). In reinforcement learning, on the other hand, the learning mechanism is based on an association between a behavior and its consequences; in other words, the behavior changes because of the resulting consequences. Reinforcement learning is particularly relevant when a DM is statistically unsophisticated, i.e. when he or she does not have the statistical ability to process and quantify forecast performance (Duffy, 2006; Millner, 2009). In this study, we use reinforcement learning to model how individuals learn from experience.

The learning mechanism in reinforcement learning is based on *reward* and *punishment*: if an action leads to a positive outcome, there is a higher chance that that action is chosen in the next time step; similarly, actions that result in negative outcomes are more likely to be avoided. One of the first mathematical models of reinforcement learning was developed by Bush and Mosteller (1951, 1953). The Bush-Mosteller model is a stochastic learning model in which choice behavior is described using a probabilistic distribution of alternatives rather than a binary choice framework, and the probability associated with each action is updated during the learning process using a simple linear rule. Cross (1973) placed the Bush-Mosteller model in an economic context and extended it to account for rewards of differing strengths. Brenner (1999, 2006) further generalized the Bush-Mosteller model by defining reinforcement strength in such a way that all rewarding (punishing) outcomes are reflected by positive (negative) reinforcement strengths. Roth and Erev (1995) also developed a reinforcement-type learning algorithm to track experimental data across various multi-player games that are analyzed in the experimental economics literature. In the Roth-Erev model, instead of directly updating the probability of choosing an action, an intermediate variable called *propensity towards an action* is employed. This variable is updated once an action is performed and is used to calculate the probability associated with that action. Here, we use a reinforcement-learning algorithm based on both the Bush-Mosteller and Roth-Erev models. The algorithm used here has two important features. First, it is *memoryless*, which corresponds to real-world behavior that is motivated by spur-of-the-moment decisions (Rahimian & Jadbabaie, 2017). Second, it captures the *spontaneous recovery* phenomenon (Rescorla, 2004; Thorndike, 1932), which makes it possible for nearly abandoned behaviors (or actions) to quickly increase in frequency if they result in positive outcomes (Millner, 2009).

We now formulate reinforcement learning mathematically. Suppose that DMs have costless access to a probabilistic drought forecast when crop-allocation decisions are being made. The forecast, denoted by p_d , indicates the probability that a drought will occur; $p_d \in \mathcal{F}$, where \mathcal{F} is a finite set of possible forecasts. Let $z_{i,t}$ be DM i 's binary forecast-adoption decision at time step t , where $z_{i,t} = 1$ indicates that the DM follows (or adopts) the forecast in making the crop-allocation decision. Following the Roth-Erev reinforcement-learning algorithm, we define $h_{i,t} \in [0,1]$ as the DM's *propensity* or *tendency* towards adopting the forecast at time t . As there are only two decision alternatives, the tendency towards not adopting the forecast (or using climatological information) is $1 - h_{i,t}$. The reinforcement-learning framework determines how the *adoption tendency*, $h_{i,t}$, evolves as a function of the DM's past decisions and outcomes. Using an updating rule based on a generalized form of the Bush-Mosteller model from Brenner (2006), the adoption tendency in the next time step, $h_{i,t+1}$, follows Equation 5:

$$h_{i,t+1} = h_{i,t} + \begin{cases} L(S_{i,t}, \tau_i) \cdot (1 - h_{i,t}) & \text{if } S_{i,t} \geq 0 \\ L(S_{i,t}, \tau_i) \cdot h_{i,t} & \text{if } S_{i,t} < 0 \end{cases}, \quad (5)$$

where $L(\cdot)$ is the learning function, $S_{i,t}$ is the *reinforcement strength* (expressed in unit u), and τ_i is the *learning rate* (expressed in unit u^{-1}). We follow convention and use a linear learning function: $L(S_{i,t}, \tau_i) = S_{i,t} \cdot \tau_i$. To ensure that $h_{i,t+1}$ remains between 0 and 1, we restrict $\tau_i \cdot \max |S_{i,t}| \leq 1$. Given the formulation in Equation 5, if $S_{i,t} \geq 0$ (i.e. the reinforcement strength is positive), the tendency of the DM to choose the action that would have led to the positive outcome will increase in the next time step. Thus, at each time step, the past is implicitly contained in the current value of $h_{i,t}$ (Brenner, 2006).

The choice of reinforcement strength is critical in this learning framework (Millner, 2009). In our case, the most natural choice for $S_{i,t}$ is the *ex post* value of the forecast, denoted by V^{exp} :

$$S_{i,t} = V_{i,t}^{exp} = \pi(x_{i,t}^{*,f}, \varphi_{i,t}) - \pi(x_{i,t}^{*,c}, \varphi_{i,t}). \quad (6)$$

Variables $x_{i,t}^{*,f}$ and $x_{i,t}^{*,c}$ are optimal crop-allocation decisions when a DM does or does not use the forecast information, respectively, at time t . If the forecast is not adopted, the DM makes the decision based on his or her own belief about drought occurrence, i.e. $p_{i,1}$; $\varphi_{i,t} \in \Phi = \{0,1\}$ is the realized state of the weather at time t , where $\varphi = 1$ indicates that a drought event has occurred. Note that V^{exp} is also expressed in the baseline unit u . The *ex post* value of the forecast denotes the value the DM would have received if he or she had made the crop-allocation decision based on the forecast. Learning occurs only when $x_{i,t}^{*,f} \neq x_{i,t}^{*,c}$ (otherwise, $S_{i,t} = 0$ and $h_{i,t+1} = h_{i,t}$). When $S_{i,t} > 0$ (hence $V_{i,t}^{exp} > 0$), the decision to adopt the forecast is reinforced; whereas when $S_{i,t} < 0$ the probability that the forecast is adopted in the next time step declines. As such, $S_{i,t}$ can be interpreted as a measure of *regret* or *happiness* regarding forecast adoption (Millner, 2009).

In our formulation, the adoption choice at each time step (i.e. z_t) is independent of adoption choices in previous time steps, and agents treat adoption and *discontinuance* decisions symmetrically. This diverges from the common approach in modeling technology adoption, where agents are assumed to continue using a new technology forever once they decide to adopt it (Ellison & Fudenberg, 1993). The rationale for assuming symmetrical behavior in our model is that, unlike in most other technological transitions, here no cost would be incurred if agents decide to switch between the two available options, i.e. adopting or not adopting a drought forecast.

2.3 Social Learning

Consider a set $\mathcal{M} = \{1, 2, \dots, m\}$ of agents interacting over a *social network*. Suppose that the underlying structure of the social network is known and can be represented by a directed graph with m vertices. Each vertex corresponds to an agent and a directed edge is present from vertex (agent) j to vertex i only if agent j is a *neighbor* of agent i . In that case, agent i 's beliefs can be influenced by agent j 's beliefs. For each agent $i \in \mathcal{M}$, define \mathcal{N}_i as the set of agents in agent i 's *social space* (Akerlof, 1997), with $|\mathcal{N}_i| = n_i$. The social network can be summarized by

matrix $\Delta = [\alpha_{ij}]_{m \times m}$, defined as the *matrix of social interaction* (Jadbabaie et al., 2012), where for each agent i , $\alpha_{ij} \geq 0$ determines the *weight* that agent i assigns to the beliefs of agent j , and the weights must satisfy $\sum_{j=1}^m \alpha_{ij} = 1$. Note that $\alpha_{ij} = 0$ if agent j is not a neighbor of agent i (or $j \notin \mathcal{N}_i$). α_{ii} is the weight that agent i assigns to his or her own belief, which is referred to as *self-reliance*, and $\sum_{j \in \mathcal{N}_i} \alpha_{ij} = 1 - \alpha_{ii}$. Therefore, matrix Δ determines both *social connections* and the extent of *social interactions* (Molavi et al., 2018).

The social-learning component of our model is based largely on the belief-aggregation model of DeGroot (1974). In DeGroot-style models, agents update their beliefs as a convex combination (i.e. weighted average) of the beliefs of their neighbors. The *weights* determine the *trust* that agents have for their neighbors (Acemoglu & Ozdaglar, 2011; DeGroot, 1974). Let $h_{i,t}$ be agent i 's tendency to adopt the forecast at time step t , as in Section 2.2. Using the DeGroot model of social learning, agent i updates the likelihood that he or she will adopt the forecast (i.e. $h_{i,t+1}$) as follows:

$$h_{i,t+1} = \alpha_{ii} \cdot h_{i,t} + \sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot h_{j,t}. \quad (7)$$

In the next section, we discuss the dynamics of the model and provide a framework within which we embed individual reinforcement-based learning into a DeGroot-style social-learning component.

2.4 An Agent-based Modeling Framework

We now integrate the components presented in Sections 2.1–2.3 into an agent-based model of forecast adoption and diffusion. Figure 2 shows the flowchart of the proposed model. For all agents ($i \in \mathcal{M}$), the risk attitude (represented by the coefficient of risk aversion, r_i), adoption threshold (h_i^*), initial wealth level (ω_i), learning rate (τ_i), belief about drought occurrence ($p_{1,i,t}$), and initial adoption tendency ($h_{i,1}^{initial}$) are taken as given. The structure of the social network ($\Delta = [\alpha_{ij}]_{m \times m}$) is also known. Time steps are indexed by $t = 1, 2, \dots, T$. Each time step represents a crop season. Let $h_{i,t}^{initial}$ and $W_{i,t}$ be agent i 's adoption tendency and wealth at the *beginning* of time step t , i.e. before crop-allocation decisions are made; note that $W_{i,1} = \omega_i$. At the beginning of each time step t , each agent receives a probabilistic drought forecast ($p_{d,i,t}$). Agents then learn from their neighbors' forecast adoption tendencies and update their own beliefs according to Equation 8:

$$\forall i: h_{i,t} = \alpha_{ii} \cdot h_{i,t}^{initial} + \sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot h_{j,t}^{initial}. \quad (8)$$

To convert agent i 's stochastic choice behavior (i.e. $h_{i,t}$) into deterministic behavior (i.e. $z_{i,t}$), a threshold (or cut-off value) defined as h_i^* is used, as indicated in Equation 9:

$$z_{i,t} = \begin{cases} 1 & \text{if } h_{i,t} \geq h_i^* \\ 0 & \text{if } h_{i,t} < h_i^* \end{cases} \quad (9)$$

Note that, in our formulation, we impose the cut-off value as an exogenous parameter. Alternatively, the cut-off value could be derived endogenously by comparing the expected utilities of forecast adoption and non-adoption (i.e. $z_t = 1$ if $E[U[\pi(x_{i,t}^{*,f}, \theta)]] > E[U[\pi(x_{i,t}^{*,c}, \theta)]]$), e.g. as shown in Ellison and Fudenberg (1993) and

Adhvaryu (2014). We cannot, however, derive an explicit relationship between an agent's adoption tendency and the cut-off value because of the functional forms of the utility function and the payoff function.

Once their adoption decisions are made, agents make their crop-allocation decisions ($x_{i,t}^*$) following Equation 10:

$$x_{i,t}^* = \begin{cases} x_{i,t}^{*,f} & \text{if } z_{i,t} = 1 \\ x_{i,t}^{*,c} & \text{if } z_{i,t} = 0 \end{cases}. \quad (10)$$

After the actual state of the weather is realized (i.e. $\varphi_{i,t}$), Equation (6) is used to calculate the reinforcement strength ($S_{i,t}$), and the adoption tendency at the end of time t will be calculated according to Equation (11):

$$h_{i,t}^{final} = h_{i,t} + \begin{cases} S_{i,t} \cdot \tau_i \cdot (1 - h_{i,t}) & \text{if } S_{i,t} \geq 0 \\ S_{i,t} \cdot \tau_i \cdot h_{i,t} & \text{if } S_{i,t} < 0 \end{cases}. \quad (11)$$

We assume that agriculture is the main economic activity that contributes to the wealth of each agent; as such, the consequence of agricultural decision-making at each time step directly affects agents' wealth. Agent i 's wealth at time $t+1$ ($W_{i,t+1}$) can be written as follows:

$$W_{i,t+1} = W_{i,t} + \pi(x_{i,t}^*, \varphi_{i,t}). \quad (12)$$

At the end of time step t , the cumulative *ex post* payoff or cumulative gain (π_t^{cum}) is calculated as follows:

$$\pi_t^{cum} = \sum_{t'=1}^t \pi(x_{i,t'}^*, \varphi_{i,t'}). \quad (13)$$

At this point, time step t is completed and period $t + 1$ begins.

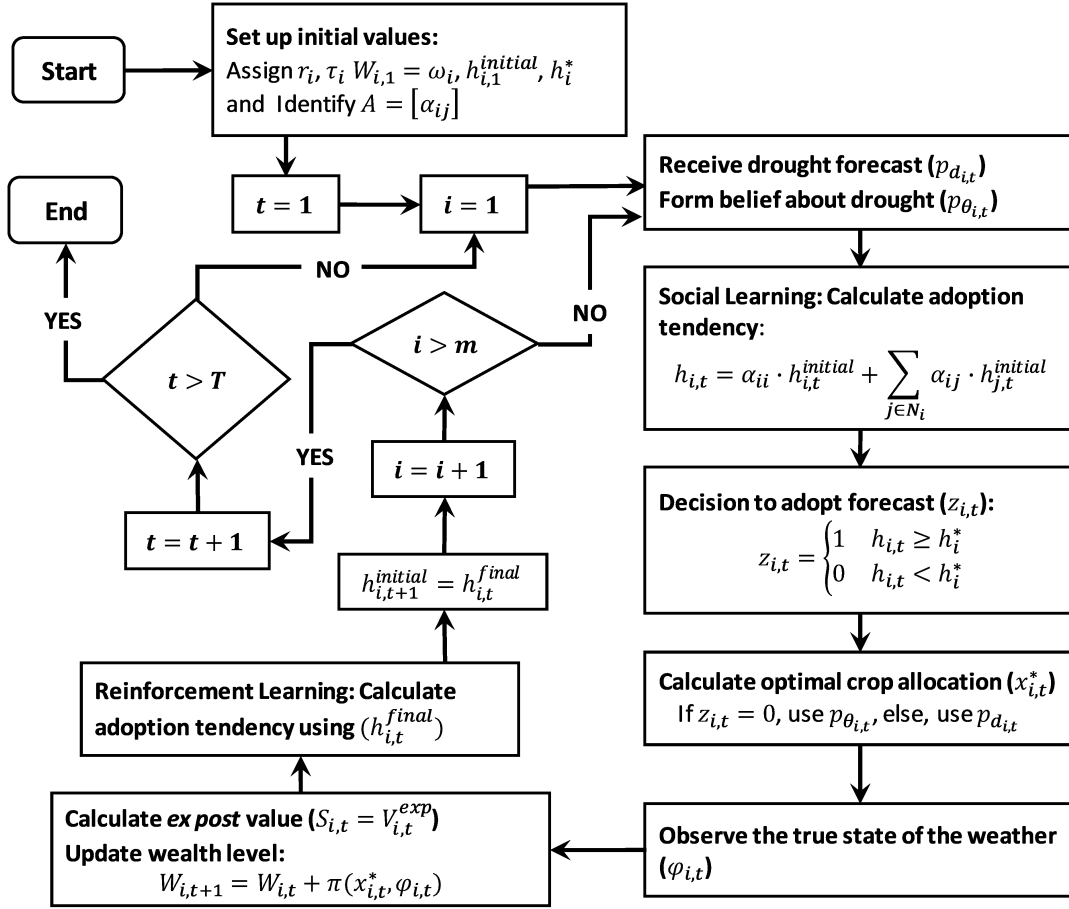


Figure 2. Flowchart of the proposed ABM simulating forecast adoption and diffusion.

3 Experimental Set-up and Assumptions

We design an experiment to demonstrate how various factors related to DMs' characteristics and their social network structure influence the dynamics of the forecast-diffusion process. The hypothetical case-study area, as shown in Figure 3, is an agriculture-dominated region consisting of 25 clusters (representing counties, communities, or villages), with 25 agents (representing farmers) in each cluster. Agents may interact with other agents in their social spaces (also referred to as *social neighborhoods*) and learn from their experiences. This interaction, which stimulates social learning, takes the form of communication of adoption beliefs, $h_{i,t}$. For all agents, the decision-making problem follows the one introduced in Section 2.1 with $y^A(0) = 0.06$, $y^A(1) = 0.03$, $y^B(0) = 0.08$, $y^B(1) = 0.01$, $c(1) = 0.04$, and $c(0) = 0.05$, all expressed in the baseline unit u .

We assume that the social neighborhood for each agent is dictated by his or her inter- and intra-county social ties, which are represented by two binary variables: $SI_{in} \in \{0,1\}$ for intra-county ties, and $SI_{out} \in \{0,1\}$ for inter-county ties. The extent of social interactions (i.e. the weights assigned to neighbors' beliefs) is represented by the matrix $\Delta = [\alpha_{ij}]_{m \times m}$, where $\alpha_{ij} = 0$ if $j \notin N_i$; $\sum_{j \in N_i} \alpha_{ij} = 1 - \alpha_{ii}$, and α_{ii} is each agent's self-reliance. When both SI_{in} and SI_{out} are zero, there is no social interaction and agents rely only on their own experience (i.e.

$\alpha_{ii} = 1$). When either of SI_{in} and SI_{out} is one, it is assumed that the agents are equally influenced (i.e. equal weights) by their own and other agents' beliefs that circulate in their social neighborhoods, unless stated otherwise.

We assume that the climatological probability of a drought event in the case-study area is 30 percent; that all agents share the same belief about the possibility that a drought event will occur, which is assumed to be equal to the climatological probability of drought in the area; and that this belief remains unchanged throughout the entire simulation, i.e. $\forall i, t: p_{1i,t} = 0.3$. Although these assumptions may not be necessarily accurate, they do not impact the purpose of this study and we leave investigating them to future work. Finally, we assume that the same time series of drought events is observed by all agents in the case-study area. Based on the time series of drought events, an approach similar to ensemble forecasting is used to generate probabilistic drought forecasts at a specified accuracy. Specifically, we assume that a probabilistic drought forecast at time t (p_{d_t}) is generated by a system that produces deterministic forecasts that have

an accuracy of κ . As such, p_{d_t} is defined as $p_{d_t} = \frac{\sum_{i=1}^N I_{\{\eta_i=1\}}}{N}$, where N is the total number of ensemble members and η_i is the deterministic forecast produced by ensemble member i (see Appendix B for additional details). Unless otherwise noted, we assume that $\kappa = 0.7$.

The key parameters of each DM are: the initial adoption tendency, the adoption threshold, the coefficient of risk aversion, initial wealth, and the learning rate. We assume that agents initially do not adopt the forecast by setting $h_{i,1}^{initial} = 0.5$. We set the adoption threshold at 0.65 for all agents (i.e. $\forall i, t: h_{i,t}^* = 0.65$). For the other three parameters as well as for parameters related to the topology of social networks (e.g. inter- and intra-county ties), we conduct sensitivity analyses.

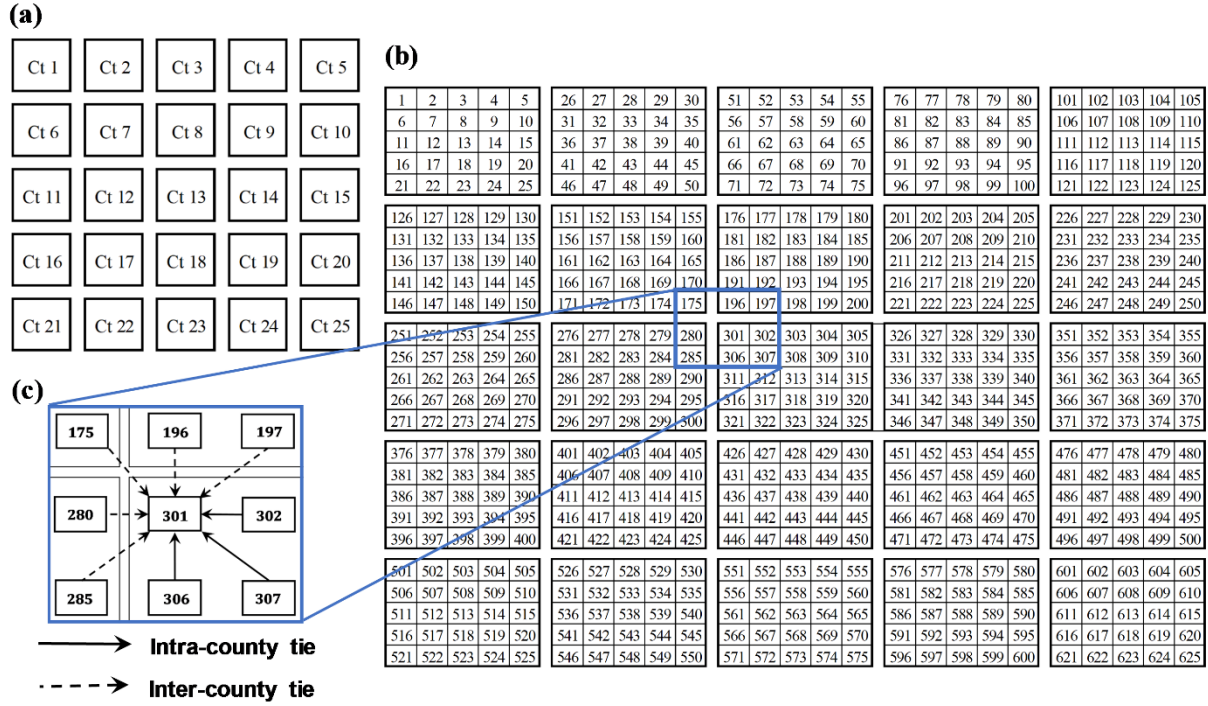


Figure 3. Hypothetical case study used to demonstrate the agent-based model of forecast adoption. (a) County locations. (b) Locations of the agents within each county. (c) Topology of the social network for agent #301.

4 Results

We first illustrate how various social-psychological and economic factors (i.e. risk aversion, r , initial wealth level, ω , and the learning rate, τ) influence an individual's learning and adoption behavior. We then explore how these factors influence the aggregate rate of forecast adoption and diffusion.

4.1 Reinforcement Learning and Belief Evolution

A DM's learning from the consequences of past adoption (or non-adoption) decisions is reflected in his or her tendency to follow the forecast (h), which depends on the reinforcement strength (S) and the learning rate (τ). Figure 4 shows one possible trajectory of the adoption tendency for a DM with a given risk-aversion coefficient (r), initial wealth (ω), and learning rate (τ). This trajectory is based on the time series of drought events and forecasts shown in Figure 4a. The Brier skill score (Wilks, 2006) for the forecasts shown in the figure is $BSS = 0.43$, which indicates that forecasts are on average more accurate than the climatological information.

When the learning rate is constant, the change in the DM's adoption tendency depends only on the current value of the reinforcement strength. This behavior reflects the key feature of the reinforcement-learning mechanism where the entire relevant history of the DM's behavior is implicitly contained in the current value of his or her adoption tendency (Brenner, 2006). No learning occurs when reinforcement strength is zero (i.e., $S = 0$). Figure 4b shows that, at first ($t < 22$), there is only one instance without learning. As a result, the DM's adoption tendency

changes frequently in this period, as shown in Figure 4c. Because there are more instances with $S > 0$, the adoption tendency exhibits an increasing trend. In the second portion ($t > 22$), the reinforcement strength is mostly zero; in those instances, the adoption tendency remains unchanged. The non-zero values of reinforcement strength are, however, relatively large and mostly positive, leading to the overall increasing trend in the DM's adoption tendency.

To further explain the two patterns, it is important to consider the decision-making context and the parameters that influence decisions under uncertainty. At first, the DM relies on climatological information (i.e. $p_1 = 0.3$) because the tendency to use the forecast remains under the adoption threshold (i.e. $h_t < h^* \rightarrow x_t^* = x_t^{*,c}$). The combination of initially low wealth and high risk aversion results in conservative crop-allocation decisions that are intended to minimize the potential impact of drought; i.e. a large fraction of land is allocated to crop A, which is a more drought-tolerant crop with lower yield variability (see Appendix A). For instance, the DM allocates 44 percent of the land to crop A at $t = 1$ (i.e. $x_1^* = x_1^{*,c} = 0.44$). For this DM, a forecast with $p_d > p_1 = 0.3$ that is followed by a drought event will result in a positive *ex post* value, thereby increasing the DM's tendency to adopt the forecast. This means that, if the DM had relied on such a forecast (instead of using the climatological information), a larger fraction of land would have been allocated to crop A (i.e. $x_1^{*,f} > 0.44$), which would have led to higher profit under drought conditions and consequently larger $S_t = V_t^{exp}$. A similar argument holds true for a forecast of $p_d < p_1 = 0.3$ that is followed by normal climatological conditions. On the other hand, the *ex post* value associated with a forecast of $p_d < p_1$ that is followed by a drought event or a forecast of $p_d > p_1$ that is followed by normal conditions is negative, which decreases the DM's tendency to adopt the forecast. Because forecasts are on average more accurate than climatological information ($BSS = 0.43$), instances with $V^{exp} > 0$ occur more frequently, which leads to an increasing trend in the adoption tendency.

As the DM's wealth increases over time (see Figure S1), his or her treatment of uncertainty approaches that of a risk-neutral DM (even though $r = 10$); as a result, for $t \geq 28$, the optimal crop land allocation would be to plant only crop B (i.e. $x_t^{*,c} = 0$) if the decision is based on climatological information. As a result, the forecast triggers learning (i.e. $V_t^{exp} > 0$) only if it leads to a decision that includes crop A in the mix of crop allocation (i.e. $x_t^{*,f} \neq 0$), which happens when p_d is relatively large. Even though these learning instances are not frequent in the remainder of the simulation (i.e. $t \geq 28$), because the corresponding reinforcement strength is positive in most cases the adoption tendency usually increases when learning is triggered, except when $S < 0$ (see Figure 4b and Figure 4c). For smaller values of p_d , decisions made with and without forecasts are similar (i.e. $x^{*,c} = x^{*,f}$). Therefore, those instances do not contribute to learning. In the scenario shown in Figure 4, the DM's tendency to adopt the forecast exceeds the threshold at $t = 39$ for the first time, and the DM continues following the forecasts until the end of the simulation, which results in a 23 percent higher cumulative economic gain than if he or she had maintained the business-as-usual practice (i.e. relying on climatological information). The DM's economic gain would however have been 37 percent higher than in a business-as-usual scenario had the forecasts been adopted from the beginning (see Figure S1).




Figure 4. Evolution of an agent's tendency to adopt drought forecast information. (a) Shows one possible time series of drought forecasts, (b) shows the corresponding time series of *ex post* forecast values, and (c) shows the corresponding trajectory of the forecast adoption tendency.

Here, $r = 10$, $\omega = 0.5$, and $\tau = 3$.

Figure 5 shows that the forecast-adoption tendency among DMs with higher learning rates, higher wealth, and lower risk aversion follows a steeper trajectory. As Equation 5 indicates, the learning rate determines the extent of a DM's response to the stimulus provided by the consequences of forecast adoption (or non-adoption) decisions. Higher values of the learning rate indicate that the DM is more susceptible to being *triggered* by the consequences of his or her past decisions, thereby representing a rapid learning behavior (Figure 5a); as such, DMs whose learning rates are higher begin following the forecasts earlier. Figure 5a also shows that, when the learning rate is high (i.e. $\tau = 5$), the adoption tendency drops significantly after only one *punishing* outcome (for example the drop at $t = 94$). This behavior, which is known in the reinforcement-learning literature as *spontaneous recovery*, implies that low-probability actions (in this case, not following the forecast) that have been abandoned by the DM could be quickly reinforced after a positive (*rewarding*) outcome (note that a punishing outcome for adoption is a rewarding outcome for non-adoption) (Brenner, 2006; Millner, 2009). On the other hand, when the learning rate is low (i.e. $\tau = 1$), even though the tendency to adopt the forecast continues to increase monotonically, it takes much longer for the DM to begin using it.

Figure 5b shows that the forecast-adoption tendency of a less risk-averse DM grows more rapidly. This is because, at first, crop-allocation decisions are made based on climatological information (i.e. $p_1 = 0.3$) and, therefore, for a DM with low risk aversion (i.e. $r = 0.5$) the

fraction of land allocated to crop A is zero because crop B has a higher expected yield. Therefore, the *ex post* value of forecasts is non-negative when $x^{*,f} \neq 0$, which corresponds to situations with relatively high values of p_d . Because droughts occur 30 percent of the time on average (i.e. $p_1 = 0.3$) and forecasts have high accuracy (i.e. $\kappa = 0.7$), instances with high p_d are not frequent. Because such instances are followed by a drought event in most cases, they provide high value to the DM, resulting in a greater increase in the forecast-adoption tendency. On the other hand, for a highly risk-averse DM (i.e. $r = 10$), even though learning occurs more frequently at first, there are more instances where $V^{exp} < 0$, and the reinforcement strength is lower than it is in the case of a DM with $r = 0.5$ (see Figure S2). As a result, such a DM's tendency to adopt forecasts increases at a slower pace. As the DM's wealth increases over time, however, the impact of risk aversion declines and the adoption tendency follows a very similar trend to the one observed where $r = 0.5$.

Similar arguments can be used to explain how the DM's initial wealth influences his or her learning pattern. The combination of $r = 10$ and $\omega = 0.25$ represents an extreme case of conservative decision-making to minimize the potential impact of drought, and the DM will decide to plant crop A regardless of whether he or she relies on climatological information or forecasts. As such, the *ex post* value of forecasts is very low at first, leading to small changes in the adoption tendency. This is the opposite of what is observed where $\omega = 1$ (see Figure 5c and Figure S2). As in the previous case, as the DM's wealth increases the decisions, *ex post* values, and consequently the adoption tendency for DMs with lower initial wealth become similar to that of a DM with large initial wealth (i.e., $\omega = 1$).




Figure 5. Sensitivity of the forecast adoption tendency to a DM's parameters based on the time series of forecasts and drought events shown in Figure 4 with varying values of (a) learning rate τ (with fixed $\omega = 0.5$ and $r = 10$), (b) risk aversion r ($\omega = 0.5$ and $\tau = 3$), and (c) initial wealth ω ($r = 10$ and $\tau = 3$).

4.2 Agent-based Model of Forecast Diffusion

In this section, we investigate learning and the dynamics of forecast diffusion in a social setting. Figure 6 shows the diffusion curve under three scenarios of social interaction: 1) full interaction, in which agents interact with all their neighbors, both inside and outside of their

counties (i.e. $SI_{in} = 1, SI_{out} = 1$); 2) intra-county interactions only, in which agents interact only with neighbors inside their counties (i.e. $SI_{in} = 1, SI_{out} = 0$); and 3) no interactions, in which agents learn based only on their own experience (i.e. $SI_{in} = 0, SI_{out} = 0$). We assign agents' learning rates (τ), risk aversion (r), and initial wealth (ω) randomly, assuming that each is normally distributed with $\tau \sim N(1.5, 0.1)$, $r \sim N(10, 1)$, and $\omega \sim N(0.5, 0.05)$ (see Figures S4–S6 for additional details). The diffusion curve for the full-interaction scenario is generally S-shaped, exhibiting logistic-type growth, which is consistent with the typical adoption path suggested in the diffusion-of-innovation literature (Mansfield, 1961; Rogers, 2003; Stoneman, 1983). Instead of exhibiting a typical monotonically increasing trend, though, the results exhibit a fluctuating trend. This is because we consider *discontinuance* in our model; in other words, agents may decide to discontinue using the forecast and base their decisions on climatological information despite having adopted the forecast earlier. The fluctuations are more frequent at first because the forecast-adoption tendency is, on average, closer to the adoption threshold of $h^* = 0.65$ during this period.

This S-shaped pattern we observe regarding adoption can be explained as follows. At first, agents exhibit a low forecast-adoption tendency, and decisions are therefore made based on climatological information. Because forecasts are on average more accurate than climatological information, however, agents gradually learn from their own and their neighbors' experience and form a greater tendency to use forecasts. This learning process varies across agents because of the heterogeneity in agents' behaviors and neighborhoods. As such, some agents adopt forecasts earlier than others. These agents are called *early adopters*. Because the agents' social network is strongly connected in this scenario, they adopt the forecasts at a higher pace as the number of adopters increases. This phase of the diffusion process ($50 < t < 70$) is known as the *take-off* phase. As the number of potential adopters decreases, the rate of adoption decreases until an adoption ceiling or equilibrium is reached.

One of the key elements in the diffusion of innovations is the social system within which diffusion occurs (Rogers, 2003). Figure 6 demonstrates how the diffusion pattern is influenced by the structure of a social network. The diffusion curve is almost linear when there is no interaction between agents (i.e. $SI_{in} = 0, SI_{out} = 0$), which can be attributed to the linear form of the learning function selected for the reinforcement-learning mechanism. When agents interact with each other, particularly in the full-interaction scenario, both individual and social learning mechanisms contribute to the forecast adoption and diffusion process, and therefore the forecast diffusion curve becomes non-linear.

Figure 6 also shows that, in the absence of social interaction, the number of adopters is higher at first than in either of the other two scenarios; with full interaction, on the other hand, agents begin adopting the forecast later than in either of the other two scenarios. In addition, the final adoption rate is highest when there is full interaction between agents, whereas almost 30 percent of the population decides not to follow the forecasts at the end of the simulation in the no-interaction scenario. These patterns can be explained by considering the spatial and temporal dynamics of diffusion. First, because risk aversion and initial wealth are randomly assigned, agents' crop-allocation decisions vary; hence, the *ex post* values of forecasts vary across agents (see Movie S1). Similarly, the learning rate is also randomly assigned. Therefore, the same forecast can lead to varying adoption tendencies. When agents interact with their neighbors, their forecast-adoption tendencies are essentially a weighted average of their own tendencies and those of their neighbors. As such, the forecast-adoption tendencies are balanced or smoothed by

neighbors' tendencies, particularly when both inter- and intra-county interactions are present. When there is no interaction, though, an agent's belief about adoption is influenced only by his or her own experience (i.e. individual learning). In this case, there is no continuity or specific spatial pattern in the way the forecast is adopted by agents (see Movie S1). When intra-county ties or full interactions exist, however, there is a strong spatial correlation in forecast adoption, and it spreads from early adopters to the entire population.




Figure 6. Diffusion curves under multiple interaction scenarios: full interaction (i.e. $SI_{in} = 1, SI_{out} = 1$), intra-county interaction only (i.e. $SI_{in} = 1, SI_{out} = 0$), and no interaction (i.e. $SI_{in} = 0, SI_{out} = 0$).

Because forecasts are more accurate on average than climatological information (the Brier skill score varies between 0.19 and 0.45 across the 25 counties), forecast adoption is expected to produce an economic gain. If forecasts were adopted by all the agents from the beginning ($Z_{i,1} = 1$), the total economic gain will on average be 29 percent (± 7 percent) higher than in the case of relying only on climatological information the entire time, which we call the *baseline scenario* hereafter. These results indicate that forecast adoption is a dynamic process and the timing and rate of adoption depends not only on agents' characteristics but also on the structure of the social network. As a result, the average increase in total economic gain with respect to the baseline scenario is 9 percent (± 7 percent) in the full interaction scenario.

Figure 7 demonstrates how forecast adoption and diffusion are influenced by the learning rate, initial wealth, and risk aversion in the full interaction scenario. Figure 7a shows that adoption starts earlier and reaches its maximum level more quickly as the learning rate increases. When the learning rate is low ($\mu_\tau = 1.5$), adoption does not occur until $t = 68$, and at the end of the simulation the adoption rate is only around 40 percent. This is because it takes a long time for an agent to form a *positive opinion* (i.e. an opinion that leads to choosing adoption) about forecasts when the learning rate is low. When $\mu_\tau = 1.5$ and there is no interaction between agents, adoption starts earlier (around $t = 45$) but remains under 40 percent at the end of the

simulation (see Figure S7). As the learning rate increases, the diffusion curves in various social interaction scenarios begin to resemble one another (see Figure S7), implying that the social structure becomes less important for the diffusion of forecasts when agents learn quickly from the consequences of their own actions. The results shown in Figure 7b and Figure 7c show that diffusion curves shift leftward as initial wealth increases or risk aversion decreases. In other words, higher values of ω or lower values of r result in earlier adoption and quicker diffusion, as in Figure 5b and Figure 5c, because increasing ω or lowering r increases an agent's willingness to adopt forecasts.




Figure 7. Diffusion curves in various scenarios of (a) the learning rate, (b) initial wealth, and (c) risk aversion. $\omega \sim N(0.5, 0.05)$ and $r \sim N(7.5, 1.5)$ in (a), $\tau \sim N(3, 1)$ and $r \sim N(7.5, 1.5)$ in (b), and $\omega \sim N(0.5, 0.05)$ and $\tau \sim N(3, 1)$ in (c).

5 Discussion

In this section, we use the ABM of forecast adoption to investigate how behavioral, contextual, and technological factors influence adoption behavior and the diffusion of forecasts.

5.1 Impact of Change Agents and Strategic Targeting

It seems possible to facilitate forecast adoption by educating agents about the potential value of forecasts, for example through local extension services, crop advisors, or boundary organizations (Buizer et al., 2016; Mase & Prokopy, 2014; Templeton et al., 2018). The impact of such educational programs can be modeled as an increase in agents' initial propensity towards forecast adoption (i.e. $h_{i,1}$). Because it may not be feasible to target the entire population with an educational program, here we consider an extension program like the *Training and Visit Extension System* (Feder & Slade, 1986; Munshi, 2004), where extension agents target only a portion of farmers in each designated region (Feder & Slade, 1984). Those farmers are referred to as *change agents* or *contact farmers*. To illustrate this phenomenon, we consider the entire case-study area as one extension region and treat all agents in county 13, located in the middle of the case-study area, as change agents.

For change agents ($i = 301, 302, \dots, 325$), we set $h_{i,1} = h_i^*$ and assume that they adopt forecasts from the beginning. Figure 8a shows that, with change agents in the system, the diffusion curve is shifted to the left and the diffusion process takes off earlier. The take-off phase develops slowly at first ($30 < t < 50$), though, as forecast adoption first spreads among agents located in change agents' neighborhoods. A stronger tendency towards adoption among these agents together with an increase in the adoption tendency among other agents based on individual learning results in a rapid increase in the adoption rate during the period $50 < t < 60$. After $t = 60$, there is a small percentage of non-adopter agents left. As a result, the adoption rate decreases, and an adoption ceiling is reached (see Movie S2).

Figure 8b demonstrates the impact of the learning rate on forecast diffusion in the presence of change agents. When the learning rate is low (i.e. $\mu_\tau = 1.5$), change agents have a significant impact on the diffusion process: the adoption rate reaches the maximum of 76 percent when change agents are present compared with the maximum of 48 percent without change agents. When the learning rate is high (i.e. $\mu_\tau = 4$), change agents have a smaller impact on the diffusion process: the diffusion curves with and without change agents are almost identical. This pattern confirms our earlier finding that the structure of the social network becomes less important in the diffusion process as the learning rate increases.




Figure 8. Impact of change agents on the diffusion process when all agents in county 13 are targeted. (a) Full interaction scenario with $\tau \sim N(3,1)$ and (b) full interaction scenario with $\tau \sim N(1.5,0.5)$ and $\tau \sim N(4,1)$. Note that $\omega \sim N(0.5,0.05)$ and $r \sim N(7.5,1.5)$

Our model can also be used to test the effectiveness of extension programs and design more efficient targeting strategies. For instance, selecting change agents is a key factor in the success of such extension methods as *training & visit* (Feder & Slade, 1984). While *opinion leaders* in farming communities are often selected as change agents, when information flows less smoothly in a social system it may be necessary to rely on less subjective measures to select change agents. To demonstrate this effect, we select agents whose initial wealth is above the 90th percentile of the wealth distribution as change agents. The rationale for this selection is that, as shown in Figure 5c and Figure 7c, agents with greater initial wealth exhibit a higher forecast adoption tendency early in the simulation. Figure 9a shows that targeting wealthier agents influences the diffusion process only slightly because, unlike in the previous example where all agents in one county were targeted, wealthy agents, scattered throughout the study area, are equally influenced by their neighbors, who have weaker adoption tendencies (recall that in the full interaction scenario, $\alpha_{ii} = \alpha_{ij}$ for $j \in \mathcal{N}_i$, i.e. self-reliance equals the weights given to each neighbor). As change agents become more self-reliant (i.e. as α_{ii} increases), however, they continue influencing their neighbors while being influenced by their neighbors to a lesser extent.

As a result, they facilitate the diffusion process: the diffusion curve shifts further to the left, implying a quicker take-off and a higher adoption rate in the short and medium runs (see Figure 9b and Movie S3).

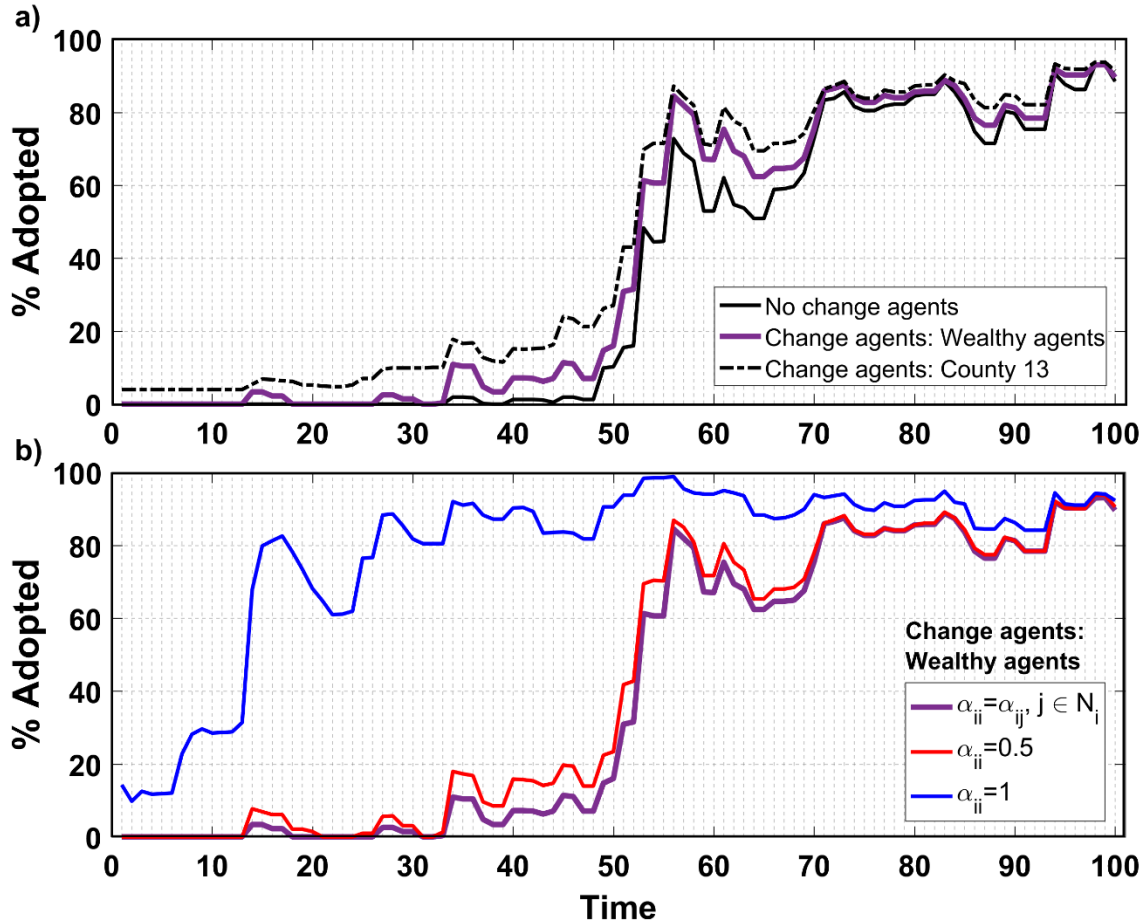


Figure 9. (a) Impact of wealthy agents as change agents on the diffusion process. (b) Impact of change agents' self-reliance (i.e., α_{ii}) on the diffusion process where wealthy agents are targeted as change agents. $\tau \sim N(3,1)$, $\omega \sim N(0.5,0.05)$ and $r \sim N(7.5,1.5)$.

5.2 Impact of Asymmetrical Learning

In the reinforcement-learning algorithm used in this study (Equation 5), we use a single learning rate to represent learning from both rewarding and punishing outcomes. Yet behavioral studies suggest that rewarding and punishing outcomes may not have symmetric impacts on decision-making (Cazé & Van Der Meer, 2013; Frank et al., 2004, 2007; Gershman, 2015). In particular, most studies have found that a negative learning rate (corresponding to punishing outcomes) is generally higher than a positive learning rate (corresponding to rewarding outcomes) (e.g. Rasmussen and Newland (2008), Niv et al. (2012), and Gershman (2015)), although some studies have found evidence of optimistic reinforcement learning, which is known as *optimism bias* (Lefebvre et al., 2017). Here, we modify the reinforcement-learning algorithm (Equation 5) to exhibit such asymmetrical updating, also known as *asymmetry in the law of effect* (Rasmussen & Newland, 2008). To do so, we use a parameter called the *asymmetric learning*

coefficient (denoted by γ) to amplify the impact of punishing outcomes: $L(S, \tau) = S \cdot \gamma \cdot \tau$, where $\gamma = 1$ if $S \geq 0$ and $\gamma > 1$ if $S < 0$.

Figure 10 shows how asymmetrical learning influences the diffusion process. In the case of symmetric learning (i.e. $\gamma = 1$), the rate of learning is the same for rewarding and punishing outcomes. Therefore, the diffusion curve is the same as the one shown in Figure 6. However, as γ increases, adoption starts later and the diffusion occurs at a slower pace (or, in the case of $\gamma = 3$, it never occurs) (see Movie S4). As γ increases, punishing outcomes (e.g. when a drought event is preceded by a low p_d), which exert negative reinforcement strength, will have a greater impact on a DM's learning. That is, when $S < 0$, the forecast-adoption tendency decreases to a greater extent for a DM with higher γ . As can be seen, when $\gamma > 2$, the adoption tendency almost never exceeds the adoption threshold, which implies that forecasts are never adopted over the simulation period.

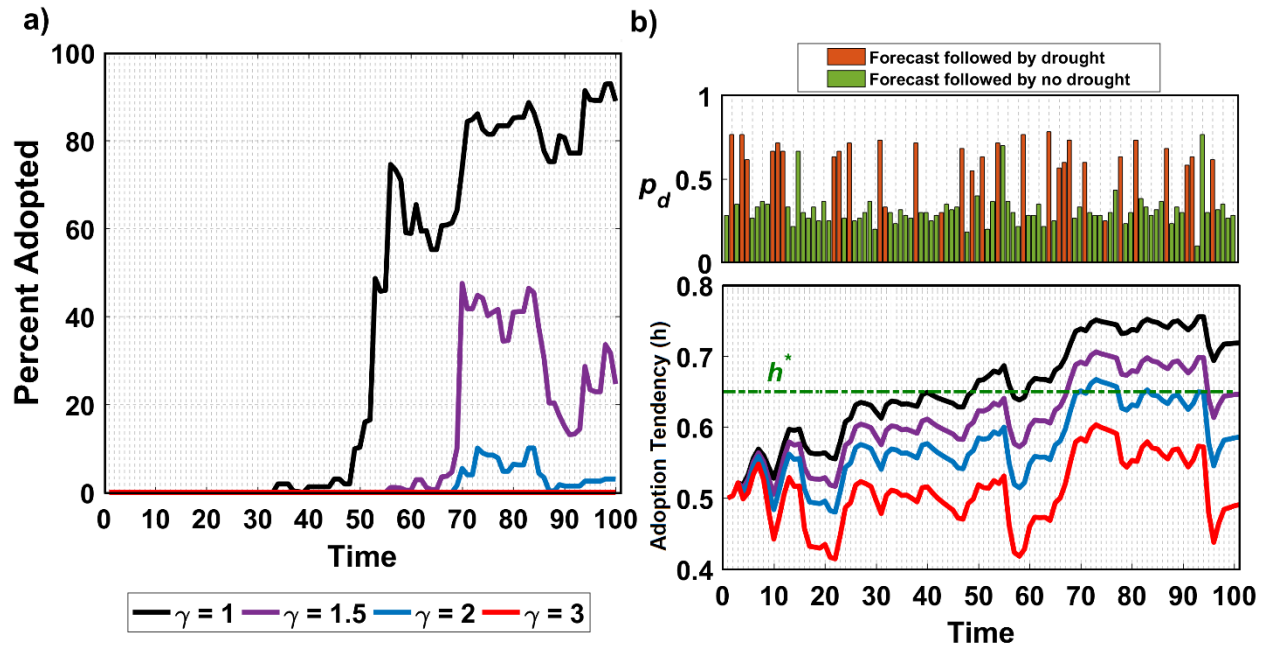


Figure 10. Impact of asymmetrical learning on forecast adoption. (a) Diffusion curves for the system in several asymmetrical learning scenarios, where $\omega \sim N(0.5, 0.05)$, $r \sim N(7.5, 1.5)$, and $\tau \sim N(3, 1)$. (b) Time series of forecasts and the evolution of the forecast-adoption tendency for a representative agent.

5.3 Impact of Forecast Accuracy

Figure 11 shows the impact of forecast accuracy on forecast diffusion. The diffusion curves are averaged across 15 realizations of a forecast time series (also see Figure S8 for the ensemble envelope). When $\kappa < 0.65$, instances with $V^{exp} < 0$ occur rather frequently; as a result, a DM's tendency to adopt forecasts never exceeds the threshold (see Figure S9). This implies that forecasts below 65 percent accuracy may never be adopted. This finding is consistent with those reported in other studies that have found that accuracy of at least 65 percent is required for seasonal forecasts to achieve long-term trust and adoption (see Ash et al.

(2007) for a review). As forecast accuracy (κ) increases, though, the take-off phase of the diffusion occurs earlier, and the adoption rate reaches its ceiling more quickly.




Figure 11. Impact of drought-forecast accuracy on the diffusion process. $\omega \sim N(0.5, 0.05)$, $r \sim N(7.5, 1.5)$, and $\tau \sim N(3, 1)$.

6 Conclusion

We develop an agent-based model to study the dynamic aspects of forecast adoption and demonstrate the impacts of farmers' characteristics and social network structure on forecast diffusion. To address forecast users' imperfect knowledge of forecasts, we model their forecast adoption as a stochastic choice and show that users' forecast-adoption tendencies evolve over time as a function of the consequences of their past decisions as well as the decisions of their neighbors. In addition, we show the influence of multiple factors on learning processes, including risk attitude, wealth, and the learning rate. We find that users with lower risk aversion, greater wealth, and higher learning rates exhibit a stronger tendency to use forecasts and therefore adopt forecasts more quickly than others.

The ABM provides a flexible tool that helps us better understand how a range of economic, behavioral, social, and forecast-related parameters influence forecast adoption and diffusion. Results derived from numerical experiments yield important insights into the effects of social interactions and social networks on the dynamics of forecast diffusion. In particular, when social interactions between agents take place, forecast diffusion follows a typical S-shaped curve, as suggested in the diffusion-of-innovation literature. In contrast, when social learning is ignored, the adoption pattern is (mainly) linear (Figure 6). Our results also show that, in a no-interaction scenario, the diffusion process starts earlier, reflecting the heterogeneities associated with farmers' characteristics but reaches a lower adoption ceiling compared with what occurs in a full interaction scenario. Moreover, our results show that asymmetrical learning reflecting the

asymmetry of reinforcement and punishment in human choice could significantly slow the diffusion process and lower the equilibrium adoption rate. On the other hand, we find that social structure has a limited impact on the diffusion process when the learning rate is high. Finally, we find that forecasts must be at least 65 percent accurate to be widely adopted and diffused in the system, which is consistent with findings reported by other studies in the literature.

Despite several constraining assumptions made in developing the ABM (e.g. discrete drought states), this model can provide valuable insights that enrich our understanding of the parameters that influence the adoption of drought forecasts, which can in turn be used to positively affect the adoption and diffusion of high-quality forecasts. In addition, once the model is tested and verified using fieldwork studies, it can be used to test the effectiveness of various intervention and targeting strategies and, ultimately, to develop more effective strategies and policies for overcoming impediments to forecast adoption. Several complementary methods could provide the necessary information for model validation: descriptive field studies, highly structured interviews, and laboratory or decision experiments. In a controlled laboratory experiment of the type that is traditionally employed in experimental economics (Kagel & Roth, 2015), researchers could observe how decision-makers respond to forecasts in stylized but reasonably realistic experiments (Millner, 2009; Sonka et al., 1988). Finally, given the demonstrated importance of social network structure for the diffusion process, field-based studies could also be used to represent a social network and its properties more realistically by extracting and mapping social and information networks and empirically analyzing the impacts of social networks and various social processes on the diffusion of forecasts.

Appendix A: Crop-Allocation Decision-Making

According to Equation 4, the optimal crop-allocation decision (x^*) depends on several factors, including the yield distribution (y), the cost function (c), initial wealth (ω), the coefficient of risk aversion (r), and beliefs about drought (p_θ). We make the following assumptions throughout: $y^A(0) = 0.06$, $y^A(1) = 0.03$, $y^B(0) = 0.08$, $y^B(1) = 0.01$, $c(1) = 0.04$, and $c(0) = 0.05$. Figure A1 shows how the optimal decision changes with r , p_θ , and ω . For a risk-neutral DM (i.e., $r = 0$), the optimal decision is to plant only crop B (i.e. $x_0^* = 0$) if $p_1 < 0.5$. As risk aversion increases, for any given p_1 or ω , a greater fraction of the land is allocated to crop A because crop A yields exhibit much less weather-related variation, helping risk-averse DMs minimize their risk exposure. Holding r constant, the fraction of land allocated to crop A increases as a DM's belief about a drought occurrence (p_1) increases. This is because crop A has a higher yield than crop B in drought conditions. Finally, more land is allocated to crop B as ω increases, as wealthier farmers' treatment of uncertainty more closely resembles that of a risk-neutral DM.

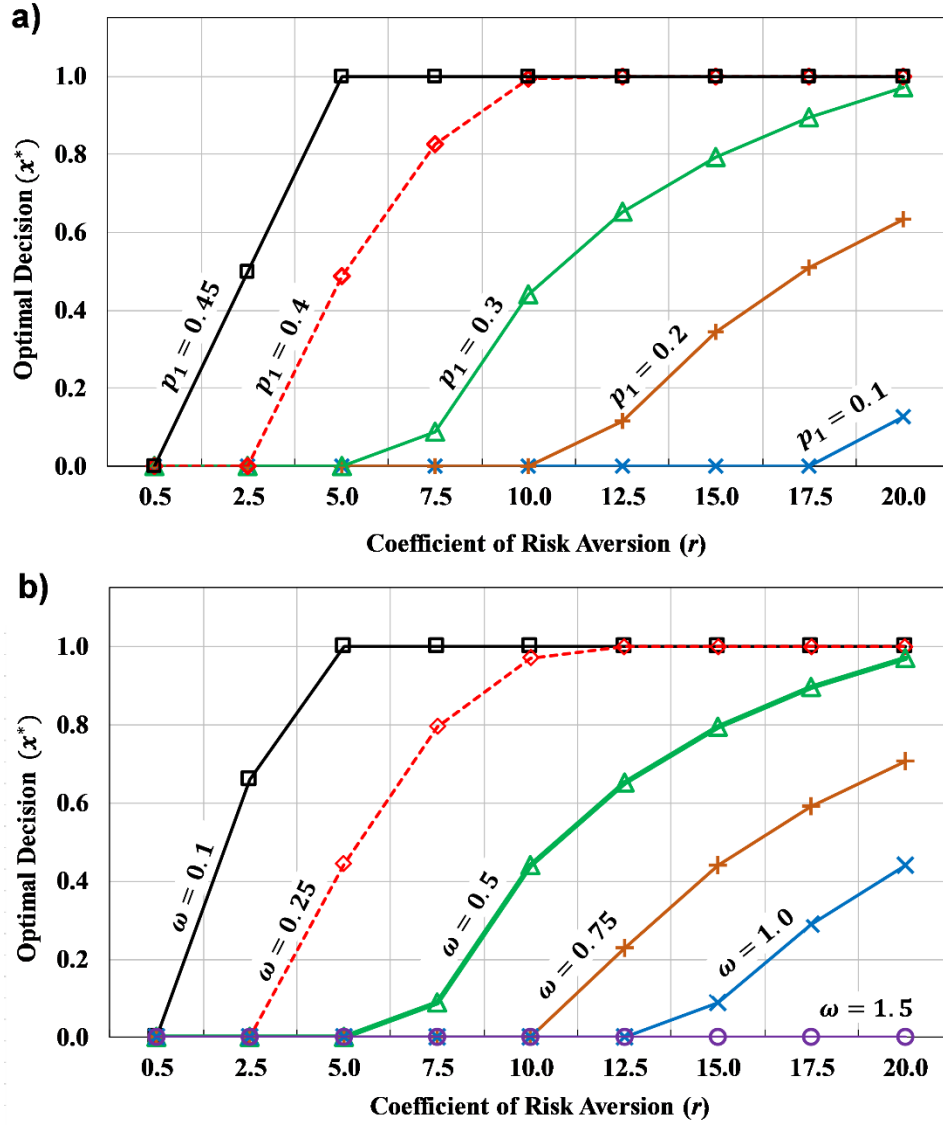


Figure A1. Optimal crop-allocation decisions: (a) sensitivity analysis for p_θ when $\omega = 0.5$, (b) sensitivity analysis for ω when $p_\theta = 0.3$.

Appendix B: Probabilistic Drought Forecast Generation

To compute the *ex post* value of forecasts, it is necessary to specify the time series of forecasts and drought realizations. One way to generate these time series is by using a joint distribution of forecasts and droughts, i.e. $f(\phi, p_d)$. We use an approach similar to ensemble forecasting to generate probabilistic drought forecasts based on a specified accuracy.

Assume that the time series of dichotomous drought events (φ_t) is known ($\varphi_t = 1$ indicates drought, $\varphi_t = 0$ indicates no drought). The system generates N deterministic forecasts of the dichotomous event at each time step t at accuracy κ . Each deterministic forecast (η_{i_t}) is referred to as an ensemble member, where $i \in [1, N]$. We assume that η_{i_t} is a Bernoulli process, defined as follows:

$$\begin{aligned} \eta_{i_t} &\sim Be(1, \kappa) & \varphi_t &= 1 \\ & & \text{if} & \\ \eta_{i_t} &\sim Be(1, 1 - \kappa) & \varphi_t &= 0 \end{aligned} \quad (B1)$$

where $Be(1, \kappa)$ indicates a binomial distribution with one trial and probability κ of success. Once N ensemble members are produced, p_d is calculated as:

$$p_{d_t} = \frac{\sum_{i=1}^N I_{\{\eta_{i_t}=1\}}}{N} \quad (B2)$$

where $\begin{cases} I_{\{\eta_{i_t}=1\}} = 1 \\ I_{\{\eta_{i_t}=1\}} = 0 \end{cases} \text{ if } \begin{cases} \eta_{i_t} = 1 \\ \eta_{i_t} = 0 \end{cases}$. The above definition has an undesirable property in that p_{d_t} could become zero or one, especially if N is small. Therefore, several post-processing methods have been suggested to account for finite ensemble size (Katz & Ehrendorfer, 2006; Roulston & Smith, 2002). We use:

$$p_{d_t} = \frac{\left(\sum_{i=1}^N I_{\{\eta_{i_t}=1\}}\right) + 0.5}{N + 1} \quad (B3)$$

In generating synthetic probabilistic drought forecasts, we also consider the possibility of low-probability events with no or limited predictability, such as the 2012 flash drought in the U.S. Midwest (Hoerling et al., 2014), by drawing a random number from a Bernoulli distribution $Be(1, 0.01)$ and flipping κ (i.e. using $1 - \kappa$ instead of κ) in Equation B1 if that random number equals 1.

Appendix C: List of Symbols

A list of mathematical notations used in the study is presented in Table C1.

Table C1. *Glossary of Notations*

Symbol	Definition
$c(\theta)$	non-land cost of crop production (e.g. fertilizers, see, labor) as a function of state of the weather, expressed in unit u
$E[\bullet]$	expectation operator
h	Forecast-adoption tendency; $h \in [0, 1]$
h^*	adoption threshold or cut-off, $h^* = 0.65$
$L(\bullet)$	learning function in the reinforcement-learning framework
\mathcal{M}	set of agents
m	total number of agents; $m = 625$
\mathcal{N}_i	set of neighbors of agent i
n_i	total number of neighbors of agent i
$p(\theta)$	user's belief about the occurrence of state θ of the random weather event

p_1	user's belief about the occurrence of drought: $p_1 = p(\theta = 1)$
p_d	probabilistic drought forecast: $p_d \in [0,1]$
r	coefficient of risk aversion, $r \geq 0$
S	reinforcement strength, expressed in unit u
SI_{in}	binary parameter indicating whether agents have social connections with neighbors <i>in</i> their counties; a connection exists if $SI_{in} = 1$.
SI_{out}	binary parameter indicating whether agents have social connections with neighbors <i>outside</i> of their counties; a connection exists if $SI_{out} = 1$.
t	time index
T	total number of time steps; $T = 625$
$U(\bullet)$	utility function
u	baseline unit used for $y(\theta)$, π , ω_0 , V^{exp} , W
V^{exp}	<i>ex post</i> value of forecast, expressed in unit u
W	Wealth, at the beginning of each time step, expressed in unit u ; W_1 is initial wealth.
x	decision variable: fraction of land allocated to crop A, $x \in [0,1]$
x^*	optimal crop-allocation decision
$x^{*,c}$	optimal crop-allocation decision based on p_θ (or climatology)
$x^{*,f}$	optimal crop-allocation decision based on p_d (forecast)
$y(\theta)$	crop yield as a function of the weather, expressed in unit u
y_0	crop yield in normal conditions ($\theta = 0$)
y_1	crop yield in drought conditions ($\theta = 1$)
z	forecast adoption decision, $z \in \{0,1\}$
α_{ij}	extent/strength of social interaction between agents i and j ; weight assigned by agent i to agent j 's belief
γ	coefficient of asymmetric learning
$\Delta = [\alpha_{ij}]$	m -by- m matrix of social interaction
η	deterministic drought forecast: $\eta \in \{0,1\}$
Θ	a set of possible states of the random weather event; of the binary drought event: $\Theta = \{0,1\}$
θ	random variable representing the uncertain weather event, $\theta \in \Theta$; $\theta = 0$: no drought, $\theta = 1$: drought
κ	forecast accuracy
$\pi(\bullet)$	normalized payoff function, expressed in unit u
τ	the learning rate
ϕ	a set of possible realized states of the event; for a binary drought event: $\phi = \{0,1\}$
φ	observation of the event, $\varphi \in \Phi$; $\varphi = 1$: drought, $\varphi = 0$: no drought

Acknowledgments, Samples, and Data

We acknowledge funding support from the NOAA Climate Program Office's Coping with Drought Initiative (grant number NA18OAR4310257). The data used in this study, including drought event time series, forecast time series, and agents' parameters, are provided in Data Set S1 and Data Set S2.

References

- Acemoglu, D., & Ozdaglar, A. (2011). Opinion Dynamics and Learning in Social Networks. *Dynamic Games and Applications*, 1(1), 3–49. <https://doi.org/10.1007/s13235-010-0004-1>
- Adhvaryu, A. (2014). Learning, misallocation, and technology adoption: Evidence from new malaria therapy in Tanzania. *Review of Economic Studies*, 81(4), 1331–1365. <https://doi.org/10.1093/restud/rdu020>
- Agrawala, S., & Broad, K. (2002). Technology Transfer Perspectives on Climate Forecast Applications. In M. de Laet (Ed.), *Research in Science and Technology Studies: Knowledge and Technology Transfer* (Vol. 13, pp. 45–69). Elsevier Science Ltd.
- Akerlof, G. A. (1997). Social Distance and Social Decisions. *Econometrica*, 65(5), 1005. <https://doi.org/10.2307/2171877>
- Arrow, K. J. (1962). The Economic Implications of Learning by Doing. *The Review of Economic Studies*, 29(3), 155. <https://doi.org/10.2307/2295952>
- Baerenklau, K. A. (2015). Toward an Understanding of Technology Adoption: Risk, Learning, and Neighborhood Effects. *Land Economics*, 81(1), 1–19. <https://doi.org/10.3368/le.81.1.1>
- Banerjee, A. V. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107(3), 797–817. <https://doi.org/10.2307/2118364>
- Barron, G., & Ursino, G. (2013). Underweighting rare events in experience based decisions: Beyond sample error. *Journal of Economic Psychology*, 39, 278–286. <https://doi.org/10.1016/j.joep.2013.09.002>
- Berger, T. (2001). Agent-based spatial models applied to agriculture: a simulation tool for technology diffusion, resource use changes and policy analysis. *Agricultural Economics*, 25(2–3), 245–260. <https://doi.org/10.1111/j.1574-0862.2001.tb00205.x>
- Besley, T., & Case, A. (1993). Modeling Technology Adoption in Developing Countries. *American Economic Review Papers and Proceedings*, 83(2), 396–402. <https://doi.org/10.1088/1757-899X/297/1/012024>
- Bharwani, S., Bithell, M., Downing, T. E., New, M., Washington, R., & Ziervogel, G. (2005). Multi-agent modelling of climate outlooks and food security on a community garden scheme in Limpopo, South Africa. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1463), 2183–2194. <https://doi.org/10.1098/rstb.2005.1742>
- Block, P. (2011). Tailoring seasonal climate forecasts for hydropower operations. *Hydrology and Earth System Sciences*, 15(4), 1355–1368. <https://doi.org/10.5194/hess-15-1355-2011>
- Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 3, 7280–7. <https://doi.org/10.1073/pnas.082080899>
- Brenner, T. (1999). *Modelling Learning in Economics*. Cheltenham, United Kingdom: Edward Elgar Publishing Limited.
- Brenner, T. (2006). Chapter 18 Agent Learning Representation: Advice on Modelling Economic Learning. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of Computational Economics* (Vol. 2, pp. 895–947). Elsevier. [https://doi.org/10.1016/S1574-0021\(05\)02018-6](https://doi.org/10.1016/S1574-0021(05)02018-6)

- Buizer, J., Jacobs, K., & Cash, D. (2016). Making short-term climate forecasts useful: Linking science and action. *Proceedings of the National Academy of Sciences*, 113(17), 4597–4602. <https://doi.org/10.1073/pnas.0900518107>
- Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, 58(5), 313–323. <https://doi.org/10.1037/h0054388>
- Bush, R. R., & Mosteller, F. (1953). A Stochastic Model with Applications to Learning. *The Annals of Mathematical Statistics*, 24(4), 559–585. <https://doi.org/10.1214/aoms/1177728914>
- Cai, J., Janvry, A. De, & Sadoulet, E. (2015). Social Networks and the Decision to Insure. *American Economic Journal: Applied Economics*, 7(2), 81–108. <https://doi.org/10.1257/app.20130442>
- Cazé, R. D., & Van Der Meer, M. A. A. (2013). Adaptive properties of differential learning rates for positive and negative outcomes. *Biological Cybernetics*, 107(6), 711–719. <https://doi.org/10.1007/s00422-013-0571-5>
- Crane, T. a., Roncoli, C., Paz, J., Breuer, N., Broad, K., Ingram, K. T., & Hoogenboom, G. (2010). Forecast Skill and Farmers' Skills: Seasonal Climate Forecasts and Agricultural Risk Management in the Southeastern United States. *Weather, Climate, and Society*, 2(1), 44–59. <https://doi.org/10.1175/2009WCAS1006.1>
- Cross, J. G. (1973). A Stochastic Learning Model of Economic Behavior. *The Quarterly Journal of Economics*, 87(2), 239. <https://doi.org/10.2307/1882186>
- DeGroot, M. H. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, 69(345), 118. <https://doi.org/10.2307/2285509>
- Duffy, J. (2006). Chapter 19 Agent-Based Models and Human Subject Experiments. In K. L. Tesfatsion, L. Judd (Ed.), *Handbook of Computational Economics* (Vol. 2, pp. 949–1011). Elsevier. [https://doi.org/10.1016/S1574-0021\(05\)02019-8](https://doi.org/10.1016/S1574-0021(05)02019-8)
- van Duinen, R., Filatova, T., Geurts, P., & Veen, A. van der. (2015). Empirical Analysis of Farmers' Drought Risk Perception: Objective Factors, Personal Circumstances, and Social Influence. *Risk Analysis*, 35(4), 741–755. <https://doi.org/10.1111/risa.12299>
- Ellison, G., & Fudenberg, D. (1993). Rules of Thumb for Social Learning. *Journal of Political Economy*, 101(4), 612–643. <https://doi.org/10.1086/261890>
- Feder, G., & Slade, R. (1984). Contact farmer selection and extension visits: the training and visit extension system in Haryana, India. *Quarterly Journal of International Agriculture*, 23(1), 6–21.
- Feder, G., & Slade, R. (1986). A Comparative Analysis of Some Aspects of the Training and Visit System of Agricultural Extension in India. *The Journal of Development Studies*, 22(2), 407–428. <https://doi.org/10.1080/00220388608421987>
- Feder, G., Just, R. E., & Zilberman, D. (1985). Adoption of Agricultural Innovations in Developing Countries: A Survey. *Economic Development and Cultural Change*, 33(2), 255–298. <https://doi.org/10.1086/451461>
- Foster, A. D., & Rosenzweig, M. R. (1995). Learning by Doing and Learning from Others:

- Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, 103(6), 1176–1209. <https://doi.org/10.1086/601447>
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, 306(5703), 1940–1943. <https://doi.org/10.1126/science.1102941>
- Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences of the United States of America*, 104(41), 16311–16316. <https://doi.org/10.1073/pnas.0706111104>
- Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic Bulletin and Review*, 22(5), 1320–1327. <https://doi.org/10.3758/s13423-014-0790-3>
- Gollier, C. (2001). *The Economics of Risk and Time. The Economics of Risk and Time*. The MIT Press. <https://doi.org/10.7551/mitpress/2622.001.0001>
- Golub, B., & Jackson, M. O. (2010). Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2(1), 112–149. <https://doi.org/10.1257/mic.2.1.112>
- Hallstrom, D. G. (2004). Interannual Climate Variation, Climate Prediction, and Agricultural Trade: the Costs of Surprise versus Variability. *Review of International Economics*, 12(3), 441–455. <https://doi.org/10.1111/j.1467-9396.2004.00460.x>
- Hansen, J. W. (2005). Integrating seasonal climate prediction and agricultural models for insights into agricultural practice. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 360(1463), 2037–47. <https://doi.org/10.1098/rstb.2005.1747>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Is reading about the kettle the same as touching it? Decisions from experience and the effects of rare events in risky choice. *Psychological Science*, 15(8), 534–539. <https://doi.org/10.1093/geronb/gbt081>
- Hoerling, M., Eischeid, J., Kumar, a., Leung, R., Mariotti, a., Mo, K., et al. (2014). Causes and Predictability of the 2012 Great Plains Drought. *Bulletin of the American Meteorological Society*, 95(2), 269–282. <https://doi.org/10.1175/BAMS-D-13-00055.1>
- Holloway, G., & Lapar, M. L. A. (2007). How big is your neighbourhood? Spatial implications of market participation among filipino smallholders. *Journal of Agricultural Economics*, 58(1), 37–60. <https://doi.org/10.1111/j.1477-9552.2007.00077.x>
- Hu, Q., Zillig, L. M. P., Lynne, G. D., Tomkins, A. J., Waltman, W. J., Hayes, M. J., et al. (2006). Understanding Farmers' Forecast Use from Their Beliefs, Values, Social Norms, and Perceived Obstacles. *Journal of Applied Meteorology and Climatology*, 45(9), 1190–1201. <https://doi.org/10.1175/JAM2414.1>
- Jadbabaie, A., Molavi, P., Sandroni, A., & Tahbaz-Salehi, A. (2012). Non-Bayesian social learning. *Games and Economic Behavior*, 76(1), 210–225. <https://doi.org/10.1016/j.geb.2012.06.001>
- Johnson, S. R., & Holt, M. T. (1997). The value of weather information. In R. W. Katz & A. H. Murphy (Eds.), *Economic Value of Weather And Climate Forecasts* (pp. 75–108). Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511608278.004>

- Kagel, J. H., & Roth, A. E. (2015). *Handbook of Experimental Economics: Vol2*. Princeton, New Jersey: Princeton University Press.
- Katz, R. W., & Ehrendorfer, M. (2006). Bayesian Approach to Decision Making Using Ensemble Weather Forecasts. *Weather and Forecasting*, 21(2), 220–231. <https://doi.org/10.1175/WAF913.1>
- Lawrence, D. B. (1999). *The Economic Value of Information*. New York, NY: Springer. <https://doi.org/10.1007/978-1-4612-1460-1>
- Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S., & Palminteri, S. (2017). Behavioural and neural characterization of optimistic reinforcement learning. *Nature Human Behaviour*, 1(4), 1–9. <https://doi.org/10.1038/s41562-017-0067>
- Lindner, R., Fischer, A., & Pardey, P. (1979). The time to adoption. *Economics Letters*, 2(2), 187–190. [https://doi.org/10.1016/0165-1765\(79\)90171-X](https://doi.org/10.1016/0165-1765(79)90171-X)
- Luseno, W. K., McPeak, J. G., Barrett, C. B., Little, P. D., & Gebru, G. (2003). Assessing the value of climate forecast information for pastoralists: Evidence from Southern Ethiopia and Northern Kenya. *World Development*, 31(9), 1477–1494. [https://doi.org/10.1016/S0305-750X\(03\)00113-X](https://doi.org/10.1016/S0305-750X(03)00113-X)
- Mansfield, E. (1961). Technical Change and the Rate of Imitation. *Econometrica*, 29(4), 741. <https://doi.org/10.2307/1911817>
- Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3), 531. <https://doi.org/10.2307/2298123>
- Marx, S. M., Weber, E. U., Orlove, B. S., Leiserowitz, A., Krantz, D. H., Roncoli, C., & Phillips, J. (2007). Communication and mental processes: Experiential and analytic processing of uncertain climate information. *Global Environmental Change*, 17(1), 47–58. <https://doi.org/10.1016/j.gloenvcha.2006.10.004>
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (2012). *Microeconomic Theory*. Oxford University Press.
- Mase, A. S., & Prokopy, L. S. (2014). Unrealized Potential: A Review of Perceptions and Use of Weather and Climate Information in Agricultural Decision Making. *Weather, Climate, and Society*, 6(1), 47–61. <https://doi.org/10.1175/WCAS-D-12-00062.1>
- Millner, A. (2009). What Is the True Value of Forecasts? *Weather, Climate, and Society*, 1(1), 22–37. <https://doi.org/10.1175/2009WCAS1001.1>
- Molavi, P., Tahbaz-Salehi, A., & Jadbabaie, A. (2018). A Theory of Non-Bayesian Social Learning. *Econometrica*, 86(2), 445–490. <https://doi.org/10.3982/ECTA14613>
- Munshi, K. (2004). Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution. *Journal of Development Economics*, 73(1), 185–213. <https://doi.org/10.1016/j.jdeveco.2003.03.003>
- Ng, T. L., Eheart, J. W., Cai, X., & Braden, J. B. (2011). An agent-based model of farmer decision-making and water quality impacts at the watershed scale under markets for carbon allowances and a second-generation biofuel crop. *Water Resources Research*, 47(9), n/a-n/a. <https://doi.org/10.1029/2011WR010399>

- 964 Nidumolu, U., Lim-Camacho, L., Gaillard, E., Hayman, P., & Howden, M. (2018). Linking
 965 climate forecasts to rural livelihoods: Mapping decisions, information networks and value
 966 chains. *Weather and Climate Extremes*, (June), 100174.
 967 <https://doi.org/10.1016/j.wace.2018.06.001>
- 968 Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a
 969 risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*,
 970 32(2), 551–562. <https://doi.org/10.1523/JNEUROSCI.5498-10.2012>
- 971 Rahimian, M. A., & Jadbabaie, A. (2017). Bayesian Learning Without Recall. *IEEE*
 972 *Transactions on Signal and Information Processing over Networks*, 3(3), 592–606.
 973 <https://doi.org/10.1109/TSIPN.2016.2631943>
- 974 Rasmussen, E. B., & Newland, M. C. (2008). Asymmetry of Reinforcement and Punishment in
 975 Human Choice. *Journal of the Experimental Analysis of Behavior*, 89(2), 157–167.
 976 <https://doi.org/10.1901/jeab.2008.89-157>
- 977 Rescorla, R. A. (2004). Spontaneous Recovery. *Learning & Memory*, 11(5), 501–509.
 978 <https://doi.org/10.1101/lm.77504>
- 979 Rogers, E. M. (2003). *Diffusion of Innovations* (Fifth). New York, NY: Free Press.
- 980 Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple
 981 dynamic models in the intermediate term. *Games and Economic Behavior*, 8(1), 164–212.
 982 [https://doi.org/10.1016/S0899-8256\(05\)80020-X](https://doi.org/10.1016/S0899-8256(05)80020-X)
- 983 Roulston, M. S., & Smith, L. A. (2002). Evaluating Probabilistic Forecasts Using Information
 984 Theory. *Monthly Weather Review*, 130(6), 1653–1660. [https://doi.org/10.1175/1520-0493\(2002\)130<1653:EPFUIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1653:EPFUIT>2.0.CO;2)
- 986 Rubas, D. J., Hill, H. S. J., & Mjelde, J. W. (2006). Economics and climate applications:
 987 exploring the frontier. *Climate Research*, 33, 43–54. <https://doi.org/10.3354/cr033043>
- 988 Rubas, D. J., Mjelde, J. W., Love, H. A., & Rosenthal, W. (2008). How adoption rates, timing,
 989 and ceilings affect the value of ENSO-based climate forecasts. *Climatic Change*, 86(3–4),
 990 235–256. <https://doi.org/10.1007/s10584-007-9293-9>
- 991 Sampson, G. S., & Perry, E. D. (2019). The role of peer effects in natural resource appropriation
 992 - The case of groundwater. *American Journal of Agricultural Economics*, 101(1), 154–171.
 993 <https://doi.org/10.1093/ajae/aay090>
- 994 Sherrick, B. J., Sonka, S. T., Lamb, P. J., & Mazzocco, M. A. (2000). Decision-maker
 995 expectations and the value of climate prediction information: Conceptual considerations and
 996 preliminary evidence. *Meteorological Applications*, 7(4), 377–386.
 997 <https://doi.org/10.1017/S1350482700001584>
- 998 Sonka, S. T., Changnon, S. A., & Hofing, S. (1988). Assessing Climate Information Use in
 999 Agribusiness. Part II: Decision Experiments to Estimate Economic Value. *Journal of*
 1000 *Climate*, 1(8), 766–774. [https://doi.org/10.1175/1520-0442\(1988\)001<0766:ACIUIA>2.0.CO;2](https://doi.org/10.1175/1520-0442(1988)001<0766:ACIUIA>2.0.CO;2)
- 1002 Stoneman, P. (1983). *The Economic Analysis of Technological Change*. Oxford University Press.
- 1003 Tarnoczi, T. J., & Berkes, F. (2010). Sources of information for farmers' adaptation practices in

- Canada's Prairie agro-ecosystem. *Climatic Change*, 98(1–2), 299–305.
<https://doi.org/10.1007/s10584-009-9762-4>
- Templeton, S. R., Hooper, A. A., Aldridge, H. D., & Breuer, N. (2018). Farmer Interest in and Uses of Climate Forecasts for Florida and the Carolinas: Conditional Perspectives of Extension Personnel. *Weather, Climate, and Society*, 10(1), 103–120.
<https://doi.org/10.1175/WCAS-D-16-0057.1>
- Tesfatsion, L. (2006). Chapter 16 Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of Computational Economics* (Vol. 2, pp. 831–880). Elsevier. [https://doi.org/10.1016/S1574-0021\(05\)02016-2](https://doi.org/10.1016/S1574-0021(05)02016-2)
- Thorndike, E. L. (1911). *Animal Intelligence*. New York, NY: Hafner Publishing.
- Thorndike, E. L. (1932). *The fundamentals of learning*. New York, NY: Teachers college, Columbia University.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Wakker, P. P. (2008). Explaining the characteristics of the power (CRRA) utility family. *Health Economics*, 17(12), 1329–1344. <https://doi.org/10.1002/hec.1331>
- Whateley, S., Palmer, R. N., & Brown, C. (2015). Seasonal Hydroclimatic Forecasts as Innovations and the Challenges of Adoption by Water Managers. *Journal of Water Resources Planning and Management*, 141(5), 04014071.
[https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000466](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000466)
- Wilks, D. S. (2006). *Statistical Methods in the Atmospheric Sciences*. Academic Press. Retrieved from [http://danida.vnu.edu.vn/cpis/files/Books/Statistical methods in the atmospheric sciences, D. Wilks \(2ed., IGS 91, Elsevier, 2006\)\(ISBN 0127519661\)\(649s\).pdf](http://danida.vnu.edu.vn/cpis/files/Books/Statistical%20methods%20in%20the%20atmospheric%20sciences,%20D.%20Wilks%20(2ed.,%20IGS%2091,%20Elsevier,%202006)(ISBN%200127519661)(649s).pdf)
- Ziervogel, G. (2004). Targeting seasonal climate forecasts for integration into household level decisions: the case of smallholder farmers in Lesotho. *The Geographical Journal*, 170(1), 6–21. <https://doi.org/10.1111/j.0016-7398.2004.05002.x>
- Ziervogel, G., & Downing, T. E. (2004). Stakeholder Networks: Improving Seasonal Climate Forecasts. *Climatic Change*, 65(1/2), 73–101.
<https://doi.org/10.1023/B:CLIM.0000037492.18679.9e>
- Ziervogel, G., Bithell, M., Washington, R., & Downing, T. (2005). Agent-based social simulation: A method for assessing the impact of seasonal climate forecast applications among smallholder farmers. *Agricultural Systems*, 83(1), 1–26.
<https://doi.org/10.1016/j.agsy.2004.02.009>

A Behavioral Social Learning Model for Studying the Dynamics of Forecast Adoption

Majid Shafiee-Jood¹, Tatyana Deryugina², and Ximing Cai^{1,3}

¹Ven Te Chow Hydrosystems Laboratory, Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA, ²Department of Finance, Gies College of Business, University of Illinois at Urbana-Champaign, Champaign, IL, USA, ³DOE Center for Advanced Bioenergy and Bioproducts Innovation, University of Illinois at Urbana-Champaign, Urbana, IL, USA.

Contents of this file

Figure S1.....	2
Figure S2.....	3
Figure S3.....	4
Figure S4.....	5
Figure S5	6
Figure S6	7
Figure S7	8
Figure S8	9
Figure S9	10

Additional Supporting Information (Files uploaded separately)

Data Sets S1 to S2
Movies S1 to S4

Introduction

The supporting information we provide includes Figures S1 through S9, Data Sets S1 and S2, and Movies S1 through S4. The figures provide supplementary information pertaining to the results and discussions presented in the paper. Data sets include time series of drought events and forecasts as well as the parameters used in the agent-based model. The animations show the spatial and temporal diffusion dynamics in the various scenarios considered in the main text.




Figure S1. Cumulative economic gains when decisions are informed by climatological information, forecasts, and the adoption mechanism shown in Figure 4.




Figure S2. Time series of *ex post* values of forecasts corresponding to belief trajectories shown in Figure 5b for (a) $r = 0.5$ and (b) $r = 10$.




Figure S3. Time series of *ex post* values of forecasts corresponding to belief trajectories shown in Figure 5c for (a) $\omega = 0.25$, (b) $\omega = 0.5$, and (c) $\omega = 1$.




Figure S4. Agents' risk aversion (r), randomly generated according to $r \sim N(7.5, 1.5)$. See Data Set S1 for the parameter values.




Figure S5. Agents' learning rate (τ), randomly generated according to $\tau \sim N(3.0, 1.0)$. See Data Set S1 for the parameter values.




Figure S6. Agents' initial wealth (ω), randomly generated according to $\omega \sim N(0.5, 0.05)$.
See Data Set S1 for the parameter values.




Figure S7. Diffusion curves for learning-rate scenarios: (a) Full social interaction, (b) intra-county ties only, and (c) no interaction. Risk aversion and initial wealth are randomly assigned based on $r \sim N(7.5, 1.5)$ and $\omega \sim N(0.5, 0.05)$. See Data Set S1 for the parameter values.




Figure S8. Impact of forecast accuracy on forecast-adoption diffusion. (a) $\kappa = 0.8$, (b) $\kappa = 0.75$, (c) $\kappa = 0.725$, (d) $\kappa = 0.7$, (e) $\kappa = 0.675$, and (f) $\kappa = 0.65$. Shaded areas represent the envelopes of diffusion curves for the 15 realizations of the forecast time series considered, and solid black lines represent mean values.




Figure S9. Evolution of the probability of forecast adoption averaged across all agents and 15 realizations of forecast time series in forecast-accuracy scenarios.

Data Set S1. Time series of drought events and drought forecasts as well as randomly generated model parameter values, including learning rate (τ), the coefficient of risk aversion (r), and initial wealth (ω).

Data Set S2. Time series of drought forecasts generated based on the following forecast accuracy values ($\kappa = 0.6, 0.625, 0.65, 0.675, 0.7, 0.725, 0.75, 0.8$). For each case, 15 realizations are generated.

Movie S1. The impact of social interaction on forecast-adoption diffusion in the case-study area. Two social interaction scenarios are considered: Full interaction (i.e. $SI_{in} = 1$ and $SI_{out} = 1$) and no interaction (i.e. $SI_{in} = 1$ and $SI_{out} = 0$).

Movie S2. The impact of change agents on forecast-adoption diffusion in the case-study area. All agents in the middle county (county 13) are considered change agents. Learning rate (τ) values are randomly assigned based on a normal distribution $N(3,1)$.

Movie S3. The impact of change agents on forecast-adoption diffusion in the case-study area. Agents whose initial wealth is above the 90th percentile are considered change agents. Two self-reliance scenarios are considered: $\alpha_{ii} = 0.5$ and $\alpha_{ii} = 1.0$.

Movie S4. The impact of asymmetrical learning on forecast-adoption diffusion in the case-study area. Two scenarios are considered: $\gamma = 1$ and $\gamma = 1.5$.