Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Managing Weather Risk with a Neural Network-Based Index Insurance

Zhanhui Chen, Yang Lu, Jinggong Zhang, Wenjun Zhu

informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–22
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Managing Weather Risk with a Neural Network-Based Index Insurance

Zhanhui Chen,[a,*] **Yang Lu,**[b] **Jinggong Zhang,**[c] **Wenjun Zhu**[c]

[a] Department of Finance, School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong SAR, China; [b] Department of Mathematics & Statistics, Concordia University, Montreal, Quebec H3G 1M8, Canada; [c] Division of Banking and Finance, Nanyang Business School, Nanyang Technological University, Singapore 639798, Singapore

*Corresponding author

Contact: chenzhanhui@ust.hk, https://orcid.org/0000-0003-2766-271X (ZC); yang.lu@concordia.ca, https://orcid.org/0000-0002-7141-1048 (YL); jgzhang@ntu.edu.sg, https://orcid.org/0000-0002-3287-4648 (JZ); wjzhu@ntu.edu.sg, https://orcid.org/0000-0002-7251-9579 (WZ)

Received: August 21, 2021
Revised: March 8, 2022; July 12, 2022; October 5, 2022
Accepted: November 15, 2022
Published Online in Articles in Advance: August 28, 2023

https://doi.org/10.1287/mnsc.2023.4902

Abstract. Weather risk affects the economy, agricultural production in particular. Index insurance is a promising tool to hedge against weather risk, but current piecewise-linear index insurance contracts face large basis risk and low demand. We propose embedding a neural network-based optimization scheme into an expected utility maximization problem to design the index insurance contract. Neural networks capture a highly nonlinear relationship between the high-dimensional weather variables and production losses. We endogenously solve for the optimal insurance premium and demand. This approach reduces basis risk, lowers insurance premiums, and improves farmers' utility.

History: Accepted by Agostino Capponi, finance.

Funding: This work was supported by the Research Grants Council, University Grants Committee [Grants GRF 16502020, GRF 16504522, and T31-603/21-N], Singapore Ministry of Education Academic Research Fund Tier 1 [Grants RG143/19 and RG55/20], the Natural Sciences and Engineering Research Council of Canada [Grants RGPIN-2021-04144 and DGECR-2021-00330], the Research Database Matching Fund, and the School of Business and Management, Hong Kong University of Science and Technology.

Supplemental Material: The data files and online appendices are available at https://doi.org/10.1287/mnsc.2023.4902.

Keywords: neural networks • weather risk • index insurance • basis risk • utility maximization

1. Introduction

Climate change and weather risk affect the economy and livelihood over large scales (Nordhaus 2019, Hong et al. 2020), especially for agricultural production, which depends on weather conditions (Fisher et al. 2012, Lesk et al. 2016). For example, the USDA (2014) estimates that 70%–90% of agricultural production loss can be attributed to adverse weather. In practice, weather index insurance can be used to hedge against weather risk (Turvey 2001). The payoff of weather index insurance is exclusively based on some prespecified weather indices instead of the actual losses incurred to the insureds. Hence, it avoids the high administration costs, adverse selection, and moral hazard issues associated with conventional indemnity-based insurance. Although promising, current index insurance faces low demand (Cole et al. 2013, Clarke 2016).

Large basis risk, the risk that the underlying indices and actual losses are mismatched, contributes to the low insurance demand (Cummins et al. 2004, Clarke 2016, Jensen et al. 2016).[1] Most current index insurances are based on low-dimensional weather indices and use linear-type payoff functions (Mahul and Skees 2007). Figure 1(a) plots the payoffs of a currently used linear index insurance based on a temperature index, exhibiting large basis risk. Moreover, given the large basis risk, the remaining uninsured risk could be high; therefore, the current insurance appears to be not cost effective for farmers (Jensen et al. 2016), which further hinders farmers' participation. This paper attempts to address these two issues to design better index insurance contracts.

We aim to reduce the basis risk with two considerations. First, crop production depends on high-dimensional weather conditions. We should carefully select and include a sufficiently large number of weather variables when constructing the contract. Second, for biological reasons, crop production depends on weather conditions in a highly nonlinear way (Schlenker and Roberts 2009, Rigden et al. 2020). Tapping into recent advances in machine learning, we propose a neural network–based (NN-based) index insurance design. Neural networks help to capture high-dimensional, nonlinear, and complex interactions between weather indices and production losses. Specifically, we embed an NN-based scheme into an expected utility maximization problem with budget constraints to capture farmers' insurance demand.

Figure 1. Payoffs of Index Insurance and Actual Losses

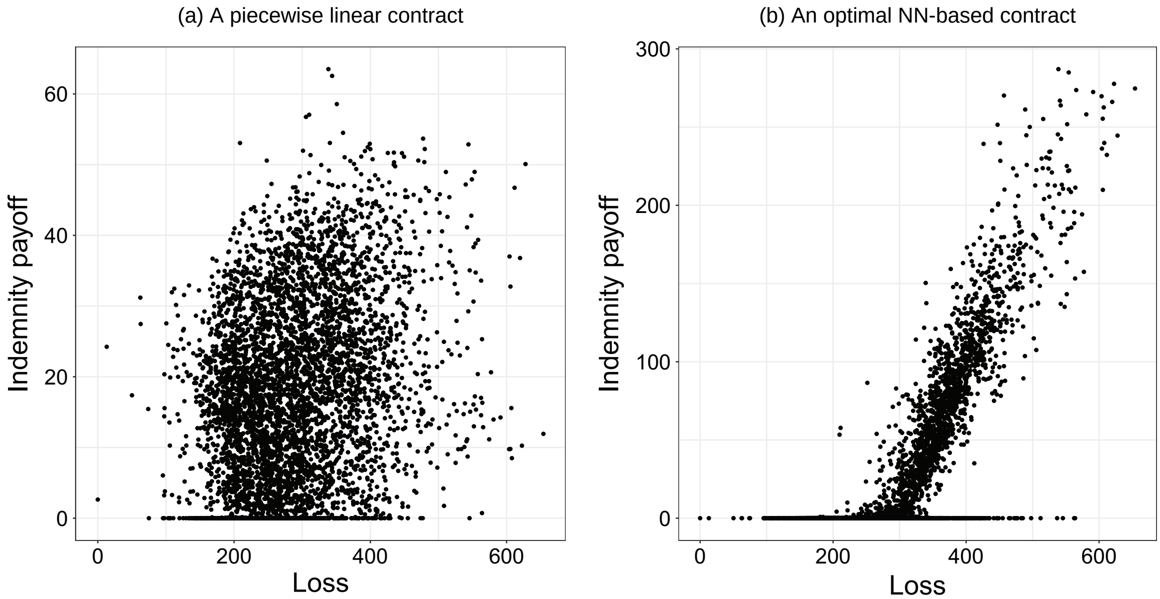

Notes. These panels plot the insurance payoffs against actual losses to illustrate the basis risk. (a) Piecewise-linear insurance contract based on a temperature index, which is commonly used in the current practice. (b) NN-based index insurance contract.

Moreover, in addition to crop yield prediction, we endogenously determine the equilibrium insurance premium and demand, while the current literature typically assumes exogenously given insurance premium (using burning cost or via an ad hoc risk loading). This economic mechanism incorporates strategic interaction between farmers and insurers, which is important in two ways. First, it allows for a fair evaluation of welfare improvement for farmers. Second, once we consider the responses of both insureds and insurers in an equilibrium setting, it improves the cost effectiveness and feasibility of the designed insurance contract.

We apply this NN-based index insurance to corn production in Illinois. We first explore various NN structures to identify the optimal neural network, which has three hidden layers (64, 64, and 16 neurons in each layer). Results show that the NN-based index insurance improves farmers' utility because this contract reduces their exposure to wealth variations, especially the extreme downside risk, and it is cost effective. For example, using a constant absolute risk aversion (CARA) utility, we show that the optimal NN-based index insurance improves farmers' utility by 14.35% for the test sample. Certainty equivalent wealth (CEW) improves by $22,767.36 for the test sample for an average farmer, which is about 4.98% of the farmer's wealth. The optimal NN-based index insurance has a premium of $28.72/acre, which is 63% of the current corn insurance premium ($45.50/acre) in Illinois, making the contract very attractive. Figure 1(b) plots the payoffs of the optimal NN-based index insurance policy. Comparing panels (a) and (b), we see that the optimal NN-based contract effectively reduces basis risk. In fact, the payoff function of our proposed NN-based index insurance is very close to a conventional indemnity-based insurance with a deductible.

We dig deeply in two ways to understand the economic mechanism of the superior performances of the NN-based index insurance contract. First, following Cong et al. (2021a) and Ribeiro et al. (2016), we perform sensitivity analyses to improve model transparency and interpretability. Interestingly, we find that some weather indices outside growing seasons (which are from May to October for corns in Illinois) are important, for example, dew point temperature and maximum temperature in December, whereas these variables are overlooked by the linear insurance contract in current practice. This demonstrates the power of neural networks to extract important information from a large set of input factors.

Second, we compare the performances of NN-based index insurance with other insurances that are based on polynomial payoff functions, such as linear polynomials (the common practice nowadays), quadratic polynomials, and cubic polynomials. We also consider contracts using different numbers of weather indices as input factors. We find that including more weather variables and nonlinear terms improve farmers' utility. These results suggest the impressive performances of NN-based contract are indeed due to its ability to extract complicated information from a large set of inputs, which are often nonlinear and nonmonotonic.

We further verify the robustness of NN-based index insurance in several ways. First, we perform an "out-of-state" tests, using major corn producing states adjacent to Illinois that have similar latitudes, namely, Indiana, Iowa, and Missouri. That is, we use Illinois data as the training sample and data from the three adjacent states as the test sample. Second, we consider the impacts of overinsuring constraint. Third, we examine the impacts of weather predictability on insurance design. Last, we consider the impacts of insurers' choices (such as various insurers' supply curves and exogenously given risk loading) and farmers' characteristics (such as different coverage levels, different risk aversion levels, time-varying risk aversion, and alternative utility functions).

Although effective, the NN-based contract is more complicated than linear contracts. Complexity aversion might deter farmers' participation. We address this complexity concern in four ways. First, we quantify the impacts of complexity aversion on index insurance performance. We find that NN-based model still outperforms polynomial models after considering the perceived value reduction by complexity averse farmers as suggested in Ceballos and Robles (2020). Second, improving the interpretability of NN-based contracts and farmers' insurance literacy can boost farmers' trust in this product (Gaurav et al. 2011, Cai et al. 2020). Third, government subsidies can significantly increase farmers' participation (Cai et al. 2020). Last, we propose a hybrid contract that provides payoffs as the maximum of a linear contract and an NN-based contract as an initial attempt. Although suboptimal, such a hybrid contract enjoys easy interpretability of a linear contract and large basis risk reduction of the NN-based contract.

Our framework can be easily extended to designing other weather risk management solutions. For example, we extend our framework to the case of revenue insurance which protects both the production and corn price risks. We also apply our approach to the case where a risk-neutral agent optimizes the tail risk, for example, the value-at-risk (VaR), instead of expected utility maximization. This can be useful for corporate farms and internal risk management of insurance companies.

This paper contributes to the growing literature applying machine learning in finance (see Capponi and Lehalle (2023) for a recent synthesis). Various machine learning algorithms have been examined to construct optimal portfolio, identify return factors, and evaluate model performances (Rossi and Timmermann 2015; Feng et al. 2020, 2021; Gu et al. 2020, 2021; Bianchi et al. 2021; Bryzgalova et al. 2021; Cong et al. 2021a, b, 2022; Li and Rossi 2021; Avramov et al. 2022; Huang and Shi 2022; Chen et al. 2023). Machine learning is also applied to risk management, including default risk, credit risk, and mortgage risk (Sirignano et al. 2021). Recently, Alsabah et al. (2021), Rossi and Utkus (2021), and Capponi et al. (2022) investigate how robo-advising helps learn investors' risk preference and improve their investment portfolios. Our paper applies neural networks to design index insurance for better managing weather risk, which also can be a robo-advising product.

This paper also contributes to the index insurance literature. Several field studies investigate the demand and efficiency of index insurance (Cole et al. 2013, 2014; Casaburi and Willis 2018). Clarke (2016) and Jensen et al. (2016) show that basis risk causes the low demand of index insurance. Assa and Wang (2021) design index insurance on agricultural goods that provides a high Sharpe ratio. Our paper illustrates the promise of an NN-based approach to reduce the basis risk and improve the demand of index insurance.

Our paper is closely related to the literature that uses machine learning techniques to improve crop yield predictions. For example, You et al. (2017), Crane-Droesch (2018), and Newlands et al. (2019) use deep learning models to predict agricultural yields. Complementing the literature, our paper applies the NN model to predict crop yields and importantly, it considers the endogenous pricing of the insurance contract, which affects insurance demand and basis risk. This is an important economic mechanism overlooked in the previous research.

This paper broadly belongs to the climate finance literature (see Hong et al. (2020) for a comprehensive overview). It is critical to design new tools to financially manage weather risks. Recently, Engle et al. (2019) consider climate risk hedging with mimicking portfolios. Our paper adds to the literature by proposing a new design of the weather-related insurance contract to manage weather risk.

The rest of this paper proceeds as follows. Section 2 introduces the model and assumptions. Section 3 discusses the NN-based approach for solving the optimal index insurance contract. Section 4 presents the empirical performance of the NN-based insurance policy. Section 5 provides robustness checks. Section 6 discusses several extensions of the model. Section 7 concludes.

2. Model

Consider a typical farmer who would like to hedge the exposure to weather risk, subject to certain budget constraints. Suppose the farmer has an initial wealth of w_0 and faces a random production loss of Y during a year. We assume there is an index insurance contract with a payoff determined by a p-dimensional random vector of indices, $\boldsymbol{X} = (X_1, X_2, \ldots, X_p)$. More specifically, the index insurance payoff is $I(\boldsymbol{X})$, where $I : \mathbb{R}^p \longmapsto \mathbb{R}^+$ is the nonnegative payoff function. Denote the premium of the index insurance contract by $\pi(I)$, which is a functional of the indemnity function I. Then the terminal wealth of the farmer at the presence of the index insurance is $w_0 - Y + I(\boldsymbol{X}) - \pi(I)$. Our objective for the index insurance design is to select the optimal indemnity function I so that the policyholder's expected utility, $\mathrm{E}(U)$, is maximized under certain budget constraints. In this paper, we focus on designing index insurance against production losses, without taking into account crop price risk, that is, a yield insurance.[2] Specifically, the index insurance design problem can be formulated as the following expected utility maximization problem:

$$\begin{cases} \sup_{I\in\mathcal{I}} & \mathrm{E}(U[w_0 - Y + I(\boldsymbol{X}) - \pi(I)]) \\ \text{s.t.} & P_L \le \pi(I) \le P_U, \end{cases} \tag{1}$$

where $\mathcal{I} := \{I : \mathbb{R}^p \longmapsto \mathbb{R}^+ \mid I \text{ is measurable}\}$ defines the feasible set of indemnity function I. The budget constraint of the farmer is given by P_L and P_U, that is, the lower and upper bounds of the premium level, which are given exogenously. The premium bounds indicate the price range that the farmer is willing to accept, and they might be chosen as zero and the farmer's highest affordable price, respectively. The premium bounds also prevent abusive and speculative use of insurances in practice, especially when there are government subsidies.

In this paper, we employ the most commonly used premium principle both in the literature and in practice to determine the premium of the index insurance contract, the expectation premium principle. That is, the index insurance premium $\pi(I)$ is proportional to the risk premium as follows:

$$\pi(I) = \lambda \mathrm{E}[I(\boldsymbol{X})], \tag{2}$$

where λ is the risk loading parameter and $\lambda \ge 1$. When $\lambda = 1$, $\pi(I)$ is called the actuarially fair premium. λ affects the insurance premium and insurer's profits and impacts farmers' index insurance demand. In this paper, we endogenously determine the equilibrium λ from the supply and demand curves of insurance contracts, which will be discussed in more details in Section 4.2.1. Such endogenous insurance premium takes into account the strategic interaction between farmers and insurers while the model solves for the utility maximization problem of farmers.

We need to compute the expected utility to solve Problem (1). Instead of modeling the joint distribution of the loss Y and indices $\boldsymbol{X}$, we replace the expected utility with its empirical counterparts and directly search for the optimal index insurance. Specifically, for a random sample of $(\boldsymbol{X}, Y)$: $\{(\boldsymbol{x}_j, y_j)\}_{j=1,\dots,n}$, where $\boldsymbol{x}_j = (x_{j1}, x_{j2}, \dots, x_{jp})$, after replacing quantities with sample statistics, Problem (1) can be reformulated as the following minimization problem:

$$\begin{cases} \min_{I\in\mathcal{I}} & -\frac{1}{n}\sum_{j=1}^{n} U(w_0 - y_j + I(\boldsymbol{x}_j) - \pi_e(I)]) \\ \text{s.t.} & P_L \le \pi_e(I) \le P_U, \end{cases} \tag{3}$$

where $\pi_e(I)$ is the empirical counterpart of $\pi(I)$ and is given by $\pi_e(I) := (\lambda/n)\sum_{j=1}^{n} I(\boldsymbol{x}_j)$.

3. NN-Based Solution

It is challenging to search for the optimal functional form of the insurance payout I in the feasible set $\mathcal{I}$ for the optimization problem (3). As a result, the existing literature has considered a certain restrictive feasible set $\tilde{\mathcal{I}}_0 \subset \mathcal{I}$. Step functions (where index insurance payment is triggered by some predefined events) and piecewise linear functions (an excess-of-loss-type contract where the triggered payment is a linear function of the index) are commonly used functional forms in the literature and in practice (Mahul and Skees 2007). However, a restrictive feasible set generates contracts with large basis risk, which leads to market failure of index insurance (Clarke 2016). Hence, it is necessary to allow for a larger set of feasible functional forms. However, if the set of feasible functional forms is too large, we will run into "overfitting" issues (see Online Appendix A for an example). Therefore, we need to choose a feasible set that balances flexibility and stability.

In this paper, we consider a feasible set which allows for nonlinear and nonmonotonic relationships. Specifically, we apply neural networks to search for the optimal contract in the expanded feasible set $\mathcal{I}_0 \subset \mathcal{I}$ (see Online Appendix B for comparisons of different feasible sets). NN-based models are attractive to our research question, for a few reasons.[3] First, NNs are designed to capture high dimensionality and nonlinearity. In particular, NN-based models can incorporate complex interactions among input factors via linear combinations of features and activation functions (Hastie et al. 2009). This is important for our research question because weather variables have complicated interactions that affect production losses. Second, NN-based models generate smooth outputs that are desirable for designing insurance payoff functions. Third, NNs are sensitive to critical details and insensitive to idiosyncratic outliers in the data and can potentially be useful in achieving good flexibility-stability balance in solving for the optimal contract.

3.1. Neural Networks and Feasible Set

We consider a standard multilayer fully connected NN (see Online Appendix C for an illustration of the NN structure); $\boldsymbol{X}_{Input}$ in the input layer includes all indices used to construct the index insurance contract, whereas I in the output layer represents the final payoff function solved by this system. In addition to an input layer and an output layer, there are $H (H \ge 1)$ additional hidden layers. The hth $(h = 1, 2, \dots, H)$ hidden layer $\boldsymbol{Z}^{(h)}$ contains p_h neurons, which are obtained by transforming the linear combination of the neurons from the previous layer through an activation function, f_h, elementwise; $\boldsymbol{\alpha}^{(h)}$ and $\boldsymbol{\omega}^{(h)}, h = 1, 2, \dots, H$, are parameters of the linear combination; $\boldsymbol{\alpha}^{(h)}$ is a $(p_h \times 1)$ vector of bias units that captures the intercepts in the model; and $\boldsymbol{\omega}^{(h)}$ is a $(p_h \times p_{h-1})$-dimensional weight matrix. $\boldsymbol{\alpha}^{(h)}$ and $\boldsymbol{\omega}^{(h)}$ $(h = 1, 2, \dots, H)$ are learned through stochastic gradient descent (SGD), where their gradients are derived by backpropagation. This network is fully connected, in the sense that neurons between two adjacent layers are fully pairwise connected, but neurons within a layer have no connections. In summary, an NN structure is defined by its number of hidden layers H, the number of neurons in each hidden layer p_h, activation functions f_h, and the parameters $\boldsymbol{\alpha}^{(h)}, \boldsymbol{\omega}^{(h)}$ $(h = 1, 2, \ldots, H)$. A specific NN structure defines the feasible set $\mathcal{I}_0 \subset \mathcal{I}$ in the optimization problem (3).

This framework conveniently guarantees that I is nonnegative, which is necessary for insurance indemnity payoffs, for example, by making the last layer activation function, f_H, nonnegative. Throughout this paper, we use the rectified linear unit (RELU), defined by $f(x) = \max(x, 0)$, as the nonlinear activation function in our empirical analysis.

One might be concerned with the overfitting issue of neural networks. To avoid overfitting, we impose complexity constraints on function I by using early stopping. We also try other complexity constraints, including regularization methods (e.g., L^1 and L^2 regularization and their combination), elastic net, and the dropout method, and we find similar results.

3.2. Solving for the Optimal Policy Using a Penalty Method

Two issues remain before solving for the optimal index insurance policy. First, the objective function in Problem (3) is a utility function and not a typical loss function in machine learning. Second, there is a budget constraint in Problem (3), that is, $P_L \le \pi_e(I) \le P_U$. The first issue can be solved by defining a customized loss function when we formulate an NN program.[4] To maneuver the second issue, we propose a penalty method.

Let's consider a sequence of unconstrained problems $\{\Phi_k\}_{k\ge 0}$:

$$\Phi_k = \min_{I\in\mathcal{I}_0}\left(-\frac{1}{n}\sum_{j=1}^{n} U(w - y_j + I(\boldsymbol{x}_j) - \pi_e(I)) + \phi_k \cdot g(I)\right), \tag{4}$$

where $\{\phi_k\}_{k\ge 0}$ is a sequence of increasing nonnegative real numbers tending toward infinity (i.e., $\phi_k \ge 0, \phi_{k+1} \ge \phi_k, \lim_{k\to\infty}\phi_k = +\infty$), and $g(I)$ is the penalty function. The penalty function is defined as a sum of squared errors (Luenberger and Ye 1984), as follows:

$$g(I) = [\max(\pi_e(I) - P_U, 0)]^2 + [\max(P_L - \pi_e(I), 0)]^2. \tag{5}$$

Our index insurance design problem (3) then can be connected to the unconstrained problem (4), using the following theorem.

Theorem 1 (Luenberger and Ye 1984). *Let I_k^* be a sequence of solutions solving the corresponding sequence of unconstrained problems $\Phi_k(I)$, as defined in Problem (4), where $\{\phi_k\}_{k\ge 0}$ is an increasing sequence tending toward infinity. Then any limit of $\{I_k^*\}_{k\ge 0}$ is a minimizer of the constrained problem (3).*

Therefore, based on Theorem 1, we construct a sequence of increasing penalty coefficients, $\{\phi_k\}_{k\ge 0}$, and specify the penalty function, $g(I)$, according to Equation (5). The optimal contract for Problem (3) can be solved from the unconstrained problem (4) iteratively. Online Appendix D summarizes the algorithm.

4. Performances of the NN-Based Index Insurance

In this section, we evaluate the empirical performances of the proposed NN-based design of a weather index insurance policy. To proceed, we consider a representative corn farmer with a negative exponential utility function $U(w) = -(1/\alpha)e^{-\alpha w}$, where $\alpha > 0$ is the absolute risk aversion coefficient. Using a CARA utility provides some advantages. First, under this CARA utility, the optimal index insurance contract in Problem (3) is independent of farmers' initial wealth level. This property is realistic from the practical viewpoint of designing insurance policies. It also makes results comparable among different policyholders. Second, the negative exponential utility function conveniently handles negative wealth (i.e., bankruptcy) cases, whereas other utility functions, such as power utility functions, may not be well defined when wealth is negative, which could occur under some catastrophic events. Therefore, we use a negative exponential utility function in our main results. As a robustness check, we consider other utility specifications, for example, logarithmic and power utility functions, in Online Appendix L.5.

Section 4.1 summarizes the data used in this paper. In Section 4.2, we focus on the baseline case of a representative farmer when budget constraints are barely binding.[5] In other words, the farmer is primarily interested in maximizing the expected utility and accepts the equilibrium coverage level and its corresponding insurance premium. We study the reduction in basis risk and the utility improvement of the optimal policy in the baseline case. Section 4.3 explores the interpretability of the optimal policy. Finally, we compare the NN-based optimal policy with some simpler contracts in Section 4.4.

4.1. Data

4.1.1. Production Loss and Farmer's Wealth Data. We use a data set consisting of county-level annual corn production experience for Illinois, obtained from the National Agricultural Statistics Service (NASS). Corn is the single most valuable agricultural commodity in the United States, and composes more than 45% of cropland by acreage in Illinois.[6] Another advantage of using Illinois corn data are its long history. We use a sample period from 1925 to 2018. To ensure stationarity of the loss experience over the long sample period, we first detrend the crop yields data with a second-order polynomial, estimated with a robust regression technique. Next, similar to Deng et al. (2007) and Harri et al. (2011), we adjust historical yields data to the 2018 level and other data heteroscedasticity. The detrending procedure also helps remove the impacts of long-run climate change or technology progress, such as improved genetics (cultivars), improved crop management, and other technological advances like the use of advanced farming equipment.[7] As the long sample of corn yield data are only available at annual frequency, the number of historical yield observations is quite limited, especially compared with the high-dimensional weather-related covariates we want to use for each area. We circumvent this difficulty by assuming that crop yield losses are both time and space homogeneous after detrending.[8] This simplification expands the size of our data to 7,869 county-years (84 counties × 94 years, with 27 missing data points removed from the sample).

In the main results of this paper, we focus on hedging fluctuations in crop yields (i.e., production losses) and assume that the corn price is a constant of \$3.5 per bushel (\$/bu), which is the five-year average of the inflation-adjusted corn price from 2014 to 2018.[9] That is, the dollar wealth is indexed on the most recent five-year price level. We further consider the price risk in Online Appendix M. The corn production loss data measured in bushels per acre (bu/acre) are then multiplied by the corn price to arrive at the monetary loss incurred by each farmer in dollars per acre.

The 2014–2018 USDA Agricultural Resource Management Survey and NASS show that net farm assets (i.e., farm assets minus farm debts) have a five-year average of \$456,977.[10] In addition, according to Illinois Farm Business Farm Management, the average size of Illinois grain farms is 1,176 acres.[11] To make the numerical results comparable and interpretable, we normalize all monetary quantities for the representative farmer by farm size. In other words, wealth levels, incomes, and other monetary quantities used in our analysis are all measured in dollars per acre, which is also consistent with the unit of our production loss data. Dividing the net farm assets by the farm size yields the initial wealth of $w_0 = \$389/\text{acre}$.

4.1.2. Climate and Weather Index Data. We collect weather data from the PRISM Climate Group.[12] PRISM publishes monthly meteorological information on six climate variables, including precipitation, max/min temperatures, max/min vapor pressure deficit, and dew points, from 1895 to present for the conterminous United States at a 4-km resolution. We use all of the six monthly climate variables over 1925–2018. This gives a 72-dimensional weather index matrix for our optimal insurance policy design. Table 1 describes the weather variables from the PRISM data set. Online Appendix E provides summary statistics for the weather variables.

The impacts of weather indices on crop production losses can be nonmonotonic and highly nonlinear due to biological reasons (Schlenker and Roberts 2009, Rigden et al. 2020). Figure 2(a) through (c), illustrates the nonlinear patterns of three selected weather indices (see Online Appendix F for scatterplots of all 72 weather indices). Furthermore, different weather indices might interact with each other and jointly affect production losses. Figure 2(d) and (e), shows that the min vapor pressure deficit in October and precipitation in October do not individually affect production losses. However, Figure 2(f) shows that they jointly affect production losses in a highly nonlinear way. These complexities cannot be captured by the linear models used by most current index insurance contracts, and, thus, more sophisticated contract design models are needed.

4.1.3. Insurance Market Data and the Supply Curve. We obtain insurance market data from the USDA National Summary of Business (SOB) Reports (1989–2017).[13] The variables that we collect include the total acres insured (*Acres*), premiums (*Prem*), liabilities (*Liab*), and indemnities (*Indem*). Liabilities stand for insurance guaranteed crop production levels. If the farmer's actual yields are lower than the liability level, the farmer will receive insurance payments. Indemnities correspond to realized insurance payments. The USDA Risk Management Agency computes the loss cost ratio (*LCR*) as the ratio of indemnity over liability: $LCR = Indem/Liab$. The premium is then calculated as the expected *LCR* times the liability, multiplied by the loading parameter λ, that is, $Prem = \lambda \cdot \mathrm{E}(LCR) \cdot Liab$. Therefore, we can use this relationship to infer the realized λ every year from the market data. Insurance coverage per acre is calculated as $Cov = Indem/Acres$ (adjusted by inflation), representing the unit acre insurance supply at market equilibrium each

Table 1. Weather Indices

Variable	Description
pcpn*k*	Total precipitation (rain+melted snow) for month *k* (mm)
tmax*k*	Daily maximum temperature averaged over all days in month *k* (°C)
tmin*k*	Daily minimum temperature averaged over all days in month *k* (°C)
dpt*k*	Daily mean dew point temperature averaged over all days in month *k* (°C)
vpdmax*k*	Daily maximum vapor pressure deficit averaged over all days in month *k* (hPa)
vpdmin*k*	Daily minimum vapor pressure deficit averaged over all days in month *k* (hPa)
k	Calendar month, that is, $k = 1, 2, \ldots, 12$ for Jan–Dec

Notes. This table summarizes the weather variables available from the PRISM data set. The sample period is 1925–2018.

Figure 2. (Color online) Relationship Between Crop Losses and Weather Indices

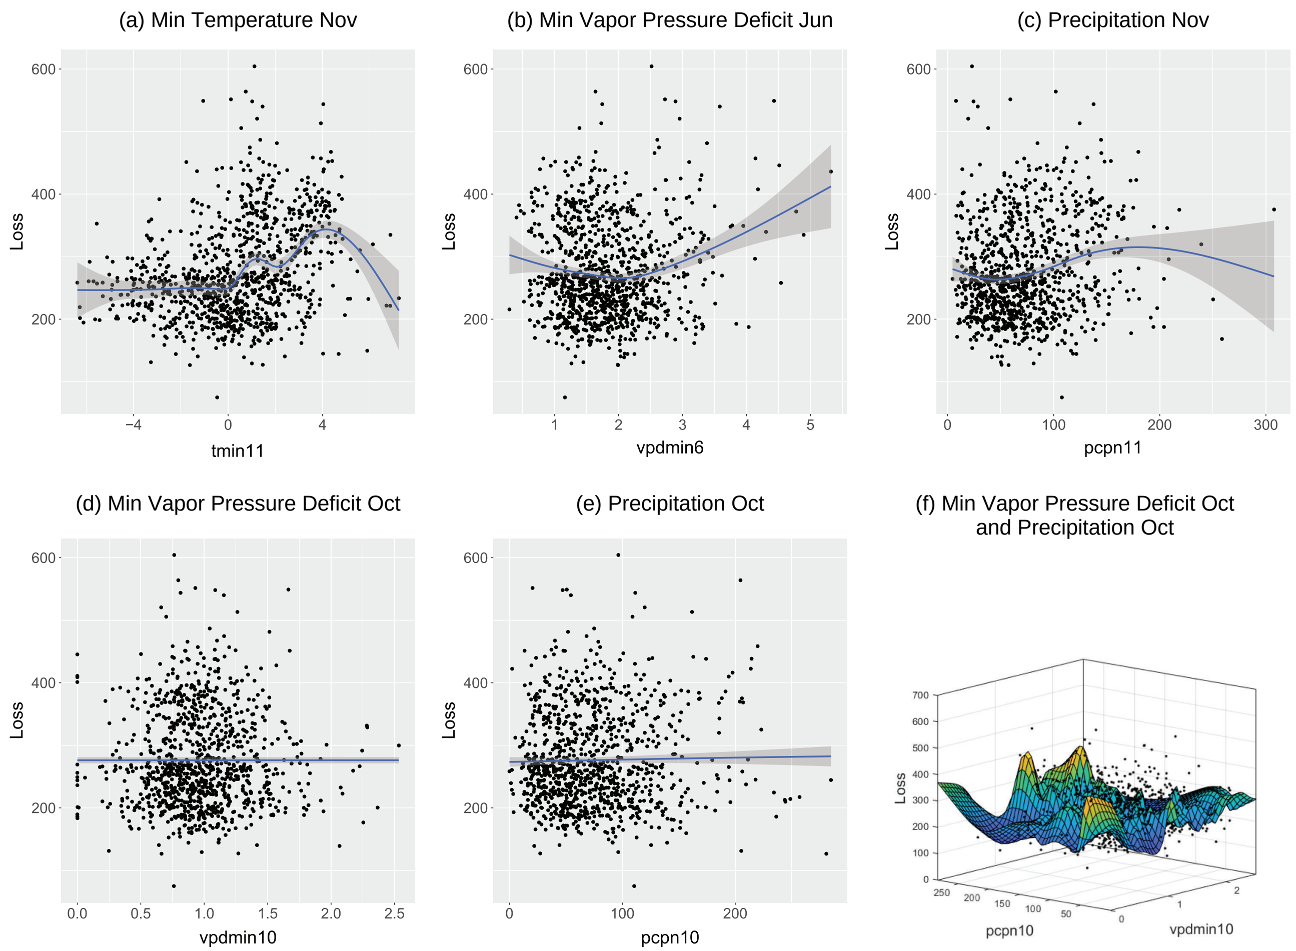

Notes. (a)–(e) visualize crop losses with five selected weather indices, based on 1,000 random draws from the sample. The solid curve is fitted by a generalized additive model. The shadow area indicates a 95% confidence interval. (f) Joint effects of two weather indices used in (d) and (e) on crop losses. The surface is fitted by a thin-plate smoothing spline model.

year. We use the historical data to fit the supply curve as a power function of λ, using nonlinear least squares method. The fitted supply curve is $Cov = f_S(\lambda) = 7.04\lambda^{2.92} + 12.83$. Because of the data limitation, one might worry about the simultaneity bias of fitting the supply curve from historical data. We perform robustness checks by using alternative supply curves in Online Appendix L.1.

4.2. Optimal NN-Based Index Insurance: The Baseline Case

We split the sample into three disjoint time periods, preserving the data temporal order, with the earliest 70% data as training set, and then the next 15% as validation set, and the last 15% as test set. NN parameters are trained on training set for each given set of hyperparameters, including the number of neurons and hidden layers. Then the validation set is used to choose the optimal set of hyperparameters. Finally, we use the test set to evaluate the performance of the selected optimal NN-based index insurance. We adopt the ensemble method to improve the robustness of the analysis. Specifically, all models are averaged over 10 training runs.

We first consider the optimal index insurance whereby farmers are barely bound by any budget constraint. For example, P_L is set to zero and P_U is set to the farmer's total wealth. In other words, the farmer would like to pay any affordable price to maximize the utility. We will discuss the effects of budget constraints in Online Appendix L.3. We set the absolute risk aversion, α, to 0.008, which corresponds to a relative risk aversion of 3.1.[14] In Section 4.2.1, we discuss how to endogenously determine the loading parameter λ at equilibrium. Then we discuss the performance of the optimal index insurance of the baseline case in Section 4.2.3.

4.2.1. Determining the Loading Parameter at Market Equilibrium.

We determine the endogenous loading parameter at market equilibrium via a reduced-form

Figure 3. (Color online) Supply and Demand Curves of the Index Insurance

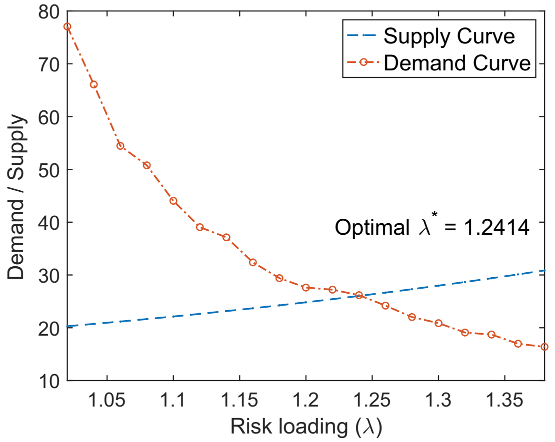

Notes. The insurance supply curve is fitted from the USDA SOB Reports data with a power function. The demand curve is for the NN-based optimal index insurance with a 64-64-16 neurons structure and is fitted with a piecewise cubic Hermite interpolating polynomial. The intersection gives the loading parameter $\lambda^* = 1.2414$ at market equilibrium.

approach, as follows. First, we use the USDA SOB Reports market data to fit the supply curve, $Cov = f_S(\lambda)$, as described in Section 4.1.3. Second, for a given value of $\lambda \in [1.02, 1.5]$,[15] we obtain the corresponding optimal coverage of index insurance policies, using the NN-based approach. Given the pairs of λ and the optimal coverage, we fit the demand curve, $Cov = f_D(\lambda)$, using the piecewise cubic Hermite interpolating polynomial. Finally, we equate the supply and demand functions, $f_S(\lambda) = f_D(\lambda)$, to solve for the loading parameter at market equilibrium, λ^*. Therefore, farmers and insurers jointly determine the insurance premium in equilibrium. Figure 3 illustrates an example of finding λ^* for the NN-based optimal index insurance with a 3-hidden-layer (64-64-16 neurons) structure, corresponding to the baseline model that we will discuss in Section 4.2.3. The equilibrium loading parameter is $\lambda^* = 1.2414$.

4.2.2. Selecting the Optimal NN-Based Model. We consider and compare different neural network structures. For each neural network structure, we use the procedures described in Section 4.2.1 to compute the equilibrium loading parameter, λ^*. As we endogenously determine λ^*, it would be very computationally costly to endogenize the hyperparameters of the model. Instead, we examine a wide range of candidate models with different number of hidden layers and different number of neurons in each layer, and then select the optimal model based on the utility improvement in the validation set. Panel A of Table 2 shows the results without using index insurance. Panels B–E present the results from various NN-based index insurance contracts with different layer and neuron structures. For example, "$i - j - k$" indicates a three-hidden-layer NN model with i, j, and k neurons in the first, second, and third hidden layers, respectively. We compute the insurance premium and compare expected utilities with and without index insurance. For ease of interpretation, we also compute the CEW with and without index insurance.

Panel B shows that a simple NN with only one hidden layer and two neurons significantly improves farmers' expected utility by 6.54% and CEW by \$8.45/acre in the validation sample. In addition, adding more neurons or adding more hidden layers generally improves the model performances, as shown in Panels B–D. However, increasing the number of hidden layers from three to four (moving from Panel D to Panel E) does not necessarily improve the NN performance. We see that the 3-hidden-layer NN with a 64-64-16 neurons structure yields the largest utility improvement and CEW improvement in the validation set. This NN-based contract provides a utility improvement of 10.60% in the training sample and 9.31% in the validation sample. We choose this as the baseline model (*NN72*, with a 64-64-16 neurons structure) and discuss its forecasting performance in the test set in Section 4.2.3.

4.2.3. Performances of the Optimal Index Insurance. The performance of the optimal index insurance in the test sample is reported in Table 4. The baseline contract (*NN72*) provides a utility improvement of 14.35% in the test sample. With this insurance contract, the CEW improves by \$19.36/acre in the test set. Given an average farm size of 1,176 acres, the wealth improvement of the NN-based index insurance for a typical farm is economically significant, for example, \$22,767.36 in the test sample, which is about 4.98% of her wealth.

The optimal baseline index insurance policy has a premium of \$28.72 in the test sample. The corresponding equilibrium loading parameter is $\lambda^* = 1.2414$, implying a coverage level of \$23.13 in the test sample. Compared with the average insurance premium for Illinois corn farmers, which is around \$45.50/acre in 2017 (Smith 2017), this optimal premium level is only 63% of the current premium level in practice. This lower premium for the proposed optimal index insurance would substantially increase the demand of index insurance in practice.

It is also shown that the proposed insurance achieves very similar expected utilities over training and test samples, suggesting that "overfitting" is of little concern. For example, the difference between policyholders' expected utilities with insurance in the training and test sample is less than 0.01, and the difference between CEW with insurance in the training and test sample is \$0.03, both of which are not sizable.

To further illustrate how the proposed NN-based insurance contract helps farmers improve their welfare, Figure 4 compares the wealth distributions between the

Table 2. Selecting an NN Model

Panel A: Without insurance	Training	Validation
U w/o insurance	−3.99	−3.99
CEW w/o insurance	430.63	430.41

Panel B: 1-hidden-layer NN	2 ($\lambda^* = 1.1956$)		8 ($\lambda^* = 1.2231$)		64 ($\lambda^* = 1.2340$)	
	Training	Validation	Training	Validation	Training	Validation
U with insurance	−3.65	−3.73	−3.59	−3.71	−3.54	−3.65
U improvement	8.37%	6.54%	9.86%	7.18%	11.17%	8.69%
CEW with insurance	441.57	438.86	443.60	439.73	445.44	441.77
CEW improvement	10.93	8.45	12.97	9.32	14.81	11.36
CEW improvement (%)	2.54%	1.96%	3.01%	2.16%	3.44%	2.64%
Premium	29.81	22.48	30.82	22.93	32.07	29.06

Panel C: 2-hidden-layers NN	16-8 ($\lambda^* = 1.2336$)		64-8 ($\lambda^* = 1.2428$)		64-16 ($\lambda^* = 1.2351$)	
	Training	Validation	Training	Validation	Training	Validation
U with insurance	−3.58	−3.67	−3.58	−3.66	−3.57	−3.70
U improvement	10.16%	8.01%	10.31%	8.27%	10.50%	7.39%
CEW with insurance	444.02	440.85	444.23	441.20	444.50	440.00
CEW improvement	13.39	10.44	13.60	10.79	13.87	9.59
CEW improvement (%)	3.11%	2.43%	3.16%	2.51%	3.22%	2.23%
Premium	26.44	26.69	26.25	28.88	26.94	28.80

Panel D: 3-hidden-layers NN	64-16-8 ($\lambda^* = 1.2400$)		64-16-16 ($\lambda^* = 1.2428$)		64-64-16 ($\lambda^* = 1.2414$)	
	Training	Validation	Training	Validation	Training	Validation
U with insurance	−3.57	−3.62	−3.56	−3.63	−3.57	−3.62
U improvement	10.53%	9.27%	10.64%	9.14%	10.60%	9.31%
CEW with insurance	444.54	442.56	444.70	442.40	444.64	442.63
CEW improvement	13.91	12.15	14.07	11.99	14.00	12.22
CEW improvement (%)	3.23%	2.82%	3.27%	2.78%	3.25%	2.84%
Premium	27.92	28.57	28.00	27.95	28.44	28.99

Panel E: 4-hidden-layers NN	64-16-8-8 ($\lambda^* = 1.2212$)		64-32-16-8 ($\lambda^* = 1.2445$)		64-64-16-16 ($\lambda^* = 1.2285$)	
	Training	Validation	Training	Validation	Training	Validation
U with insurance	−3.55	−3.63	−3.57	−3.68	−3.56	−3.68
U improvement	11.07%	9.09%	10.52%	7.98%	10.85%	7.78%
CEW with insurance	445.30	442.32	444.53	440.80	444.99	440.53
CEW improvement	14.67	11.92	13.90	10.39	14.35	10.12
CEW improvement (%)	3.41%	2.77%	3.23%	2.41%	3.33%	2.35%
Premium	29.32	30.07	28.91	27.04	27.63	26.77

Notes. This table summarizes the performances of several NN-based index insurance policies with various neuron structures. Panel A shows expected utility and certainty equivalent wealth (CEW) without index insurance, over the training and validation samples. Panels B–E present results with different neuron structures. For example, "$i-j-k$" indicates i, j, and k neurons in the first, second, and third hidden layers, respectively. Each panel reports the expected utility, percentage utility improvement, CEW, CEW improvement in dollars, CEW improvement in percentage, and insurance premium for each optimal insurance policy, for both training and validation samples. The risk loading parameter at equilibrium (λ^*) is reported in parentheses. The optimal NN structure is the one with the largest utility improvement in the validation set (i.e., the 64-64-16 neuron structure).

"with insurance" case and the "without insurance" case. Without insurance, farmers' wealth distribution has a larger variation with a heavy left tail, that is, the downside risk. In contrast, with the proposed index insurance, farmers' wealth distribution becomes less dispersed, and the left tail is significantly reduced.[16] We observe similar patterns in both training and test samples. Figure 4 demonstrates that our proposed index insurance policy

Figure 4. (Color online) Farmers' Wealth Distribution: With Index Insurance vs. Without Insurance

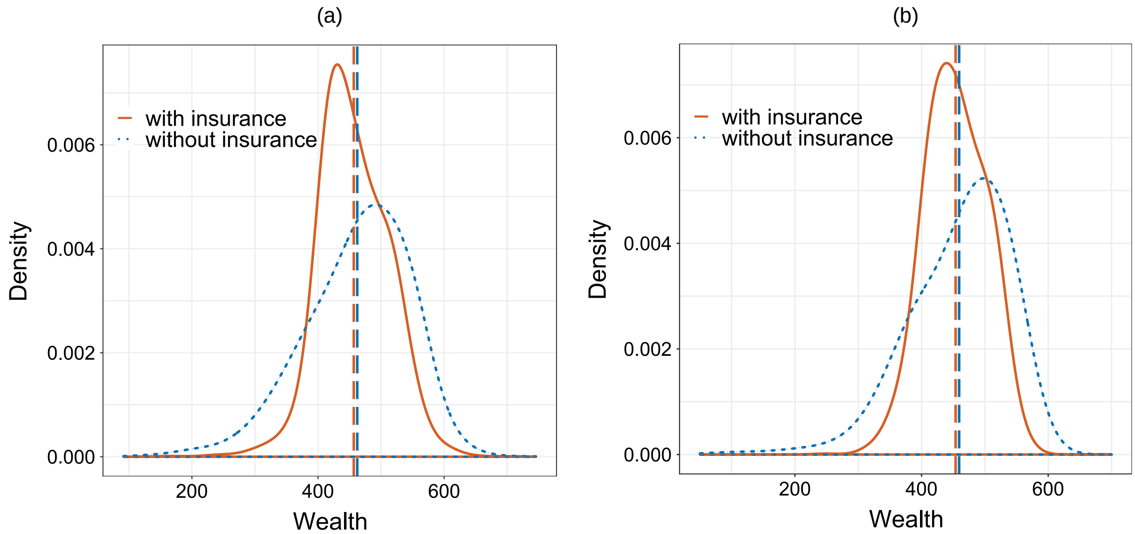

Notes. The dashed curve represents the probability density of wealth without using insurance, and the solid curve represents the probability density of wealth using the optimal NN-based index insurance. (a) Training sample. (b) Test sample.

effectively hedges the downside weather risk for farmers and substantially improves their utilities.

4.3. Interpreting the NN-Based Index Insurance

Neural network results are often difficult to explain while economic interpretability is crucial for economics (Cong et al. 2020, 2021a). In this section, we perform a battery of sensitivity analyses to interpret the baseline case of NN-based index insurance contract.

4.3.1. Global Interpretability. Following Cong et al. (2021a), we adopt a polynomial feature sensitivity analysis, which involves using the gradient-based characteristic importance method to examine contributions of feature inputs and extracts important features. Hence, it provides "economic distillation" of the NN-based index insurance and increases interpretation of the "black box" insurance policy. This provides a global viewpoint because it is based on the whole sample. Given the optimal NN-based index insurance, $I(\boldsymbol{X})$, where $\boldsymbol{X}$ is the vector of weather indices, the sensitivity of $I(\boldsymbol{X})$ to the weather index X_k can be measured by its partial derivative with respect to X_k:

$$\delta_{X_k}(\boldsymbol{X}) = \frac{\partial I(\boldsymbol{X})}{\partial X_k} = \lim_{\Delta x_k \to 0} \frac{I(\boldsymbol{X}) - I(X_k + \Delta x_k, \boldsymbol{X}_{-k})}{\Delta x_k}, \tag{6}$$

where $\boldsymbol{X}_{-k}$ is the vector of weather indices with the kth index removed. Then the gradient-based sensitivity of insurance payoff with respect to X_k, $\mathcal{S}^g_{X_k}$, can be expressed as the average influence of index X_k to the optimal index insurance payoff, calculated as

$$\mathcal{S}^g_{X_k} = \mathrm{E}(\delta_{X_k}) = \int_\Omega \delta_{X_k}(\boldsymbol{X}) d\mathbb{P}(\boldsymbol{X}), \tag{7}$$

where $\mathbb{P}(\boldsymbol{X})$ is the joint probability distribution of $\boldsymbol{X}$, Ω is the sample space, and the superscript g indicates the gradient-based sensitivity analysis. Empirically, $\mathcal{S}^g_{X_k}$ can be estimated by

$$\bar{\mathcal{S}}^g_{X_k} = \frac{1}{N}\sum_{i=1}^{N} \delta_{X_k}(\boldsymbol{x}_i), \tag{8}$$

where $\boldsymbol{x}_i, i = 1, \ldots, N$, are the sample data of $\boldsymbol{X}$ in the training set.[17]

We rank the importance of weather indices by the absolute value of $\bar{\mathcal{S}}^g_{X_k}$. Table 3 lists the five most important weather indices, as shown in the right panel.[18] For comparison, we also rank weather indices based on the absolute correlations between insurance payoffs and the indices, and the top five indices are shown in the left panel. We see that all five indices are different between these two panels, and most of them have very different ranks under these two ranking methods. Although the correlations capture linear dependence, sensitivities from NN-based insurance provide a more general nonlinear dependence measure to study the relationship between insurance payoffs and weather indices. Those weather indices with a large absolute value of correlations are not necessarily the most important ones, whereas those variables with small correlations (tmin6, dpt6, dpt11, dpt12) may be very useful for the insurance. Finally, we notice that the five most important weather indices based on absolute correlations are all within the corn growing season, but the NN-based insurance contract picks up three weather indices outside the growing season, that is, dew point temperature in November (dpt11), maximum temperature in December (tmax12), and dew point temperature in December (dpt12). In fact,

Table 3. Identifying Important Weather Indices

$Rank^{corr}$	Index	Absolute correlations	Rank difference $Rank_g^{NN} - Rank^{corr}$	$Rank_g^{NN}$	Index	NN policy absolute $\bar{S}^g$	Rank difference $Rank_g^{NN} - Rank^{corr}$
1	tmax8	0.49	20	1	tmin6	10.30	−27
2	tmax7	0.46	11	2	dpt6	10.03	−37
3	vpdmax8	0.45	3	3	dpt11	9.58	−42
4	tmin7	0.41	27	4	tmax12	8.70	−3
5	vpdmax7	0.41	6	5	dpt12	7.76	−25

Notes. This table shows the five most important weather indices ranked by the absolute correlations between insurance payoffs and weather indices (left panel) or the gradient-based sensitivities of an NN-based insurance payoffs to the weather indices (right panel). Column "Rank Diff" shows the difference between two ranks for a weather index. See Table 1 for the index variable descriptions.

agronomic studies show that hydrological conditions over nongrowing seasons might be important for corn production for two reasons. First, water accumulation in winter will affect the soil moisture that becomes available for corn in the next growing season.[19] Second, soil water accumulation during the nongrowing season affects cover crops, which are grown to improve soil quality and influence corn growth later.[20] This further illustrates the ability of neural networks to extract important nongrowing season information from the data, which is often overlooked by the linear functions in current practice.

4.3.2. Local Interpretability. From the marketing perspective, sometimes a local interpretation of the insurance contract may be more helpful, because a farmer often wants to know what weather variables will contribute to the farmer's own insurance payoff in a certain

Table 4. Performances of the Baseline Insurance Contract and Alternative Contracts

Contract	Linear1 ($\lambda^* = 1.0255$)	Linear5 ($\lambda^* = 1.0730$)	Quadradic5 ($\lambda^* = 1.0785$)	Cubic5 ($\lambda^* = 1.0840$)	NN5 ($\lambda^* = 1.1778$)	Linear72 ($\lambda^* = 1.2133$)	NN72 (BL, $\lambda^* = 1.2414$)
Panel A: Utility improvement							
U with insurance	−4.14	−4.03	−4.03	−4.00	−3.78	−3.76	−3.57
U w/o insurance	−4.16	−4.16	−4.16	−4.16	−4.16	−4.16	−4.16
U improvement (%)	0.55%	3.08%	3.28%	3.84%	9.11%	9.60%	14.35%
Panel B: CEW improvement							
CEW with insurance	425.94	429.16	429.43	430.15	437.19	437.87	444.61
CEW w/o insurance	425.26	425.26	425.26	425.26	425.26	425.26	425.26
CEW improvement	0.69	3.91	4.17	4.90	11.94	12.61	19.36
CEW improvement (%)	0.16%	0.92%	0.98%	1.15%	2.81%	2.97%	4.55%
Panel C: Policy characteristics							
Premium	24.48	26.94	28.55	26.21	27.94	28.56	28.72
Coverage	23.87	25.10	26.47	24.18	23.72	23.54	23.13
Insurer profit	0.61	1.83	2.08	2.03	4.22	5.02	5.59
Panel D: Risk reduction measured by standard deviation							
Std	75.13	71.05	70.01	69.69	62.86	59.36	47.49
Std w/o insurance	78.92	78.92	78.92	78.92	78.92	78.92	78.92
Std reduction	4.80%	9.97%	11.29%	11.69%	20.34%	24.78%	39.82%
Panel E: Risk reduction measured by value-at-risk (VaR)							
$VaR_{5\%}$	332.92	338.48	339.68	338.01	339.00	345.93	379.64
$VaR_{5\%}$ w/o insurance	325.91	325.91	325.91	325.91	325.91	325.91	325.91
$VaR_{5\%}$ improvement	7.02	12.58	13.77	12.10	13.09	20.03	53.73

Notes. This table summarizes the performances of seven insurance contracts in the test sample, including (1) a linear insurance contract with one weather index (*Linear1*), which corresponds to a currently used contract; (2) a linear insurance contract with five weather indices (*Linear5*); (3) a quadratic insurance contract with five weather indices (*Quadratic5*); (4) a cubic insurance contract with five weather indices (*Cubic5*); (5) an NN-based contract with five weather indices (*NN5*); (6) a linear insurance contract with 72 weather indices (*Linear72*); and (7) the baseline model (*NN72*, an NN-based contract with 72 weather indices). Panel A summarizes expected utilities with and without (w/o) index insurance and the percentage of utility improvement. Panel B reports certainty equivalent wealth (CEW) with and without (w/o) index insurance and the CEW improvement in dollars and as a percentage. Panel C summarizes policy characteristics including policy premium, coverage, and profits of the insurer. Panel D summarizes the risk reduction effect of an index insurance policy, measured by the standard deviation of wealth. Panel E summarizes the tail risk reduction, measured by the 5%-level value-at-risk (VaR). The risk loading parameter at equilibrium (λ^*) for each contract is reported in parentheses.

year. Therefore, we apply the local interpretable model-agnostic explanations (LIME; Ribeiro et al. 2016) to provide local interpretations. LIME approximates the underlying black box model with a simpler and more interpretable local surrogate model, typically a linear regression model, to investigate a local region of the data. With LIME, farmers can understand the rationale behind the NN-based contract, assisting them to decide whether to trust the insurance payoffs.

Figure 5 displays some representative examples of LIME results. For example, all the top five weather indices suggest a good harvest, and hence the farmer in Warren county would receive an insurance payoff of zero in 2005. By contrast, the farmer in Edwards county would receive a large insurance payoff of $288.15 per acre in 2012. The large payoff is due to a high vpdmax8, tmin11, vpdmax7, and vpdmin6, whereas a high vpdmin5 reduces the payoff. Such explanations will help improve the farmers' understanding and trust of the index insurance contract.

4.4. Comparison with Other Simple Insurance Contracts

Previously, we demonstrated the effectiveness of the NN-based index insurance design and assessed its interpretability. In this section, we further compare the NN-based insurance contract with other commonly used simple contracts, that is, contracts with fewer indices and/or less nonlinear structures. This comparison highlights the benefits of using the NN-based insurance contract. In addition to the NN-based insurance contract, we consider some polynomial-based contracts, including (1) a linear insurance contract; (2) a quadratic insurance contract; and (3) a cubic insurance contract.[21] Depending on the contract structure, we use either all 72 weather indices if manageable or the most important one or five weather indices identified by gradient-based sensitivity analysis in the NN-based contract, discussed in Section 4.3.[22] Specifically, we consider the following seven index insurance contracts:

- *Linear contract with one input (Linear1)*: a linear contract written on a single temperature index, which is a currently used index insurance contract in practice;
- *Linear contract with five inputs (Linear5)*: only linear polynomials are used, with the top-five weather indices;
- *Quadratic contract with five inputs (Quadratic5)*: quadratic polynomials are used, with the top-five weather indices;
- *Cubic contract with five inputs (Cubic5)*: cubic polynomials are used, with the top-five weather indices;
- *NN with five inputs (NN5)*: NN with a 64-64-16 neuron structure, using the top-five weather indices;
- *Linear contract with 72 inputs (Linear72)*: only linear polynomials are used, with all 72 weather indices;
- *Baseline model (NN72)*: NN with a 64-64-16 neuron structure, using all 72 weather indices. This is the baseline optimal contract discussed in Sections 4.2 and 4.3.

For each contract, the loading parameter is endogenously determined via the reduced-form approach described in Section 4.2.1. Table 4 presents the performances of these seven contracts in the test set. We compare the contract performances mainly by examining their utility improvements, CEW improvements, and risk reductions. We measure risk reductions by the standard deviation, or the 5% value-at-risk (VaR) of policyholders' wealth.

First, adding more weather indices as inputs significantly improves the contract performance of *Linear72*,

Figure 5. (Color online) Local Explanation of the NN-Based Index Insurance with LIME

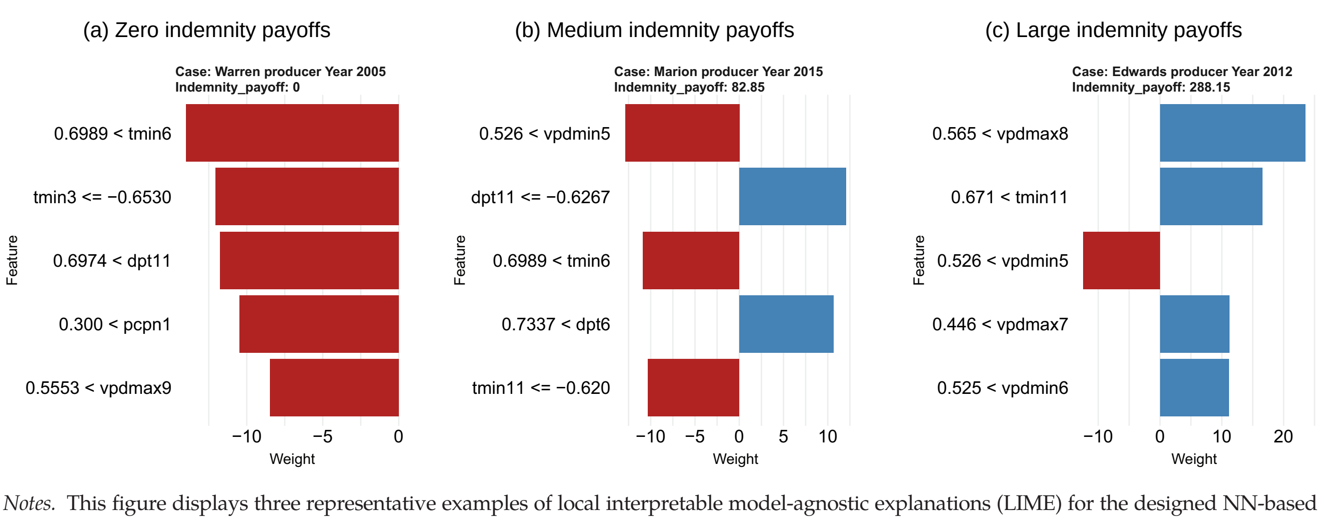

Notes. This figure displays three representative examples of local interpretable model-agnostic explanations (LIME) for the designed NN-based index insurance. (a)–(c) Farmer-year pairs that receive indemnity payoffs of zero, medium, and large amounts, respectively. Five most important weather indices for each farmer-year pair are illustrated. Bars with positive (negative) weights indicate that the corresponding indices contribute positively (negatively) to the indemnity payments, that is, bad (good) weather conditions. Longer bars imply higher impacts that the indices have on indemnity payments.

relative to *Linear1* and *Linear5*. Second, comparing *Linear5*, *Quadratic5*, and *Cubic5*, we see that utility improves a little when adding nonlinear terms to the insurance contract in Panel A. We see similar improvements in CEW in Panel B, standard deviation of wealth in Panel D, and tail risk in Panel E. Third, using the same five weather indices, we see *NN5* outperforms *Linear5*, *Quadratic5*, and *Cubic5*. Fourth, the baseline model (*NN72*) delivers the best performance. *NN72* provides the highest expected utility and CEW to farmers, and it also achieves the highest risk reduction measured by standard deviation reduction and tail risk reduction. Furthermore, from the insurers' perspective, *NN72* is also the best contract because it provides the largest profits. Quantitatively, *Linear1* improves farmers' utility by 0.55%; *Linear72* improves their utility by 9.60%; and *NN72* improves their utility by 14.35%. In terms of CEW, *Linear1*, *Linear72*, and *NN72* improve it by \$0.69, \$12.61, and \$19.36, respectively. Insurers profits are \$0.61/acre, \$5.02/acre, and \$5.59/acre, for the *Linear1*, *Linear72*, and *NN72*, respectively. This is because of higher market demand and hence a higher equilibrium loading, λ^*.

Basis risk often dissuades farmers from buying index insurance and makes insurance less cost effective. Policyholders are most concerned with the worst scenario when they suffer a huge loss in crop yields, but only receive little payments from the insurance. The opposite case is also unappealing: although it does not hurt when farmers get a good harvest and receive some insurance payments at the same time, insurance companies will suffer losses in this case. This means that the insurance premiums are not effectively allocated to the worst scenarios when insurance payoffs are more needed. Online Appendix I compares the basis risk of these seven index insurance contracts. We see that except *NN72*, there is a notably large mismatch between losses and insurance payoffs, especially for the test set. Overall, our results demonstrate that, by reducing basis risk and solving for the endogenous insurance premium, the NN-based solution can significantly improve the demand of index insurance and increase the profit margin of insurers.

5. Robustness of NN-Based Index Insurance

We investigate the robustness of the NN-based index insurance contract in several ways, using an NN

Table 5. Out-of-State Tests Using Adjacent States

	Training set	Test set		
	Illinois	Indiana	Iowa	Missouri
Panel A: Utility improvement				
U with insurance	−3.55	−2.88	−3.95	−2.55
U w/o insurance	−4.02	−3.15	−4.21	−2.78
U improvement (%)	11.63%	8.69%	6.16%	8.53%
Panel B: CEW improvement				
CEW with insurance	445.21	471.48	431.77	486.74
CEW w/o insurance	429.75	460.11	423.82	475.59
CEW improvement	15.46	11.37	7.95	11.15
CEW improvement (%)	3.60%	2.47%	1.88%	2.34%
Panel C: Policy characteristics				
Premium	29.09	17.83	10.54	17.89
Coverage	23.43	14.36	8.49	14.41
Insurer profit	5.66	3.47	2.05	3.48
Panel D: Risk reduction measured by standard deviation				
Std	51.29	51.17	53.39	64.20
Std w/o insurance	80.60	67.93	66.26	79.35
Std reduction	36.36%	24.66%	19.43%	19.10%
Panel E: Risk reduction measured by value-at-risk (VaR)				
$\text{VaR}_{5\%}$	381.12	396.90	354.55	394.52
$\text{VaR}_{5\%}$ w/o insurance	317.70	361.01	321.92	362.57
$\text{VaR}_{5\%}$ improvement	63.42	35.90	32.63	31.95

Notes. We perform out-of-state tests for the NN-based index insurance. We use Illinois data as the training set to estimate the contract parameters. We use data from three states adjacent to Illinois, namely, Indiana, Iowa, and Missouri, as the test sets. The NN uses a 3-hidden-layer (64-64-16 neurons) structure. Panel A summarizes utilities with and without (w/o) different index insurance policies and the percentage of utility improvement. Panel B summarizes certainty equivalent wealth (CEW) with and without (w/o) index insurance policies and the CEW improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage, and profits for the insurer. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation of wealth. Panel E summarizes risk reduction at the tail, measured by the 5%-level value-at-risk (VaR).

structure of 64-64-16 neurons. First, we use data from three states adjacent to Illinois to perform the out-of-state tests. Second, we consider the impacts of over-insuring constraint. Third, we examine the impacts of weather predictability on insurance design. Online Appendix L provides additional robustness checks, including (1) the impacts of insurers' choices (various insurers' supply curves and exogenously given risk loading) and (2) the impacts of farmers' characteristics (such as different coverage levels, different risk aversion levels, time-varying risk aversion, and alternative utility functions).

5.1. Out-of-State Tests

As a more stringent out-of-sample test, we run some out-of-state tests for states adjacent to Illinois. Specifically, we use Illinois data as the training set to estimate the parameters in the NN-based index insurance policy and use data from the adjacent states as the test set to evaluate the model performance. We select three states adjacent to Illinois and that have similar latitudes: Indiana, Iowa, and Missouri, all of which are important corn producers.[23] Table 5 summarizes the results. We see that the NN-based index insurance performs reasonably well in these states. For example, it improves the expected utility by 8.69% in Indiana, 6.16% in Iowa, and 8.53% in Missouri. As for CEW, improvements are \$11.37/acre, \$7.95/acre, and \$11.15/acre in Indiana, Iowa, and Missouri, respectively. These results demonstrate the power of our proposed NN-based index insurance.

5.2. Overinsuring Constraint

The overinsuring constraint usually is necessary for traditional optimal (re)insurance design:

$$I(\boldsymbol{X}) \leq \text{Loss}.$$

This constraint ensures that the indemnity payment from the insurance contract cannot exceed actual losses, which rules out the "overinsuring" issue. Indeed, previous studies find that there is an incentive for policyholders to over insure at the presence of basis risk (Doherty and Schlesinger 1983). This constraint is less of a concern in our setting for two reasons. First, there is little moral hazard and adverse selection in index insurance, because insurance payments cannot be manipulated by policyholders. Second, our baseline model exhibits very small basis risk and over-indemnifying payments. In this section, we further compare the numerical results for the optimal NN-based index insurance contract with and without the overinsuring constraint. The results are

Table 6. Overinsuring Constraint

	With overinsuring constraint		Without overinsuring constraint (BL)	
	Training	Test	Training	Test
Panel A: Utility improvement				
U with insurance	−3.56	−3.57	−3.57	−3.57
U w/o insurance	−3.99	−4.16	−3.99	−4.16
U improvement (%)	10.61%	14.34%	10.60%	14.35%
Panel B: CEW improvement				
CEW with insurance	444.65	444.60	444.64	444.61
CEW w/o insurance	430.63	425.26	430.63	425.26
CEW improvement	14.02	19.34	14.00	19.36
CEW improvement (%)	3.25%	4.55%	3.25%	4.55%
Panel C: Policy characteristics				
Premium	28.37	28.54	28.44	28.72
Coverage	22.85	22.99	22.91	23.13
Insurer profit	5.52	5.55	5.53	5.59
Panel D: Risk reduction measured by standard deviation				
Std	54.07	47.63	54.05	47.49
Std w/o insurance	81.94	78.92	81.94	78.92
Std reduction	34.02%	39.65%	34.04%	39.82%
Panel E: Risk reduction measured by value-at-risk (VaR)				
$VaR_{5\%}$	383.01	378.56	382.89	379.64
$VaR_{5\%}$ w/o insurance	316.28	325.91	316.28	325.91
$VaR_{5\%}$ improvement	66.73	52.66	66.61	53.73

Notes. We consider the optimal insurance performance with and without over-insuring constraint. Panel A summarizes utilities with and without (w/o) different index insurance policies and the percentage of utility improvement. Panel B summarizes CEW with and without (w/o) index insurance policies and certainty equivalent wealth (CEW) improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage, and profits for the insurers. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation of wealth. Panel E summarizes the risk reduction at the tail, measured by the 5%-level value-at-risk (VaR). "BL" indicates the baseline case studied in Section 4.2.

summarized in Table 6. We can see that the overinsuring constraint does not have significant impacts.

5.3. Weather Predictability and Conditional Contracts

Our baseline model uses only current year weather indices as inputs. However, one might wonder whether weather conditions are predictable and historical weather conditions affect the insurance contract design, suggesting the discrepancy between conditional and unconditional probability measures of weather indices. We address this concern in two ways.

First, although short-term weather technology has been improving dramatically over the last decades, long-term (e.g., several months or one-year ahead) weather prediction is still unlikely, given current technology (Alley et al. 2019, Voosen 2019). For example, currently the best forecast at the European Centre for Medium-Range Weather Forecasts runs out to around 10 days only (see Online Appendix K). Given the unpredictable nature of annual weather conditions, and the fact that index insurance contract is typically renewed annually, we expect the difference between the conditional and unconditional contracts to be small.

Second, to provide more direct evidence, in Table 7, we compare contracts that use the current year, the previous year, and the most recent two years (i.e., current year plus previous year) weather conditions as inputs. Examining the performances over the test sample in Table 7, we see that training the model with the previous year weather conditions generates significantly worse results than the model trained with the current year weather conditions. Also, using two-year weather conditions as inputs generates results similar to the model trained by the current year weather conditions. This suggests using weather conditions in the previous year adds little information to the current contract, and hence the potential discrepancy between the conditional and unconditional contracts is likely negligible.

6. Extensions

In this section, we explore several extensions to the baseline model. First, we discuss the complexity issue associated with the NN-based insurance contract and propose

Table 7. Contracts Conditioning on Various Weather Information

	Current year weather (BL)		Previous year weather		Last two years weather	
	($\lambda^* = 1.2414$)		($\lambda^* = 1.2131$)		($\lambda^* = 1.2469$)	
	Training	Test	Training	Test	Training	Test
Panel A: Utility improvement						
U with insurance	−3.57	−3.57	−3.68	−3.94	−3.57	−3.59
U w/o insurance	−3.99	−4.16	−4.02	−4.15	−4.02	−4.15
U improvement (%)	10.60%	14.35%	8.36%	5.26%	11.19%	13.57%
Panel B: CEW improvement						
CEW with insurance	444.64	444.61	440.52	432.28	444.45	443.76
CEW w/o insurance	430.63	425.26	429.61	425.52	429.61	425.52
CEW improvement	14.00	19.36	10.91	6.76	14.84	18.24
CEW improvement (%)	3.25%	4.55%	2.54%	1.59%	3.45%	4.29%
Panel C: Policy characteristics						
Premium	28.44	28.72	31.71	42.45	28.96	28.28
Coverage	22.91	23.13	26.14	34.99	23.23	22.68
Insurer profit	5.53	5.59	5.57	7.46	5.74	5.60
Panel D: Risk reduction measured by standard deviation						
Std	54.05	47.49	58.20	64.14	51.58	50.23
Std w/o insurance	81.94	78.92	81.17	78.97	81.17	78.97
Std reduction	34.04%	39.82%	28.30%	18.78%	36.46%	36.39%
Panel E: Risk reduction measured by value-at-risk (VaR)						
$VaR_{5\%}$	382.89	379.64	355.26	341.52	385.05	371.49
$VaR_{5\%}$ w/o insurance	316.28	325.91	315.86	325.90	315.86	325.90
$VaR_{5\%}$ improvement	66.61	53.73	39.40	15.63	69.19	45.59

Notes. This table compares contracts that use the current year, the previous year, and both years weather conditions as inputs. Panel A summarizes utilities with and without (w/o) index insurance policies and the percentage of utility improvement. Panel B summarizes certainty equivalent wealth (CEW) with and without (w/o) index insurance policies and CEW improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage, and profits for the insurers. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation of wealth. Panel E summarizes the risk reduction at the tail, measured by the 5%-level value-at-risk (VaR). "BL" indicates the baseline case studied in Section 4.2.

potential ways to address it. Next, we extend our setting to corporate farming where the insureds are risk neutral. We consider an objective function of minimizing tail risk, instead of maximizing expected utility.

6.1. Contract Complexity

In practice, contract complexity often deters insureds' participation in insurance markets, that is, complexity aversion (Sonsino et al. 2002, Bernheim and Sprenger 2020).[24] Admittedly, the NN-based contract is much more complicated than a simpler linear contract to farmers, even though the former provides much better protections. In this section, we first quantify the impacts of contract complexity on our index insurance performance in Section 6.1.1.

Clearly, it is important to alleviate farmers' concerns over contract complexity. First, we can increase the interpretability of the NN-based index insurance to improve farmers' understanding and trust in this product (Section 4.3). The literature also finds that education and insurance literacy can effectively improve insurance demand (Gaurav et al. 2011, Cai et al. 2020). Second, agricultural insurance often involves a public-private partnership (PPP). Government subsidies could be used to improve the communication and trust between the insurance companies and farmers, and we discuss this in Section 6.1.2. Last, we consider a hybrid index insurance that provides payoffs as the maximum of a linear contract and an NN-based contract in Section 6.1.3. This hybrid contract provides payoffs at least as good as a linear contract, which might help reduce the complexity concern of farmers.

6.1.1. Quantifying the Impacts of Contract Complexity. Concerns and distrust generated by contract complexity may affect farmers' perception of index insurance and deter their participation. However, it is difficult to directly measure contract complexity and complexity aversion. We follow Ceballos and Robles (2020) and consider the perceived value of insurance contract to the complexity averse farmers.[25] Ceballos and Robles (2020) show that the perceived value of index insurance could be reduced by about 40% relative to its true value if policyholders cannot understand the products. In Table 8,

Table 8. Impacts of Complexity Aversion

	(1) BL		(2) Value reduction of 20%		(3) Value reduction of 40%	
	($\lambda^* = 1.2414$)		($\lambda^* = 1.1253$)		($\lambda^* = 1.0376$)	
	Training	Test	Training	Test	Training	Test
Panel A: Utility improvement						
U with insurance	−3.57	−3.57	−3.69	−3.70	−3.81	−3.84
U w/o insurance	−3.99	−4.16	−3.99	−4.16	−3.99	−4.16
U improvement (%)	10.60%	14.35%	7.56%	11.16%	4.57%	7.66%
Panel B: CEW improvement						
CEW with insurance	444.64	444.61	440.46	440.05	436.48	435.22
CEW w/o insurance	430.63	425.26	430.63	425.26	430.63	425.26
CEW improvement	14.00	19.36	9.82	14.79	5.85	9.96
CEW improvement (%)	3.25%	4.55%	2.28%	3.48%	1.36%	2.34%
Panel C: Policy characteristics						
Premium	28.44	28.72	21.57	19.06	13.64	12.82
Coverage	22.91	23.13	19.17	16.94	13.15	12.36
Insurer profit	5.53	5.59	2.40	2.12	0.49	0.46
Panel D: Risk reduction measured by standard deviation						
Std	54.05	47.49	58.25	55.47	64.87	64.17
Std w/o insurance	81.94	78.92	81.94	78.92	81.94	78.92
Std reduction	34.04%	39.82%	28.91%	29.72%	20.83%	18.69%
Panel E: Risk reduction measured by value-at-risk (VaR)						
$VaR_{5\%}$	382.89	379.64	374.58	363.85	356.26	344.10
$VaR_{5\%}$ w/o insurance	316.28	325.91	316.28	325.91	316.28	325.91
$VaR_{5\%}$ improvement	66.61	53.73	58.30	37.95	39.98	18.19

Notes. This table quantifies the impact of complexity aversion on insurance contracts. Contract complexity affects farmers' perceived value of the insurance contract. Cases (2) and (3) assume farmers' perceived value is reduced by 20% and 40% relative to the true value of insurance, respectively. Case (1) is the baseline model without value reduction. Panel A summarizes utilities with and without (w/o) index insurance policies and the percentage of utility improvement. Panel B summarizes certainty equivalent wealth (CEW) with and without (w/o) index insurance policies and CEW improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage, and profits for the insurers. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation of wealth. Panel E summarizes the risk reduction at the tail, measured by the 5%-level value-at-risk (VaR). "BL" indicates the baseline case studied in Section 4.2.

we consider two different levels of value reduction in the *NN72* model. We consider a severe case (Case (3) in Table 8) when farmers do not understand the index insurance and their perceived value of the insurance is reduced by 40% relative to the true value of the insurance, and another moderate case when farmers better understand the index insurance; that is, their perceived value of the insurance is reduced by 20% relative to its true value (Case (2) in Table 8). We see that CEW improvement and purchased insurance coverage in Cases (2) and (3) decrease to about two thirds and one half of those in the baseline model (Case (1)), respectively. Nonetheless, Cases (2) and (3) still better enhance farmers' utility and CEW than simpler contracts, such as *Linear1*, *Linear5*, *Quadratic5*, and *Cubic5* in Table 4. Cases (2) and (3) perform similarly to *NN5* and *Linear72* in Table 4. This suggests the effectiveness of the NN-based contracts, even after reasonably capturing the complexity aversion.

6.1.2. Role of Government Subsidies.

Government subsidies help launch insurance programs and improve insurance market participation. For example, Cai et al. (2020) show that government subsidies together with education programs enhance long-term adoption of index insurance programs. In the United States, the Federal Crop Insurance Program (FCIP) costs the federal government about \$6.26 billion in 2019, with a subsidy ratio of about 60% (Rosa 2018, USDA 2019).

In this section, we quantitatively evaluate the impacts of government subsidies on the NN-based index insurance. Assume that the index insurance could be subsidized by the government for a proportion of θ. Then the terminal wealth of the farmer at the presence of the index insurance is $w_0 - Y + I(\boldsymbol{X}) - (1-\theta)\pi(I)$. The maximization problem is

$$\begin{cases} \min\limits_{I\in\mathcal{I}} & -\dfrac{1}{n}\sum\limits_{j=1}^{n} U(w_0 - y_j + I(\boldsymbol{x}_j) - (1-\theta)\pi_e(I)]), \\ \text{s.t.} & P_L \le \pi_e(I) \le P_U. \end{cases} \tag{9}$$

We consider various subsidy rates, $\theta = \{0, 5\%, 10\%, 15\%, 19.45\%\}$, in Table 9. Note that $\theta = 19.45\%$ is the subsidy level such that farmers pay the actuarially fair

Table 9. Impacts of Government Subsidies

	θ = 0 (BL) (λ^* = 1.2414)		θ = 0.05 (λ^* = 1.2948)		θ = 0.1 (λ^* = 1.3359)		θ = 0.15 (λ^* = 1.3911)		θ = 0.1945 (λ^* = 1.4400)	
	Training	Test	Training	Test	Training	Test	Training	Test	Training	Test
Panel A: Utility improvement										
U with insurance	−3.57	−3.57	−3.56	−3.56	−3.54	−3.54	−3.52	−3.52	−3.51	−3.50
U w/o insurance	−3.99	−4.16	−3.99	−4.16	−3.99	−4.16	−3.99	−4.16	−3.99	−4.16
U improvement (%)	10.60%	14.35%	10.79%	14.52%	11.26%	15.01%	11.62%	15.41%	12.00%	15.84%
Panel B: CEW improvement										
CEW with insurance	444.64	444.61	444.91	444.87	445.56	445.59	446.07	446.17	446.61	446.82
CEW w/o insurance	430.63	425.26	430.63	425.26	430.63	425.26	430.63	425.26	430.63	425.26
CEW improvement	14.00	19.36	14.28	19.61	14.93	20.33	15.43	20.91	15.97	21.56
CEW improvement (%)	3.25%	4.55%	3.32%	4.61%	3.47%	4.78%	3.58%	4.92%	3.71%	5.07%
Panel C: Policy characteristics										
Premium	28.44	28.72	29.90	29.79	32.77	33.09	36.36	36.29	37.11	39.54
Coverage	22.91	23.13	23.09	23.01	24.53	24.77	26.14	26.09	25.77	27.46
Insurer profit	5.53	5.59	6.81	6.78	8.24	8.32	10.22	10.20	11.34	12.08
Panel D: Risk reduction measured by standard deviation										
Std	54.05	47.49	53.91	47.58	53.02	46.34	52.10	45.44	52.36	44.56
Std w/o insurance	81.94	78.92	81.94	78.92	81.94	78.92	81.94	78.92	81.94	78.92
Std reduction	34.04%	39.82%	34.22%	39.71%	35.30%	41.28%	36.42%	42.42%	36.10%	43.54%
Panel E: Risk reduction measured by value-at-risk (VaR)										
$\text{VaR}_{5\%}$	382.89	379.64	382.92	378.89	383.70	380.40	384.96	384.57	385.29	386.77
$\text{VaR}_{5\%}$ w/o insurance	316.28	325.91	316.28	325.91	316.28	325.91	316.28	325.91	316.28	325.91
$\text{VaR}_{5\%}$ improvement	66.61	53.73	66.64	52.99	67.42	54.50	68.69	58.67	69.02	60.86

Notes. We consider an NN-based index insurance with various government subsidy rates, that is, $\theta = 0, 5\%, 10\%, 15\%$, or 19.45%. Panel A summarizes utilities with and without (w/o) different index insurance policies and the percentage of utility improvement. Panel B summarizes certainty equivalent wealth (CEW) with and without (w/o) index insurance policies and the CEW improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage and profits for the insurer. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation. Panel E summarizes the risk reduction at the tail, measured by the 5%-level value-at-risk (VaR). "BL" represents the baseline case studied in Section 4.2. The risk loading parameter at equilibrium (λ^*) for each contract is reported in parentheses.

premium, and the government pays related costs and loadings, which is much lower than the rate of 60% in crop insurance practice (Rosa 2018). With subsidies, the demand curve of policyholders moves upward, resulting in a higher loading parameter at equilibrium. Producer's utility, CEW, and coverages increase slightly. The insurer's profits increase significantly. This is because the insurer has a small price elasticity, whereas farmers are more sensitive to prices. Therefore, subsidies cause larger changes to insurance demand curve. A higher profit margin will encourage more insurance companies to participate in the weather index insurance market. Therefore, government subsidies can significantly increase insurance participation and improve social welfare.

6.1.3. Hybrid Contracts with Guaranteed Linear Payoffs. In this section, we propose a hybrid insurance contract as a temporary solution to the complexity concern. This hybrid insurance provides payoffs (I^{hybrid}) as the maximum of a linear contract and an NN-based contract, defined as follows:

$$I^{hybrid} = \max(\text{Linear}i, \text{NN}j) = \text{Linear}i + \max(0, \text{NN}j - \text{Linear}i),$$

where Linear$i$ and NNj denote generic payoff functions for a linear contract with i inputs and an NN-based index insurance with j inputs, respectively. With this hybrid contract, a farmer is guaranteed with a linear payoff that could be well understood and trusted and meanwhile enjoys the upside potential from the NN-based contract. That is, this hybrid insurance can simultaneously benefit from easy interpretability of a linear contract and large basis risk reduction of the NN-based contract.

We consider three hybrid contracts where an NN-based contract (*NN1*, *NN5*, or *NN72*) is combined with *Linear1* (the simplest linear contract). Table 10 summarizes the performances of these three hybrid contracts in the test sample. We can see that the hybrid contracts perform very close to the corresponding NN-based contracts. For example, the utility improvement of max(*Linear1*, *NN72*) is 14.02%, which is only slightly lower than that of the *NN72* contract (14.35%). Also, the hybrid contracts generally have higher premiums due to higher indemnity payments. Although hypothetical, such analysis provides important and direct evidence that our proposed NN-based index insurance is very effective in providing protections to farmers even after considering complexity aversion. Moreover, this analysis also sheds

Table 10. Performances of the Linear Contract, NN-Based Contracts, and Hybrid Contracts

	Linear1 ($\lambda^* = 1.0255$)	NN1 ($\lambda^* = 1.0512$)	NN5 ($\lambda^* = 1.1778$)	NN72 (BL, $\lambda^* = 1.2414$)	max(*Linear*1, *NN*1) ($\lambda^* = 1.0353$)	max(*Linear*1, *NN*5) ($\lambda^* = 1.0899$)	max(*Linear*1, *NN*72) ($\lambda^* = 1.1114$)
Panel A: Utility improvement							
U with insurance	−4.14	−4.14	−3.78	−3.57	−4.14	−3.74	−3.58
U w/o insurance	−4.16	−4.16	−4.16	−4.16	−4.16	−4.16	−4.16
U improvement (%)	0.55%	0.47%	9.11%	14.35%	0.47%	10.18%	14.02%
Panel B: CEW improvement							
CEW with insurance	425.94	425.84	437.19	444.61	425.84	438.68	444.14
CEW w/o insurance	425.26	425.26	425.26	425.26	425.26	425.26	425.26
CEW improvement	0.69	0.58	11.94	19.36	0.59	13.42	18.88
CEW improvement (%)	0.16%	0.14%	2.81%	4.55%	0.14%	3.16%	4.44%
Panel C: Policy characteristics							
Premium	24.48	27.24	27.94	28.72	30.90	40.24	41.87
Coverage	23.87	25.91	23.72	23.13	29.84	36.92	37.67
Insurer profit	0.61	1.33	4.22	5.59	1.05	3.32	4.20
Panel D: Risk reduction measured by standard deviation							
Std	75.13	72.92	62.86	47.49	73.07	61.84	52.06
Std w/o insurance	78.92	78.92	78.92	78.92	78.92	78.92	78.92
Std reduction	4.80%	7.60%	20.34%	39.82%	7.41%	21.64%	34.03%
Panel E: Risk reduction measured by value-at-risk (VaR)							
$\text{VaR}_{5\%}$	332.92	335.30	339.00	379.64	334.96	346.23	371.79
$\text{VaR}_{5\%}$ w/o insurance	325.91	325.91	325.91	325.91	325.91	325.91	325.91
$\text{VaR}_{5\%}$ improvement	7.02	9.39	13.09	53.73	9.05	20.32	45.88

Notes. This table compares the performances of three hybrid contracts in the test sample, including a contract with the maximum payment between *Linear1* and *NN1*, a contract with the maximum payment between *Linear1* and *NN5*, and a contract with the maximum payment between *Linear1* and *NN72*. For ease of comparison, we also list the results of *Linear1*, *NN1*, *NN5*, and *NN72*. Panel A summarizes expected utilities with and without (w/o) index insurance and the percentage of utility improvement. Panel B reports certainty equivalent wealth (CEW) with and without (w/o) index insurance and the CEW improvement in dollars and as a percentage. Panel C summarizes policy characteristics including policy premium, coverage, and profits of the insurer. Panel D summarizes the risk reduction effect of an index insurance policy, measured by the standard deviation of wealth. Panel E summarizes the tail risk reduction, measured by the 5%-level value-at-risk (VaR). The risk loading parameter (λ^*) for each contract is reported in parentheses.

Table 11. Tail Risk Optimization

	$\text{VaR}_{5\%}$ maximization		Utility maximization (BL)	
	Training	Test	Training	Test
Panel A: Utility improvement				
U with insurance	−3.58	−3.60	−3.57	−3.57
U w/o insurance	−3.99	−4.16	−3.99	−4.16
U improvement (%)	10.19%	13.43%	10.60%	14.35%
Panel B: CEW improvement				
CEW with insurance	444.07	443.28	444.64	444.61
CEW w/o insurance	430.63	425.26	430.63	425.26
CEW improvement	13.44	18.02	14.00	19.36
CEW improvement (%)	3.12%	4.24%	3.25%	4.55%
Panel C: Policy characteristics				
Premium	46.37	45.17	28.44	28.72
Coverage	37.36	36.39	22.91	23.13
Insurer profit	9.02	8.79	5.53	5.59
Panel D: Risk reduction measured by standard deviation				
Std	51.65	43.45	54.05	47.49
Std w/o insurance	81.94	78.92	81.94	78.92
Std reduction	36.97%	44.94%	34.04%	39.82%
Panel E: Risk reduction measured by value-at-risk (VaR)				
$\text{VaR}_{5\%}$	388.57	386.91	382.89	379.64
$\text{VaR}_{5\%}$ w/o insurance	316.28	325.91	316.28	325.91
$\text{VaR}_{5\%}$ improvement	72.30	61.01	66.61	53.73

Notes. We consider a risk-neutral agent with a tail risk minimization objective. The tail risk measure considered is $\text{VaR}_{5\%}$. Panel A summarizes utilities with and without (w/o) different index insurance policies and the percentage of utility improvement. Panel B summarizes certainty equivalent wealth (CEW) with and without (w/o) index insurance policies and the CEW improvements in dollars and as a percentage. Panel C summarizes policy characteristics, including premiums, coverage, and profits for the insurer. Panel D summarizes the risk reduction effect of different index insurance policies, measured by the standard deviation of wealth. Panel E summarizes risk reduction at the tail, measured by the 5%-level value-at-risk (VaR).

light on the implementation of the NN-based index insurance in practice. Initially, insurance companies could collaborate with the government and provide some subsidized programs of such hybrid index insurance contracts, together with certain types of insurance education programs, to gain trust from farmers and improve insurance demand. Once the idea of this hybrid insurance program becomes well accepted by farmers, the hybrid contract may be gradually phased into the cheaper and more efficient NN-based index insurance.

6.2. Corporate Farming

The framework discussed in this paper can be easily extended to corporate farming or internal risk management of insurance companies. Corporate farms usually are risk neutral and interested in reducing tail risk to achieve solvency in extreme scenarios. Therefore, policyholders aim to minimize their tail risk instead of maximizing expected utility. In particular, we consider the following tail risk minimizing problem:

$$\min_{I\in\mathcal{I}} -\text{VaR}_{5\%}(\{w_0 - y_j + I(\boldsymbol{x}_j) - \pi_e(I)\}_{j=1,\ldots,n}), \tag{10}$$

where $\text{VaR}_{5\%}$ is the 5%-level VaR that is a commonly used tail risk measure. Table 11 summarizes the results. We see that, compared with the utility maximization baseline framework, the improvements in $\text{VaR}_{5\%}$ are larger when farmer maximizes $\text{VaR}_{5\%}$, whereas the improvements in expected utility and CEW are slightly lower. It is also shown that the designed contract is substantially more expensive than the baseline contract, indicating that it is more costly to manage tail risk. However, such a premium level is still affordable in the case of corporate farming.

7. Conclusion

Index insurance could effectively manage systemic weather risk. However, the current insurance has large basis risk and is less cost effective, which leads to low insurance demand. In this paper, we formulate a neural network-based design of an index insurance contract which helps to reduce basis risk. Moreover, we consider endogenous insurance premium and demand to further improve the cost effectiveness of index insurance.

We illustrate the superior performance of this framework by applying it to corn farmers in Illinois. Results show that the NN-based index insurance contract effectively reduces basis risk and greatly outperforms other contracts (e.g., piecewise linear contract, quadratic contract, or cubic contract). This NN-based insurance contract significantly increases market demand, mitigates downside risk, and improves farmers' utility, even when we consider some frictions in practice (e.g., overinsuring constraints, regulatory costs, and contract complexity aversion). Our framework can be easily extended to other settings, for example, the revenue index insurance and corporate farming. Overall, our results suggest the promise of using the NN-based model to design a broad class of financial products. Nevertheless, the marketing and operating perspectives of NN-based insurance contracts need further investigations, which we leave for future work.

Acknowledgments
The authors are grateful for helpful comments from Agostino Capponi (the editor), the associate editor, three anonymous referees, Utpal Bhattacharya, Darwin Choi, Lin William Cong, Guanhao Feng, Yan Ji, Wenxi Jiang, Simone Krummaker (discussant), Kai Li, Weikai Li, Yingying Li, Peter Mackay, Abhiroop Mukherjee, Yoshio Nozawa, George Panayotov, Yang Shen, Ke Tang, Qihe Tang, Jialin Yu, Chu Zhang, and the seminar participants at World Risk and Insurance Economics Congress 2020, 24th International Congress on Insurance: Mathematics and Economics 2021, AI in Finance Symposium 2022, 2022 CORS/INFORMS International Conference, East China Normal University, Hong Kong University of Science and Technology, Peking University, Renmin University, Southwestern University of Finance and Economics, Southern University of Science and Technology, Sun Yat-sen University, Tsinghua University, and University of New South Wales.

Endnotes

[1] Other factors also affect the demand of index insurance, including asymmetric information (Hartman-Glaser and Hébert 2020), ambiguity aversion (Bryan 2019), farmers' past insurance payout experiences (Cole et al. 2014, Cai et al. 2020), and insurance literacy (Gaurav et al. 2011, Cai et al. 2020).

[2] We extend this framework to revenue protection in Online Appendix M, which considers crop price risk.

[3] Other machine learning approaches could be feasible. For example, one could suggest reinforcement learning (Alsabah et al. (2021) use reinforcement learning to learn investors' risk preference based on their trading experiences), but we expect this approach offers limited benefits in our setting. Because current atmospheric models cannot forecast weather conditions and production losses one-year ahead (Voosen 2019), farmers cannot time the index insurance market at the annual frequency, which implies limited gains from reinforcement learning. Nevertheless, reinforcement learning could be helpful if we consider farmers with bounded rationality, for example, farmers are short-sighted and only learn from the recent experiences. We avoid designing the insurance contract based on behavioral bias as this might introduce some litigation risks in practice. Tree-based models (Rossi and Timmermann 2015, Gu et al. 2020, Li and Rossi 2021, Rossi and Utkus 2021, Cong et al. 2022) and support vector machines (SVMs) are also popular machine learning approaches that could address high-dimensionality and nonlinearity. In Online Appendix O, we compare the performance of NN-based model with different tree-based models, including a simple regression tree, tree bagging, random forest, and tree boosting, and a SVM model with the radial kernel function. We find that the NN-based models perform best for our problem.

[4] In a standard NN-based statistical learning problem, the objective function is usually a loss function defined as certain distance measure (e.g., L^1 norm for mean absolute error, L^2 norm for mean squared error, the Huber loss, or the cosine similarity, all of which are commonly used built-in functions in Keras). For our problem, we create a customized loss function (4) using the backend functions in Keras.

[5] We will test the impact of budget constraints in the robustness analysis in Online Appendix L.3.

[6] *Source:* USDA's National Agricultural Statistics Service, Economic Research Service.

[7] One potential issue of the raw crop yield data is data contamination. For example, in early years, farmers were insured with conventional insurance and might be prone to moral hazard. Hence, losses might have been exaggerated. This issue could be partially mitigated via the detrending process but still inevitably introduces noises in our results.

[8] We discuss the data homogeneity in Online Appendix G.

[9] Specifically, we obtain the commodity prices for corn at harvest from the USDA Economic Research Service, which are estimated using data from USDA's Agricultural Resource Management Survey and other sources (see https://www.ers.usda.gov/data-products/commodity-costs-and-returns/). Inflation is adjusted for by the annual consumer price index (CPI) from the Bureau of Labor Statistics.

[10] Available at https://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/Ag_Resource_Management/.

[11] See "Farm Income and Production Costs for 2017: Advance Report", available at http://www.fbfm.org/pdfs/AdvanceReport17.pdf.

[12] PRISM is the USDA's official climatological data, available at http://prism.oregonstate.edu/.

[13] Available at: https://www.rma.usda.gov/en/Information-Tools/Summary-of-Business.

[14] This is in line with the literature. For example, Chavas and Holt (1990) find that maize and soybean farmers have a relative risk aversion ranging from 1.41 to 7.62.

[15] This interval covers the feasible loading parameter for index insurance in practice.

[16] The mean wealth is slightly lower with insurance, because of the insurance premium paid.

[17] We find similar results over the test set.

[18] Online Appendix H ranks all indices by the gradient-based sensitivity analysis.

[19] Suyker and Verma (2008) find that cumulative evapotranspiration during the "nongrowing" seasons contributes to 20%–25% of the annual evapotranspiration totals for corn. Li et al. (2019) show that water storage over nongrowing seasons affects corn yields.

[20] Cover crops, such as cereal rye, which help build and improve soil fertility and quality, control diseases and pests, and promote biodiversity, are commonly integrated into corn production in Illinois. In fact, in Illinois, farmers who adopt cover crops may be eligible to receive an insurance premium discount in the following year through the Illinois Department of Agriculture (IDOA) Cover Crop Premium Discount Program (Illinois Department of Agriculture 2020).

[21] Polynomial terms might not be orthogonal, so using polynomials of a higher degree would introduce multicollinearity/robustness issues.

[22] Instead of using 72 weather indices, one might selectively use fewer weather indices and still achieve reasonable insurance

contracts. Online Appendix J further compares NN-based models with different number of weather indices.

[23] In Online Appendix Q, we further perform a distant out-of-state test. That is, we use North Dakota, a state in corn belt but geographically distant from Illinois, as a negative test sample. As expected, the NN-based contract trained with Illinois data does not work well for North Dakota due to their dissimilarities in weather patterns.

[24] In Online Appendix N, we consider the regulatory costs associated with contract complexity.

[25] Contract complexity might effectively increase farmers' perceived uncertainty about the insurance payout. In Online Appendix P, we consider this alternative way to capture contract complexity and find similar results.

References

Alley RB, Emanuel KA, Zhang F (2019) Advances in weather prediction. *Science* 363(6425):342–344.

Alsabah H, Capponi A, Ruiz Lacedelli O, Stern M (2021) Robo-advising: Learning investors' risk preferences via portfolio choices. *J. Financial Econom.* 19(2):369–392.

Assa H, Wang M (2021) Price index insurances in the agriculture markets. *North Amer. Actuary J.* 25(2):286–311.

Avramov D, Cheng S, Metzker L (2022) Machine learning vs. economic restrictions: Evidence from stock return predictability. *Management Sci.* 69(5):2587–2619.

Bernheim BD, Sprenger C (2020) On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. *Econometrica* 88(4):1363–1409.

Bianchi D, Büchner M, Tamoni A (2021) Bond risk premia with machine learning. *Rev. Financial Stud.* 34(2):1046–1089.

Bryan G (2019) Ambiguity aversion decreases the impact of partial insurance: Evidence from African farmers. *J. Eur. Econom. Assoc.* 17(5):1428–1469.

Bryzgalova S, Pelger M, Zhu J (2021) Forest through the trees: Building cross-sections of stock returns, Working paper, Stanford University, Stanford, CA.

Cai J, De Janvry A, Sadoulet E (2020) Subsidy policies and insurance demand. *Amer. Econom. Rev.* 110(8):2422–2453.

Capponi A, Lehalle CA, eds. (2023) *Machine Learning and Data Sciences for Financial Markets: A Guide to Contemporary Practices* (Cambridge University Press, Cambridge, UK).

Capponi A, Olafsson S, Zariphopoulou T (2022) Personalized robo-advising: Enhancing investment through client interaction. *Management Sci.* 68(4):2485–2512.

Casaburi L, Willis J (2018) Time vs. state in insurance: Experimental evidence from contract farming in Kenya. *Amer. Econom. Rev.* 108(12):3778–3813.

Ceballos F, Robles M (2020) Demand heterogeneity for index-based insurance: The case for flexible products. *J. Developement Econom.* 146:102515.

Chavas JP, Holt MT (1990) Acreage decisions under risk: The case of corn and soybeans. *Amer. J. Agricultural Econom.* 72(3):529–538.

Chen L, Pelger M, Zhu J (2023) Deep learning in asset pricing. *Management Sci.*, ePub ahead of print February 20, https://doi.org/10.1287/mnsc.2023.4695.

Clarke DJ (2016) A theory of rational demand for index insurance. *Amer. Econom. J. Microeconom.* 8(1):283–306.

Cole S, Stein D, Tobacman J (2014) Dynamics of demand for index insurance: Evidence from a long-run field experiment. *Amer. Econom. Rev.* 104(5):284–290.

Cole S, Giné X, Tobacman J, Topalova P, Townsend R, Vickery J (2013) Barriers to household risk management: Evidence from India. *Amer. Econom. J. Appl. Econom.* 5(1):104–135.

Cong LW, Liang T, Zhang X (2020) Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. Working paper, Cornell University, Ithaca, NY.

Cong LW, Feng G, He J, He X (2022) Asset pricing with panel trees under global split criteria. Working paper, Cornell University, Ithaca, NY.

Cong LW, Tang K, Wang J, Zhang Y (2021a) AlphaPortfolio: Direct construction through deep reinforcement learning and interpretable AI. Working Paper, Cornell University, Ithaca, NY.

Cong LW, Tang K, Wang J, Zhang Y (2021b) Deep sequence modeling: Development and applications in asset pricing. *J. Financial Data Sci.* 3(1):28–42.

Crane-Droesch A (2018) Machine learning methods for crop yield prediction and climate change impact assessment in agriculture. *Environment. Res. Lett.* 13(11):114003.

Cummins JD, Lalonde D, Phillips RD (2004) The basis risk of catastrophic-loss index securities. *J. Financial Econom.* 71(1):77–111.

Deng X, Barnett BJ, Vedenov DV (2007) Is there a viable market for area-based crop insurance? *Amer. J. Agricultural Econom.* 89(2): 508–519.

Doherty NA, Schlesinger H (1983) Optimal insurance in incomplete markets. *J. Political Econom.* 91(6):1045–1054.

Engle RF, Giglio S, Kelly BT, Lee H, Stroebel J (2019) Hedging climate change news. *Rev. Financial Stud.* 33(3):1184–1216.

Feng G, Giglio S, Xiu D (2020) Taming the factor zoo: A test of new factors. *J. Finance* 75(3):1327–1370.

Feng G, Polson N, Xu J (2021) Deep learning in characteristics-sorted factor models. Working paper, University of Chicago, Chicago, IL.

Fisher AC, Hanemann WM, Roberts MJ, Schlenker W (2012) The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *Amer. Econom. Rev.* 102(7):3749–3760.

Gaurav S, Cole S, Tobacman J (2011) Marketing complex financial products in emerging markets: Evidence from rainfall insurance in India. *J. Marketing Res.* 48(SPL):S150–S162.

Gu S, Kelly B, Xiu D (2020) Empirical asset pricing via machine learning. *Rev. Financial Stud.* 33(5):2223–2273.

Gu S, Kelly B, Xiu D (2021) Autoencoder asset pricing models. *J. Econometrics* 222(1):429–450.

Harri A, Coble KH, Ker AP, Goodwin BJ (2011) Relaxing heteroscedasticity assumptions in area-yield crop insurance rating. *Amer. J. Agricultural Econom.* 93(3):707–717.

Hartman-Glaser B, Hébert B (2020) The insurance is the lemon: Failing to index contracts. *J. Finance* 75(1):463–506.

Hastie T, Tibshirani R, Friedman JH, Friedman JH (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer, Berlin).

Hong HG, Karolyi GA, Scheinkman J (2020) Climate finance. *Rev. Financial Stud.* 33(3):1011–1023.

Huang JZ, Shi Z (2022) Machine-learning-based return predictors and the spanning controversy in macro-finance. *Management Sci.* 69(3): 1780–1804.

Illinois Department of Agriculture (2020) Cover crops premium discount program. Technical report, Illinois Department of Agriculture, Springfield, IL.

Jensen ND, Barrett CB, Mude AG (2016) Index insurance quality and basis risk: Evidence from northern Kenya. *Amer. J. Agricultural Econom.* 98(5):1450–1469.

Lesk C, Rowhani P, Ramankutty N (2016) Influence of extreme weather disasters on global rop production. *Nature* 529:84–88.

Li B, Rossi AG (2021) Selecting mutual funds from the stocks they hold: A machine learning approach, Working paper, Georgetown University, Washington, DC.

Li Y, Guan K, Schnitkey GD, DeLucia E, Peng B (2019) Excessive rainfall leads to maize yield loss of a comparable magnitude to extreme drought in the United States. *Global Change Biology* 25(7):2325–2337.

Luenberger DG, Ye Y (1984) *Linear and Nonlinear Programming*, 5th ed. (Springer, Berlin).

Mahul O, Skees J (2007) *Managing Agricultural Risk at the Country Level: The Case of Index-Based Livestock Insurance in Mongolia* (The World Bank).

Newlands N, Ghahari A, Gel YR, Lyubchich V, Mahdi T (2019) Deep learning for improved agricultural risk management. Bui TX, ed. *Proc. 52nd Hawaii Internat. Conf. on System Sci.* (HICSS, Honolulu, HI), 1033–1042.

Nordhaus W (2019) Climate change: The ultimate challenge for economics. *Amer. Econom. Rev.* 109(6):1991–2014.

Ribeiro MT, Singh S, Guestrin C (2016) "Why should I trust you?" Explaining the predictions of any classifier. *Proc. 22nd ACM SIGKDD Internat. Conf. on Knowledge Discovery and Data Mining*, 1135–1144.

Rigden A, Mueller N, Holbrook N, Pillai N, Huybers P (2020) Combined influence of soil moisture and atmospheric evaporative demand is important for accurately predicting US maize yields. *Nature Food* 1(2):127–133.

Rosa I (2018) Federal crop insurance: Program overview for the 115th congress. Technical report, Congressional Research Service.

Rossi AG, Timmermann A (2015) Modeling covariance risk in Merton's ICAPM. *Rev. Financial Stud.* 28(5):1428–1461.

Rossi AG, Utkus SP (2021) Who benefits from robo-advising? Evidence from machine learning. Working paper, Georgetown University, Washington, DC.

Schlenker W, Roberts MJ (2009) Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proc. National Acad. Sci. USA* 106(37):15594–15598.

Sirignano JA, Sadhwani A, Giesecke K (2021) Deep learning for mortgage risk. *J. Financial Econom.* 19:313–368.

Smith VH (2017) Achieving rational farm subsidy rates. Technical Report 113, R Street Policy Study, R Street Institute, Washington, DC.

Sonsino D, Benzion U, Mador G (2002) The complexity effects on choice with uncertainty: Experimental evidence. *Econom. J. (London)* 112(482):936–965.

Suyker AE, Verma SB (2008) Interannual water vapor and energy exchange in an irrigated maize-based agroecosystem. *Agricultural Forest Meteorology* 148(3):417–427.

Turvey CG (2001) Weather derivatives for specific event risks in agriculture. *Appl. Econom. Perspective Policy* 23(2):333–351.

USDA (2014) World agricultural supply and demand estimates report (WASDE). Office of the chief economist (OCE). U.S. Department of Agriculture, Washington, DC.

USDA (2019) State/county/crop summary of business. Accessed August 18, 2023, https://rma.usda.gov/en/Information-Tools/Summary-of-Business/State-County-Crop-Summary-of-Business.

Voosen P (2019) A 2-week weather forecast may be as good as it gets. *Science* 363(6429):801–801.

You J, Li X, Low M, Lobell D, Ermon S (2017) Deep Gaussian process for crop yield prediction based on remote sensing data. *Proc. 31st AAAI Conf. on Artificial Intelligence*, 4559–4565.