Received: 29 November 2021 | Accepted: 1 April 2022

DOI: 10.1002/jaa2.9

JOURNAL OF THE AGRICULTURAL AND APPLIED ECONOMICS ASSOCIATION
Open Access

ARTICLE

Risk, crop yields, and weather index insurance in village India

Jeffrey D. Michler[1] | **Frederi G. Viens**[2] | **Gerald E. Shively**[3]

[1]Department of Agricultural and Resource Economics, University of Arizona, Tucson, Arizona, USA

[2]Department of Statistics and Probability, Michigan State University, East Lansing, Michigan, USA

[3]Department of Agricultural Economics, Purdue University, West Lafayette, Indiana, USA

Correspondence

Jeffrey D. Michler, Department of Agricultural and Resource Economics, University of Arizona, Tucson, AZ, USA.
Email: jdmichler@arizona.edu

Funding information

We acknowledge funding from the Actuarial Science Program at Michigan State University, the United States Agency for International Development's Feed the Future Innovation Lab for Nutrition, and partial support from the National Science Foundation (DMS 1811779).

Abstract

We investigate the sources of variability in crop yields and their relative importance in the context of weather index insurance for smallholder farmers in India. Using parcel-level panel data, multilevel modeling, and Bayesian methods we measure how large a role seasonal variation in weather plays in explaining yield variance. Seasonal variation in weather accounts for 14%–22% of total variance in crop yields, a value that has rarely been convincingly estimated in the economics literature. These calculations shed light on the relatively low rates of uptake of index insurance by resource-constrained farmers. The results also provide direction for designing more suitable index insurance products.

KEYWORDS

agricultural production, Bayesian analysis, index insurance, India, multilevel models, weather risk

JEL CLASSIFICATION

C11, D81, G22, O12, O13, Q16, Q12

1 | INTRODUCTION

Agricultural production is complex and risky. Weather is just one potential cause of yield variability, along with the quantity and quality of inputs, the characteristics of parcels, farmer ability, the policy environment, and changes in technologies (Hardaker et al., 1997).

In this paper we use parcel-level panel data from India to measure the sources of variability in crop yields and assess their relative importance. Ascertaining weather-induced variation in yields is essential when designing and implementing weather index insurance, a micro-level risk management strategy promoted in developing countries. However, pilot projects and randomized control trials in Asia have found limited farmer uptake of index insurance (Bjerge & Trifkovic, 2018; Cai et al., 2015; Cole et al., 2013;

Giné et al., 2012; Hazell et al., 2010; Matsuda & Kurosaki, 2019; Mobarak & Rosenzweig, 2013; Shirsath et al., 2019; Stein, 2018). Using a multilevel modeling approach and Bayesian estimation we find that a small, but not insubstantial, fraction of the variability in yields can be attributed to seasonal variation in weather. Our Bayesian results reveal the presence of infrequent but potentially costly extreme weather events. However, we find that, on average, weather variability plays a minority role in contributing to yield variability. This result may help to explain the low rates of uptake of weather index insurance by farmers.

The goal of weather index insurance is to assist farmers in managing covariate risk. In developing country agriculture, covariate risk is particularly difficult to informally insure against since, by its very nature, it affects neighboring households and limits the effectiveness of traditional risk sharing mechanisms. Several early studies looked at the role of weather risk on crop yields and insurance in India. Townsend (1994) found that household consumption moves with village-level consumption and is not affected much by idiosyncratic shocks, which households cover via borrowing, gifts, and asset sales. He concludes that while households can insure against individual risks, covariate risk remains a problem. Rosenzweig and Binswanger (1993) find that village-level rainfall variables explain only a small proportion of household-level profit variability. And, similar to Townsend (1994), they find households insuring against non-covariate risk. In part, the recognition that households are less successful in insuring against covariate weather risk than in insuring against non-covariate risk motivates recent attempts to design and implement weather-based index insurance for poor farmers (Mobarak & Rosenzweig, 2013).

This apparent need for weather index insurance has been met by surprisingly little uptake of the product by farmers in India (Giné et al., 2008, 2012, 2017). There are many potential barriers to uptake of insurance, including trust in the insurance company, household liquidity constraints, and lack of financial literacy of household decision makers (Cole et al., 2013).[1] In addition to these consumer-centric explanations for low uptake, flaws in the product itself may be a cause, specifically the imperfect correlation between crop yield and farm profit, the outcome of ultimate interest to farmers (Binswanger-Mkhize, 2012; Shirsath et al., 2019). This uninsured exposure, or basis risk, can be a significant deterrent to purchasing insurance by reducing the utility gain for households. There are two potential sources of basis risk. One is the loss caused by an event not measured by the index. This could include, for example, low temperatures that retard crop growth or high winds that cause crop damage, when the insurance index is based on rainfall. Another potential source of basis risk is low correlation between the index and covered losses. This can occur, for example, when measurement stations that are used to trigger payouts are spatially distant from insured crops and provide poor congruence with local conditions. A less common example is potentially insured crops that are sufficiently resilient to the phenomena which forms the basis for the insurance, such as drought-tolerant varieties covered by rainfall-based insurance or pest-resistant cotton to pest infestation (Liu, 2013).

The most common type of weather index insurance in India is rainfall insurance (Akter et al., 2009; Barnett & Mahul, 2007; Matsuda & Kurosaki, 2019; Stein, 2018). However, surprisingly little economic research has been conducted to confirm the underlying relationship between rainfall or seasonal weather variability and yield variability. Barnett and Mahul (2007) suggest that rainfall variability accounts for 50% of yield variability and Giné et al. (2012) assert that as much as 90% of variation in Indian crop yields is driven by rainfall volatility, citing Parchure (2002) as their source. But Parchure relies on a 1976 report by the Indian National Commission on Agriculture (NCA), the details of which are somewhat opaque.[2] Finally, to our knowledge, no data exist on covariate risk's share in explaining variance in crop yield, which is where we focus our attention. Previous studies that have attempted such measurements

[1] For example, Banerjee et al. (2014) find no demand among Indian households for health insurance bundled with microfinance, even among those for whom there was clear value. They attribute this low uptake to poor understanding of the insurance product and poor support for enrolling by insurance underwriters.

[2] According to Parchure (2002), the 90% variation is actually for cotton and groundnuts while variation in yield due to rainfall variability is 45% and 47% for wheat and barley, respectively. Parchure (2002) actually cites the 1976 NCA report as stating "rainfall variations accounted for 50% of the variability in agricultural yields." This may be where Barnett and Mahul (2007) get their number. Examining the report, specifically pages 47–8, these numbers were obtained from a linear regression "with yield as the dependent variable and total rainfall during the five crop growth phases as the independent variable" (National Commission on Agriculture, 1976). The 50% number is apparently the R^2 on the regression.

were constrained by data, econometric techniques, and computing power so that they were unable to net out measurement error and fully describe the sources of output variability. We attempt to rectify these shortcomings and fill this information gap.

We examine the different sources of yield variability using a multilevel/hierarchical regression framework. This approach more fully accounts for the covariance structure of the data than a standard regression framework and allows us to control for inputs at the parcel-level and also to isolate the amount of yield variance due to differences between parcels, households, seasonal weather, villages, and time. Using Bayesian methods, we draw the underlying distribution of the random error term corresponding to different sources of variability, thereby providing a quantitative measure of potentially insurable risk. Bayesian methods are particularly useful to this task because the underlying distributions for several of the error terms turn out to be highly skewed and non-normal. Considering all sources of yield variance, we find that, depending on crop type, seasonal variation in weather accounts for 14%–22% of total variance in crop yields. This value is substantially below the 50%–90% figures commonly cited in the economics literature, though our estimate is more in line with evidence from agronomy and environmental science (Porter & Semenov, 2005; Ray et al., 2015).

By using a multilevel approach to measuring the roles of idiosyncratic and covariate risk in crop yields in village India we contribute to two separate streams of research. The first is the literature focused on production risk and insurance in the developing world. Our aim is to connect two strands of this study. The first, originating with Townsend (1994), Rosenzweig and Binswanger (1993), Rosenzweig and Wolpin (1993), Kochar (1999), and Mobarak and Rosenzweig (2013) has focused on the theoretical and empirical relationship between production, risk, and insurance. The second focuses on interventions to help insure production against risk (Akter et al., 2009; Barnett & Mahul, 2007; Chantarat et al., 2013; Clarke et al., 2012a; Cole et al., 2013; Gaurav et al., 2011; Giné et al., 2007, 2008; Skees et al., 2007). We empirically estimate the role different sources of variability play in determining the mean and variance of crop yield. This information is informative for determining the value and pricing of index insurance products.

The second body of literature to which we contribute is the empirical estimation of production and yield functions (Griliches, 1963; Ker et al., 2016; Ramsey, 2020; Tack et al., 2012; Woodard & Sherrick, 2011). Until recently it has been computationally difficult to estimate yield functions with multiple nested levels of data. While panel data fixed effects methods mimic a multilevel model with a single level, few alternatives exist for models with a large number of nested levels. Production data are often collected at the farm or factory level, which suggests at minimum a two-level model. Adding additional levels to account for and distinguish among temporal and/or spatial variation can provide more efficient estimation of production by more accurately modeling the data generating process. We contribute to this strand of literature by comparing yield function parameter estimates obtained by a standard linear model to multilevel regression estimates obtained by maximum likelihood and Bayesian estimation techniques.

2 | DATA

To conduct our empirical analysis, we use household data from villages in India. These data were collected as part of the Village Level Studies/Village Dynamics Study of South Asia (VDSA, 2015). The data set combines high and low-frequency household data from 30 Indian villages. The villages include the three studied by Townsend (1994) and the 10 studied by Rosenzweig and Binswanger (1993). However, while much of the previous research has relied on the low-frequency data, we utilize a newly available high-frequency data set covering the years 2009 through 2013. These data include monthly household observations on input purchases and labor expenditure for on-farm activities and crop production. The added value of a high-frequency data set is that, with multiple crops on multiple parcels for multiple seasons in a single year, it provides much more detailed and accurate farm production information than the low-frequency data.

We utilize monthly parcel-level data aggregated to the seasonal level. We focus on five crops: paddy rice, sorghum, wheat, maize, and cotton. Together these crops account for 60% of total crop observations in the VDSA and cover 78% of the total parcel-level observations.[3] This provides us with 11,942 parcel-level observations. Rice is the most common crop, accounting for 46% of total observations. Next most common is sorghum, accounting for 23% of observations. Wheat, maize, and cotton account for 14%, 9%, and 8% of observations.

The locations of VDSA villages provides us with a great deal of heterogeneity in climate, weather, crop choice, and cultivation practices. Statistics describing these data, by crop, are presented in Table 1. Twenty-seven villages cultivate at least two crops, generally a single crop in each of the two growing seasons. Only two villages cultivate all five crops. Cotton is the most input-intensive crop, using more labor, fertilizer, mechanization, and pesticide than any other crop. Sorghum is the least labor intensive while wheat uses the least fertilizer, mechanization, and pesticide. In total, our 11,942 parcel-level observations come from 5100 unique parcels operated by 1079 distinct households, in 30 villages, farming across 10 seasons. We exploit this nested data structure in our empirical analysis.

One final source of variation in the data structure, which plays an important role in our analysis, is the number of time observations for each crop. Three crops (rice, wheat, and maize) are grown in both the *Kharif* monsoon season and the post-monsoon *Rabi* season. Sorghum is only grown in *Rabi* while cotton is only grown in *Kharif*. The limited time series for these crops may affect our estimation of seasonality's effects on yield variability, albeit in an unknown way.

3 | ECONOMETRIC FRAMEWORK

3.1 | The linear model

We begin by estimating a simple linear regression for yield via ordinary least squares (OLS). Let y_{is} denote the log of yield for parcel i in season s. We estimate

$$y_{is} = X_{is}\beta + \alpha_s + \epsilon_{is}, \tag{1}$$

where X_{is} is a matrix of data and β is a vector of regression coefficients associated with various crops. To account for the spatial dimension and allow our proxy for weather to vary across both location and time, the seasonal indicator, α_s, is a fixed or constant effect for each time period $(t = 10)$ in each village $(v = 30)$ on yield. We assume the error term is $\epsilon_{is} \sim \mathcal{N}(0, \sigma^2)$ so that $y_{is} \sim \mathcal{N}(X_{is}\beta + \alpha_s, \sigma^2)$, where β and σ^2 are regression and variance parameters which are season independent while α_s depends on the season.

We note several drawbacks associated with this linear estimation of the yield function. First is that it precludes one from discerning the role of weather in yield variability. Equation (1) can estimate the impact of parcel-level inputs on yield at each point in time as well as the impact of a seasonal weather indicator on yield. But because these variables impact mean yield, the specification does not allow us to measure the share of yield variability (σ^2) represented by seasonal variability in weather. A second drawback is that the linear model limits our ability to account for additional group effects, such as differences between parcels, households, villages, or time. While changes in season clearly impact the effectiveness of parcel-level inputs, equally relevant variation may exist at these other levels. Some households may be more efficient in their application of labor compared to

[3]With the exception of pigeon pea, a perennial crop, the five crops used in our analysis are the five most frequently observed crops in the data set. We include cotton instead of pigeon pea for two reasons. First, pigeon pea is a perennial crop and therefore may be treated by farmers differently than annual crops when considering insurance. Second, cotton is one of the crops considered by Clarke et al. (2012b) in their product design and ratemaking for insurance contracts in Gujarat. Thus, despite it being a nonfood crop, we include it to bring our analysis more closely in line with existing literature.

TABLE 1 Descriptive statistics by crop

	Rice	Sorghum	Wheat	Maize	Cotton	Total
Yield (kg/ha)	3356	2368	1537	2596	1383	2658
	(3252)	(1429)	(4038)	(2397)	(870.4)	(2956)
Labor (hr/ha)	880.1	301.6	760.5	831.7	1156	749.8
	(825.8)	(265.9)	(2777)	(769.6)	(644.4)	(1239)
Fertilizer (kg/ha)	215.0	289.3	174.6	252.3	333.9	239.2
	(213.6)	(191.9)	(724.0)	(224.6)	(210.0)	(333.4)
Mechanization (Rs/ha)	1362	2891	805.6	878.7	2927	1716
	(2781)	(4887)	(1289)	(1755)	(3778)	(3369)
Pesticide (Rs/ha)	61.70	124.5	57.20	59.80	2066	235.0
	(266.7)	(566.3)	(345.9)	(319.0)	(2748)	(1012)
Parcel area (ha)	0.355	0.597	0.674	0.317	0.904	0.494
	(0.501)	(0.833)	(0.733)	(0.334)	(0.662)	(0.652)
Rainfall (mm)	539.1	395.0	286.5	456.3	426.3	468.4
	(253.3)	(166.6)	(83.1)	(192.2)	(178.1)	(234.8)
Number of observations	5572	2720	1625	1073	952	11,942
Number of parcels	2678	1636	1151	549	600	5100
Number of households	612	491	383	268	254	1079
Number of seasons	106	81	69	86	51	240
Number of villages	21	20	20	20	12	30
Number of time periods	10	5	10	10	5	10

Note: Means of data by crop with standard deviations in parenthesis are displayed. All monetary values are in real 2010 Indian Rupees (Rs).

others, while some parcels may be of better quality, resulting in less need for, say, fertilizer. Some households may live in villages in states with favorable agricultural policy, resulting in better access to inputs. Finally, while we do not expect much technological change between 2009 and 2013, the quality of inputs continues to evolve, which may impact yields over time. A third drawback, and perhaps the most important for us, is that OLS requires the variance terms to be normally distributed. While in many applications this assumption is appropriate, we should not expect weather events to be normally distributed. If time-specific variation in yield has a highly skewed distribution, OLS will produce biased parameter estimates and estimates of risk severity may ignore extreme events.

3.2 | The multilevel model

A multilevel or hierarchical modeling strategy addresses the first two drawbacks associated with the standard linear approach to estimation. First, multilevel models offer a natural way to assess the role of seasonal changes in weather on variation in yields by explicitly modeling the variance, not just the mean of the data. This allows us to measure the different sources of variance in yields for particular groups of data but also across groups, which would be difficult in classical linear regressions

(Gelman & Hill, 2007). Second, a multilevel approach allows us to account for each grouping of the data without adding to the computational burden and without violating independence assumptions. While this is possible in classical linear regression by adding indicators for each individual or for each group, it is often infeasible when there are many groups or few observations within a group. Multilevel models offer a natural way to model group-level variation in the uncertainty for individual coefficients (Gelman & Hill, 2007). In our case, a multilevel approach also allows us to disaggregate total variance in yields into its multiple sources, so as to measure the relative contribution of seasonal weather risk in crop yield.

For expository purposes we start with an illustrative example of a simple two-level model in which realizations of yields are grouped within seasons. Let y_n denote the log of an observed yield, n, realized in season s. We estimate

$$y_n = X_n\beta + \alpha_s + \epsilon_n, \tag{2a}$$

$$\alpha_s = \mu + \nu_s, \tag{2b}$$

where X_n and β are as previously defined, and α_s is a season index that is a function of an overall mean, μ, and a random disturbance term, ν_s. Note that unlike the OLS model, where α_s was the fixed or constant effect of season on mean yield, α_s in the MLM is a random variable, a component of the error term, that allows the variance of yields to change from season to season. We assume that $\epsilon_n \sim \mathcal{N}(0, \sigma_\epsilon^2)$, $\nu_s \sim \mathcal{N}(0, \tau_\nu^2)$, and the ϵ_n are mutually independent, as are the ν_s.

To make explicit that our interest lies not in mean-shifting fixed effects but in the variability of yields, we can rewrite Equations (2a) and (2b) in terms of a probability distribution so that

$$y_n \sim \mathcal{N}(X_n\beta + \mu, u_n), \tag{3}$$

with $u_n \equiv \tau_\nu^2 + \sigma_\epsilon^2$. The above distribution is obtained by substituting (2b) into (2a) and using the independence of ϵ_n and ν_s. Defining the regression equation in this way highlights the very specific dispersion structure of the residual. What the formula for the variance u_n does not say is that, for two data points n and n' which belong to the same seasons s, the variables y_n and $y_{n'}$ are not independent, while they are indeed independent when using two different seasons s and s'. From here we can easily define the intraclass correlation coefficient (ICC),

$$\rho = \frac{\tau_\nu^2}{\tau_\nu^2 + \sigma_\epsilon^2}, \tag{4}$$

which measures the proportion of yield variance that is attributed to uncertainty at the season level.

The value of a multilevel model becomes obvious as we add additional levels. In our analysis each observation on yield, y_n is a unique data point that corresponds to the index n, each data point corresponds to a unique parcel i, each parcel corresponds to a unique household h, each household experiences a unique season s, each season occurs in a unique village v, and all villages exists at unique time t. Thus i, h, s, v, and t can be understood as functions of n. We can then write the multilevel model for yields as

$$y_n = \mu + X_n\beta + \gamma_i + \delta_h + \theta_s + \xi_v + \omega_t + \epsilon_n, \tag{5}$$

with the following error structure:

$$\text{Yields} \;\; (n = 11{,}942) : \epsilon_n \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{6a}$$

$$\text{Parcels}\quad (i = 5{,}100) : \gamma_i \sim \mathcal{N}\left(0, \tau_\gamma^2\right), \tag{6b}$$

$$\text{Households}\quad (h = 1{,}079) : \delta_h \sim \mathcal{N}(0, \tau_\delta^2), \tag{6c}$$

$$\text{Seasons}\quad (s = 240) : \theta_s \sim \mathcal{N}(0, \tau_\theta^2), \tag{6d}$$

$$\text{Villages}\quad (v = 30) : \xi_v \sim \mathcal{N}(0, \tau_\xi^2), \tag{6e}$$

$$\text{Time}\quad (t = 10) : \omega_t \sim \mathcal{N}(0, \tau_\omega^2), \tag{6f}$$

where again X_n is a matrix of input data and β is a vector of regression coefficients associated with the various crops.

The estimating equation models the log of yield as a function of inputs, similar to that in Equation (1), with yields being a function of a specific parcel i and an idiosyncratic error term $\epsilon_n \sim \mathcal{N}(0, \sigma_\epsilon^2)$, where σ_ϵ^2 is a constant variance parameter which we assume does not depend on i, h, s, v, or t.

Each n observation (data point) comes from a parcel group i which we assign a unique index value, γ_i, which is distributed $\gamma_i \sim \mathcal{N}(0, \tau_\gamma^2)$. Though similar to intercepts or indicators in a fixed effects regression context, in the multi-level model γ_i is not a constant but rather a mean-zero normal variable. Said another way, the parcel-level index is an error term (random variable) which changes from item to item within the parcel level. This allows us to specify a relationship between inputs and yield that differs across parcels depending on parcel-level characteristics. While some parcel characteristics can be observed, many are difficult to measure or costly to observe. Such characteristics include soil micro-nutrients, grade, and aeration or composition. By incorporating a unique index value for each parcel we can partition out the portion of the error term that is due to differences in these parcel characteristics.

Parcels are grouped within households, with each household, h, being given a unique index value, δ_h, which is distributed $\delta_h \sim \mathcal{N}(0, \tau_\delta^2)$. The household-level index allows variation in parcel-level yield to vary by household characteristics. In most applications, the analyst attempts to control for unobserved household ability through proxy variables such as age or education. The multilevel approach allows us to partition out the portion of the error term that is due to differences in any unobserved household-level characteristics by assigning each household a unique index value without the need to rely on proxies. Modeling variance in this way across all data points n corresponding to a given parcel group allows for this by imposing a covariance structure which is consistent with variation at the parcel group level.

Households are grouped within seasons, which vary across villages as well as within a village across time. Here each season, s, has a unique index value, θ_s, which is distributed $\theta_s \sim \mathcal{N}(0, \tau_\theta^2)$. As with the indices at every other level, θ_s is not a constant but a mean-zero normal variable. Each unique season-level index value allows for season-on-season variation of average yields within each village. By making seasonality a function of both village and time, we are able to account for both the spatial and temporal nature of weather. The unique season index value allows us to partition out the portion of the error term that is due to differences in weather that occurs in a given village at a given time and is not attributable to variation in yields coming from other levels of the model.

The final two sources of variation in yields come from the time-invariant village-level and the spatial-invariant time-level. Here the set of village-level index values, ξ_v, are distributed $\xi_v \sim \mathcal{N}(0, \tau_\xi^2)$. These index values allow us to partition out the portion of the error term that is due to cross village differences in culture, politics, society, or policy, which vary across space but not time. The set of time-level index values, ω_t, are distributed $\omega_t \sim \mathcal{N}(0, \tau_\omega^2)$ and allow for variation across time. While our data cover only 5 years, we cannot rule out national changes in policy, technology, or both during the period that results in variation in yields. The time-level index allows us to determine what portion of the error term is due to differences across time but not space.

Having partitioned out all other potentially relevant sources of yield variability, we interpret seasonal variation as coming solely from weather events. The main goal in uncertainty quantification is to estimate the disturbance terms. Indeed, this is the key to evaluating the share of variance in yield corresponding to each level in the hierarchy.

Given the various sources of variability in the MLM, there are different ICCs that one could define for the model. We follow Lambert et al. (2004) in defining the ICC as the percentage of the total variance that is explained by the variance within groups or clusters of groups. Analagous to the two-level model, in our MLM the total variance is $u_n = \tau_\gamma^2 + \tau_\delta^2 + \tau_\theta^2 + \tau_\xi^2 + \tau_\omega^2 + \sigma_\epsilon^2$. So, the correlation between realizations of yield within the same parcel is

$$\rho(\text{intra-parcel}) = \frac{\tau_\gamma^2}{u_n}, \tag{7}$$

which is a measure of the similarity between observations on parcels across all parcels. The correlation between realizations of yield on parcels within the same household is

$$\rho(\text{intra-household}) = \frac{\tau_\gamma^2 + \tau_\delta^2}{u_n}, \tag{8}$$

which is a measure of the similarity between observations in a households across all households. The correlation between realizations of yield within the same season is

$$\rho(\text{intra-season}) = \frac{\tau_\gamma^2 + \tau_\delta^2 + \tau_\theta^2}{u_n}, \tag{9}$$

which is a measure of the similarity between observations in a season across all seasons. The correlation between realizations of yield within the same village is

$$\rho(\text{intra-village}) = \frac{\tau_\gamma^2 + \tau_\delta^2 + \tau_\theta^2 + \tau_\xi^2}{u_n}, \tag{10}$$

which is a measure of the similarity between observations in a village across all villages in India. And, the correlation between realizations of yield during the same time period is

$$\rho(\text{intra-temporal}) = \frac{\tau_\gamma^2 + \tau_\delta^2 + \tau_\theta^2 + \tau_\xi^2 + \tau_\omega^2}{u_n}, \tag{11}$$

which is a measure of the similarity between observations at a given time across all time in India. By construction the ICC increases as we move to higher levels of aggregation. Thus, we also calculate each level's contribution to total variance in the model. This is simply the variance at each level divided by u_n, the total variance in yields.

3.3 | Bayesian estimation of the multilevel model

While multilevel models address the first two drawbacks of linear OLS models, they still rely on the standard, though potentially unsupported, assumption of normality of the disturbance term at each level.

We can verify if the normality assumption placed on each of the random disturbance terms is reasonable using a likelihood ratio (LR) test. We visualize the results of these tests using zeta profile plots, which plot the sensitivity of the model fit to changes in values of particular parameters. While these plots are not equivalent to drawing the underlying distributions of the estimators, they represent a similar idea and can be interpreted as representing the underlying distributions. First, we estimate the model in Equation (5) using maximum likelihood. Then we hold a single parameter fixed and vary the other parameters, assessing the fit of each new iteration compared to the globally optimal fit using the LR as our comparison statistic. We then apply a signed square root transformation to the LR statistic and plot the absolute value of the resulting function, $|\zeta|$, in comparison to the estimated parameter values. The zeta profile plots resulting from the model fit to Equation (5) is presented in Figure 1. Parameters with underlying normal distributions have straight line zeta profile plots. When this is the case these parameters provide a good approximation to the normal distribution and standard confidence intervals can be used for inference. When zeta profile plot lines are not straight, the normal distribution is a poor approximation of the underlying distribution. For parameter estimates on data, a nonnormal distribution simply requires an adjustment of the relevant test statistics. But, if the underlying distributions of our disturbance parameters are non-normal, we are faced with a violation of the normality assumptions of the error term. In Figure 1, zeta profile plots for the data, the parcel variance term $\left(\tau_{\gamma}^{2}\right)$, and the idiosyncratic variance term (σ_{ϵ}^{2}) come from distributions that are good approximations of normality. The zeta profile plot for the weather variance term (τ_{θ}^{2}) and the village variance term (τ_{ξ}^{2}) are slightly skewed. The zeta profile plot for the household variance term (τ_{δ}^{2}) and the time variance term (τ_{ω}^{2}) are highly skewed, meaning they cannot be assumed to come from a normal distribution. Thus, we cannot assume normality of the individual error terms nor the composite error term. Since we do not, a priori, know the distribution of ϵ_n, maximum likelihood methods cannot be used to reliably estimate the model.

We address this drawback to classical estimation of multilevel models by adopting a Bayesian framework and using Markov Chain Monte Carlo methods to obtain posterior estimates. For estimation of posterior distributions we use the Gibbs sampler, which iteratively constructs a sequence of samples from the univariate random values of each response variable, and of each model parameter, conditional on all other parameters and variables. This method allows us to compute all features of the marginal and joint distributions because the marginal samples are

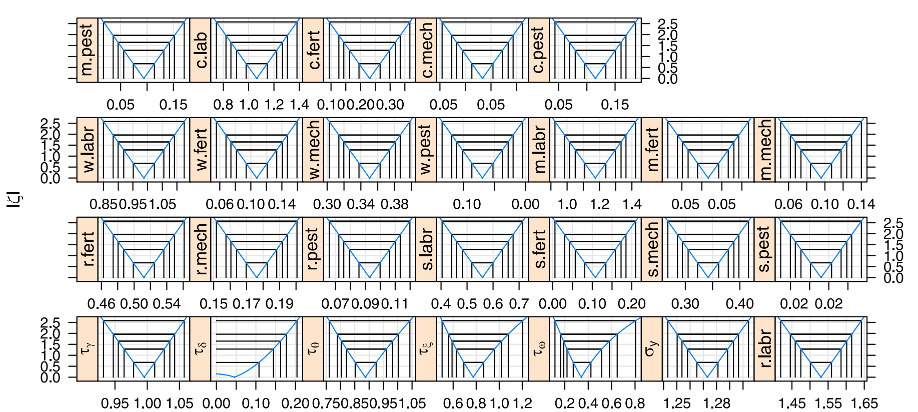

FIGURE 1 Zeta profile plots for multilevel model *Note*: The figure presents the zeta profile plots, representing the underlying distribution of the parameters resulting from the model fit to Equation (5). Plots are calculated by first estimating the model. Then, holding a single parameter fixed, the other parameters are varied to assess the fit of each new iteration compared to the globally optimal fit. The comparison statistic is the likelihood ratio test. Finally, a signed square root transformation is applied to the LR statistic. The plots are the absolute value of the resulting function, $|\zeta|$, in comparison to the estimated parameter values. Parameters with underlying normal distributions have straight line zeta profile plots. The vertical lines delimit 50%, 80%, 90%, 95%, and 99% confidence intervals. For zeta profile plot lines that are not straight, the normal distribution is a poor approximation of the underlying distribution

iteratively fed back into the conditional posterior densities of all other parameters and variables for each sampling. This allows us to calculate unbiased point estimates and confidence intervals for all variables without recourse to normality assumptions.

We define the Bayesian estimator using our illustrative two-level model and then provide parameter definitions for the specific model defined by Equation (5) and the variance structure in Equations (6a)–(6f).[4] For a general two level model, let y_{km} be a unique observation $k = 1, \ldots, K$ from the group $m = 1, \ldots, M$. Within each group m, the data are distributed according to a particular distribution G with parameter γ such that $y_{km} \sim G(\gamma_m)$. We assume that parameter γ_m comes from a distribution L with parameter λ such that $\gamma_m \sim L(\lambda)$. Finally, we assign a distribution to the hyperparameter λ so that $\lambda \sim Q(a, b)$, where a is the mean and b is the variance of the distribution. The joint posterior distribution of all unknown parameters is derived using Bayes' theorem:

$$p(\lambda, \gamma | y) \propto p(y|\gamma, \lambda) p(\gamma|\lambda) p(\lambda), \tag{12}$$

where $y = (y_{11}, \ldots, y_{1K}, \ldots, y_{M1}, \ldots, y_{MK})$ is all the data, $\gamma = (\gamma_1, \ldots, \gamma_M)$ is the group level parameter, and λ is the population parameter. The density of the data is obtained through

$$p(y|\gamma, \lambda) = \prod_{m=1}^{M} \prod_{k=1}^{K} p(y_{km}|\gamma_m, \lambda). \tag{13}$$

This is an independence assumption, and the individual density, $p(y_{km}|\gamma_m, \lambda)$, is assumed to be known. By the independence assumption, the prior for the group level effect is assumed to be

$$p(\gamma|\lambda) = \prod_{m=1}^{M} p(\gamma_m|\lambda). \tag{14}$$

From the joint posterior distribution we can derive the conditional posterior distributions.[5]

The multilevel model estimated by Bayesian techniques is exactly the same model as estimated via MLE and specified in Equations (5) and (6a)–(6f). But Bayesian estimation requires that we define a set of hyperpriors, or the prior distributions of the hyperparameters:

$$\begin{aligned} \text{Hyperpriors:} \quad \mu &\sim \mathcal{N}(a, b) \\ \sigma_\epsilon^2 &\sim IG(r_\epsilon, q_\epsilon) \\ \tau_\gamma^2 &\sim IG(r_\gamma, q_\gamma) \\ \tau_\delta^2 &\sim IG(r_\delta, q_\delta) \\ \tau_\theta^2 &\sim IG(r_\theta, q_\theta) \\ \tau_\xi^2 &\sim IG(r_\xi, q_\xi) \\ \tau_\omega^2 &\sim IG(r_\omega, q_\omega), \end{aligned} \tag{15}$$

where IG is the inverse gamma distribution. The assumption that the priors are distributed inverse gamma is a common assumption in Bayesian econometrics to facilitate the mathematical analysis

[4]Cameron and Trivedi (2005) and Hanmaker and Klugkist (2011) outline Bayesian estimation of multilevel models.

[5]In terms of the independence assumptions made for the MLE multilevel model, assumption (13) corresponds exactly to assuming that the error terms (ϵ_n) are independent of each other while assumption (14) corresponds exactly to assuming that the elements in each group-level are independent of each other. The Bayesian multilevel setup in which each group-level parameter is assumed to have its own prior distribution is consistent with assuming in the MLE multilevel model that ϵ_n and the group indices are independent.

via the Gibbs sampler. We select values for the distributions (0, 10) so that priors are very weakly informative. We then use a burn-in period of 5,000 iterations and an additional 5,000 iterations to ensure convergence of the iterative simulations and sufficient mixing of the Gibbs sampler.[6]

4 | ECONOMETRIC RESULTS

We present the results from a large complement of regressions in Tables 2–4. Table 2 presents a trio of results from fitting the yield function using OLS, MLE, and Bayes. Table 3 presents estimated variances, ICCs, and variance shares from the MLE and Bayesian versions of the yield function. Table 4 presents both MLE and Bayesian estimation results from crop-by-crop yield functions.[7] All models are estimated using the log of yield as the dependent variable and log values of inputs as independent variables.[8]

4.1 | Yield function results

The three yield functions presented in Table 2 rely on the same sample and contain the same set of inputs. To account for heterogeneous input response across crops, we allow all slope and intercept estimates to vary by crop. Model 1 is the classical linear model estimated via OLS as represented by Equation (1). This regression model contains fixed effects for seasons but does not account for the multilevel structure of the data. Models 2 and 3 are the MLE and Bayesian regressions and explicitly account for the indexing of observations at the parcel, household, season, village, and temporal levels. The point estimates from the Bayesian regression are based on posterior density estimates derived from iterating from our weakly informative priors. Results from these regressions point to a fairly robust set of basic patterns that are repeated with few exceptions. These include (i) positive and significant production relationships between yields and measured inputs (labor, fertilizer, mechanization and pesticides), and (ii) diminishing returns to inputs other than labor. We observe only one instance in which the point estimate for an input is negative and significant (for pesticides in the case of wheat). Returns to scale appear to be increasing for all crops. The Bayesian estimation of the yield function generates point estimates that are broadly similar in sign, magnitude, and significance to those of the OLS and MLE multilevel regressions.

Table 3 reports estimated variance parameters (Panel A), ICCs (Panel B) and variance shares (Panel C) for the multilevel MLE and Bayes regressions. These statistics establish the key findings that inform our insights into the potential role of weather index insurance.[9] We focus attention on Panel C of Table 3, which compactly summarizes the data expressed in the upper panels of the table. Reading down the rows of the table allows us to assess the decomposition of variance and, by extension, the relative importance of each level in explaining overall variance in yields. We find for the MLE specification (Model 2 reported in Table 2) that 23% of the total variance in yields comes from between-parcel differences, 19% is attributed to between-season differences, 15% comes from between-village differences, 3% is attributed to differences across time, and 40% of the total residual is idiosyncratic noise. In the Bayesian specification (Model 3 reported in Table 2) 24% of total variance in yields comes from between-parcel differences, 20% is attributed to between-season differences, 15% comes from between-village differences, 6% is attributed to differences across time, and 35% is idiosyncratic noise.

[6]Diagnostic trace plots of convergence for each of the variance terms are presented in Supporting Information Appendix Figure A1.

[7]In the interests of parsimony we focus on the residual estimates for the crop specific models. The point estimates on inputs are presented in Supporting Information Appendix Table A1.

[8]Given the prevalence of zero values in the input data, and to a lesser extent in the output data, we use the inverse hyperbolic sine transformation to convert levels to logarithmic values.

[9]An expanded set of results from models estimated using different hierarchical structures is provided in Supporting Information Appendix Tables B1 and B2.

TABLE 2 Results of yield function regressions

	(1)	(2)	(3)
Rice			
Log labor	1.635***	1.531***	1.127***
	(0.043)	(0.043)	(0.032)
Log fertilizer	0.598***	0.512***	0.542***
	(0.019)	(0.019)	(0.018)
Log mechanization	0.206***	0.176***	0.147***
	(0.010)	(0.009)	(0.009)
Log pesticides	0.123***	0.093***	0.088***
	(0.011)	(0.010)	(0.010)
Sorghum			
Log labor	0.585***	0.559***	0.987***
	(0.063)	(0.060)	(0.042)
Log fertilizer	0.087**	0.108***	0.214***
	(0.042)	(0.040)	(0.038)
Log mechanization	0.321***	0.342***	0.491***
	(0.030)	(0.029)	(0.026)
Log pesticides	0.009	0.012	0.009
	(0.018)	(0.016)	(0.016)
Wheat			
Log labor	0.993***	0.988***	1.217***
	(0.054)	(0.053)	(0.036)
Log fertilizer	0.110***	0.108***	0.073***
	(0.021)	(0.020)	(0.036)
Log mechanization	0.386***	0.350***	0.366***
	(0.019)	(0.018)	(0.036)
Log pesticides	−0.059**	−0.069***	−0.100***
	(0.026)	(0.025)	(0.025)
Maize			
Log labor	1.126***	1.174***	1.584***
	(0.100)	(0.097)	(0.043)
Log fertilizer	0.048	0.016	−0.005
	(0.041)	(0.039)	(0.040)
Log mechanization	0.111***	0.096***	0.109***
	(0.018)	(0.017)	(0.017)

TABLE 2 (Continued)

	(1)	(2)	(3)
Log pesticides	0.086***	0.094***	0.098***
	(0.033)	(0.030)	(0.029)
Cotton			
Log labor	1.047***	1.064***	1.291***
	(0.129)	(0.124)	(0.057)
Log fertilizer	0.220***	0.230***	0.230***
	(0.057)	(0.053)	(0.053)
Log mechanization	0.072**	0.033	0.036
	(0.031)	(0.031)	(0.029)
Log pesticides	0.128***	0.115***	0.095***
	(0.029)	(0.028)	(0.027)
Observations	11,942	11,942	11,942
R^2	0.964		
Log Likelihood		−22,268	
Akaike Inf. Crit.		44,598	
Deviance Inf. Crit.		44,828	42,578

Note: Dependent variable is log of yield. All specifications include statistically significant crop-specific intercept terms as controls. Column (1) is a classical OLS regression that includes season fixed effects but does not account for the multilevel structure of the data. Column (2) is the maximum likelihood estimate of the multilevel model and contains covariates, and data indexed by parcel, household, season, village, and time. Column (3) is the Bayesian estimation of the multilevel model. Bayesian calculations use a burn-in period of 5000 iterations and an additional 5000 iterations to ensure convergence. Standard errors are reported in parentheses (** $p < 0.05$; *** $p < 0.01$).

An intuitive interpretation of these results is that much of the differences observed in yields reflects differences between parcels, such as soil quality. Household or farmer capability, relative to other sources of variability, is unimportant in explaining differences in yields. In other words, good farmers cannot make up for bad soil but bad farmers can still prosper if they have good soil. Similar to Townsend (1994) and Rosenzweig and Binswanger (1993) we find that idiosyncratic sources of risk play a much larger role in determining observed yields than covariate sources. Unlike Townsend (1994) and Rosenzweig and Binswanger (1993), we are able to quantify these differences. Considering all sources of variance in yields, only 37%–41% comes from what could be considered covariate sources while the remaining 59%–63% comes from what could be considered idiosyncratic sources. Thinking in terms of insurable weather risk, 19%–20% of the variability in crop yield is due to seasonal weather variation. While this share is by no means tiny, it is a minority source of variation in yields and much lower than the amount typically cited in the literature to justify the need for weather index insurance. These basic patterns highlight the relatively small importance of between-season yield variance compared with other sources of yield variance.[10]

[10]Note that the share of variance that remains unexplained is not the appropriate measure of model fit. The Akaike Information Criterion (AIC) and Deviance Information Criterion (DIC) reported in Table 2 provides a measure of model fit. In both the AIC and DIC, lower values correspond to better model fit. The DIC is a generalization of the more common Bayesian Information Criterion (BIC) and is used here to facilitate comparison between the ML and Bayesian estimations of the same models.

TABLE 3 Estimated variance, ICCs, and variance shares from multilevel regressions

	MLE		*Bayes*
Panel A: Variance parameter estimates			
Parcel $\left(\tau_{\gamma}^{2}\right)$	0.933		1.098
Household (τ_{δ}^{2})	0.002		0.004
Season (τ_{θ}^{2})	0.790		0.903
Village (τ_{ξ}^{2})	0.623		0.682
Time (τ_{ω}^{2})	0.126		0.261
Idiosyncratic (σ_{ϵ}^{2})	1.623		1.613
Panel B: Intraclass correlation coefficients			
Parcel	0.228		0.241
Household	0.228		0.242
Season	0.421		0.440
Village	0.573		0.589
Time	0.604		0.646
Panel C: Shares of variance from each level			
Parcel	23%		24%
Household	00%		00%
Season	19%		20%
Village	15%		15%
Time	03%		06%
Idiosyncratic	40%		35%
Observations		11,942	

Note: In Panel A, estimates of the variance parameters on each level's residuals come from estimation of models reported in Columns (2) and (3) in Table 2. Variances τ_{γ}^{2}, τ_{δ}^{2}, τ_{θ}^{2}, τ_{ξ}^{2}, τ_{ω}^{2} represent the variance in crop yield that comes from the corresponding level. The final variance parameter (σ_{ϵ}^{2}) corresponds to the idiosyncratic or unexplained portion of the model. Intraclass correlation coefficients (ICC) in Panel B are calculated using the formulas in Equations (7)–(11). Panel C decomposes the ICC into percent of variance accorded to each level. For comparison, Supporting Information Appendix Tables B1 and B2 reports alternative specifications of the models presented here.

Abbreviation: ICC, intraclass correlation coefficient.

Strictly speaking, the non-normality of the error term means that the OLS and MLE regressions are misspecified. Accordingly, we would expect estimates from these models to be biased. However, in our application, only the disturbance terms associated with households and time are highly skewed and non-normal (recall Figure 1). Furthermore, these two sources of variance make up a very small share of total variance. Thus, any bias that is introduced through model misspecification is mitigated by the small impact of the skewed noise term on the overall estimation. We highlight that this is an artefact of the current application and data. If non-normal variation accounted for more of the total variance we would expect to find greater bias in our OLS and MLE estimates.

Figure 2 reports distribution profiles based on the Bayesian regression results for the associated sources of variance measured by Model 3 in Table 2. These histograms are constructed from the posterior distributions estimated using the Gibbs sampler. We can draw several conclusions from this graph. First, several of the variance terms are not normally distributed, which violates the

TABLE 4 Estimated variance, ICCs, and variance shares by crop

	Rice		Sorghum		Wheat		Maize		Cotton	
	MLE	*Bayes*	*MLE*	*Bayes*	*MLE*	*Bayes*	*MLE*	*Bayes*	*MLE*	*Bayes*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Variance parameter estimates										
parcel $\left(\tau_{\gamma}^{2}\right)$	1.858	1.852	0.000	0.002	1.092	1.059	0.564	0.557	0.000	0.018
Household (τ_{δ}^{2})	0.000	0.003	0.048	0.046	0.046	0.034	0.000	0.007	0.013	0.010
Season (τ_{θ}^{2})	0.832	0.865	0.166	0.174	2.323	2.573	0.234	0.266	0.367	0.401
Village (τ_{ξ}^{2})	1.098	1.266	0.133	0.156	5.527	6.411	0.185	0.177	0.000	0.041
Time (τ_{ω}^{2})	0.203	0.298	0.000	0.011	0.612	0.694	0.000	0.013	0.521	0.689
Idiosyncratic (σ_{ϵ}^{2})	1.918	1.918	0.686	0.684	1.598	1.626	0.517	0.518	0.785	0.766
Panel B: Intraclass correlation coefficients										
Parcel	0.314	0.299	0.000	0.002	0.098	0.085	0.376	0.362	0.000	0.009
Household	0.314	0.299	0.046	0.045	0.102	0.088	0.376	0.366	0.008	0.015
Season	0.455	0.439	0.207	0.207	0.309	0.296	0.532	0.539	0.225	0.223
Village	0.641	0.643	0.336	0.352	0.803	0.813	0.655	0.654	0.225	0.244
Time	0.675	0.691	0.336	0.363	0.857	0.869	0.655	0.663	0.534	0.602
Panel C: Shares of variance from each level										
Parcel	31%	30%	00%	00%	10%	09%	38%	36%	00%	01%
Household	00%	00%	05%	04%	00%	00%	00%	00%	01%	01%
Season	14%	14%	16%	16%	21%	21%	16%	17%	22%	21%
Village	19%	20%	13%	15%	49%	52%	12%	12%	00%	02%
Time	03%	05%	00%	01%	05%	06%	00%	01%	31%	36%
Idiosyncratic	32%	31%	66%	64%	14%	13%	34%	34%	47%	40%
Observations	5572		2720		1625		1073		952	

Note: In Panel A, estimates of the variance parameters on each level's residuals come from estimation of models in corresponding columns in Supporting Information Appendix Table A1. Variances τ_{γ}^{2}, τ_{δ}^{2}, τ_{θ}^{2}, τ_{ξ}^{2}, τ_{ω}^{2} represent the variance in crop yield that comes from the corresponding level. The final variance (σ_{ϵ}^{2}) parameter corresponds to the idiosyncratic or unexplained portion of the model. Intraclass correlation coefficients (ICC) in Panel B are calculated using the formulas in Equations (7)–(11). Panel C decomposes the ICC into percent of variance accorded to each level. Bayesian calculations use a burn-in period of 5000 iterations and an additional 5000 iterations to ensure convergence.

Abbreviation: ICC, intraclass correlation coefficient.

normality assumption in OLS and non-Bayesian multilevel regression. Second, normal confidence interval calculation for the parameters will be inaccurate. Third, the long right-hand side tails in the distributions for households and time, and to a lesser extent season, indicates that these levels by themselves are insufficient to explain the observed yield variability in a classical setting. Said differently, our model relies on indices to determine which portion of the error term varies due to differences in household ability, occurrences of weather, and technical change but does not take into account their severity. The skewed output of our Bayesian analysis hints that while most households are similar in ability there are a few outlying farmers with exceptional ability. Similarly, while weather events account for only 20% of variance in yields, severe weather events remain important in considering tools to mitigate risk.

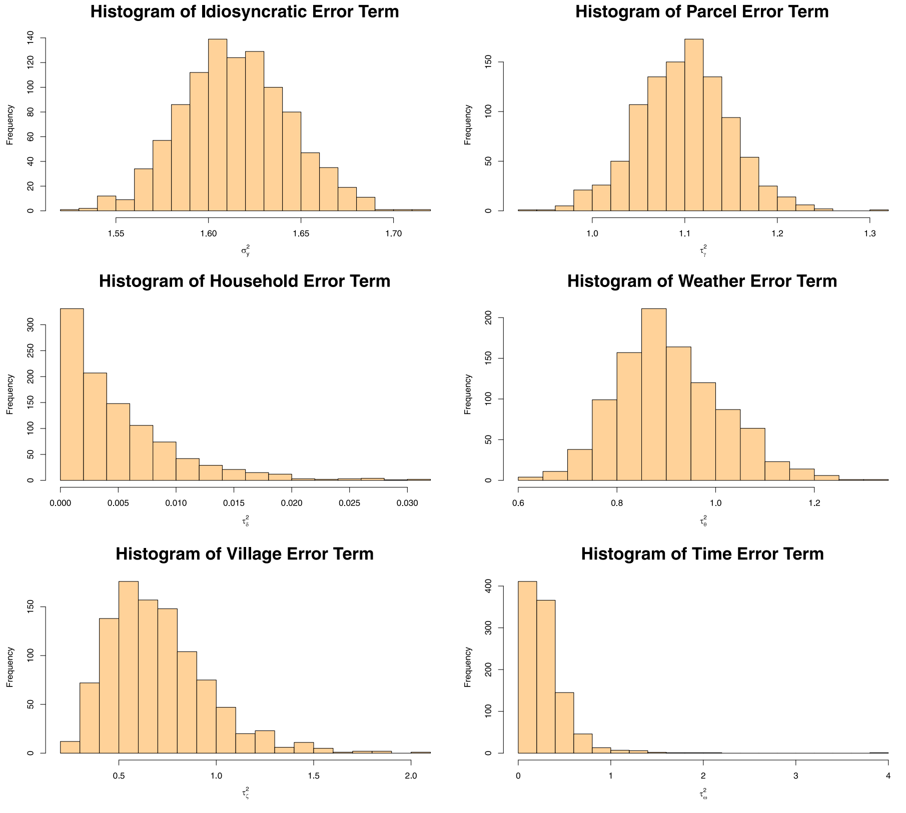

FIGURE 2 Histograms of level variance from Bayesian estimation. Histograms are drawn from posterior distributions of variance terms estimated from Column (3) reported in Table 2

4.2 | Variance results by crop type

Up to this point, our analysis has accommodated heterogeneity in the input-response curves of the crops under consideration by allowing slopes and intercepts to vary across crops. By pooling these data, however, the variance results reported in Table 3 are implicitly derived under the assumption that the same variance structure of the hierarchical model applies to all crops under consideration and that, by extension, the same weather risk profile applies to each crop. We now relax that assumption and re-estimate the MLE and Bayesian regressions, in each case using five separate, crop-specific sub-samples of the data. Our goal is to allow for a more comprehensive assessment of the variance structure of yields, to gain additional insights into the nature of weather risk exposure among farmers in the sample. The new samples vary in size and coverage, which also permits us to limit the influence of rice in our results, since rice dominates the pooled sample by contributing roughly half of all observations. In three cases (rice, wheat, and maize) we observe data across all 10 seasons. In the case of sorghum and cotton we observe data for only five seasons.

To complement the results reported in Table 3, we present in Table 4 a full set of estimated variance parameters, ICCs, and variance shares for the crop-specific versions of our regressions.[11] As in the case of the regressions with the pooled samples, the MLE and Bayesian regressions explicitly account for the indexing of observations at the parcel, household, season, village, and

[11]We relegate the discussion of the point estimates for parcel-level inputs to Supporting Information Appendix A and Table A1.

temporal levels. The point estimates from the Bayesian regression are based on posterior density estimates derived from iterating from our weakly informative priors.

As before, we focus attention on Panel C of Table 4. Once again we see a close correspondence between the MLE and Bayesian results. Reading across the table, however, we observe substantial variation in the sources and relative importance of each level in explaining overall variance in crop-specific yields. For example, over 30% of variance comes from between-parcel differences for rice and maize, compared with less than 2% for sorghum and cotton. Of particular importance for index insurance, we observe only moderate variation in the contribution of between-season weather differences, which range from 14% of variance in the case of rice, to as much as 22% in the case of wheat and cotton. We also find more notable differences between the estimated variances from MLE and Bayesian regressions, with implications for the role of non-normality and the inability of the MLE approach to properly account for skewness in several of the disturbance terms. As noted above, to the extent that nonnormal errors account for more of the total variance one would expect to find greater bias in the OLS and MLE point estimates vis-à-vis the Bayes estimates. This is evident when level variance terms tend to be small, resulting in MLE estimates that collapse to zero while Bayesian estimates are small but non-zero.

5 | IMPLICATIONS FOR WEATHER INDEX INSURANCE

The results from our estimation of the sources of variability in yields has implications for both the demand for and the supply of weather index insurance, particularly in village India. Much of the previous literature on index insurance has focused on consumer-centric barriers to demand, such as financial literacy or household liquidity (Cole et al., 2013). Results from both the MLE and Bayesian estimation of the multilevel model show that, relative to previous statements in the economics literature about the role of weather in yield variability, weather is a minority source of variation in yield. At least in our data, weather accounts for 14%–22% of variance in yield, a substantial amount but far short of the 50%–90% figures that are frequently cited to justify the need for farmers to purchase weather index insurance.

The implication for the demand for index insurance is clear. With so many other sources of variation in yields, farmers may rationally choose to expend limited resources to help improve yields and ensure against the variety of risks they face. This could include purchases of improved inputs to try and reduce the variability in yields due to parcel-level differences. It could also include purchases of stress tolerant seed varieties, which have been shown to improve yields, reduce risk to abiotic stress, and crowd in the purchase of other inputs (Emerick et al., 2016). That weather makes up a minority source of variability in yields means that weather index insurance carries a substantial amount of basis risk, at least in terms of insuring the outcome of ultimate interest to farmers—profit. While indemnity insurance contracts can be written to cover losses from a variety of sources, index insurance contracts payout contingent on a change in the index. With a majority of variation in yield coming from non-weather sources, that means an insurance contract based on weather can only ever insure against a minority of potential risks. While from an insurer's point of view, these other sources of risk are not meant to be covered by the index insurance product and are therefore not a relevant part of the "basis risk" of the contract, from a farmer's perspective any loss not covered by the insurance contract is basis risk, that is, uninsured exposure (Binswanger-Mkhize, 2012).

The implication for the supply of index insurance is not that insurers should stop writing index insurance contracts but rather that they need to develop better insurance products. While weather variability accounts for a much smaller share of yield variability than has been previously argued, it still makes up a substantial source of total variation. The likely importance of extreme and infrequent weather events in yield variability, evidenced by the skewness in the posterior distribution of τ_θ^2, highlights the need for farmers to have access to affordable risk-management tools. In other words, despite the relatively low share of seasonal weather variation in yield variability, our data still suggest that potential insurance purchasers need to consider weather risk, since extreme weather events, while rare, are nevertheless potentially costly over the long run. This motivates us to ask whether tools for managing this type of risk are adequately accessible. As we show in Supporting Information Appendix C, current index insurance in the market in India from the major insurers is overpriced, with extremely high loading factors, when compared to index insurance in developed or other developing countries.

We conclude that these compounding factors—weather being less important than previously assumed and current contracts being overpriced—have led to the observed low uptake of weather index insurance. The practical outcome is that farmers operate in an environment where they experience risk from a variety of sources but have few good options for insuring against that risk. Evidence shows that farmers adapt their crop portfolios and inputs based on the need to manage risk across seasons and within seasons (Jagnani et al., 2021; Mobarak & Rosenzweig, 2013), meaning farmers likely operate below their production possibility frontier. Weather index insurance, as currently marketed, is unlikely to be an attractive risk management tool on its own, given the various sources of variability in yields. But, improved products may be able to play an important role in a risk management portfolio that include tools such as stress tolerant seed varieties (Emerick et al., 2016), improved forecasting (Rosenzweig & Udry, 2013), microfinance (Assah & Sberro-Kessler, 2017), and a variety of digital technologies (Benami & Carter, 2021).

6 | CONCLUSION

Despite long standing evidence that rural households are unable to fully insure covariate risk, few studies have attempted to measure just how large a role covariate events, such as weather, play in agricultural yields. We address this information gap using agricultural production data covering 11,942 parcel level observations from India. Using a multilevel/hierarchical regression framework, we estimate the different sources of yield variance. This approach controls for inputs at the parcel-level and also isolates the amount of yield variability attributed to parcel-level effects, household-level effects, seasonal weather effects, village-level effects, and time effects. Adopting Bayesian estimation techniques allows us to account for the highly skewed distribution of several of the disturbance terms.

Overall, we find that variability in weather makes up a small but not insubstantial share of the total variance in crop yields. This suggests that basis risk on the production side (i.e., low correlation between the weather index and yield loss) may be large. The majority of variation in yields does not come from seasonal variability, which is our proxy for rainfall or weather variability. Rather, the majority of variation in yields comes from differences in parcels and from the random disturbance term which captures idiosyncratic events which would be, by definition, not covered by index insurance. Given the long history of evidence that farmers are effective in using self-insurance and mitigation measures to protect against idiosyncratic risk, and given the small role covariate seasonal risk plays in crop variability, farmers may rationally prefer to forgo weather index insurance to focus their risk management choices on minimizing other sources of risk.

Although there are many potential impediments to insurance uptake, this study also provides evidence for an obvious but until now unsupported explanation, at least in the case of India: existing insurance contracts are overpriced from the perspective of the farm household. Product design and ratemaking of index insurance in India has generally taken a long view, utilizing rainfall data extending as far back as four decades to calculate actuarial rates. From the perspective of the insurance company, such a long time horizon may support the company's longevity in managing it's own financial risk. However, it may not easily align with the planning horizons of rural households.

ACKNOWLEDGMENTS

We greatly appreciate the helpful comments and suggestions from Philip Abbott, Hans Binswanger-Mkhize, Michael Roberts, and Neda Trifkovic along with those from seminar participants at the Econometric Society World Congress in Montreal, the Midwest International Economics Development Conference in Madison, the Agricultural and Applied Economics Association Annual Meeting in San Francisco, and the International Conference of Agricultural Economists in Milan. This paper and our thinking have been shaped by conversations with Steve Boucher, Garth Holloway, Mario Miranda, and Robert Townsend. We are solely responsible for any errors or misunderstandings.

OPEN RESEARCH BADGES

This article has earned Open Data and Open Materials badges. Data and materials are available at 10.5281/zenodo.6476466.

DATA AVAILABILITY STATEMENT

The data that supports the findings of this study are available at the link in the section above.

REFERENCES

Akter, S., Roy Brouwer, Saria Choudhury, and Salina Aziz. 2009. "Is There a Commercially Viable Market for Crop Insurance in Rural Bangladesh?" *Mitigation and Adaptation Strategies for Global Change* 14(3): 215–29.

Assah, Fatou C., and Rachel Sberro-Kessler. 2017. "Microfinance and Index Insurance: Testing Business Models in the Agricultural Sector." *Global Index Insurance Facility Knowledge Notes*. Washington, D.C.: World Bank.

Banerjee, Abhijit, Esther Duflo, and Richard Hornbeck. 2014. "Bundling Health Insurance and Microfinance in India: There Cannot Be Adverse Selection if There is No Demand." *American Economic Review* 104(5): 291–7.

Barnett, Barry J., and Olivier Mahul. 2007. "Weather Index Insurance for Agriculture and Rural Areas in Lower-Income Countries." *American Journal of Agricultural Economics* 89(5): 1241–7.

Benami, Elinor, and Michael R. Carter. 2021. "Can Digital Technologies Reshape Rural Microfinance? Implications for Savings, Credit, and Insurance." *Applied Economics Perspectives and Policy* 43(3): 1196–220.

Binswanger-Mkhize, Hans P. 2012. "Is There Too Much Hype About Index-Based Agricultural Insurance?" *Journal of Development Studies* 48(2): 187–200.

Bjerge, Benedikte, and Neda Trifkovic. 2018. "Extreme Weather and Demand for Index Insurance in Rural India." *European Review of Agricultural Economics* 45(3): 397–431.

Cai, Hongbin, Yuyu Chen, Hanming Fang, and Li-An Zhou. 2015. "The Effect of Microinsurance on Economic Activities: Evidence from a Randomized Field Experiment." *Review of Economics and Statistics* 97(2): 287–300.

Cameron, A. Colin, and Pravin K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.

Chantarat, Sommarat, Andrew G. Mude, Christopher B. Barrett, and Michael R. Carter. 2013. "Designing Index-Based Livestock Insurance for Managing Asset Risk in Northern Kenya." *Journal of Risk and Insurance* 80(1): 205–37.

Clarke, Daniel J., Olivier Mahul, Kolli N. Rao, and Niraj Verma. 2012a. "Weather Based Crop Insurance in India." Policy Research Working Paper no. 5985. Washington, D.C.: World Bank.

Clarke, Daniel J., Olivier Mahul, and Niraj Verma. 2012b. "Index Based Crop Insurance Product Design and Ratemaking: The Case of Modified NAIS in India." Policy Research Working Paper no. 5986. Washington, D.C.: World Bank.

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert M. Townsend, and James Vickery. 2013. "Barriers to Household Risk Management: Evidence from India." *American Economic Journal: Applied Economics* 5(1): 104–35.

Cole, Shawn, Daniel Stein, and Jeremy Tobacman. 2014. "Dynamics of Demand for Index Insurance: Evidence from a Long-Run Field Experiment." *American Economic Review* 104(5): 284–90.

de Nicola, Francesca. 2015. "Handling the Weather: Insurance, Savings, and Credit in West Africa." Policy Research Working Paper no. 7187. Washington, D.C.: World Bank.

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H. Dar. 2016. "Technological Innovations, Downside Risk, and The Modernization of Agriculture." *American Economic Review* 106(6): 1537–61.

Gaurav, Sarthak, Shawn Cole, and Jeremy Tobacman. 2011. "Marketing Complex Financial Products in Emerging Markets: Evidence from Rainfall Insurance in India." *Journal of Marketing Research* 48(SPL): S15–62.

Gelman, Andrew, and Jennifer Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. New York: Cambridge University Press.

Giné, Xavier, Robert M. Townsend, and James Vickery. 2007. "Statistical Analysis of Rainfall Insurance Payouts in Southern India." *American Journal of Agricultural Economics* 89(5): 1248–54.

Giné, Xavier, Robert M. Townsend, and James Vickery. 2008. "Patterns of Rainfall Insurance Participation in Rural India." *World Bank Economic Review* 22(3): 539–66.

Giné, Xavier, Robert M. Townsend, and James Vickery. 2017. "Forecasting When it Matters: Evidence from Semi-Arid India." Manuscript.

Giné, Xavier, James Vickery, Lev Menand, and Robert M. Townsend. 2012. "Microinsurance: A Case Study of the Indian Rainfall Index Insurance Market." In *Handbook of the Indian Economy*, edited by C. Ghate. New York: Oxford University Press.

Griliches, Zvi. 1963. "Estimates of the Aggregate Agricultural Production Function from Cross-Sectional Data." *Journal of Farm Economics* 45(2): 419–28.

Hanmaker, Ellen L., and Irene Klugkist. 2011. "Bayesian Estimation of Multilevel Models." In *Handbook of Advanced Multilevel Analysis*, edited by Joop Hox and J. Kyle Roberts, 137–61. Hoboken: Taylor and Francis.

Hardaker, J. B., R. B. M. Huirne, and J. R. Anderson. 1997. *Coping With Risk in Agriculture*. Wallingford: CAB International.

Hazell, P., J. Anderson, N. Balzer, A. Hastrup Clemmensen, U. Hess, and E. Rispoli. 2010. *Potential for Scale and Sustainability in Weather Index Insurance for Agriculture and Rural Livelihoods*. Rome: International Fund for Agricultural Development and World Food Program.

Jagnani, Maulik, Christopher B. Barrett, Yanyan Liu, and Liangzhi You. 2021. "Within-Season Producer Response to Warmer Temperatures: Defensive Investments by Kenyan Farmers." *Economic Journal* 131(633): 392–419.

Ker, Alan P., Tor N. Tolhurst, and Yong Liu. 2016. "Bayesian Estimation of Possibly Similar Yield Densities: Implications for Rating Crop Insurance Contracts." *American Journal of Agricultural Economics* 98(2): 360–82.

Kochar, Ajini. 1999. "Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India." *Review of Economics and Statistics* 81(1): 50–61.

Lambert, Lambert C., Lambert R. Burton, Keith R. Abrams, and Adrian M. Brooke. 2004. "The Analysis of Peak Expiratory Flow Data Using a Three-Level Hierarchical Model." *Statistics in Medicine* 23: 3821–39.

Liu, Elaine M. 2013. "Time to Change What to Sow: Risk Preferences and Technology Adoption Decisions of Cotton Farmers in China." *Review of Economics and Statistics* 95(4): 1386–403.

Matsuda, Ayako, and Takashi Kurosaki. 2019. "Demand for Temperature and Rainfall Index Insurance in India." *Agricultural Economics* 50(3), 353–66.

Mobarak, Ahmed Mushfiq, and Mark M. R. Rosenzweig. 2013. "Informal Risk Sharing, Index Insurance, and Risk Taking in Developing Countries." *American Economic Review* 103(3): 375–80.

National Commission on Agriculture. 1976. "Report of the National Commission on Agriculture." Technical report, Ministry of Agriculture, New Delhi.

Parchure, Rajas. 2002. "Varsha Bonds and Options: Capital Market Solutions for Crop Insurance Problems." Manuscript.

Porter, John R., and Mikhail A. Semenov. 2005. "Crop Responses to Climate Variation." *Philosophical Transactions of the Royal Society B: Biological Sciences* 360(1463): 2021–35.

Ramsey, A. Ford. 2020. "Probability Distributions of Crop Yields: A Bayesian Spatial Quantile Regression Approach." *American Journal of Agricultural Economics* 102(1): 220–39.

Ray, Deepak K., J. S. Gerber, Graham K. MacDonald, and Paul C. West. 2015. "Climate Variation Explains a Third of Global Crop Yield Variability." *Nature Communications* 6: 5989.

Rosenzweig, Mark, and Christopher R. Udry. 2013. "Forecasting Profitability." Working Paper no. 19334. Cambridge: NBER.

Rosenzweig, Mark R., and Hans P. Binswanger. 1993. "Wealth, Weather Risk, and the Consumption and Profitability of Agricultural Investments." *Economic Journal* 103(416): 56–78.

Rosenzweig, Mark R., and Kenneth I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy* 101(2): 223–44.

Shirsath, Paresh, Shalika Vyas, Pramod Aggarwal, and Kolli N. Rao. 2019. "Designing Weather Index Insurance of Crops for the Increased Satisfaction of Farmers, Industry and the Government." *Climate Risk Management* 25: 100189.

Skees, Jerry J. R., Jason J. Hartell, and Anne G. Murphy. 2007. "Using Index-Based Risk Transfer Products to Facilitate Micro Lending in Peru and Vietnam." *American Journal of Agricultural Economics* 89(5): 1255–61.

Stein, Daniel. 2018. "Dynamics of Demand for Rainfall Index Insurance: Evidence From a Commercial Product in India." *World Bank Economic Review* 32(3): 692–708.

Tack, Jess, Ardian Harri, and Keith Coble. 2012. "More Than Mean Effects: Modeling the Effect of Climate on the Higher Order Moments of Crop Yields." *American Journal of Agricultural Economics* 94(5): 1037–54.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62(3): 539–91.

VDSA. 2015. "Village Dynamics in South Asia (VDSA) Database." Generated by ICRISAT/IRRI/NCAP in Partnership With National Institutes in India and Bangladesh. https://vdsa.icrisat.ac.in

Woodard, Joshua D., and Bruce J. Sherrick. 2011. "Estimation of Mixture Models using Cross-Validation Optimization: Implications for Crop Yield Distribution Modeling." *American Journal of Agricultural Economics* 93(4): 968–82.

World Bank. 2011. "Weather Index Insurance for Agriculture: Guidance for Development Practitioners." Agriculture and Rural Development Discussion Paper 50. Washingtons, D.C.: World Bank.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Michler, Jeffrey D., Frederi G. Viens, and Gerald E. Shively. 2022. "Risk, Crop Yields, and Weather Index Insurance in Village India." *Journal of the Agricultural and Applied Economics Association* 1: 61–81. https://doi.org/10.1002/jaa2.9