Technology, Risk, and Climate Adaptation: Corn Specialization in the U.S. Corn Belt

Xiaodong Du* Xueying Sun† Yuhan Wang‡

Abstract

We study how technological change facilitates climate adaptation in agricultural land use by examining the evolution of corn specialization in the U.S. Corn Belt. We construct a county-level Technical Change Index (TCI) using agro-ecological yield potential data and combine it with a long-run panel of land use, weather, and crop insurance outcomes. Reduced-form evidence shows that higher TCI increases corn specialization on both the intensive and extensive margins and is associated with lower crop insurance loss ratios, consistent with greater resilience to adverse weather shocks. To identify the mechanisms behind these patterns, we estimate a structural discrete-choice model in which counties choose specialization by trading off expected returns, downside risk, and adjustment costs. The estimates show that technological change raises the returns to specialization, lowers adjustment frictions, and reduces exposure to downside weather risk. These channels help explain both the intensification of specialization within the traditional Corn Belt and the more limited expansion at its margins. Counterfactual simulations show that modern technology mitigates the welfare costs of greater climate volatility, generating adaptation gains of about $0.85 per acre. The results imply that technological progress is an important adaptation margin in agriculture, reshaping the risk-return profile of production and allowing specialized systems to persist under changing climatic conditions.

Keywords: crop insurance, downside risk, land use, technological change.

JEL Classification: Q54, Q55, Q12, C35.

*Agricultural and Applied Economics, University of Wisconsin-Madison. xdu23@wisc.edu.

†Agricultural and Resource Economics, University of Maryland. xueying@umd.edu.

‡Agricultural and Applied Economics, University of Wisconsin-Madison. ywang2558@wisc.edu.

Author names are listed in alphabetical order; all three authors contributed equally.

1 Introduction

Agricultural production is increasingly exposed to climate risk, yet the mechanisms through which producers adapt through land use remain imperfectly understood. A large literature documents the effects of temperature and precipitation on crop yields, but we know less about how technological change shapes producers' ability to respond to shifting climatic conditions. In particular, it remains unclear whether technological progress facilitates reallocation across crops and locations and, in turn, mitigates the aggregate economic costs of climate change.

The U.S. Corn Belt provides a useful setting to study these issues. Its historical emergence was initially driven by declining transportation costs and market integration, as railroad expansion and westward settlement promoted specialization in grain production (Donaldson and Hornbeck, 2016; Solakoglu and Goodwin, 2005). A second transformation followed in the postwar period, with the diffusion of hybrid seeds (Griliches, 1957), chemical fertilizer, and mechanization (Anderson, 2008). More recent advances in biotechnology (Klümper and Qaim, 2014), precision agriculture (Basso et al., 2019; Schimmelpfennig, 2016), and input management have further raised both the level and stability of corn yields.

These technological changes interact with climate to reshape the spatial distribution of production. Because crop yields respond nonlinearly to temperature and moisture, climate change alters both expected returns and the risk profile of land use (Schlenker and Roberts, 2009). At the same time, technological progress expands the feasible production set by improving resilience to adverse weather and enabling cultivation in previously marginal environments (Beddow et al., 2015). Consistent with this view, recent evidence documents a shifting geographic footprint for the Corn Belt, especially along its northern and western margins (Green et al., 2018). Similar long-run adjustments in North American wheat production suggest that geographic reallocation has historically been an important adaptation margin in agriculture (Olmstead and Rhode, 2011).

This paper studies how technological change facilitates climate adaptation in agricultual land use. We focus on county-level corn specialization in the U.S. Corn Belt and pursue two objectives. First, we construct and validate a county-level Technical Change Index (TCI) using agro-ecological yield potential data. Second, we estimate how TCI affects land-use decisions, risk exposure, and welfare under changing climate conditions.

We begin by constructing a new county-level measure of technological capacity based on the gap between high-input and low-input yield potential in the FAO Global Agro-Ecological Zones (GAEZ) database (Fischer et al., 2021). The index captures the productivity advantage of modern production systems across space and time, following the idea that technological change can be inferred from differences in attainable productivity across production regimes (Bustos et al., 2016). We then show that higher TCI is associated with greater corn specialization and lower insurance loss ratios, consistent with technology improving resilience to adverse weather. To interpret these patterns, we develop and estimate a structural discrete-choice model of land use in which counties trade off expected returns, downside risk, and adjustment costs. The model allows us to identify the channels through which technology facilitates adaptation and to quantify welfare under counterfactual climate scenarios.

The structural estimates highlight several mechanisms. Technological change increases the return to specialization: for a given revenue differential between corn and alternative crops, higher-TCI counties derive greater utility from specialization. It also reduces adjustment frictions, lowering the cost of reallocating land, although these frictions remain substantial in highly specialized regions. In addition, risk enters the objective function primarily through downside exposure, measured by insurance loss ratios, rather than overall yield volatility. Technological change increases the value of mitigating this downside risk, consistent with evidence that farmers place particular weight on downside protection (Emerick et al., 2016). Overall, these results show that technological progress reshapes both the profitability and the risk profile of agricultural production.

We then use the structural model to evaluate welfare under greater climate volatility. Counterfactual simulations show that technological advancement mitigates the welfare losses associated with increased weather variability, generating adaptation gains of about \$0.85 per acre. These gains arise mainly through risk mitigation rather than greater land reallocation, suggesting that technology strengthens the resilience of existing specialization patterns rather than simply inducing more adjustment.

This paper relates to three strands of literature. First, it contributes to work on climate impacts in agriculture by emphasizing adaptation through land-use reallocation rather than focusing only on yield responses (e.g., Costinot et al., 2016; Mendelsohn et al., 1994). Second, it contributes to the literature on technological change and structural transformation by proposing a new measure of technological capacity and documenting its effects on both returns and risk (Bustos et al., 2016; Emerick et al., 2016). Third, it contributes to the literature on spatial and environmental adaptation by linking reduced-form evidence to a structural model of decision-making under risk and adjustment costs (Balboni and Shapiro, 2025; Fezzi and Bateman, 2011; Hendricks et al., 2014).

The remainder of the paper proceeds as follows. Section 2 describes the data and the construction and validation of the Technical Change Index. Section 3 presents the conceptual framework linking climate, technological capacity, and corn specialization. Section 4 presents reduced-form evidence on specialization, risk, and adaptation mechanisms. Section 5 develops the structural model and estimation strategy. Section 6 reports the structural estimates. Section 7 presents the welfare and counterfactual analysis. Section 8 concludes.

2 Data and Variable Construction

In this section, we describe the data sources and the construction of the main variables used in the analysis, including land-use shares, weather and climate measures, and the Technical Change Index (TCI).

2.1 Land Use and Weather

We obtain county-level data on harvested corn and soybean acreage from the USDA National Agricultural Statistics Service (NASS).[1] We normalize these measures by total cropland acreage from the Census of Agriculture to construct land-use shares.[2] Our main outcome is the county-level corn share of total cropland, which captures the intensity of corn specialization.

Our study area consists of 22 states in the "Extended Corn Belt" (ECB), which accounts for over 95% of total U.S. corn and soybean production (USDA NASS, 2023) and roughly one-third of global corn and soybean output (Zhong et al., 2016). The ECB includes Arkansas, Colorado, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana, Michigan, Minnesota, Mississippi, Missouri, Nebraska, North Carolina, North Dakota, Ohio, Oklahoma, Pennsylvania,

[1]Acreage data are available at: https://quickstats.nass.usda.gov/.

[2]Census data are available at https://www.nass.usda.gov/AgCensus/.

South Dakota, Tennessee, Texas, and Wisconsin. This broad geographic scope allows us to examine adaptation across diverse agro-ecological conditions, especially along the climate margins of the corn-producing region.

Figure 1 illustrates the spatial evolution of the Corn Belt by comparing county-level specialization in 1955 and 2023. We classify counties as "Core" if their corn land share is at least 15% and as "Edge" otherwise. The figure shows a substantial northward and westward expansion of the high-specialization core over seven decades. In 1955, the core was concentrated in central Iowa and Illinois. By 2023, it extended into northern Minnesota, the Dakotas, and western Nebraska.

Figure 2 summarizes these changes quantitatively. Panel A shows that the geographic footprint of the Core nearly doubled. Panel B shows the average corn shares increased from about 25% to nearly 40%. Panel C follows counties classified as Edge in 1955 and shows that roughly 70% has transitioned into the specialized Core by 2023. This high transition rate points to substantial long-run adjustment along the extensive margin.

To control for climatic determinants of land use, we construct weather variables following the standard empirical literature (e.g., Du et al., 2015; Schlenker and Roberts, 2009). All weather data come from the U.S. Historical Climatology Network (USHCN) maintained by the National Oceanic and Atmospheric Administration (NOAA).[3]

Using daily temperature data over the March–September growing season, we calculate growing degree days (GDD) and overheating degree days (ODD). We define GDD as the sum of daily temperatures within the $[10°C, 30°C]$ ($[50°F, 86°F]$) range, and ODD as cumulative degree days above $32°C$ ($90°F$), which captures heat stress. We also include total precipitation, measured as cumulative rainfall over the same growing season. These variables capture the nonlinear effects of temperature and moisture on expected returns and production risk.

To capture moisture extremes, we also use the Palmer Z-Index (PZ). Following Gong et al. (2023), we define severe drought as months with $PZ < -2$ and severe wetness as months with $PZ > 2.5$. Our main moisture measures, D_c and W_c, denote the number of growing-season months in which county c experiences severe drought or severe wetness, respectively.

[3]The USHCN data are available at https://www.ncei.noaa.gov/products/land-based-station/us-historical-climatology-network.

2.2 Technical Change Index (TCI)

We construct the Technical Change Index (TCI) using data from the FAO Global Agro-Ecological Zones (GAEZ) database.[4] The GAEZ framework simulates crop yield potential as a function of climate, soil, and terrain conditions and reports yield estimates under alternative management regimes. Following Bustos et al. (2016), we define TCI as the difference between attainable yields under high-input and low-input management.

In the GAEZ framework, these management regimes differ in irrigation, fertilizer application, pest control, mechanization, and improved seed varieties. TCI therefore captures the productivity advantage of a modern production system relative to a traditional one. Spatial and temporal variation in TCI reflects both the returns to improved technology and the extent to which technologies mitigate environmental constraints. Appendix B provides additional details on the GAEZ yield calculations and the distinction between low- and high-input scenarios. We aggregate the gridded GAEZ data to the county level by averaging potential yield within each county under each management scenario.

Validation and Imputation

Table A1 shows that TCI is positively and significantly associated with realized corn yields. With county and year fixed effects, a one-unit increase in TCI (scaled by 1,000) is associated with an 8.7% increase in county corn yields. This relationship is robust to first-differencing and to the inclusion of state-by-year fixed effects, which absorb time-varying state-level shocks (Table A2). TCI also interacts positively with heat and drought shocks, consistent with the index capturing both productivity gains and greater resilience to climatic stress (Table A3).

Because the GAEZ potential yield data are available only for the 1990, 2000, and 2010 anchor years, we construct an annual series using supervised learning. We train the model on the anchor years using spatial coordinates and five-year lags of growing-season weather variables, including GDD, ODD, precipitation, and Palmer dry and wet indicators, as predictors. We compare four imputation methods: histogram-based gradient boosting (Ke et al., 2017), random forests (Breiman, 2001), multi-layer perceptron neural networks (Goodfellow et al., 2016), and regularized linear benchmarks (Zou and Hastie, 2005). We use histogram-based gradient

[4] Available at `https://gaez.fao.org/`.

boosting in the baseline analysis because it provides the best fit on the log-TCI scale used in the structural estimation, achieving the highest out-of-sample R^2 (0.842) with a mean absolute error of 0.113 (Table A4). Although random forests perform slightly better on some alternative metrics, gradient boosting better matches the transformed outcome relevant to our application and more flexibly captures the nonlinear relationship between environmental conditions and the technological frontier. This approach preserves meaningful cross-sectional and intertemporal variation in the latent TCI series and helps reduce attenuation from measurement error in the structural analysis.

We further validate the annual imputed series by regressing realized county corn yields on predicted TCI over the full panel. Table A5 reports the results. In a specification with county and year fixed effects, the coefficient on log predicted TCI is 0.116 and statistically significant. After adding weather controls, the coefficient falls to 0.065, consistent with the correlation between technological capacity and environmental conditions. In the most demanding specification, which includes county fixed effects, state-by-year fixed effects, and weather controls, the coefficient remains positive at 0.020 and is marginally significant.

Finally, we compare the machine-learning-based TCI series with a simple linear interpolation across the GAEZ anchor years. Table A5 shows that the machine-learning series significantly predicts realized yields both on its own and when entered jointly with the linear interpolator, with coefficients of 0.070 and 0.066, respectively. By contrast, the linear interpolator is not statistically significant in either specification. These results suggest that the imputed annual series captures meaningful nonlinear variation in technological capacity that is missed by a simple interpolation procedure.

3 Conceptual Model

We model county-level corn specialization as the outcome of a dynamic optimization problem in which producers trade off the returns to specialization against climate-related risk and land-use adjustment frictions.

Let $q_{ct} \in [0, 1]$ denote the corn share of cropland in county c and year t. In each county, a representative producer chooses q_{ct} to maximize the expected present value of net returns. We represent this problem using the following per-period value function:

$$\mathcal{V}_{ct}(q) = \underbrace{\Pi_{ct}(q; W_{ct}, TCI_{ct})}_{\text{Expected surplus}} - \underbrace{\mathcal{R}_{ct}(q; W_{ct}, TCI_{ct})}_{\text{Risk cost}} - \underbrace{\mathcal{C}_{ct}(q, q_{c,t-1}; TCI_{ct})}_{\text{Adjustment cost}}. \tag{1}$$

The first term, Π_{ct}, denotes expected operating surplus from corn relative to alternative crops. In locations with a comparative advantage in corn production, we assume $\partial\Pi/\partial q > 0$. Climate conditions, summarized by W_{ct}, affect the marginal return to specialization, so that $\partial^2\Pi/\partial q\partial W$ may be positive in cooler northern margins, where warming improves suitability, and negative in hotter southern areas, where heat stress reduces returns. Technological capacity, measured by TCI_{ct}, shifts these returns upward, implying $\partial^2\Pi/\partial q\partial TCI > 0$.

The second term, $\mathcal{R}_{ct}$, captures the risk cost of specialization. Greater specialization may increase exposure to weather-sensitive downside outcomes if corn is more vulnerable to extreme realizations than a more diversified crop mix. We therefore allow $\partial\mathcal{R}/\partial q > 0$. We further hypothesize that technological capacity attenuates this exposure, so that $\partial^2\mathcal{R}/\partial q\partial TCI < 0$. In this sense, technology reduces the marginal risk cost of specializing.

The third term, $\mathcal{C}_{ct}$, captures adjustment costs in land use. Reallocating acreage is costly because it requires crop-specific investments in machinery, storage, and agronomic knowledge. We assume a quadratic adjustment cost:

$$\mathcal{C}_{ct} = \frac{\phi(TCI_{ct})}{2}\left(q_{ct} - q_{c,t-1}\right)^2, \tag{2}$$

where $\phi(TCI_{ct})$ governs the magnitude of adjustment frictions. If technology facilitates reallocation through improved seed availability, more flexible input use, or more modular management practices, then $\phi'(TCI_{ct}) < 0$.

The optimal specialization level, q^*_{ct}, satisfies the first-order condition

$$\frac{\partial\Pi}{\partial q} - \frac{\partial\mathcal{R}}{\partial q} - \phi(TCI_{ct})\left(q_{ct} - q_{c,t-1}\right) = 0. \tag{3}$$

This condition states that the optimal corn share equalizes the marginal gain in expected surplus with the marginal increase in risk and the marginal cost of adjustment relative to the previous period's land use.

Totally differentiating the first-order condition yields the response of optimal specialization to a change in climate conditions:

$$\frac{dq^*_{ct}}{dW_{ct}} = \frac{\frac{\partial^2\Pi}{\partial q\partial W} - \frac{\partial^2\mathcal{R}}{\partial q\partial W}}{\phi(TCI_{ct}) - \left(\frac{\partial^2\Pi}{\partial q^2} - \frac{\partial^2\mathcal{R}}{\partial q^2}\right)}. \tag{4}$$

This expression highlights two channels through which technology shapes adaptation. First, technology affects how climate changes the return to specialization and the associated risk exposure, captured in the numerator. Second, technology lowers adjustment frictions, captured by $\phi(TCI_{ct})$ in the denominator, which makes land use more responsive to changing conditions.

The framework yields several testable implications for the empirical analysis. First, higher technological capacity should increase corn specialization by raising the return to specialization and reducing adjustment frictions. Second, higher technological capacity should reduce downside risk, particularly in counties more exposed to heat and moisture stress. Third, climate shocks should have larger effects on land-use adjustment where technological capacity is higher, because producers in those counties face lower effective adjustment costs and weaker risk penalties.

4 Reduced-Form Evidence

This section presents reduced-form evidence on the relationship between technological change, corn specialization, and agricultural risk. The results serve two purposes. First, they establish the empirical relevance of TCI for land-use patterns along both the intensive and extensive margins. Second, they show that technology is associated with greater resilience to climatic stress, which motivates the structural analysis that follows.

4.1 TCI and County-Level Specialization

Table 1 reports the relationship between technological capacity and corn specialization. We begin with the intensive margin, using the continuous county corn share as the dependent variable. Column (1) of Panel A shows that the coefficient on log TCI is 0.009 and precisely estimated. The estimate implies that higher technological capacity is associated with deeper specialization even within established corn-producing regions.

We next examine the extensive margin using indicator variables for whether a county's corn share exceeds 15%, 20%, or 30%. Across these thresholds, the linear fixed-effects estimates are positive and statistically significant, with coefficients around 0.023. This pattern indicates that higher TCI is associated with a greater likelihood that a county enters a specialized "Core" regime. Nonlinear specifications in Panels B and C, estimated using logit and probit models, yield similar conclusions. For example, the logit coefficient for the 30% threshold is 1.855, consistent with technological change playing an important role in reaching high levels of specialization.

All specifications include county and year fixed effects, with standard errors clustered at the county level. To ensure that the estimates do not simply reflect contemporaneous weather realizations or long-run climate trends, we control for growing-season GDD, ODD, precipitation, and Palmer dry and wet indices, each decomposed into a 30-year rolling mean ("climate") and a deviation from that mean ("shock"). Panel D reports robustness checks for the 15% threshold under alternative weather control sets. The estimates are stable when using log climate controls and when excluding either precipitation or Palmer measures, which suggests that the TCI coefficient is not driven by a particular weather specification.

4.2 Specialization, Technology, and Downside Risk

A central implication of the conceptual framework is that technological progress facilitates adaptation by reducing downside risk. We examine this hypothesis using county-level crop insurance loss ratios for corn and soybean policies from 1989 to 2022. Table 2 shows that more specialized regions exhibit systematically lower loss ratios. In Panel A, the coefficient on the Corn Belt indicator is negative and statistically significant across specifications, ranging from -0.131 to -0.229. Panel B shows a similar pattern when specialization is measured using the continuous corn share, with coefficients of roughly -2.29. These results indicate that more specialized corn-producing regions experience lower insured losses on average.

Panel C provides more direct evidence on the technological channel by interacting TCI with specialization and climate shocks. The coefficient on log TCI is negative and highly significant, indicating that higher technological capacity is associated with lower insurance losses. The interaction between TCI and corn specialization is positive and statistically significant. One interpretation is that technological progress narrows the risk gap between highly specialized and less-specialized counties by improving resilience in marginal areas. More broadly, the estimates are consistent with the view that technological change lowers the downside risk associated with corn-intensive land use.

4.3 Mechanisms of Adaptation

To connect the reduced-form results to the structural framework, we examine three mechanisms: climate–technology interactions, land-use adjustment dynamics, and spatial heterogeneity. First, we test whether technological capacity modifies the effect of climate on specialization. Panel A of Table 3 shows that the interaction between TCI and growing-season GDD normals is positive and statistically significant. This result suggests that higher-TCI counties respond more strongly to favorable climatic conditions. The interaction between TCI and extreme heat (ODD) is negative, consistent with technology dampening the adverse effect of excessive heat on specialization.

Second, we examine whether technology facilitates land-use reallocation using a first-difference specification. In Panel B of Table 3, we regress changes in corn share on lagged TCI and weather shocks, including ΔGDD and ΔODD. The coefficients on TCI and its interactions with these shocks are small and imprecisely estimated. This pattern suggests limited short-run responsiveness of land use to annual variation in technology and is consistent with the presence of substantial adjustment frictions.

Finally, we examine spatial heterogeneity in the TCI effect. Table 4 shows that the association between TCI and specialization is stronger in historically "Core" counties than in "Edge" counties, with the interaction term indicating a weaker response at the climatic margins. This result suggests that technological change contributes both to Corn Belt expansion and, more strongly, to intensification within established production regions.

Overall, the reduced-form evidence indicates that TCI is associated with greater specialization, lower downside risk, and differential responses across climatic and spatial environments. Because these estimates cannot separately identify the relative roles of returns, risk, and adjustment frictions, the next section develops a structural discrete-choice model to quantify these channels explicitly.

5 Structural Model and Estimation

We model county-level corn specialization as a discrete choice over a finite set of corn-share alternatives. In each county c and year t, the decision-maker selects a specialization level $q^j \in \mathcal{Q}$

to maximize utility. To preserve tractability while retaining economically meaningful heterogeneity, we discretize the continuous corn-share variable into a set of bins, $\mathcal{Q} = \{q^1, \ldots, q^J\}$, and represent each alternative by its bin midpoint. This yields a conditional logit model defined on an expanded county-year choice set and allows the data to discipline the trade-offs among returns, risk, and adjustment frictions.

5.1 Choice Environment and Utility

Let j index the corn-share bins, and let q^j denote the midpoint of bin j. Utility for county c in year t from choosing alternative j is

$$\begin{aligned} U_{ctj} = {} & \beta_R \, \mathrm{Rev}_{ctj} + \beta_{R,TCI} \, \mathrm{Rev}_{ctj} \, \widetilde{\mathrm{TCI}}_{ct} + \beta_V \, \mathrm{Risk}_{ctj} + \beta_{V,TCI} \, \mathrm{Risk}_{ctj} \, \widetilde{\mathrm{TCI}}_{ct} \\ & + \beta_L \, \mathrm{Loss}_{ctj} + \beta_{L,TCI} \, \mathrm{Loss}_{ctj} \, \widetilde{\mathrm{TCI}}_{ct} - C_{ctj} + \alpha_j + \varepsilon_{ctj}, \end{aligned} \tag{5}$$

where α_j are alternative-specific constants and ε_{ctj} follows a Type I extreme value distribution. The variable $\widetilde{\mathrm{TCI}}_{ct}$ is log TCI centered at its sample mean, so the interaction coefficients capture how technological capacity changes the marginal utility of expected returns, revenue volatility, and downside risk. The term C_{ctj} denotes the cost of moving from last period's specialization level, $q_{c,t-1}$, to the current alternative q^j. This specification allows technological change to affect specialization through four channels: expected returns, revenue volatility, downside exposure, and adjustment frictions.

5.2 Payoff Components

Revenue premium. The revenue component captures the expected return to corn specialization relative to soybeans:

$$\mathrm{Rev}_{ctj} = q^j \left(\mathbb{E}[R^c_{ctj}] - \mathbb{E}[R^s_{ctj}] \right), \tag{6}$$

where $\mathbb{E}[R^k_{ctj}]$ denotes predicted per-acre revenue for crop $k \in \{c, s\}$ under choice j. This term measures the incremental return from allocating a larger share of cropland to corn and is scaled by 1,000 for estimation.

We construct expected revenues in two steps. First, we estimate semi-log yield equations for corn and soybeans with county and year fixed effects, allowing yields to vary with TCI, climatic conditions, and the intensity of corn specialization. In the prediction step, we replace the observed corn share with the midpoint q^j of each candidate choice bin so that predicted yields reflect the agronomic implications of each specialization level. Second, we multiply predicted yields by state-year producer prices from USDA NASS:

$$\mathbb{E}[R^k_{ctj}] = \exp\left(\ln \widehat{\text{Yield}}^k_{ctj}\right) \times \text{Price}^k_{st}. \tag{7}$$

Using predicted rather than realized yields ensures that the revenue premium reflects the producer's ex ante information set at the time of the land-use decision. The underlying yield equations include a quadratic term in specialization, which allows for diminishing returns at high corn intensity. As documented in Table A6, these nonlinear effects capture the agronomic penalty associated with continuous corn production, including higher pest pressure, nutrient depletion, and lower average yields. Incorporating this nonlinearity prevents the structural model from overstating the incentive to specialize by accounting for the productivity trade-offs inherent in intensive monoculture.

Portfolio risk. We measure revenue volatility using the variance–covariance structure of crop revenues:

$$\text{Risk}_{ctj} = (q^j)^2\sigma^2_{c,ct} + (1-q^j)^2\sigma^2_{s,ct} + 2q^j(1-q^j)\sigma_{cs,ct}, \tag{8}$$

where $\sigma^2_{c,ct}$, $\sigma^2_{s,ct}$, and $\sigma_{cs,ct}$ denote the variance of corn revenue, the variance of soybean revenue, and their covariance, respectively. These moments are estimated from residualized county-level revenue series using a 15-year trailing window. We first construct realized per-acre revenue as yield times price and residualize the series on county and year fixed effects to isolate local idiosyncratic volatility from persistent productivity differences and common price shocks. We then estimate the variance–covariance matrix using the 15 years immediately preceding year t, so that the risk measure reflects historical volatility within the producer's information set.

We interpret the producer's portfolio problem as a choice between corn specialization, q, and its primary hedging alternative, $1-q$, which we view as a soybean-intensive rotation. To reduce mechanical correlation with the revenue premium, we orthogonalize the risk measure within each choice set and scale it by 10,000. The interaction coefficient $\beta_{V,TCI}$ captures whether technological capacity changes the effective cost of revenue volatility. Appendix C shows that an analogous portfolio-risk measure based on observed crop shares is systematically related to climatic volatility, especially moisture stress and baseline precipitation.

Projected loss ratios. To capture insured downside exposure, we construct a projected loss-ratio measure that maps each discrete alternative q^j into an implied insurance outcome:

$$\text{Loss}_{ctj} = L_{ct}^{\text{obs}} + \delta\left(q^j - q_{ct}^{\text{obs}}\right), \tag{9}$$

where L_{ct}^{obs} is the observed county-level loss ratio and $\delta = -2.289$ is the estimated marginal effect of corn specialization from the reduced-form loss-ratio regressions.[5]

Because $\delta < 0$, greater specialization is associated with lower projected insurance losses in the reduced-form evidence. This projection allows the structural model to incorporate the downside-risk benefits of specialization. The variable is scaled by 100 for estimation, and $\beta_{L,TCI}$ captures whether technology amplifies the gains from downside protection.

Adjustment costs. We model persistence in land use through an adjustment cost function that depends on the distance between the current alternative q^j and the previous period's specialization, $q_{c,t-1}$:

$$\begin{aligned} C_{ctj} = &\left(\kappa_F + \kappa_{F,TCI}\widetilde{\text{TCI}}_{ct}\right)\mathbf{1}(q^j \neq q_{c,t-1}) \\ &+ \left(\kappa_L + \kappa_{L,TCI}\widetilde{\text{TCI}}_{ct} + \kappa_{L,q}\widetilde{q}_{c,t-1}\right)|q^j - q_{c,t-1}| \\ &+ \left(\kappa_Q + \kappa_{Q,TCI}\widetilde{\text{TCI}}_{ct} + \kappa_{Q,q}\widetilde{q}_{c,t-1}\right)(q^j - q_{c,t-1})^2, \end{aligned} \tag{10}$$

where $\mathbf{1}(\cdot)$ is an indicator for any change in specialization bin, $|\cdot|$ is the absolute adjustment distance, and the quadratic term captures convex costs of large reallocations. The variable $\widetilde{q}_{c,t-1}$ denotes lagged corn share centered at 0.16, the sample mean, allowing frictions to vary with the county's initial level of specialization.

This specification allows adjustment frictions to vary with both technological capacity and initial specialization. Interactions with $\widetilde{\text{TCI}}_{ct}$ test whether technology lowers the cost of reallocation, while interactions with $\widetilde{q}_{c,t-1}$ allow frictions to differ between historically specialized and marginal counties. Because C_{ctj} enters utility with a negative sign, larger coefficients imply greater persistence in land use.

5.3 Estimation

We estimate the model as a conditional logit. Each county-year observation (c, t) is expanded into J alternatives corresponding to the discrete specialization bins in $\mathcal{Q}$. Under the assumption

[5]The parameter δ is obtained from a county-level panel regression of insurance loss ratios on corn share, controlling for weather variables and county and year fixed effects: $L_{ct} = \gamma_0 + \delta q_{ct} + X_{ct}'\gamma + \eta_c + \tau_t + \epsilon_{ct}$. Here L_{ct} denotes the observed loss ratio, q_{ct} is corn share, X_{ct} includes weather controls, and η_c and τ_t are county and year fixed effects. Results are reported in Table A7.

that ε_{ctj} is i.i.d. Type I extreme value, the probability that county c chooses specialization level q^j in year t is

$$P(q_{ct} = q^j \mid \mathbf{X}_{ct}) = \frac{\exp(V_{ctj})}{\sum_{k \in \mathcal{Q}} \exp(V_{ctk})}, \tag{11}$$

where V_{ctj} denotes the deterministic component of utility. We estimate the model by maximum likelihood, with choice sets defined at the county-year level. All payoff components are constructed using information available at the time of the land-use decision and are scaled for numerical stability. The estimation sample is restricted to observations with valid lagged choices and non-missing payoff components.

5.4 Identification

The model identifies three distinct determinants of corn specialization: expected returns, downside risk, and adjustment frictions. Expected returns are identified from how counties' choices vary with predicted revenue differentials across alternatives. The downside-risk channel is identified from variation in projected insurance losses across county-years and specialization bins. Adjustment costs are identified from persistence in observed land-use choices and from the response of specialization to changes in the payoff environment over time.

Climate enters the model through the construction of these payoff components. Changes in temperature, moisture, and weather volatility affect predicted yields, revenue differentials, and projected insurance losses, thereby shifting the attractiveness of alternative specialization levels. Technological capacity mediates these effects by altering the marginal utility of returns, volatility, and downside exposure, and by changing the cost of reallocating land. In this sense, observed shifts in Corn Belt participation are interpreted as optimal responses to changing climatic conditions filtered through local technology and historical persistence.

Because TCI is not randomly assigned, the interaction coefficients should not be interpreted as the causal effect of exogenous technological shocks. Instead, identification comes from within-county variation over time after controlling for county fixed effects, year fixed effects, weather realizations, and other location-specific factors embedded in the payoff construction. The stability of the estimates across specifications supports the interpretation that technological capacity is an important determinant of land-use adjustment.

6 Structural Estimates and Economic Mechanisms

We estimate the structural land-use model using county-level panel data from 1989 to 2022. This period spans major changes in agricultural technology, from the diffusion of transgenic traits to the broader adoption of precision agriculture, and therefore provides substantial variation in the Technical Change Index (TCI). Table 5 reports the maximum-likelihood estimates. The results show that technology shapes land-use decisions through multiple channels, including expected returns, downside risk, and adjustment frictions.

Returns and risk. Technological advancement substantially amplifies the return to specialization. Although the baseline coefficient on the revenue premium ($\beta_R = 34.17$) is imprecisely estimated, its interaction with TCI ($\beta_{R,TCI} = 167.30$) is large and statistically significant. This implies that a given revenue differential translates into a larger utility gain in high-TCI counties, consistent with technology strengthening the comparative advantage of corn-intensive production.

Downside risk, captured by the projected loss ratio, is an important deterrent to specialization. The baseline coefficient ($\beta_L = -97.22$) indicates that higher expected insurance losses make corn-intensive allocations less attractive. The interaction between loss exposure and TCI ($\beta_{L,TCI} = -77.09$) is also negative and statistically significant, implying that specialization decisions in high-TCI counties are more responsive to insurance-relevant downside risk. One interpretation is that technological progress raises the value of protecting against tail losses even as it improves average productivity. By contrast, the coefficients on portfolio volatility are small and statistically insignificant. Conditional on downside exposure, year-to-year revenue variance does not appear to be a primary driver of land-use reallocation.

Adjustment frictions and heterogeneity. Adjustment costs are economically large and precisely estimated, pointing to substantial persistence in land-use patterns. The fixed and linear adjustment-cost components indicate strong baseline inertia in agricultural systems. At the same time, technological change facilitates reallocation by reducing the linear cost of adjustment, as reflected in the positive and significant interaction term $\kappa_{L,TCI} = 8.43$.

The estimates also reveal substantial heterogeneity in these frictions across locations. The positive interaction between linear adjustment costs and initial specialization ($\kappa_{L,q} = 86.40$) suggests that small adjustments are less costly in counties that are already more specialized in corn. However, the negative interaction in the quadratic component ($\kappa_{Q,q} = -287.60$) implies that large reallocations are substantially more difficult in historically specialized "Core" counties. This pattern is consistent with capital lock-in: counties with highly specialized infrastructure, storage, and agronomic systems may adjust at the margin but face high costs of moving far from established corn-intensive production. Technological progress mitigates this curvature, as indicated by $\kappa_{Q,TCI} = -26.65$, thereby lowering the barriers to larger structural shifts in land use over time.

7 Welfare and Counterfactual Analysis

The structural estimates indicate that technological change affects land-use decisions through expected returns, downside risk, and adjustment frictions. We now use these estimates to assess the economic significance of these channels. We begin by translating the utility coefficients into approximate monetary equivalents, then evaluate welfare under counterfactual increases in climate volatility, and finally decompose the resulting adaptation gains across the main technological channels.

7.1 Economic Benchmarks and Monetary Equivalents

To give the structural estimates an economically interpretable scale, we translate the utility coefficients into approximate monetary equivalents measured in dollars per acre. We use the marginal utility of revenue, β_R, as the numeraire. Given the estimated revenue coefficient, $\beta_R = 34.17$ per \$1,000, the implied conversion factor is

$$M = \frac{1{,}000}{34.17} \approx 29.27,$$

which maps changes in choice-specific utility into dollars per acre. These monetary equivalents should be interpreted as local benchmarks for the value producers place on different payoff and friction components, rather than as exact welfare effects.

The structural coefficients imply that technological change materially increases the value of specialization through several channels. First, technology amplifies the return to specialization. The estimated interaction between TCI and the revenue premium, $\beta_{R,TCI} = 167.25$, indicates that technological capacity substantially increases the utility value of a given corn-versus-soy revenue advantage. Using the revenue coefficient as the numeraire, the implied magnitudes are economically meaningful: for counties facing typical corn revenue premia, even modest increases in TCI generate sizable gains through the revenue channel alone. This pattern is consistent with the view that technological progress strengthens the comparative advantage of corn-intensive production.

Second, technology reduces adjustment frictions. The positive interaction between TCI and the linear adjustment-cost term, $\kappa_{L,TCI}$, implies that higher TCI lowers the marginal cost of reallocating land. Quantitatively, our estimates imply that a one-unit increase in log TCI reduces the marginal cost of a 10 percentage-point shift in corn share by roughly \$25 per acre. This is a sizable effect. It suggests that technological progress not only raises the payoff to specialization, but also makes land-use decisions more responsive to changing climatic and economic conditions by reducing the inertia embedded in existing production patterns.

Third, technology increases the value of downside-risk mitigation. The interaction between TCI and the projected loss ratio, $\beta_{L,TCI}$, indicates that higher technological capacity raises the structural value of reducing insurance-relevant losses. For a benchmark reduction in loss-ratio exposure, a one-unit increase in TCI corresponds to approximately \$5.12 per acre for a 10-point shift in specialization. Although this magnitude is smaller than those associated with the revenue and adjustment-cost channels, it remains economically meaningful because it operates on exposure to tail-risk outcomes rather than average returns.

These comparisons suggest that the revenue and adjustment-cost channels are quantitatively the most important margins through which technology affects specialization, while the downside-risk channel provides an additional source of value by lowering exposure to adverse states. Because the baseline revenue coefficient is imprecisely estimated, these dollar equivalents should be interpreted as approximate local benchmarks rather than exact welfare measures. Even so, they show that the TCI-mediated shifts in returns, downside risk, and adjustment frictions are economically meaningful and provide a useful scale for the counterfactual analysis that follows.

7.2 Climate Risk and Welfare Resilience

We next use the structural model to evaluate welfare under increased climate risk. Whereas the previous subsection translated individual coefficients into approximate dollar equivalents, here we measure welfare using the model's expected maximum utility, or log-sum. This object captures the value of the full menu of specialization choices available to each county, including the option to adjust land use in response to changing conditions. We use it to quantify how technological change affects resilience under a counterfactual increase in climatic volatility.

Specifically, we simulate a 50 percent increase in the structural portfolio-variance term, V_q. As shown in Appendix C, this corresponds to an approximately 22.5% increase in the standard deviation of underlying weather shocks. Using the empirical sensitivities reported in Table C1, we interpret this counterfactual as a widening of the historical distribution of heat and moisture realizations. Under this high-volatility regime, extreme events become more frequent. For example, a historical 1-in-20-year drought event occurs with an annual probability of about 6.7%. This mapping anchors the structural counterfactual in a plausible intensification of observed climate risk.

For each county-year, we compute the expected log-sum under both the baseline and high-volatility environments and convert the resulting utility differences into dollars per acre using the numeraire M. Table 6 reports the results. Technological advancement mitigates the welfare losses associated with higher climatic volatility. As shown in Panel B, the implied resilience gain is \$0.91 per acre in expanding "Edge" regions, \$0.79 per acre in entrenched "Core" counties, and \$0.85 per acre on average across the region. Technology therefore does not eliminate the welfare costs of climate shocks, but it provides a meaningful buffer against them.

To understand the source of this resilience effect, we decompose the total adaptation gain into three channels: revenue enhancement, risk mitigation, and adjustment-cost reduction. We do so by selectively allowing TCI to interact with subsets of the structural parameters while holding the remaining channels at their baseline values. This exercise isolates the contribution of each economic margin to the overall welfare effect.

The results show that nearly all of the average adaptation gain of \$0.85 per acre operates through the risk-mitigation channel. This effect is driven primarily by the interaction between TCI and downside-risk exposure, captured by projected loss ratios, which reduces the effective utility penalty associated with adverse insured outcomes. By contrast, the revenue and adjustment-cost channels contribute relatively little to buffering the variance shock. This pattern suggests that, in the climate counterfactual, the main value of modern agricultural technology lies less in inducing large reallocations of land use than in strengthening the resilience of existing specialization patterns.

These welfare gains are smaller than the dollar-equivalent benchmarks associated with the revenue channel in the previous subsection, but the two quantities measure different objects. The larger revenue-based benchmark captures the value of technology in raising expected returns under typical conditions. By contrast, the \$0.85 to \$0.91 per acre estimates isolate the value of technology in cushioning a mean-preserving increase in climatic volatility. This distinction is economically important. In our counterfactual, the variance shock imposes a baseline welfare loss of \$13.54 per acre for the average county under stagnation-era technology (Table 6, Panel A). An average resilience gain of \$0.85 therefore implies that modern technology offsets roughly 6 to 7 percent of the welfare damage caused by the increase in climate risk. The primary value of technological change remains productivity enhancement, but its secondary role as a form of resilience is economically meaningful under adverse climatic conditions.

8 Conclusion

This paper studies how technological change shapes climate adaptation in agricultural land use, using the evolution of corn specialization in the U.S. Corn Belt as a laboratory. Combining a newly constructed measure of technological capacity with long-run county-level data and a structural model, we show that technology plays a central role in determining how producers respond to changing environmental conditions. Higher technological capacity encourages deeper corn specialization, supports limited expansion into marginal areas, and reduces exposure to adverse weather through lower downside risk.

Our structural results clarify the mechanisms behind these patterns. Technological progress not only increases the returns to specialization, but also lowers adjustment frictions and changes the way risk enters land-use decisions. In particular, technology makes production systems more resilient by reducing the importance of severe downside losses, allowing producers to maintain specialized land-use patterns even when climate conditions become more volatile. The welfare analysis underscores the economic importance of these mechanisms: under a counterfactual increase in climate volatility, technological advancement generates adaptation gains of about \$0.85 per acre, with most of these gains coming from improved resilience rather than from large shifts in land allocation.

The results also carry important policy implications. They suggest that innovation is not simply a source of productivity growth, but a key complement to climate adaptation. Public and private investments in agricultural R&D, seed technology, and other innovations that reduce weather-related losses can lower the economic costs of climate change and strengthen the resilience of existing production systems. At the same time, the presence of adjustment frictions implies that adaptation will remain gradual, even when technology improves. This points to a role for policies that ease land-use adjustment, reduce barriers to reallocation, and support adoption in regions that remain vulnerable to climate risk. More broadly, the framework developed here can be applied to other climate-sensitive sectors in which technology, risk, and adjustment costs jointly shape long-run adaptation.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work the author(s) used ChatGPT in order to improve the language of the manuscript. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

Anderson, J. L. (2008). *Industrializing the Corn Belt: Agriculture, Technology, and Environment, 1945–1972.* DeKalb, IL: Northern Illinois University Press.

Balboni, C. and J. S. Shapiro (2025). Chapter 9 - Spatial environmental economics. In D. Donaldson and S. J. Redding (Eds.), *Handbook of Regional and Urban Economics*, Volume 6 of *Handbook of Regional and Urban Economics*, pp. 585–652. Elsevier.

Basso, B., G. Shuai, J. Zhang, and G. P. Robertson (2019). Yield stability analysis reveals sources of large-scale nitrogen loss from the US midwest. *Scientific Reports 9*(5774).

Beddow, J. M., P. G. Pardey, Y. Chai, T. M. Hurley, D. J. Kriticos, H.-J. Braun, R. F. Park, W. S. Cuddy, and T. Yonow (2015). Research investment implications of shifts in the global geography of wheat stripe rust. *Nature Plants 1*(15132).

Breiman, L. (2001). Random forests. *Machine Learning 45*, 5–32.

Bustos, P., B. Caprettini, and J. Ponticelli (2016). Agricultural productivity and structural transformation: Evidence from Brazil. *American Economic Review 106*(6), 1320–1365.

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

Donaldson, D. and R. Hornbeck (2016). Railroads and American economic growth: A "market access" approach. *Quarterly Journal of Economics 131*(2), 799–858.

Du, X. D., C. Yu, D. A. Hennessy, and R. Miao (2015). Geography of crop yield skewness. *Agricultural Economics 46*(4), 463–473.

Emerick, K., A. de Janvry, E. Sadoulet, and M. H. Dar (2016). Technological innovations, downside risk, and the modernization of agriculture. *American Economic Review 106*(6), 1537–1561.

Fezzi, C. and I. J. Bateman (2011). Structural agricultural land use modelling for spatial agro-environmental policy analysis. *American Journal of Agricultural Economics 93*(4), 1068–1088.

Fischer, G., F. O. Nachtergaele, H. T. van Velthuizen, F. Chiozza, G. Franceschini, M. Henry, D. Muchoney, and S. Tramberend (2021). Global agro-ecological zones v4 – model documentation. Technical report, Food and Agriculture Organization of the United Nations (FAO), Rome.

Gong, X., D. Hennessy, and H. Feng (2023). Systemic risk, relative subsidy rates, and area yield insurance choice. *American Journal of Agricultural Economics 105*(3), 888–913.

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning.* MIT press.

Green, T. R., H. Kipka, O. David, and G. S. McMaster (2018). Where is the USA corn belt, and how is it changing? *Science of the Total Environment 618*, 1613–1618.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica 25*(4), 501–522.

Hendricks, N. P., A. Smith, and D. A. Sumner (2014). Crop supply dynamics and the illusion of partial adjustment. *American Journal of Agricultural Economics 96*(5), 1469–1491.

Ke, G., Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu (2017). Lightgbm: a highly efficient gradient boosting decision tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, Red Hook, NY, USA, pp. 3149–3157. Curran Associates Inc.

Klümper, W. and M. Qaim (2014). A meta-analysis of the impacts of genetically modified crops. *PLOS ONE 9*(11), e111629.

Mendelsohn, R., W. D. Nordhaus, and D. Shaw (1994). The impact of global warming on agriculture: A Ricardian analysis. *American Economic Review 84*(4), 753–771.

Olmstead, A. L. and P. W. Rhode (2011). Adapting North American wheat production to climatic challenges, 1839–2009. *Proceedings of the National Academy of Sciences 108*(2), 480–485.

Schimmelpfennig, D. (2016). Farm profits and adoption of precision agriculture. Technical Report ERR-217, U.S. Department of Agriculture, Economic Research Service.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences 106*(37), 15594–15598.

Solakoglu, E. G. and B. K. Goodwin (2005). The effects of railroad development on price convergence among the states of the USA from 1866 to 1906. *Applied Economics 37*(15), 1747–1761.

USDA NASS (2023). Crop production 2022 summary. United States Department of Agriculture, National Agricultural Statistics Service.

Zhong, L., L. Yu, X. Li, L. Hu, and P. Gong (2016). Rapid corn and soybean mapping in us corn belt and neighboring areas. *Scientific Reports 6*, 36240.

Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67*(2), 301–320.

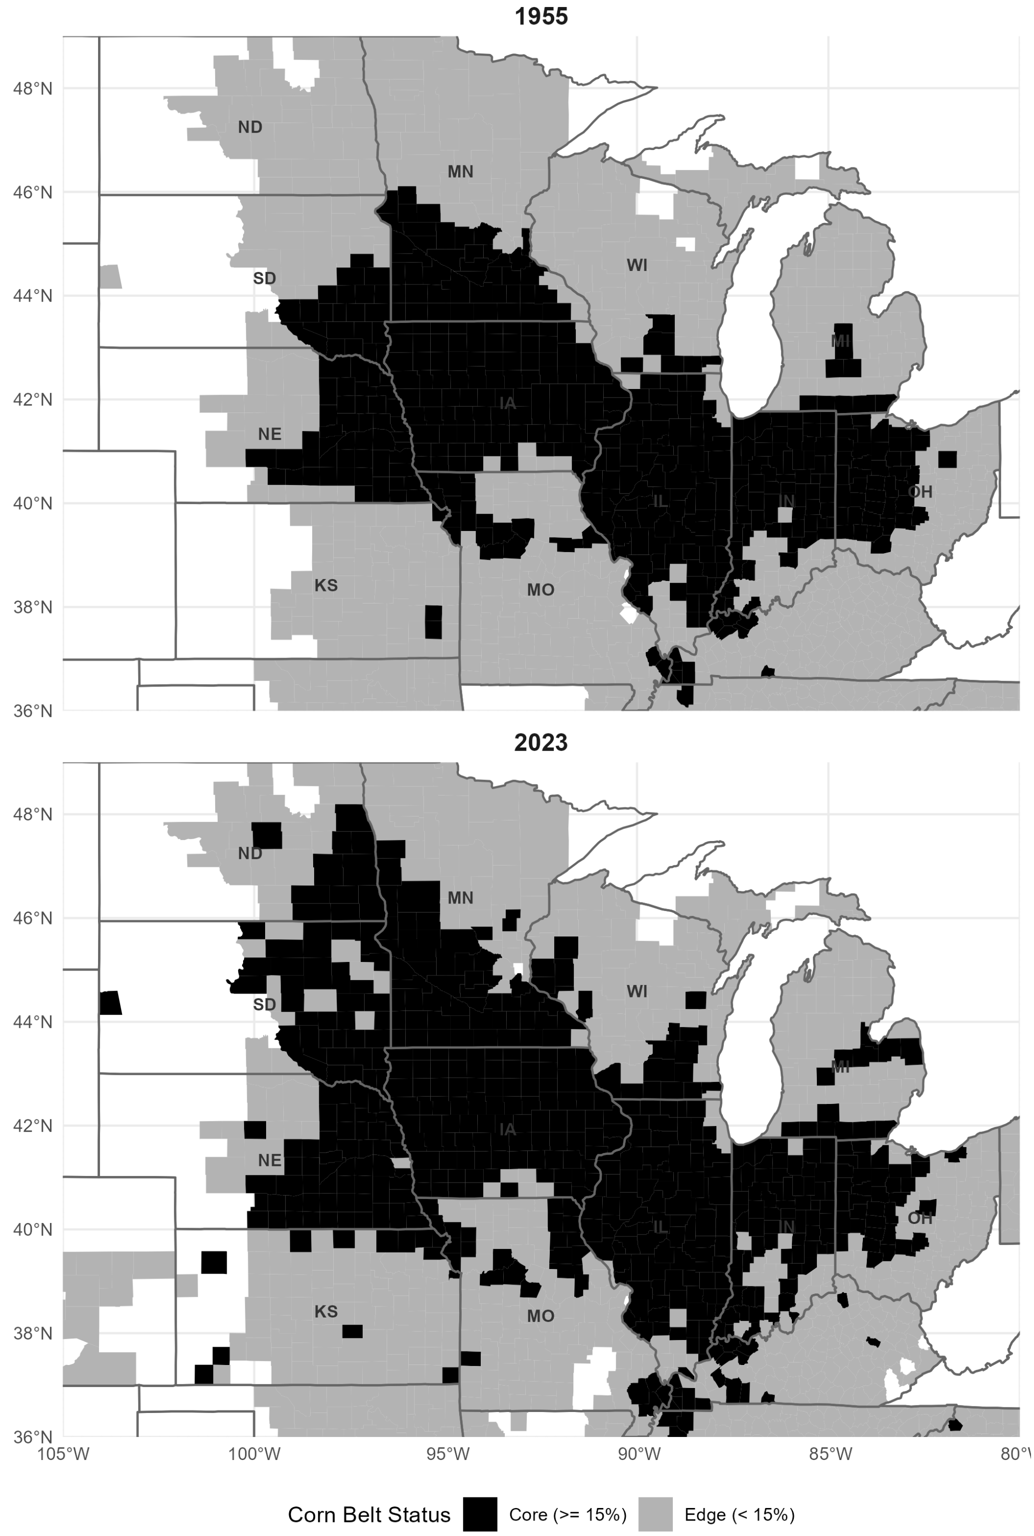

Figure 1: Historical Expansion of the Corn Belt: 1955 vs. 2023

Notes: Spatial distribution of the Corn Belt Core ($\geq$ 15% share, black) and Edge (< 15% share, gray). The comparison shows the significant northward and westward expansion of high-intensity corn specialization over seven decades of technological progress. State abbreviations provided for geographic context.

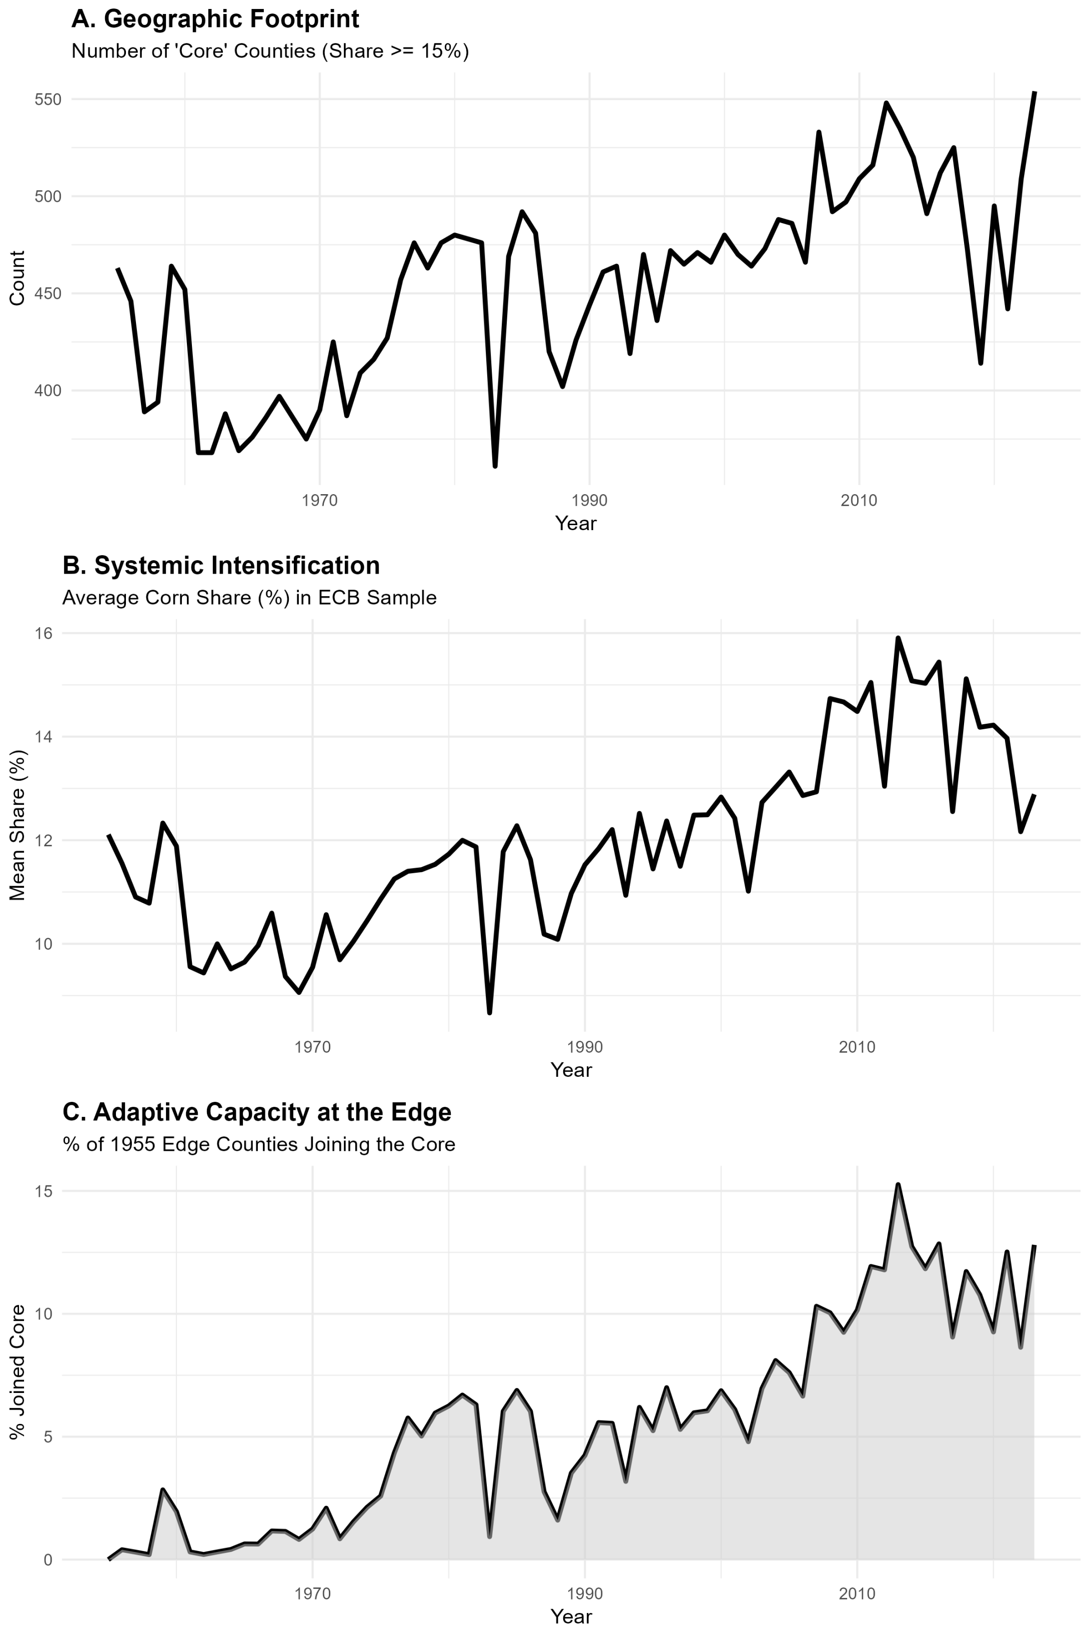

Figure 2: The Great Specialization: Core-Edge Dynamics (1955–2023)

Notes: (A) The geographic footprint of the Core has nearly doubled since the 1950s. (B) Systemic intensification reflects the 40% rise in average corn shares. (C) Evidence of adaptive capacity shows that approximately 70% of the original 1955 Edge counties have now joined the specialized Core.

Table 1: TCI and county corn specialization

	Corn share (1)	0.15 threshold (2)	0.20 threshold (3)	0.30 threshold (4)	0.15 threshold (5)
Panel A. Linear fixed-effects models					
Predicted TCI (log)	0.009***	0.023**	0.022**	0.022***	0.023**
	(0.002)	(0.010)	(0.009)	(0.008)	(0.010)
Observations	58,066	58,066	58,066	58,066	58,066
R^2	0.966	0.885	0.888	0.846	0.885
Panel B. Fixed-effects logit models					
Predicted TCI (log)	–	0.416*	0.500	1.855***	0.424**
	–	(0.218)	(0.333)	(0.664)	(0.215)
Observations	–	12,064	10,701	9,369	12,064
Pseudo-R^2	–	0.528	0.545	0.546	0.528
Panel C. Fixed-effects probit models					
Predicted TCI (log)	–	0.252**	0.340*	0.986***	0.255**
	–	(0.121)	(0.195)	(0.373)	(0.120)
Observations	–	12,064	10,701	9,369	12,064
Pseudo-R^2	–	0.526	0.543	0.544	0.525
Panel D. Robustness for the 0.15-threshold linear probability model					
Specification	Baseline	Log climate	No precipitation	No Palmer	
Predicted TCI (log)	0.023**	0.024**	0.023**	0.023**	
	(0.010)	(0.010)	(0.010)	(0.010)	
Observations	58,066	58,066	58,066	58,066	
R^2	0.885	0.885	0.885	0.885	

Notes: The sample is an annual county panel from the Extended Corn Belt spanning 1955–2024. The dependent variable in column (1) is the continuous corn share of cropland, defined as harvested corn acres divided by total cropland. In columns (2)–(5), the dependent variable is an indicator equal to one if the county corn share exceeds the indicated threshold. Predicted TCI is the annually imputed Technical Change Index derived from a gradient-boosting model trained on GAEZ potential yield gaps. All specifications include county and year fixed effects. Weather controls include growing degree days (GDD, 50–86°F), overheating degree days (ODD, > 90°F), total precipitation, and Palmer Z-index measures of severe drought and severe wetness. Each weather variable is decomposed into a 30-year rolling mean ("climate") and a deviation from that mean ("shock"). In Panel D, the baseline specification corresponds to column (2). Fixed-effects logit and probit models are estimated only on counties that exhibit within-county variation in the threshold outcome. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Corn specialization, technology, and insurance loss ratios

	Dependent variable: county crop insurance loss ratio					
	(1) Baseline	(2) Shock int.	(3) Clim.+prec.	(4) Clim.+prec. int.	(5) Clim.+Palmer	(6) Full
Panel A. Binary specialization measure						
Corn Belt indicator	-0.149***	-0.149***	-0.149***	-0.149***	-0.131**	-0.229***
	(0.056)	(0.057)	(0.057)	(0.057)	(0.055)	(0.062)
Panel B. Continuous specialization measure						
Corn share	-2.289***	–	-2.284***	–	-2.272***	–
	(0.312)		(0.312)		(0.307)	
Panel C. Technology and downside risk						
Corn share	–	-2.279***	–	-15.618***	–	–
		(0.312)		(3.325)		
Predicted TCI (log)	–	-0.275***	–	-0.508***	–	–
		(0.095)		(0.107)		
Corn share × TCI	–	0.00004	–	1.595***	–	–
		(0.0001)		(0.388)		
Observations	39,430	39,430	39,430	39,430	39,430	39,430
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Weather specification	Baseline	Shock int.	Clim.+prec.	Clim.+prec. int.	Clim.+Palmer	Full

Notes: The dependent variable is the county crop insurance loss ratio. Panel A uses a binary specialization measure based on the Corn Belt indicator. Panel B uses continuous county corn share. Panel C examines whether technological capacity attenuates downside risk by including predicted TCI and its interaction with specialization. Column (1) includes county and year fixed effects only. Column (2) adds interactions between specialization and weather shocks. Column (3) adds decomposed long-run climate and weather-shock controls. Column (4) augments column (3) with specialization–weather interactions. Column (5) adds Palmer Z-index measures of severe drought and wetness. Column (6) includes the full set of weather controls and interaction terms. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Technological change, climate interactions, and adjustment dynamics

	(1)	(2)	(3)	(4)
Panel A. Climate–technology interactions				
Predicted TCI (log)	0.008***	0.009***	0.010***	0.012***
	(0.002)	(0.002)	(0.002)	(0.003)
TCI × GDD normal	1.91e-5***	1.85e-5***	1.94e-5***	1.88e-5***
	(0.42e-5)	(0.41e-5)	(0.43e-5)	(0.41e-5)
TCI × ODD shock	-0.004**	-0.005***	-0.004**	-0.005***
	(0.002)	(0.001)	(0.002)	(0.001)
Dependent variable: corn share				
Panel B. Adjustment dynamics				
Lagged TCI (log)	0.001	0.001	0.002	0.001
	(0.001)	(0.001)	(0.002)	(0.001)
TCI ×ΔGDD	-0.12e-5	-0.10e-5	-0.14e-5	-0.11e-5
	(0.22e-5)	(0.20e-5)	(0.24e-5)	(0.21e-5)
TCI ×ΔODD	0.0003	0.0002	0.0004	0.0003
	(0.0004)	(0.0003)	(0.0004)	(0.0003)
Dependent variable: Δ corn share				
Observations	58,066	58,066	58,066	58,066
County fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Weather controls	Baseline	Shock int.	Clim.+prec.	Full

Notes: Panel A reports estimates of interactions between technological capacity (TCI) and climate variables in specifications where the dependent variable is county corn share. Panel B reports first-difference specifications in which the dependent variable is the annual change in county corn share. All models include county and year fixed effects. Weather controls follow the baseline specification in Table 1. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: TCI and Corn Specialization: Core versus Edge Counties

	Core (1)	Edge (2)	Pooled (3)
TCI (log)	0.033***	0.007***	0.033***
	(0.010)	(0.002)	(0.011)
TCI × Edge			-0.0267**
			(0.0109)
Observations	17,927	40,139	58,066
R^2	0.920	0.866	0.966
County FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: Core counties are defined based on 1955 specialization status; edge counties are those outside the historical core. All specifications include county and year fixed effects. Weather controls include growing-season GDD, ODD, precipitation, and Palmer measures of severe drought and wetness. Each weather variable is decomposed into a long-run climate component and a transitory shock component. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Structural estimates of corn specialization

Parameter	Estimate	Std. error
Panel A. Payoff channels		
Relative revenue (β_R)	34.170	36.420
$\times$ TCI ($\beta_{R,TCI}$)	167.300***	62.150
Portfolio risk (β_V)	-0.040	0.260
$\times$ TCI ($\beta_{V,TCI}$)	0.560	0.470
Projected loss ratio (β_L)	-97.220***	24.170
$\times$ TCI ($\beta_{L,TCI}$)	-77.090**	35.390
Panel B. Adjustment frictions		
Fixed cost (κ_F)	-0.730***	0.110
$\times$ TCI ($\kappa_{F,TCI}$)	-0.560***	0.090
Linear cost (κ_L)	-26.950***	2.900
$\times$ TCI ($\kappa_{L,TCI}$)	8.430***	1.780
$\times$ Core ($\kappa_{L,q}$)	86.400***	3.920
Quadratic cost (κ_Q)	-46.320***	16.370
$\times$ TCI ($\kappa_{Q,TCI}$)	-26.650***	6.430
$\times$ Core ($\kappa_{Q,q}$)	-287.600***	35.380
Panel C. Choice-bin fixed effects (α_j)		
Bin 2: [0.05, 0.10)	-0.340***	0.030
Bin 3: [0.10, 0.15)	-0.520***	0.040
Bin 4: [0.15, 0.20)	-0.540***	0.050
Bin 5: [0.20, 0.30)	-0.020	0.050
Bin 6: [0.30, 0.40)	-0.020	0.050
Observations	244,678	
County-year events	34,954	
Concordance (C-index)	0.957	

Notes: Estimates are from a conditional logit model of county-level corn-share bin choice for 1989–2022. TCI is log-centered at 8.95. "Core" denotes lagged corn share centered at 0.16. Bin 1, corresponding to the interval [0, 0.05), is the omitted reference category. Payoff variables are rescaled for numerical stability: revenue by 1,000, portfolio risk by 10,000, and loss ratio by 100. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Welfare resilience and adaptation gains under climate stress

Scenario	Core counties	Edge counties	Regional average
Panel A. Welfare loss from a 50% risk shock ($/acre)			
1990 technology (stagnation)	-14.22	-12.85	-13.54
2020 technology (innovation)	-13.43	-11.94	-12.69
Panel B. Net resilience effect ($/acre)			
Adaptation gain from innovation	0.79	0.91	0.85

Notes: Entries are reported in dollars per acre relative to the historical baseline. The 50% risk shock corresponds to an approximately 22.5% increase in the standard deviation of underlying weather shocks. Panel A reports welfare losses under the high-volatility counterfactual evaluated at 1990 and 2020 technology levels. Panel B reports the difference between these two scenarios and measures the resilience gain attributable to technological progress.

Appendix

A Tables

Table A1: County-level validation of the Technical Change Index (TCI)

	Fixed-effects specifications		First-difference specifications	
	(1) TCI/1000	(2) Normalized TCI	(3) Δ TCI/1000	(4) Δ Normalized TCI
Coefficient	0.087***	0.146***	0.082***	0.132***
	(0.014)	(0.035)	(0.016)	(0.039)
Observations	5,606	5,587	3,598	3,588
R^2	0.755	0.754	0.010	0.004

Notes: The dependent variable is log county corn yield in columns (1) and (2), and the change in log county corn yield in columns (3) and (4). Columns (1) and (2) include county and year fixed effects. Columns (3) and (4) are estimated in first differences. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Additional county-level validation of TCI

	Main validation specifications		Balanced-sample specifications	
	(1) Production (log, + acres)	(2) State × year FE (log yield)	(3) Fixed effects (log yield)	(4) First differences (Δ log yield)
TCI / 1000	0.080***	0.049***	0.083***	0.072***
	(0.015)	(0.015)	(0.015)	(0.017)
Normalized TCI	0.129***	0.063*	0.141***	0.115***
	(0.035)	(0.036)	(0.036)	(0.040)
Observations	5,606	5,606	4,770	3,180
R^2	0.990	0.861	0.720	0.009

Notes: Column (1) regresses log county corn production on TCI while controlling for log harvested acres and including county and year fixed effects. Column (2) uses log county corn yield as the dependent variable and replaces year fixed effects with state × year fixed effects. Columns (3) and (4) restrict the sample to the balanced three-period panel. Column (3) reports fixed-effects estimates, and column (4) reports first-difference estimates. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Weather interactions with the Technical Change Index (TCI)

	(1) TCI / 1000	(2) Normalized TCI
Panel A. Heat and precipitation shocks		
TCI measure	0.079***	0.136***
	(0.015)	(0.038)
TCI × heat shock	9.59e-05**	3.42e-04**
	(4.07e-05)	(1.54e-04)
TCI × precipitation shock	2.02e-06	8.01e-06
	(2.08e-06)	(6.44e-06)
Observations	3,868	3,866
R^2	0.793	0.792
Panel B. Heat and drought shocks		
TCI measure	0.085***	0.165***
	(0.015)	(0.038)
TCI × heat shock	7.97e-05**	2.88e-04**
	(3.54e-05)	(1.39e-04)
TCI × drought shock	0.004**	0.018***
	(0.002)	(0.006)
Observations	3,868	3,866
R^2	0.798	0.798

Notes: The dependent variable is log county corn yield. All specifications include county and year fixed effects. Heat shock is measured using growing-season overheating degree day (ODD) shocks. Precipitation shock is measured using deviations in growing-season precipitation. Drought shock is based on the Palmer drought index. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Comparison of machine-learning algorithms for TCI imputation

	Log-scale performance		Level-scale performance	
Algorithm	R^2	MAE	R^2	Within 10% error
Histogram-based gradient boosting	0.842	0.113	0.957	80.3%
Random forest	0.825	0.108	0.965	82.6%
Neural network (MLP)	0.685	0.222	0.453	44.0%
Elastic net (linear)	0.451	0.299	0.577	32.8%

Notes: This table compares alternative machine-learning methods for imputing annual TCI values using a 20% out-of-sample validation sample. R^2 denotes the coefficient of determination, MAE denotes mean absolute error, and "Within 10% error" reports the share of predictions that fall within 10% of the observed GAEZ anchor value. We use the histogram-based gradient boosting model in the baseline analysis because it provides the best fit on the log-TCI scale used in the structural estimation, although random forests perform slightly better on some alternative metrics. MLP denotes a three-layer multi-layer perceptron.

Table A5: Full-sample validation of the annual predicted TCI series

	Annual validation regressions			Horse-race specifications		
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted TCI (log)	0.116***	0.065***	0.020*	0.070***		0.066***
	(0.015)	(0.014)	(0.010)	(0.015)		(0.014)
Linear interpolator (log)					0.014	0.010
					(0.011)	(0.011)
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	No	Yes	Yes	Yes
State × year fixed effects	No	No	Yes	No	No	No
Weather controls	No	Yes	Yes	Yes	Yes	Yes
Observations	86,929	68,815	68,814	67,690	67,690	67,690
R^2	0.809	0.778	0.879	0.779	0.779	0.779

Notes: The dependent variable is log county corn yield. Columns (1)–(3) regress realized county corn yields on the annual predicted TCI series. Column (1) includes county and year fixed effects. Column (2) adds weather controls. Column (3) replaces year fixed effects with state × year fixed effects and also includes weather controls. Columns (4)–(6) compare the machine-learning-based TCI series with a county-specific linear interpolator constructed from the 1990, 2000, and 2010 GAEZ anchor years. Column (4) includes only the predicted TCI series, column (5) includes only the linear interpolator, and column (6) includes both. Weather controls include growing-season GDD, ODD, precipitation, and Palmer measures of severe drought and wetness. Standard errors, reported in parentheses, are clustered at the county level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: First-stage reduced-form yield regressions

	Dependent variable	
	ln(Corn Yield$_{ct}$)	ln(Soybean Yield$_{ct}$)
Log TCI ($\widetilde{\text{TCI}}_{ct}$)	0.065***	0.028*
	(0.015)	(0.016)
Corn share (q)	-0.055	1.242***
	(0.178)	(0.135)
Corn share squared (q^2)	-0.534**	-1.292***
	(0.246)	(0.203)
Drought (PDSI)	-0.062***	-0.059***
	(0.002)	(0.001)
County fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Observations	44,506	49,195
Adjusted R^2	0.931	0.705

Notes: Standard errors, reported in parentheses, are clustered at the county level. The corn-share variable (q) denotes the observed specialization level in the reduced-form estimation and the bin midpoint q^j in the structural prediction step. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A7: Reduced-form impact of specialization on insurance loss ratios

	Dependent variable: insurance loss ratio (L_{ct})	
	(1)	(2)
	Baseline	Climate controls
Corn share (q_{ct})	-2.289***	-2.272***
	(0.312)	(0.307)
Drought (Palmer dry)		0.248***
		(0.010)
Excess moisture (Palmer wet)		0.151***
		(0.009)
County fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Observations	39,430	39,430
Adjusted R^2	0.004	0.026

Notes: Standard errors, reported in parentheses, are clustered at the county level. The dependent variable is the raw county insurance loss ratio. The coefficient $\delta = -2.289$ from column (1) is used as the shift parameter in the structural loss projections. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Yield Potential Calculation in the GAEZ Model

This appendix describes the Global Agro-Ecological Zones (GAEZ) methodology used to construct crop yield potential measures, following Fischer et al. (2021). The GAEZ framework combines biophysical and agronomic information to simulate attainable yields under alternative management regimes. By comparing low-input and high-input scenarios, it provides a measure of the productivity gains associated with modern agricultural technology.

Yield Potential in GAEZ

In the GAEZ framework, yield potential is defined as the maximum attainable yield under a given set of environmental conditions and management assumptions. The simulation combines four main inputs:

1. Climate: Temperature, precipitation, radiation, and evapotranspiration determine the length of the growing period, thermal suitability, and photosynthetic potential.
2. Soil and terrain: Soil texture, depth, fertility, drainage, and slope affect nutrient availability, water retention, and the feasibility of mechanized cultivation.
3. Crop physiology: Crop-specific parameters for corn and soybeans, including biomass accumulation, water requirements, and harvest indices, determine growth and yield formation.
4. Management inputs: Assumptions about seed genetics, fertilizer application, irrigation, and crop protection determine the attainable production frontier.

Yield potential is obtained by matching local agro-climatic conditions to crop-specific ecological requirements. An important distinction across scenarios is the assumed seed technology. Low-input systems are based on traditional varieties selected primarily for robustness under limited external inputs, whereas high-input systems assume improved varieties designed to respond more strongly to fertilizer, irrigation, and crop protection.

Management Scenarios

The GAEZ model distinguishes between alternative management regimes to capture differences in technological intensity.

The low-input scenario represents a production environment with limited external inputs. It assumes no chemical fertilizer, no irrigation, reliance on manual labor or simple tools, and limited pest and disease control. Yield potential in this case reflects the combined effect of environmental constraints and restricted technological support.

The high-input scenario represents modern intensive agriculture. It assumes improved nutrient management, irrigation sufficient to reduce moisture stress, mechanized cultivation, and more advanced pest and disease control. It also incorporates high-yielding varieties or hybrids. These assumptions provide an estimate of the attainable yield frontier under modern production practices in a given location.

The Technological Gap and Model Constraints

We define the Technical Change Index (TCI) as the difference between high-input and low-input yield potential. This gap serves as a proxy for the local productivity gains associated with advanced technology. In most cases, the difference is positive. However, the GAEZ framework also allows for cases in which high-input yield potential is lower than low-input yield potential, implying a negative TCI. Such cases arise when local environmental constraints disproportionately limit intensive production systems.

Several mechanisms can generate this outcome. First, mechanized cultivation may be less effective in areas with steep slopes, poor drainage, or excessive wetness, where low-input systems relying less on machinery may be more adaptable. Second, humid environments with long growing periods may generate pest and disease pressure that is difficult to control even under intensive management. Third, improved varieties may offer higher yield ceilings under normal conditions while remaining more sensitive to certain extreme events, such as frost or excessive moisture during harvest, than locally adapted traditional varieties.

These features imply that TCI captures not only the abstract value of innovation, but also the location-specific suitability of advanced technology. In this sense, the GAEZ-based measure reflects both the productivity potential of modern agriculture and the extent to which local environmental conditions permit that potential to be realized.

C Empirical Alignment of the Portfolio Risk Measure

This appendix validates the portfolio-risk measure, $V_{q,ct}$, by examining its relationship with observed climatic conditions. We define the measure as the variance of a county-level crop revenue portfolio:

$$V_{q,ct} = q_{c,ct}^2 \sigma_{c,ct}^2 + q_{s,ct}^2 \sigma_{s,ct}^2 + 2 q_{c,ct} q_{s,ct} \sigma_{cs,ct},$$

where $q_{c,ct}$ and $q_{s,ct}$ denote the observed land-use shares of corn and soybeans, respectively. This formulation captures both crop-specific revenue volatility and the risk-reduction benefits of diversification.

We estimate the underlying variance and covariance terms using rolling moments from residualized revenue series. Specifically, per-acre crop revenues, defined as yield times state-level price, are residualized using county and year fixed effects to remove persistent productivity differences and common price movements. We then compute the variance–covariance matrix from these residuals using a 15-year trailing window.

Table C1 reports the relationship between V_q and climate variables. The results indicate that moisture stress is the primary driver of portfolio risk. The drought indicator enters with a large and statistically significant coefficient of 38.03, indicating that moisture stress is a major source of revenue volatility. Precipitation normals enter negatively and significantly, suggesting that counties with higher baseline moisture levels exhibit lower overall revenue risk. By contrast, heat normals and temperature shocks are not statistically significant in this specification, implying that, for the average county, water balance is a more informative predictor of revenue risk than temperature alone.

Mapping Risk to Weather Shocks

To discipline the counterfactual simulations in Section 7, we map a 50 percent increase in the structural variance term, V_q, into changes in the physical distribution of weather shocks. Because the risk object is defined in variance units, a 50 percent increase in variance implies an approximately 22.5 percent increase in the standard deviation of the underlying shocks.

Using the empirical distribution of county-level weather shocks over 1950–2022 as a benchmark, this counterfactual widens the distributions of both overheating degree days (ODD) and growing-season precipitation. For ODD, the standard deviation rises from 52.0 to 63.7; for precipitation, it rises from 2,063 to 2,527. This scaling implies a substantial increase in the frequency of tail events. For example, a historical 1-in-20-year weather event, which occurs with a 5% annual probability, would occur with a probability of 6.7% under this counterfactual, while a 1-in-10-year event would rise from 10% to 15.5%.

This mapping links the structural risk parameter to observable changes in climate volatility. It therefore allows the counterfactual analysis to interpret changes in V_q as shifts in the physical distribution of weather realizations, rather than as abstract changes in an unobserved risk parameter. In this way, the model can evaluate how technological change affects the agricultural sector's resilience to increasingly volatile climatic conditions.

Table C1: Climate drivers of refined portfolio risk (V_q)

	Dependent variable: portfolio risk (V_q)
Weather shocks	
Drought indicator (Dry)	38.030***
	(4.636)
Excess moisture (Wet)	2.123
	(4.763)
Heat shock (ODD)	-0.006
	(0.092)
Precipitation shock	-0.005
	(0.020)
Climate normals	
Heat normals (ODD)	-0.319
	(0.655)
Precipitation normals	-0.613***
	(0.139)
County fixed effects	Yes
Year fixed effects	Yes
Observations	38,277
Within R^2	0.003

Notes: Standard errors, reported in parentheses, are clustered at the county level. The dependent variable, portfolio risk (V_q), measures the variance of the county crop revenue portfolio and is calculated as

$$V_q = q_c^2 \text{Var}_c + q_s^2 \text{Var}_s + 2 q_c q_s \text{Cov}_{cs},$$

where q_c and q_s are the observed land-use shares of corn and soybeans. Weather shocks are defined as annual deviations from the county-specific 30-year rolling mean, and climate normals are defined as 30-year rolling averages. Dry and Wet denote the number of growing-season months in which the Palmer Z-index indicates moderate-to-extreme drought or moisture surplus, respectively. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.