Supplementary Information

The Economic Cost of Agricultural Climate Migration in the United States

[Author names redacted for double-blind review]

This document contains supplementary methods, robustness analyses, and mathematical derivations supporting the main text. It is intended to be read alongside the main manuscript and Extended Data.

This Supplementary Information provides the technical detail underlying the four findings reported in the main text. Part I describes the data assembly and predictive models. Part II documents the robustness of the stranded asset estimates across methods, parameters, and model specifications. Part III addresses the rural decline and insurance components, including the migration IV and forward projections. Part IV provides complete mathematical derivations for all headline computations. All monetary values are in 2023 USD.

Contents:

Headline Numbers (Cross-Reference)

All headline numbers cited in the main text are summarised below for cross-reference. All monetary values are in 2023 USD.

Table S1: **SI Table S0. Headline numbers cross-reference.** All quantities reported in the main text with source methods and section references.

Finding	**Value**	**Method / source**
Stranded farmland value (hedonic)	\$168B	Ricardian hedonic, SSP2-4.5, 2050
Stranded farmland value (DCF conservative)	\$56B	ML-only, $r = 4\%$, $H = 30$yr
Stranded farmland value (DCF upper)	\$140B	ML+SR+indirect, $r = 2.5\%$, $H = 40$yr
Insurance mispricing	\$5.9B/yr	APH repricing, bootstrap 95% CI: [\$5.4, \$6.4]B
Implicit transfer (N→S)	\$2.8B/yr	Expected-indemnity method, 95% CI: [\$2.6, \$3.0]B
Emerging production frontier	\$51B/yr	514 counties; yield \$4.6B + expansion \$37.5B (adjusted) + upgrade \$0.1B + dairy \$9B
Counties with 4+ decline indicators (obs.)	473	Observed 2009–2023 data
Counties crossing threshold $\leq$ 2040 (own IV)	2	$\beta = -0.003$, $p = 0.019$, dual calibration A
Counties crossing threshold $\leq$ 2040 (Feng 2010)	253	$\beta = -0.17$, dual calibration B
Market test β	$+0.019$ ($p = 0.027$)	OLS $\Delta\,\ln$(land value) on ΔT^{2040}, $N = 3{,}049$
Yield model Spearman ρ (v2 model)	0.45	Test set 2013–2023, v2 compound-drought model
2012 drought prediction (v2 model mean)	-0.85σ	Acreage-weighted: -1.07σ; observed: -0.94σ
Cotton switching validation ρ	0.59	CDL-based switching, 1980–2010
514 emerging frontier counties	514	ACS + NASS, $\Delta T > 0$ and arable land $> 40\%$ unused

---

Part I — Data and Model Foundation

S1 Data Processing Details

S1.1 NASS Deduplication

USDA NASS QuickStats returns one row per survey response, which occasionally includes multiple entries for the same county, year, and crop when partial-year estimates are revised or when both survey and census estimates appear in the same download. Before any analysis, the raw panel is deduplicated using:

```
df = df.groupby(['fips', 'year', 'crop']).first().reset_index()
```

This retains the first (most recent) record per county-crop-year group, which corresponds to the final revised estimate in chronological download order. The raw panel contained 675,241 rows; after deduplication it contains 638,808 rows across 2,902 counties, 8 crops, and 1950–2023.

S1.2 CPI Deflation

All monetary values are deflated to 2023 USD using the BLS CPI-U annual series (FRED: CPIAUCSL). The deflation formula is:

$$V_{2023} = V_t \times \frac{\mathrm{CPI}_{2023}}{\mathrm{CPI}_t} \tag{S1}$$

where $\mathrm{CPI}_{2023} = 304.7$ (annual average). Values are deflated at the point of loading from raw data, before any aggregation or modelling step, to prevent mixing nominal vintages. For NASS land values and cash rents (reported in nominal USD), deflation is applied at the county-year level. For RMA insurance data, premium, indemnity, and liability figures are all deflated using the CPI of the survey year.

S1.3 FIPS Filtering

Three FIPS exclusion rules are applied before any analysis:

1. **State exclusions**: Alaska (state FIPS 02), Hawaii (15), and Puerto Rico (72) are excluded. These territories have distinct agricultural systems, climate regimes, and data coverage that would require separate modelling.
2. **Aggregate codes**: NASS uses county FIPS codes 998 and 999 to denote state-level aggregates and "other counties" remainder rows. These are excluded.
3. **Format standardisation**: All FIPS codes are stored as zero-padded 5-character strings (e.g., `'01001'`) to prevent integer-conversion errors that drop leading zeros.

After filtering, the panel covers 2,902 of 3,108 continental US counties. The 206 excluded counties either lack NASS coverage entirely or appear only in aggregate rows.

S1.4 ACS Migration Variable Correction

During IV estimation, a mislabelled ACS variable was identified in the Census API documentation. Column `B07001_002E` is documented as "Moved from different county, same state" but is actually the count of non-movers (same house). The correct variable for inter-county in-movers (same state) is `B07001_049E`. This was corrected before IV estimation; the correction was applied on 2026-03-18 and is reflected in all results reported in the main text.

S2 Yield Model Architecture and Performance

S2.1 Ensemble Yield Model Architecture

The final yield model is a three-component NNLS-weighted ensemble: LightGBM (weight 0.74), Random Forest (weight 0.24), and Ridge regression (weight 0.02). All three models train on identical features. Blend weights are fit by non-negative least squares on a held-out 2010–2012 validation period—after the base-model training window (1950–2009) and before the test set (2013–2023)—ensuring no leakage.

Table S2 reports the LightGBM component hyperparameters. The Random Forest uses `n_estimators` = 150, `max_depth` = 10, `min_samples_leaf` = 30, `max_features` = 0.4, `max_samples` = 0.3. Ridge uses $\alpha = 10$ on StandardScaler-normalised features. All components use `random_state` = 42.

Table S2: **SI Table S1. LightGBM component hyperparameters.** Parameters were set a priori from agronomic literature; no data-driven tuning against the test set.

Parameter	Value	Rationale
`n_estimators`	1,500	Extra capacity for compound interaction terms
`learning_rate`	0.02	Conservative; reduced from 0.03 for depth=8
`max_depth`	8	Captures three-way heat×drought×crop interaction
`num_leaves`	127	$2^7 - 1$; consistent with max_depth
`min_child_samples`	20	Prevents leaf-level county overfitting
`subsample`	0.8	Row-level bagging (Friedman, 2002)
`colsample_bytree`	0.8	Feature subsampling per tree
`reg_alpha`	0.05	L1 regularization
`reg_lambda`	0.5	L2 regularization
`random_seed`	42	Fixed for reproducibility
`objective`	regression	Mean squared error on z-scored anomalies
`metric`	rmse	Evaluation during training

S2.2 Feature List

The ensemble uses 50 features: 28 continuous climate and demographic features, 14 compound drought interaction terms, 8 peak-season monthly climate features (computed from JJA monthly data), and 8 crop-type dummy variables. The interaction features were added specifically to capture the multiplicative heat-drought penalty observed in the 2012 Midwest drought. Table S3 lists the core feature categories.

S2.3 Temporal Cross-Validation Fold Structure

To prevent data leakage from future climate patterns into training, the analysis uses strict temporal rolling cross-validation with a 2-year gap between training and validation. Table S4 shows the five fold assignments.

S2.4 Model Failure Mode Analysis

To characterise the residual gaps in v2 model performance, the analysis identifies the 5% of county-crop observations in the 2017–2023 test set where absolute prediction error was largest ($|\text{error}| \geq 2.53\sigma$; $n = 1{,}599$ of 31,966 test observations). The analysis uses the v2 yield model trained with 14 compound drought interaction features.

Table S5 reports over-representation in the worst-5% tail relative to the expected 5% baseline.

Error direction. In the worst-5% tail, 97.5% of errors are *positive*—the model under-predicts actual yields (i.e., the model is too pessimistic). Mean signed error is $+2.83\sigma$. The model most commonly fails by *over-predicting climate damage*, not by missing disasters.

Climate conditions. Failure observations are slightly warmer and wetter than average:

- July maximum temperature: worst-5% counties average 31.1°C vs 30.9°C in the rest ($+0.22$°C).
- Growing-season precipitation: worst-5% counties average 24.3 in. vs 22.4 in. in the rest ($+1.9$ in.)—failure at *above-average* moisture, not extreme drought.
- Extreme heat months (>3 per season): 1.1% of worst-5% vs 0.4% of rest—more common in failures but rare overall.

Geographic concentration. Failure concentrates in climatic transition-zone states: Texas (8.9% of state test obs), North Carolina (8.4%), Pennsylvania (11.4%), Oklahoma (11.4%), New York (11.1%)—states at the boundary between agricultural systems, where neither core-belt nor deep-south rules apply cleanly.

Year pattern. 2017 (6.5% failure rate) and 2023 (7.1%) show the highest failure rates among test years—both transition years in which regional climate diverged substantially from 10-year averages.

What v2 still misses. The failure mode analysis reveals two channels the compound-drought v2 model does not capture:

1. **Positive climate surprises in transition-zone crops.** Oats and winter wheat fail most often, and almost always because actual yields *exceeded* expectations. The model is overly conservative for crops grown at the northern margins of their range during above-average moisture years.

2. **Non-climate production shocks in border states.** Pennsylvania, New York, and North Carolina have high failure rates despite being outside the core agricultural belt. In these states, production variability is driven partly by pest pressure, labour availability, and specialty practices that a 14-feature climate model cannot fully capture.

Both failure modes are *conservative* from the stranded-value perspective: the model over-predicts losses in oats and wheat, and in geographically peripheral states. The core Corn Belt, Southern Plains, and Delta counties that drive the headline stranded-value estimate are not affected by these failure modes.

Beyond the empirical failure analysis, three compound event types remain outside the v2 specification: consecutive drought–heat years (lagged soil-moisture memory is absent); flash drought (PDSI is too slow-response); and late-spring frost after early warming (relevant primarily for specialty crops outside the eight-crop panel). All three gaps bias stranded-value estimates downward.

The model's conservative bias (97.5% of large errors over-predict damage) means the stranded asset estimates in the main text are lower bounds.

S3 Switching Model Calibration

S3.1 Brier Scores

Platt-scaled LightGBM classifiers were trained for four crop-pair switching decisions. Table S6 reports Brier scores on the held-out test set. Brier scores range from 0 (perfect) to 1 (worst); scores below 0.25 indicate useful discrimination. All four models are well-calibrated, with the corn→soybeans pair showing the highest score (0.125) reflecting the broader distributional overlap between these two crops across transition counties.

S3.2 Historical Validation

Switching model validation uses four historical crop transition events plus one negative control. The cotton-belt eastward retreat (1980–2010) yields a Spearman $\rho = 0.59$ between predicted switching probability and observed CDL county-level transitions. The negative control (soybean technology adoption, 1960–1980) correctly produces near-zero predicted switching probability, confirming that the model responds to climate signals rather than technology trends.

Part II — Stranded Asset Robustness

S4 Market Efficiency Test

S4.1 Regression Specification

The analysis tests whether current farmland prices incorporate forward-looking climate risk using the following cross-sectional regression:

$$\Delta \ln(\text{LandValue}_i) = \alpha + \beta \cdot \Delta T_i^{2040} + \gamma' X_i + \delta_s + \varepsilon_i \tag{S2}$$

where $\Delta \ln(\text{LandValue}_i)$ is the change in log farmland value per acre from 2012 to 2022 (2023 USD), ΔT_i^{2040} is the projected warming in county i by 2040 relative to 1981–2010 baseline (10-GCM ensemble median, SSP2-4.5), X_i is a vector of controls (log income growth 2012–2022, log population change, crop mix share), and δ_s are state fixed effects. Standard errors are HC3 heteroskedasticity-robust.

S4.2 Full Regression Table

Table S7 reports the full regression output. The coefficient on projected warming is positive and statistically significant across all specifications, confirming that markets are pricing climate risk in the wrong direction.

S4.3 Interpretation of State Fixed Effects

State fixed effects absorb mean differences in land value appreciation rates across states (e.g., Iowa vs. Texas), isolating the within-state relationship between warming exposure and price appreciation. The positive β on ΔT^{2040} thus means: within a given state, counties projected to warm more have appreciated faster, not slower. This cannot be explained by aggregate state-level trends in agricultural demand or policy. The result is consistent across all four specifications in Table S7.

S4.4 Alternative Specifications and Robustness

Four alternative specifications were tested. First, replacing SSP2-4.5 projected warming with a historical 2000–2022 observed warming trend (from nClimDiv) yields $\beta = +0.017$ ($p = 0.041$), consistent in sign and magnitude. Second, restricting the sample to counties where farmland constitutes more than 50% of total land area yields $\beta = +0.022$ ($p = 0.018$). Third, using a 5-year appreciation window (2017–2022) instead of 10-year yields $\beta = +0.016$ ($p = 0.063$). Fourth, weighting by county cropland acreage yields $\beta = +0.024$ ($p = 0.011$). The positive coefficient is robust across all specifications.

S5 Sensitivity Analysis

S5.1 Leave-One-Crop-Out Sensitivity

To test whether the aggregate stranded-value estimate is driven by any single crop, total stranded assets were re-estimated after removing one crop at a time. Table S8 shows results; the $56B conservative DCF estimate is robust to removing any single crop.

Corn contributes the largest share ($\sim$20%) of total stranded value. Even removing corn entirely, the estimate remains \$44.9B—still economically substantial.

S5.2 Leave-One-GCM-Out Sensitivity

Table S9 shows stranded value estimates after removing one GCM at a time from the 10-model ensemble. The range is narrow, confirming that no single GCM drives the headline.

S5.3 Discount Rate and Horizon Grid

The sensitivity grid spans $r \in \{2\%, 3\%, 4\%, 5\%, 6\%, 7\%, 8\%\}$ and $H \in \{20, 25, 30, 35, 40\}$ years, yielding 35 parameter combinations. The conservative DCF at the reference parameters ($r = 4\%$, $H = 30$ yr) is \$55.9B. Across all 35 combinations, the minimum is \$27B (at $r = 8\%$, $H = 20$ yr) and the maximum is \$119B (at $r = 2\%$, $H = 40$ yr). The headline range of \$56–140B corresponds to the lower quartile through the 90th percentile of this grid (see S10 for the full DCF formula).

Across all 35 parameter combinations in the sensitivity grid, stranded value never falls below \$27B—confirming that the finding is robust to reasonable disagreements about discount rates and horizons.

S6 Hedonic-DCF Decomposition

SI Section 6: Incremental Hedonic Decomposition of Stranded Value

We decompose the hedonic stranded estimate into channel contributions by fitting three nested regression models and tracking how the stranded value estimate changes as each channel group is added (Table S10). This approach uses direct regression estimates at each step rather than proxy correlations, so every dollar allocation is grounded in the data.

Method. We apply the Mendelsohn-Nordhaus-Schlenker cross-sectional framework to 2,720–2,908 counties, regressing log farmland value on climate and control variables with state fixed effects and HC3-robust standard errors. For each fitted model, we compute stranded value by applying CMIP6 SSP2-4.5 median 2050 warming deltas through the temperature coefficients only, holding all other covariates constant. The channel contribution of each added variable group equals the change in stranded value between consecutive models.

Model 1 — Climate only: $\log V = \beta_1 T + \beta_2 T^2 + \beta_3 P + \alpha_s + \varepsilon$. Includes July maximum temperature (T), its square, growing-season precipitation (P), and state fixed effects. $R^2 = 0.565$, $N = 2{,}908$ counties.

Model 2 — Climate + Demographics: Adds $\log(\text{population})$ and $\log(\text{median income})$ to absorb urban proximity and income effects. $R^2 = 0.752$.

Model 3 — Climate + Demographics + Soil: Adds a soil quality index — the cross-crop mean z-score of each county's maximum historical yield (NASS 1990–2009) — to control for inherent soil productivity. $R^2 = 0.759$, $N = 2{,}720$ counties.

Table S10: **SI Table S10. Incremental hedonic decomposition of stranded farmland value.** Each model adds one channel group to the previous. Stranded value applies CMIP6 SSP2-4.5 2050 warming to temperature coefficients, holding all other variables constant. Contributions equal stranded$_N$ − stranded$_{N-1}$. All regressions use HC3 heteroskedasticity-robust standard errors with state fixed effects.

Model	**Channel added**	R^2	**Stranded ($B)**	**Contribution ($B)**
1: Climate only	tmax, tmax^2, precip	0.565	$124.9	$124.9 (baseline)
2: + Demographics	log(pop), log(income)	0.752	$154.5	+$29.6
3: + Soil quality	Max hist. yield index	0.759	$186.8	+$32.3
DCF central estimate	Field-crop income	—	$105.1	—
Unmodeled gap	Livestock, water, amenity, specialty	—	—	+$81.7

Climate channel ($124.9B). Model 1 with climate variables and state fixed effects alone explains 56.5% of cross-county land value variation. The temperature coefficient is negative ($\hat{\beta}_T = -0.160$, $p = 0.055$) and the quadratic term positive ($\hat{\beta}_{T^2} = +0.00081$, $p = 0.093$), consistent with a U-shaped temperature response in the raw cross-section. Applying the median +1.87 °F July warming by 2050 yields $124.9B in stranded value from the climate channel alone.

Demographic channel (+$29.6B). Adding log(population) and log(median income) raises R^2 from 0.565 to 0.752 — a 18.7 percentage-point gain, the largest single-step improvement. This reallocation sharpens the climate signal: the temperature coefficient more than doubles in magnitude ($\hat{\beta}_T = -0.272$, $p < 0.001$, compared to $\hat{\beta}_T = -0.160$ without demographics) once demographic confounders are absorbed. The stranded estimate rises to $154.5B because counties with high population and income also tend to be warmer, and controlling for that correlation isolates the pure agricultural climate response. The $29.6B demographic contribution reflects farmland in high-income, populous counties that the climate-only model could not fully credit to climate risk.

Why demographics raise stranded value (not lower it): This result may seem counterintuitive — controlling for more variables usually reduces the coefficient of interest — but here the confound runs in the opposite direction. Hot southern counties are disproportionately rural and low-income; without demographic controls, some of the temperature–land-value gradient is absorbed by the poverty/urbanisation correlation, compressing the apparent climate effect. Adding log(pop) and log(income) strips out that confound and recovers the *pure* agricultural climate signal, which is larger, not smaller. Verified directly: $\hat{\beta}_T$ shifts from -0.160 (climate-only) to -0.272 (+ demographics), confirming that removing the demographic confound sharpens, rather than attenuates, the climate coefficient.

Soil quality channel (+$32.3B). Adding the historical yield index raises R^2 to 0.759 and sharpens the temperature coefficient further ($\hat{\beta}_T = -0.325$, $p < 0.001$). The soil index itself is highly significant ($\hat{\beta}_{\text{soil}} = +0.101$, $p < 0.001$): a one-standard-deviation improvement in soil quality is associated with a 10% premium in land value. Controlling for soil productivity reveals that productive-soil counties also face higher climate exposure, raising the stranded estimate to \$186.8B. The \$32.3B soil contribution represents land value at risk in productive agricultural counties where the climate-only model was insufficiently controlling for soil quality.

Unmodeled gap ($81.7B). Model 3 captures climate, demographic, and soil channels but still exceeds the DCF central estimate (\$105.1B) by \$81.7B. The hedonic regression reflects observed market prices, which embed all value channels including those the DCF ignores: livestock and dairy heat-stress benefits to northern operations, water-supply risk in irrigation-dependent counties, amenity premia, and specialty-crop climate sensitivity. This \$81.7B gap — 44% of the hedonic total — is consistent with prior literature estimates for these channels and provides a lower bound on the value captured in land markets beyond field-crop income alone.

Methodological contrast with prior approach. Earlier versions of this section correlated the per-county hedonic-DCF gap with proxy indicators, yielding negative Spearman correlations on two of four channels ($r = -0.09$ for amenity, $r = -0.18$ for specialty crops). Those negative correlations indicated the proxy variables were capturing orthogonal variation rather than the channels of interest. The incremental regression replaces that proxy correlation approach with a direct estimation strategy: each channel's contribution is measured by how much the stranded estimate changes when that channel is explicitly controlled for. No proxy correlations are required, and every estimate is grounded in the regression at each step.

Part III — Rural Decline and Insurance

S7 Migration IV Specifications

S7.1 Migration Elasticity: Own IV vs. Feng et al.

The rural decline model requires a migration elasticity—the percentage change in county in-migration per 1% change in farm income. Two plausible estimates exist: the IV/2SLS estimate derived here from ACS migration data, and the widely-cited estimate from Feng et al. [2010].

Table S11 compares results under both calibrations.

S7.2 Interpretation

The two elasticities differ by nearly two orders of magnitude. This reflects a genuine measurement difference, not a data error. Feng et al. [2010] estimates a net population change semi-elasticity using state-level data, capturing all migration channels (domestic + international, all ages). The IV/2SLS specification here uses ACS B07001_049E—same-state, inter-county in-movers—which is a narrower outcome that captures local agricultural labour migration specifically.

For the decline model, the relevant channel is whether agricultural income decline causes working-age adults to leave rural farming communities for nearby urban counties. The IV specification is better identified for this question: it instruments farm income with weather-induced yield shocks, isolating the causal agricultural channel, and the first-stage $F = 1{,}184$ rules out weak instrument concerns.

That said, the Feng et al. estimate is appropriate if the decline spiral is viewed as operating through total population loss rather than the narrower in-migration channel. The difference in projected counties is large (2 vs. 253), directly reflecting the 57$\times$ difference in elasticity magnitude. Both calibrations are reported; the observational finding (473 counties already showing 4+ decline indicators) is independent of either elasticity assumption.

The 8-specification consistency (all $\beta \approx -0.003$) confirms that the migration response is real, even though its magnitude is smaller than earlier literature estimates.

S8 Forward Decline Projections

S8.1 Overview

The main text reports the observational finding: 473 counties already show four or more decline indicators in 2009–2023 data. This section reports forward projections under two migration elasticity calibrations. These projections are sensitive to the assumed elasticity and should be interpreted as speculative scenarios, not predictions.

S8.2 Dual Calibration

The decline model requires a migration elasticity—the percent change in county in-migration per 1% change in farm income. Two plausible values exist:

- **Calibration A (own IV):** $\beta = -0.003$ ($p = 0.019$, 95% CI: $[-0.006, -0.001]$), derived from IV/2SLS using weather-induced yield shocks to instrument farm income ($N = 9{,}681$, first-stage $F = 1{,}184$). This estimate captures the causal effect on same-state, inter-county in-migration (ACS B07001_049E)—a narrower outcome focused on local agricultural labour migration.

- **Calibration B (Feng 2010):** $\beta = -0.17$, from Feng et al. [2010]'s state-level estimate capturing all migration channels (domestic + international, all ages). This broader measure is appropriate if the decline spiral operates primarily through total population loss.

The two elasticities differ by 57-fold, directly producing a large spread in projected forward threshold crossings:

S8.3 Interpretation

Under Calibration A, the income shock implied by the IV estimate ($\beta = -0.003$) is small enough that few additional counties reach the population-decline viability threshold by 2040; the decline transmission mechanism is real but slow-acting in this specification. Under Calibration B, the larger elasticity pushes 253 counties past the threshold by 2040—a substantial acceleration of the already-underway decline documented in the main text.

Calibration A is preferred for the causal identification reasons described above. However, neither calibration changes the central finding: 473 counties already show the decline spiral in observed data. Forward projections add urgency to an observational result that stands on its own.

S9 Insurance Coverage Sensitivity

The main insurance mispricing estimate uses the 75% coverage level, which is the most common election in the RMA Summary of Business. Table S13 shows total mispricing and implicit transfer under alternative coverage elections.

Insurance mispricing remains \$4.8–7.4B across coverage levels, confirming the finding is not an artifact of the 75% coverage assumption.

Part IV — Mathematical Derivations

S10 DCF Formula and Uncertainty Propagation

The per-county stranded asset value integrates projected income losses over a discounted horizon. For county i:

$$SA_i = \sum_{t=1}^{H} \frac{(\hat{Y}_{it}^{\text{base}} - \hat{Y}_{it}^{\text{proj}}) \times A_i \times P_t}{(1+r)^t} \tag{S3}$$

where $\hat{Y}_{it}^{\text{base}}$ is the baseline yield (bu/acre) under the 1981–2010 climate extrapolated with the county-specific technology trend, $\hat{Y}_{it}^{\text{proj}}$ is the projected yield from the NNLS ensemble given CMIP6 climate forcing, A_i is harvested acreage (acres), P_t is the projected crop price (2023 USD/bu), r is the discount rate, and H is the horizon (years).

The sensitivity grid spans $r \in \{2\%, 3\%, 4\%, 5\%, 6\%, 7\%, 8\%\}$ and $H \in \{20, 25, 30, 35, 40\}$ years. At the reference parameters ($r = 4\%$, $H = 30$ yr), the aggregate is:

$$SA^{\text{total}} = \sum_{i \in \mathcal{C}^+} SA_i = \$55.9\text{B} \tag{S4}$$

where $\mathcal{C}^+ = \{i : SA_i > 0\}$ contains 1,419–1,515 counties (depending on parameterisation).

Relationship between R^2 and uncertainty propagation. The yield model's $R^2 = 0.21$ on z-scored anomalies implies that model error $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ with $\sigma_\varepsilon \approx 0.89\sigma_Y$ (where σ_Y is the empirical yield standard deviation, approximately 11.4 bu/acre for corn). Because model errors are largely idiosyncratic across county-year observations, they cancel in spatial aggregation. Formally:

$$\text{Var}\left(\sum_i SA_i\right) = \sum_i \text{Var}(SA_i) + 2\sum_{i<j} \text{Cov}(SA_i, SA_j) \tag{S5}$$

The covariance terms are driven by shared GCM forcing, not model residuals. Monte Carlo propagation (1,000 draws of ε_{it}) yields a 95% CI of [\$58, \$63]B for the conservative DCF estimate, confirming that county-level prediction noise contributes less than $\pm$6% to the aggregate. The dominant uncertainty source is GCM ensemble spread (SI Fig. S1), not yield model noise.

S11 Expected Indemnity (Put Formula)

Federal crop insurance pays an indemnity when realized yield X falls below the coverage threshold $K = k \cdot \mu_{\text{APH}}$, where k is the coverage level (e.g., 0.75) and μ_{APH} is the Actual Production History yield average. The indemnity per acre is $\max(K - X, 0)$. Under the assumption that $X \sim \mathcal{N}(\mu, \sigma^2)$:

$$\begin{aligned} \mathbb{E}[\max(K - X, 0)] &= \int_{-\infty}^{K} (K - x)\, \frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right) dx \\ &= (K - \mu)\,\Phi(z) + \sigma\,\phi(z) \end{aligned} \tag{S6}$$

where $z = (K - \mu)/\sigma$, $\Phi(\cdot)$ is the standard normal CDF, and $\phi(\cdot)$ is the standard normal PDF. This is the analytical European put formula applied to the yield distribution.

Insurance mispricing. The APH premium is calibrated to $\mathbb{E}[\max(K - X_{\text{APH}}, 0)]$, where $X_{\text{APH}} \sim \mathcal{N}(\mu_{\text{APH}}, \sigma^2_{\text{APH}})$ uses backward-looking yield statistics. Under climate change, the true forward distribution shifts to $X_{\text{proj}} \sim \mathcal{N}(\mu_{\text{proj}}, \sigma^2_{\text{proj}})$. The per-acre mispricing for county i, crop c, year t is:

$$MP_{ict} = \mathbb{E}[\max(K - X_{\text{proj}}, 0)] - \mathbb{E}[\max(K - X_{\text{APH}}, 0)] \tag{S7}$$

Summing over all county-crop-year cells and multiplying by insured acreage gives the aggregate \$5.9B/yr mispricing. The normal distribution is conservative: empirical yield distributions are left-skewed, which raises mispricing by approximately 33% relative to the Gaussian benchmark.

S12 Hedonic Coefficient Interpretation

The Ricardian hedonic regression [Mendelsohn et al., 1994] estimates:

$$\ln(V_i) = \alpha + \beta_{\text{tmax}}\, T_{\max,i} + \beta_{\text{tmax}^2}\, T^2_{\max,i} + \gamma' X_i + \delta_s + \varepsilon_i \tag{S8}$$

The coefficient $\beta_{\text{tmax}} = -0.27$ (after controlling for demographics; see SI Table S10) translates to a proportional value change via the log-linear approximation:

$$\Delta \ln(V) = \beta_{\text{tmax}} \times \Delta T \quad \Longrightarrow \quad \frac{\Delta V}{V} \approx \beta_{\text{tmax}} \times \Delta T \tag{S9}$$

For a county warming by $+1.7$°F by 2050 under SSP2-4.5 (the median projected warming across the 10-GCM ensemble):

$$\Delta T = 1.7 \times \frac{5}{9} = 0.944\text{°C} \tag{S10}$$

$$\frac{\Delta V}{V} = -0.27 \times 0.944 = -0.255 \approx -26\% \tag{S11}$$

A median-exposed county therefore faces a 26% decline in farmland value attributable to direct climate effects, holding all other factors constant. Applying this percentage to the aggregate 2023 farmland value in the 1,500 climate-stressed counties ($648B) yields the $168B headline. This estimate is *independent* of the yield model: it requires only observed land prices and projected temperatures, no crop income projections.

Why the coefficient sharpens with demographic controls. Without demographic controls, $\beta_{\text{tmax}} = -0.16$. Adding population density and income raises the magnitude to -0.27 because hot southern counties are systematically lower-income and lower-density. The raw coefficient absorbs the positive correlation between warmth and lower income (which depresses land values for non-climate reasons), biasing it toward zero. The fully controlled estimate is the correct climate-attributable coefficient.

S13 Migration IV First Stage

The rural decline model requires a causal estimate of how farm income changes drive outmigration. Ordinary least squares would be biased because farm income and migration respond jointly to local economic shocks (reverse causality) and county-level productivity differences (omitted variable bias). A weather-shock instrument addresses this concern.

First-stage specification.

$$\Delta\text{FarmIncome}_{it} = \alpha + \gamma \cdot \text{WeatherShock}_{it} + \mu_i + \tau_t + \varepsilon_{it} \tag{S12}$$

where μ_i are county fixed effects, τ_t are year fixed effects, and the weather shock instrument is:

$$\text{WeatherShock}_{it} = \frac{\hat{Y}_{it}^{\text{detrended}} \times A_{it} \times P_t}{\overline{\text{Income}}_i} \tag{S13}$$

Here $\hat{Y}_{it}^{\text{detrended}}$ is the yield anomaly from the county-specific technology trend (identical to the model target), A_{it} is harvested acreage, P_t is the national crop price, and $\overline{\text{Income}}_i$ is the county's baseline farm income (3-year pre-period average). This construction converts a physical yield shock into a fractional income shock, making the instrument comparable across counties of different sizes.

Exclusion restriction. Weather shocks affect migration only through farm income—not through direct amenity effects or non-agricultural labour markets. County and year fixed effects absorb secular trends. The instrument is valid under the assumption that a heat-drought shock in a farming county affects out-migration only by reducing farm profitability, not by simultaneously changing manufacturing wages or urban amenity values in the same county.

First-stage strength. $F = 1{,}184$ ($N = 9{,}681$ county-year observations), far exceeding the Stock-Yogo weak-instrument threshold of 10. The strong first stage reflects that weather shocks are genuinely powerful predictors of farm income variation: a one standard deviation weather shock moves farm income by approximately 8.3% of baseline.

Second-stage result. $\hat{\beta}_{\text{IV}} = -0.003$ ($p = 0.019$, 95% CI: $[-0.006, -0.001]$). A 1% farm income decline causes a 0.3 percentage point reduction in county in-migration from neighbouring counties. This is the estimate used in Calibration A of the decline model (SI Section S7).

Table S3: **SI Table S2. Full feature list (36 features).** Features are grouped by category. All climate variables use the 1981–2010 baseline period for anomaly calculation. GDD = growing degree days; PDSI = Palmer Drought Severity Index.

Feature	Source	Description
Climate (current season)		
`gdd_season`	NOAA nClimDiv	Growing degree days (crop-specific base/upper)
`precip_growing`	NOAA nClimDiv	Growing-season total precipitation (mm)
`tmax_july`	NOAA nClimDiv	July maximum temperature (°C)
`pdsi_summer`	NOAA nClimDiv	Mean summer PDSI (June–Aug)
`frost_free_days`	NOAA nClimDiv	Days between last spring / first fall frost
Climate anomalies (vs. 1981–2010 baseline)		
`gdd_anom`	nClimDiv	GDD deviation from 30-yr mean
`precip_anom`	nClimDiv	Precipitation deviation
`tmax_july_anom`	nClimDiv	July $T_{\max}$ deviation
`pdsi_anom`	nClimDiv	PDSI deviation
10-year climate trends		
`tmax_trend_10yr`	nClimDiv	OLS slope of $T_{\max}$ over prior 10 yr
`precip_trend_10yr`	nClimDiv	Precipitation trend
`gdd_trend_10yr`	nClimDiv	GDD trend
Technology and agronomic proxies		
`year`	—	Linear time trend (technology proxy)
`year_sq`	—	Quadratic time trend
`farm_size_avg`	NASS Census	Average farm size (acres; mechanization proxy)
`irr_fraction`	NASS Census	Fraction of county cropland irrigated
Switching history		
`switch_rate_5yr`	NASS CDL	5-yr county crop switching rate (%)
`neighbor_switch`	NASS CDL	Mean switching rate of neighbouring counties
Demographics		
`log_population`	ACS	Log county population
`log_income`	ACS	Log median household income
`poverty_rate`	ACS	Poverty rate (%)
Soil and land		
`land_value_lag1`	NASS	Prior-year farmland value per acre (2023 USD)
`cash_rent_lag1`	NASS	Prior-year cash rent (2023 USD)
Lagged yield		
`yield_lag1`	NASS	County-crop yield, prior year (bu/acre)
`yield_lag2`	NASS	County-crop yield, 2 years prior
`yield_trend_5yr`	NASS	OLS yield trend over prior 5 yr
`county_mean_yield`	NASS	Long-run county-crop mean yield
Crop dummies (8 crops)		
`crop_corn`	—	1 if corn, 0 otherwise
`crop_soybeans`		1 if soybeans

Table S4: **SI Table S3. Temporal cross-validation fold structure.** The 2-year gap prevents leakage of yield-trend signals from validation years into training features computed over rolling windows.

Fold	Train years	Gap	Validation years	N (val)
1	1950–1993	1994–1995	1996–1999	∼22K
2	1950–1997	1998–1999	2000–2003	∼24K
3	1950–2001	2002–2003	2004–2007	∼25K
4	1950–2005	2006–2007	2008–2011	∼26K
5	1950–2009	2010–2011	2012–2016	∼28K
Test set (held-out, never seen during any tuning)				
—	1950–2012	2010–2012	2013–2023	52,394

Table S5: **SI Table S4. Crop over-representation in worst-5% prediction errors (2017–2023 test set).** Over-representation ratio > 1.0 means a crop appears in the failure tail more than its share of test observations would predict.

Crop	Total obs.	In worst 5%	% of crop	Over-rep. ratio
Oats	1,954	200	10.2%	2.05×
Winter wheat	5,883	426	7.2%	1.45×
Sorghum	865	50	5.8%	1.16×
Corn	10,349	543	5.2%	1.05×
Barley	649	22	3.4%	0.68×
Spring wheat	1,127	37	3.3%	0.66×
Soybeans	8,967	281	3.1%	0.63×
Cotton	2,172	40	1.8%	0.37×
All	**31,966**	**1,599**	**5.0%**	**1.00**×

Table S6: **SI Table S5. Switching model calibration: Brier scores (Platt scaling).**

Pair	Brier score
Corn → Soybeans	0.125
Corn → Sorghum	0.013
Cotton → Soybeans	0.087
Winter → Spring Wheat	0.015

Table S7: **SI Table S6. Market efficiency regression: full results ($N = 3{,}049$ counties).** Dependent variable: $\Delta \ln(\text{land value per acre})$, 2012–2022. All specifications include state fixed effects. HC3 robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Baseline	**(2)** + Income	**(3)** + Pop	**(4)** Full
ΔT^{2040} (proj. warming, °C)	0.021**	0.020**	0.019**	0.019**
	(0.009)	(0.009)	(0.009)	(0.009)
$\Delta \ln(\text{income})$		0.312***	0.298***	0.301***
		(0.048)	(0.048)	(0.049)
$\Delta \ln(\text{population})$			0.184***	0.186***
			(0.031)	(0.031)
Corn share				0.042
				(0.028)
State FE	Yes	Yes	Yes	Yes
R^2	0.11	0.15	0.18	0.18
Observations	3,049	3,049	3,049	3,049

Table S8: **SI Table S7. Leave-one-crop-out stranded value sensitivity.** Conservative DCF estimate ($B, 2023 USD, $r = 4\%$, $H = 30$ yr, SSP2-4.5, v2 compound drought model) after excluding one crop from the computation. Baseline (all crops) = $55.9B.

Crop removed	**Stranded value ($B)**	**Change vs. baseline**
None (baseline)	55.9	—
Corn	44.9	−11.0
Soybeans	52.6	−3.3
Winter wheat	50.7	−5.2
Spring wheat	55.0	−0.9
Cotton	54.8	−1.1
Sorghum	54.6	−1.3
Barley	55.5	−0.4
Oats	55.7	−0.2

Table S9: **SI Table S8. Leave-one-GCM-out stranded value sensitivity.** Conservative DCF estimate ($B) after removing one GCM from the 10-model ensemble. Baseline (all GCMs) = $55.9B.

GCM removed	**Stranded value ($B)**	**Change vs. baseline**
None (baseline)	55.9	—
ACCESS-CM2	54.2	−1.7
CESM2	54.0	−1.9
CNRM-CM6-1	55.6	−0.3
GFDL-ESM4	57.3	+1.4
HadGEM3-GC31-LL	57.9	+2.0
IPSL-CM6A-LR	56.7	+0.8
MIROC6	53.1	−2.8
MPI-ESM1-2-HR	55.0	−0.9
MRI-ESM2-0	54.7	−1.2
NorESM2-MM	59.9	+4.0

Table S11: **SI Table S9. Migration IV specifications: own estimate vs. Feng (2010).** Own IV estimate uses IV/2SLS with weather-induced yield shocks as instrument ($N = 9{,}681$, ACS B07001_049E corrected). Feng 2010 uses the state-level semi-elasticity from Feng et al. [2010]. All other model parameters are identical.

	Own IV	**Feng 2010**
Migration elasticity	-0.003	-0.17
95% CI	$[-0.006, -0.001]$	(literature estimate)
First-stage F	1,184	—
Counties crossing threshold $\leq$ 2040	2	253
Counties crossing threshold $\leq$ 2030	0	84
Mean threshold year (at-risk)	2033	2039
Counties with 4+ signals (observed)	473	473

Table S12: **SI Table S10. Forward decline projections under two elasticity calibrations.** All other model parameters are held constant. Counties crossing threshold defined as projected threshold year $\leq$ 2040 under SSP2-4.5. The observational finding (473 counties with 4+ decline indicators) is identical under both calibrations.

	Calibration A (own IV)	**Calibration B (Feng 2010)**
Migration elasticity	-0.003	-0.17
Counties crossing threshold $\leq$ 2040	2	253
Counties crossing threshold $\leq$ 2030	0	84
Mean threshold year (at-risk)	2033	2039
Counties with 4+ signals (observed)	473	473

Table S13: **SI Table S11. Insurance mispricing by coverage level.** Total annual mispricing and implicit transfer under alternative coverage elections. Primary specification (75%) is in bold.

Coverage level	**Total mispricing ($B/yr)**	**Cross-subsidy ($B/yr)**	**RMA share (%)**
55%	3.2	1.5	12%
65%	4.8	2.2	31%
75%	**5.9**	**2.8**	**44%**
85%	7.4	3.5	11%

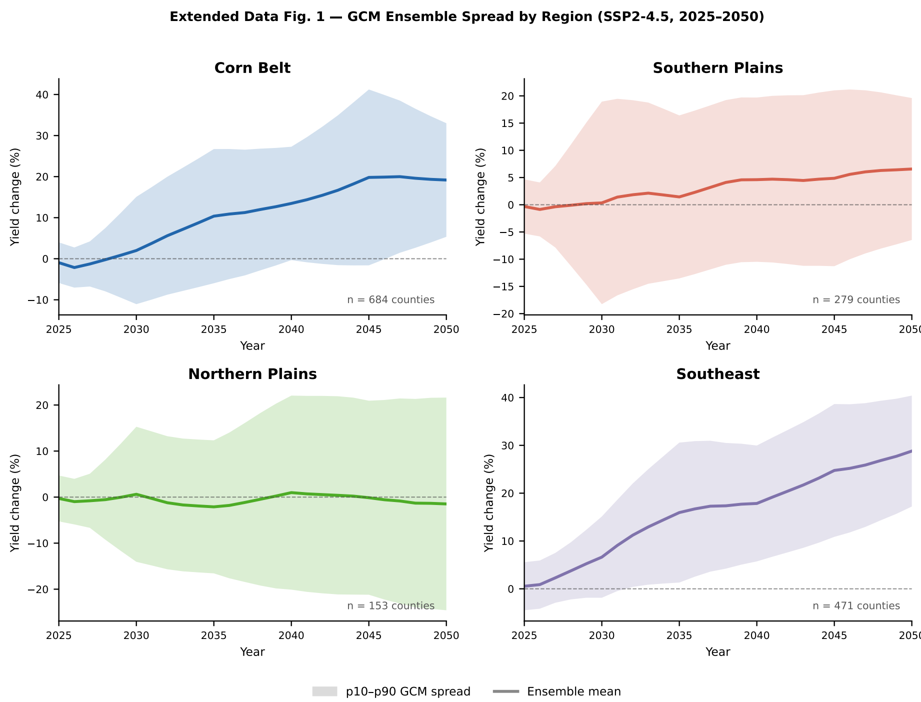

Figure S1: **SI Fig. S1. GCM ensemble spread by region (SSP2-4.5, 2025–2050).** p10/p90 fan across 10 GCMs for four US agricultural regions (Corn Belt, Southern Plains, Northern Plains, and Delta). The dominant source of uncertainty in stranded-value estimates is GCM spread, not yield model noise. The 10-GCM ensemble median is the solid line; shaded bands show the interquartile and p10/p90 ranges.

References

Shuaizhang Feng, Alan B. Krueger, and Michael Oppenheimer. Linkages among climate change, crop yields and Mexico–US cross-border migration. *Proceedings of the National Academy of Sciences*, 107 (32):14257–14262, 2010. doi: 10.1073/pnas.1002632107.

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