NBER WORKING PAPER SERIES

# DYNAMIC SELECTION AND THE PUZZLE OF SLOW CLIMATE CHANGE ADAPTATION IN U.S. AGRICULTURE

Robert Huang
Matthew E. Kahn

Working Paper 34771
http://www.nber.org/papers/w34771

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We thank Barry Goodwin, Jennifer Ifft, Juan Moreno-Cruz, Aaron Smith, and Jan Stuhler for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Dynamic Selection and the Puzzle of Slow Climate Change Adaptation in U.S. Agriculture
Robert Huang and Matthew E. Kahn
NBER Working Paper No. 34771
January 2026
JEL No. Q1, Q54

**ABSTRACT**

Long-run studies of agricultural productivity use balanced county-level panel data to document slow adaptation to climate change in U.S. agriculture. We revisit this treatment-effect puzzle using a selection model that incorporates the optimizing behavior of heterogeneous farmers. Using data from 1980 to 2020, we document three dynamic selection effects. First, counties that completely exit farming alter the composition of the balanced panel sample. Second, higher-skilled workers are more likely to leave rural areas as weather conditions worsen. Third, smaller farms consolidate into larger farms in counties affected by climate change. We document that the yield damage function's slope depends on the county's educational level and farm size. Causal estimates of weather impacts must account for selection effects along several adaptation margins of adjustment.

Robert Huang
Stanford University
Department of Economics
huangrob@stanford.edu

Matthew E. Kahn
University of Southern California
Department of Economics
and NBER
kahnme@usc.edu

## 1. Introduction

Using county-level panel data and a long-difference econometric framework, Burke and Emerick (2016) document an important puzzle: U.S. agriculture shows limited adaptation to rising heat, as benchmarked by the persistent sensitivity of corn and soy yields to extreme heat over the period 1950–2000. An extensive literature has used region aggregate data with balanced panel econometric models to study the effects of climate change on agricultural productivity (Schlenker and Roberts 2009; Tack, Barkley, and Nalley 2015; Ortiz-Bobea et al. 2021), mortality (Deschênes and Greenstone 2011), economic growth (Dell, Jones, and Olken 2012), and many other socioeconomic variables (Ranson 2014; Burke et al. 2024). More broadly, similar methods have been used to study the impacts of macro shocks such as war (Federle et al. 2024), trade (Autor, Dorn, and Hanson 2016), and technology (Acemoglu and Restrepo 2020) on economic outcomes.

This place-based balanced panel approach implicitly assumes away the importance of heterogenous workers' optimizing choices, failing to incorporate the logic of opportunity cost. For example, if climate change lowers agricultural yields and rural worker earnings, high-skilled farmers (who have better nonfarm opportunities) have an incentive to exit completely or partially from agriculture. The remaining farming population becomes less skilled, and the decline in crop yields would have been smaller in magnitude without this change in population composition.

Labor economists studying trends in the earnings gender gap have emphasized the importance of accounting for the selection issue of which women choose to work (Heckman and Honore 1990; Goldin 2014). They have been able to utilize micro longitudinal methods because they have had access to individual level panel data.

Yet, agricultural researchers seeking to study how long run climate change impacts cannot access similar data that matches land parcels and which managers own the land and which workers work on that land. Costinot, Donaldson, and Smith (2016) have compiled a field level dataset similar to what we describe above, but they do not have a long panel and do not observe characteristics of each parcel's owner. Faced with this missing data problem, agricultural researchers often adopt a representative-agent framework, which rules out composition shifts driven by sorting at the extensive margin.

Our paper focuses on the allocation of human capital across space as predictable weather trends play out. We study the implications of three selection margins—counties exiting farming, skilled rural workers moving away, and land consolidation—for interpreting the causal estimates of weather based on county-level panel data. A growing literature examines what economic parameters can be recovered from panel regressions with aggregate data. Ortiz-Bobea and Just (2013) underscore the need to structurally model adaptation and decompose the mechanisms driving observed crop yields. Ortiz-Bobea (2020) shows that the reduced-form Ricardian approach tends to overstate climate damages by ignoring nonfarm factors, such as urban development. Jones et al. (2026) show that panel regressions can generate spurious U-shaped damage functions when differential weather trends are correlated with a place's baseline temperature. We contribute to this literature by emphasizing the role of selection. Our work is in a similar spirit to Dustmann et al. (2025), who document that regional effects of immigration shocks estimated using repeated cross sections conflate pure supply-and-demand responses with compositional changes, as natives selectively out-migrate from affected regions or switch occupations.

We start by providing a micro foundation optimizing framework that helps to explain the marginal effect of heat on county level aggregate dynamics. We simulate agricultural economies in which different areas face distinct weather patterns, and heterogeneous farms optimize accordingly. Within each simulated county, we solve for the optimal exit decision, land input, and adaptation technology take-up of each farm. We aggregate up their choices to the county level to create synthetic data that we use to estimate the long difference regression commonly employed in the literature.

We argue that the long-difference coefficients represent the net effect of climate change, adaptation, and selection. Burke and Emerick (2016) conclude that U.S. agriculture has seen little adaptation because long-difference estimates do not statistically differ from first-difference estimates, which approximate the counterfactual marginal effect of warming without adaptation.[1] Our simulation exercise highlights that the gap between these estimates should be interpreted as the combined effect of adaptation and selection. Adaptation includes adopting new technologies or crop varieties, while selection arises when heterogeneous farmers decide whether to continue farming and how much land to use. When high-skilled farmers exit the sector (negative place-based selection), a more negative long-difference coefficient may not indicate a lack of

[1] Their rationale is that no adaptation can take place on an annual basis, so adaptation can be benchmarked by whether the long-difference estimates are smaller than the first-difference estimates.

adaptation.[2] As people adapt to the changing rural weather, the place (based on aggregate productivity yields) appears to not be adapting (Hornbeck 2012).

In our empirical sections, we provide descriptive evidence on different margins of selection. About 25% of the corn-producing counties in 1980 have stopped growing corn, and about 15% of the soy-producing counties have ceased soybean production. Using county-decade panel data from 1980 to 2020, we document that counties are more likely to stop producing corn or soy if they experience more extreme heat, if the county population is more educated, and if the soil quality is lower. By combining census data with natural disaster data, we show that workers, especially the educated ones, are moving away from farming counties experiencing more disasters. The opportunity cost of locating in a rural area is to urbanize. Urbanization and human capital are complements, and climate change accelerates the rural "brain drain" (Glaeser 2012).

An additional selection case offers a very different hypothesis. Comparative advantage logic predicts that farmers differ in their ability to cope with new weather conditions. If high ability farmers have an edge, then they will buy the land of farmers who fail to adapt (Sumner 2014; Bekkerman, Belasco, and Smith 2019). Such "superstar managers" have a greater incentive to build up skill and to invest in equipment if they control more land (scale economies). We document that the distribution of farm sizes has become more right-skewed in dry and warming areas, and the proportion of farms exceeding 1000 acres has risen significantly. Our discussion here highlights that selection effects can move in opposite directions. While this complicates the interpretation of aggregate findings, we contribute to understanding the mechanisms through which a key sector copes with a key emerging trend while investing under uncertainty.

In the final section of the paper, we estimate new versions of the classic Burke and Emerick (2016) long-difference regressions, updating their data sample through 2020. We find that the long difference coefficient has shrunk in the period 2000 to 2020. In addition, we partition counties based on their 1980 levels of educational attainment and average farm size, then estimate the long-difference specification on subsamples. For example, one of our data samples represents rural

[2] Burke and Emerick (2016) recognize that selection effects may explain part of their results. In appendix table 10, they use county data on the percent of farms in 1978 and 1997 with more than $20k in farmland equipment to test for selection effects. They do not explore the human capital hypothesis and do not use a Roy Model to investigate their aggregate dynamics. Our work and their work highlights that farmers and rural places have multiple margins of adjustment to increase their profits during times of weather stress. In the conclusion, we discuss even more margins of adjustment such as ecotourism. This high-dimensional set of strategies together determines whether a farmer makes an extensive margin decision to keep farming, and this determines how aggregate production changes over time.

counties with a less educated population and more small farms. We posit the overall selection effect is more negative in these counties. We report evidence consistent with the dynamic selection hypothesis. On net, counties experiencing brain drain and not undergoing farm consolidation suffer a greater marginal yield loss from negative climate trends.

This paper is organized as follows. Sections 2 and 3 present a simulation example to illustrate our points on dynamic selection. Section 4 describes our primary data sources. Section 5 provides descriptive evidence of the three margins of selection. Section 6 presents our summary regressions for testing how rural economy extensive margin decisions influence estimates of reduced form long difference regressions.

## 2. Optimizing Farmer Behavior Generates County Level Aggregate Agricultural Output Data

The reallocation of heterogeneous labor between the farming and non-farming sectors, as well as the reallocation of land across heterogeneous farmers, affects the data-generating process that agricultural economists use to benchmark place-based trends in agricultural productivity. For example, if the least productive farmers are more likely to sell their land for other uses, then the average human capital of farmers increases over time. In this case, an empirical researcher who lacks access to such dynamic manager-quality data will systematically underestimate the causal effect of weather due to this dynamic selection process. In this section, we present a stylized example of farms with heterogeneous managers to illustrate our points on selection.

### 2.1 Setup

The initial farming population has measure 1, and farmers have heterogenous ability $m \in [0,1]$ which follows a distribution $F$. The initial land endowment $L_0 = 1$ for every farmer. Farmers have the option to exit farming and earn a wage $w(m)$, an increasing function of m, in the non-agriculture sector. Farmers can sell (part or all of) their land at price $p$, the market-clearing price in the land market. The agricultural output price is normalized to 1. Climate change reduces productivity by a proportion $d$, but for a fixed cost $C$, farmers can invest in an adaptation technology that mitigates this negative effect.

The agricultural production function is given as follows:

$$Y(L, I; m) = m(1 - (1 - \phi I)d)L^{\alpha}$$

where I is an indicator of whether the farmer invests in adaptation; $\phi \in [0,1]$ is adaptation effectiveness; and $\alpha \in (0,1)$ so that there is diminishing return to land input $L$. Farmers (who continue farming) choose $L$ and $I$ to maximize the profit function:

$$\pi(L, I; m, p) = Y(L, I; m) - p(L-1) - CI$$

subject to the constraint $L > 0$ and $I \in \{0,1\}$. Denote $\pi^{stay}(m) \equiv argmax_{L,I}\, \pi(L, I; m, p)$, the maximum profit that can be attained by farmers with ability m if they choose to stay and the equilibrium land price is p. In what follows, we consider the equilibrium with and without climate change respectively.

### 2.2 Equilibrium without climate change

In the baseline case without climate change, farmers choose whether to continue farming, and if so, the land input. If they stay, the land demand can be found from the first order condition (FOC) of the profit function:

$$L(m; p) = (\frac{m\alpha}{p})^{\frac{1}{1-\alpha}}$$

which is always positive (if m>0) and increasing in m. Based on the FOC, we can rewrite $mL^{\alpha} = \frac{p}{\alpha}L$ when L is the optimal land input. The profit from staying is:

$$\pi^{stay}(m) = p + \frac{p(1-\alpha)}{\alpha}L(m; p)$$

which equals the revenue from selling their crops netting the cost of acquiring new land (or the revenue from selling part of their land). They stay if $\pi^{stay}(m) > p + w(m)$, meaning the income from farming is higher than the revenue from selling all their land plus the outside wage. We make the following assumption:

**Assumption (Complementarity between human capital and non-farming sectors)**

At the equilibrium land price, the function $g(m) = \pi^{stay}(m) - p - w(m)$ is decreasing in m, and there exists a unique $m^*$ such that $g(m^*) = 0$. This implies $m^*$ is the cutoff such that farmers with ability $m > m^*$ exit. The interpretation is that the opportunity cost of farming is higher for farmers with higher human capital.

The equilibrium is characterized by the cutoff $m^*$ and the land price $p^*$ such that the following two equations jointly hold:

$$\frac{p^*(1-\alpha)}{\alpha} L(m^*; p^*) - w(m^*) = 0 \tag{1A}$$

$$\int_0^{m^*} L(m; p^*) dF(m) = 1 \tag{1B}$$

where the land demand L is solved from the FOC of the farmer's profit-maximizing problem. Equation (1A) means the marginal farmer is indifferent between staying and exiting. Equation (1B) is the land market clearing condition. In this motivating example, we assume (and check in our simulations) that the equilibrium prices are positive. In reality, they may be zero and the land is reallocated out of farming or left idle.

### 2.3 Equilibrium with climate change

As before, we can derive farmers' land demand and profit conditional on the adaptation decision $I$. Denote $A(I) \equiv 1 - (1 - \phi I)d$.

$$L(m, I; p) = \left(\frac{mA(I)\alpha}{p}\right)^{\frac{1}{1-\alpha}}$$

$$\pi^{stay}(m, I) = p + \frac{p(1-\alpha)}{\alpha} L(m, I; p) - CI$$

and farmers adapt if $\pi^{stay}(m, 1) > \pi^{stay}(m, 0)$, equivalent to $\frac{p(1-\alpha)}{\alpha}\big(L(m, 1; p) - L(m, 0; p)\big) > C$. Because the LHS is an increasing function of m, there exists a unique cutoff $m^A$ such that farmers with ability above it choose to adapt if they stay. Then the profit when staying $\pi^{stay}(m) = \pi^{stay}(m, 1)$ if $m > m^A$ and $\pi^{stay}(m, 0)$ otherwise.

The equilibrium is characterized by the adaptation cutoff $m^A$, the exit cutoff $m^E$, and the land price $p^C$ such that the following equations jointly hold:

$$\frac{p^C(1-\alpha)}{\alpha}\big(L(m^A, 1; p^C) - L(m^A, 0; p^C)\big) = C \tag{2A}$$

If $m^E > m^A$:

$$\frac{p^C(1-\alpha)}{\alpha} L(m^E, 1; p^C) - C - w(m^E) = 0$$

$$\int_0^{m^A} L(m, 0; p^C) dF(m) + \int_{m^A}^{m^E} L(m, 1; p^C) dF(m) = 1$$

(2B)

If $m^E \leq m^A$:

$$\frac{p^C(1-\alpha)}{\alpha} L(m^E, 0; p^C) - w(m^E) = 0$$

$$\int_0^{m^E} L(m, 0; p^C) dF(m) = 1$$

(2C)

Equations (2B) and (2C) refer to the indifference of the marginal farmer and the land market clearing condition as before. The only addition is equation (2A), which states that farmers at the adaptation cutoff are indifferent about adopting the technology.

Under our assumption of complementarity between human capital and non-farming sectors, we show that the post-climate change exit cutoff and land prices are both lower than their pre-climate change equilibrium levels (see Appendix 1). Intuitively, the average farming population has lower skill levels following farmers' reoptimizing behavior induced by climate change.

**2.4 A calibration example**

We now use the model to quantify the interplay between farm size and manager quality as climate change unfolds. In the following calibration example, we demonstrate that observed changes in crop yields depend on the net effect of climate change, adaptation, and selection. If the overall effect of selection is negative, a simple test based on the change in yields may understate adaptation progress.

We calibrate the production function parameter $\alpha = 0.5$, the distribution of ability $F \sim U[0,1]$, and the non-agricultural wage function $w(m) = 2m^2 - 0.3$. The rationale is that farmers need to pay a fixed cost (e.g. get more training) when switching to another sector. Using these values, we solve for the exit cutoff $m^* = 0.719$ and the equilibrium land price $p^* = 0.176$ without climate change. In this benchmark case, the average crop yield (production per unit of land) is 0.352. The average farmer has ability 0.36, and the average farm is operated by farmers with ability 0.54.

We then introduce climate damage $d = 0.5$, adaptation effectiveness $\phi = 0.5$, and adaptation cost $C = 0.25$. Solving the model yields an adaptation cutoff $m^A = 0.493$, an exit cutoff $m^E = 0.575$, and a land price $p^C = 0.076$. Crop yield falls to 0.152, and the average farm is owned by farmers with ability 0.464. These simulation results align with our theoretical predictions: higher-ability farmers are more likely to exit, and land becomes less valuable under climate change. Also, it highlights the farmland reallocation from lower-ability to higher-ability farmers—while the ability of the farming population decreases by 20%, it decreases only by 14% when weighted by farm sizes.[3]

In Figure 1, we plot the distributions of farm size and farm income before and after climate change. Several patterns emerge. First, the average farm size increases as remaining farmers acquire land from those who exit. Second, despite this consolidation, average farm income declines due to lower aggregate production. Third, both distributions become bimodal. Two groups of farmers remain: a lower-ability group that neither acquires additional land nor invests in adaptation, and a higher-ability group (the "superstar" farms) that purchases more land and undertakes costly adaptation.

This example shows two margins of selection: farmers exiting and land reallocation across farmers who stay. Although climate damage is calibrated to 50%, equilibrium crop yield declines by 57% (from 0.352 to 0.152). This is the net effect of climate change, adaptation, and selection through both exit and land consolidation. Adaptation and land reallocation to high-skilled farmers mitigate some of the damage, but the exit of high-ability farmers, who face high opportunity cost of farming, reduces overall productivity.

To decompose these effects, we simulate a counterfactual in which adaptation technology is available, but no land can be reallocated and farmers cannot quit farming. This means farmers only decide whether to take up the technology, keeping the land input as in the pre-climate change equilibrium. In this case, crop yield declines to 0.215 (a 39% decrease from the benchmark), higher than the equilibrium yield in the presence of selection (i.e. 0.152). That is, adaptation mitigates the negative effect of climate change on yields, but it is offset by negative selection as high-skilled farmers exit.[4]

---

[3] We have also checked that the assumption of complementarity between human capital and non-farming sectors holds in all simulated equilibria. In the pre climate change equilibrium, $g(m) = -0.58m^2 + 0.3$ and in the post climate change equilibrium $g(m) = -0.15m^2 + 0.05$. Both are decreasing in m on [0,1].
We repeat this simulation exercise for different levels of damage. In Appendix 1, we show that the exit cutoff and the equilibrium land price decrease as damage increases.

[4] Our stylized example has several limitations. First, we find an equilibrium within each county instead of a spatial equilibrium. We rule out the option for farmers to move to another county and continue to farm. Second, we assume a stationary production function and abstract away from modelling endogenous innovations in the agricultural sector (Moscona and Sastry 2023). If the technology frontier shifts out over time, this lowers the marginal cost of adaptation. Third, we assume that farmers can sell land only to other farmers, which is reasonable when zoning

## 3. Bridging Theory and Empirics

The ideal data for studying long-run agricultural adaptation to climate change would be geocoded parcel-level data indicating what commodity is grown on the land, the exogenous soil quality, the manager of the parcel, and the input of other production factors, such as labor, at each point in time for decades. If adjacent parcels are purchased, then a farmer would manage a larger plot of land. The researcher should also observe the weather realizations and crop yields on each plot at each point in time. With such long panel data, the researcher could study the productivity of farmland conditional on the dynamic selection that the parcel is controlled by a given farmer at that point in time (Olley and Pakes 1992). However, due to data limitations, researchers often use aggregate data such as county/year data by crop.

### 3.1 What do reduced-form aggregate climate damage functions estimate?

In the classic Nordhaus IAM climate change economics model, the damage function plays a central role (Nordhaus 2019). This equation maps temperature deviations into lost economic activity. Over the years, empirical economists have estimated many versions of this reduced form function to explain many outcome variables, including mortality rates, GNP dynamics, worker productivity, and farm productivity (Ortiz-Bobea et al. 2021; Burke et al. 2024).

In the absence of detailed microdata, empirical research in climate change economics often uses panel data at the geographic level and estimates similar versions of equation (3) on county i and period t:

$$\Delta Y_{it} = \beta * \Delta Temperature_{it} + FEs + \varepsilon_{it} \tag{3}$$

where Y is the outcome of interest such as crop yields; temperature is a measure of climate change and can be replaced by other variables such as heat and precipitation; and all variables are differenced from the base year to the end year. We argue that this reduced-form estimate captures the net effect of climate change and farmers' reoptimization behavior.

Farmers in our model (see Section 2) have optimized to maximize their profits given their circumstances and constraints, including switching out of agriculture or selling part of their land.

---

regulations, such as California's Williamson Act, restrict conversion of farmland to other uses. Lastly, the literature has documented moral hazard in agricultural production, as farmers do not bear the full cost of climate change (Annan and Schlenker 2015). We do not incorporate crop insurance programs in the example.

Such behavior may work in the same direction as climate change, for example when high-skilled farmers leave the sector. In this section, we discuss how to interpret the slope estimates from equation (3) in the presence of dynamic selection.

The stylized example in Section 2.4 illustrates that the observed change in crop yields ($\Delta Y_i$) reflects the net effect of three forces: the direct damage from climate change, adaptation, and extensive margin selection (i.e. land reallocation across farmers and labor reallocation out of agriculture). We formalize this decomposition below.

Consider a general production function $f(m, T, A, L)$ and a cost function $C(A, L; m, T)$, both of which depend on farmers' ability m, temperature T, adaptation input A, and land input L. Define the maximized profit for farmers of ability m given temperature T as $\pi(m, T) = max_{A,L}\ f(m, T, A, L) - C(A, L; m, T)$. Their production is $y(m, T) = f(m, T, A^*, L^*)$ where $A^*$ and $L^*$ are the solutions to the profit-maximization problem. Farmers produce if $\pi(m, T) \geq 0$ and the exit cutoff $m^E(T)$ is defined by $\pi\left(m^E(T), T\right) = 0$.

We normalize the total farmland area to 1 so that the total production is equivalent to the crop yield. Denote the density of ability by $g$. The total production can bet written as:

$$Y(T) = \int_0^{m^E(T)} y(m, T) g(m) dm$$

We can apply Leibniz rule to derive the marginal change in crop yields with respect to a unit increase in temperature, $\frac{dY}{dT}$. We can then decompose the derivative into the direct effect, adaptation effect, and selection effect:

$$\frac{dY}{dT}|_{T=T_0} = \int_0^{m^E} (f_T + f_A A_T^*) g(m) dm + \left(\int_0^{m^E} f_L L_T^* g(m) dm + \frac{dm^E}{dT} y(m^E, T_0) g(m^E)\right) \tag{4}$$

where the subscripts denote partial derivatives evaluated at $T_0$ and $m^E$ is the exit cutoff at $T_0$. In the first term, $f_T$ (often negative) represents the change in crop yields that would occur if farmers neither exited the sector nor adjusted their adaptation levels or land inputs, and $f_A A_T^*$ (often positive) refers to the change in yields due to adaptation. This whole term thus shows the net effect of climate change and adaptation. The balanced panel studies implicitly assume their estimates of

$\frac{dY}{dT}$ equals this net effect (so that it must be less negative when adaptation takes place). Yet, this implicitly assumes everything in the bracket adds up to zero.

The terms in the bracket illustrate two margins of selection, the land consolidation effect and the exit effect. In the first term, $f_L L_T^*$ is the change in yields as a farmer reoptimizes land input, and the integral gives a weighted average of this over the distribution of ability. The integral is positive if high-skilled farmers acquire land from low-skilled farmers as the economy transitions to the post climate change equilibrium. Intuitively, high-skilled farmers can produce more on each unit of land (all else equal), so the average land productivity increases when low-ability farmers sell their land to high-ability farmers. This is selection through land reallocation.

The second term in the bracket captures the change in crop yields resulting from the shifting composition of farmers. This equals the product of three components: the yield of marginal farmers, $y(m^E, T)$, times the density of farmers at the margin, $g(m^E)$, times the shift in the exit cutoff, $\frac{dm^E}{dT}$. If the marginal farmers are exiting ($\frac{dm^E}{dT} < 0$), then the remaining farmers become less-skilled and exit effect is negative. If the overall selection effect (sum in the bracket) is negative, the estimated marginal effect $\frac{dY}{dT}$ is more negative than $\int_0^{m^E} (f_T + f_A A_T) g(m) dm$ and the balanced panel approach would underestimate adaptation.[5]

When equation (3) is estimated with yearly differences rather than long differences, the coefficient is arguably closer to the case without any adjustment (i.e. $\int_0^{m^E} f_T g(m) dm$), since both adaptation and selection are more likely to occur over longer horizons. Studies such as Burke and Emerick (2016) implicitly assume no dynamic selection, positing that the long-difference coefficient should be less negative than the first-difference coefficient when farmers adapt. If we drop this assumption, as illustrated in equation (4), the deviation of the long-difference coefficient from the first-difference coefficient should be interpreted as the net effect of adaptation and

---

[5] The true relationship between $\Delta Y_i$ and $\Delta T_i$ is: $\Delta Y_i = Y_{i,t+1} - Y_{it} = \int_{T_{it}}^{T_{i,t+1}} \frac{dY}{dT}(t) dt = \Delta T_i \frac{dY}{dT}(t_i)$
where the last equation follows from mean value theorem for some $t_i \in (T_{it}, T_{i,t+1})$. The long-difference coefficient can be written as: $\hat{\beta} = \frac{\sum (\Delta T_i - \overline{\Delta T})^2 \frac{dY}{dT}(t_i)}{\sum (\Delta T_i - \overline{\Delta T})^2} \to \frac{E\left[(\Delta T - \overline{\Delta T})^2 \frac{dY}{dT}(t)\right]}{E[(\Delta T - \overline{\Delta T})^2]}$
It is a weighted average of the county-specific marginal effect given by equation (4), asymptotically approaching the expected weighted average over the joint distribution of the initial temperature $T_{it}$ and the temperature change $\Delta T_i$. Equation (4) suggests that this coefficient confounds three channels: direct productivity effects of climate change, adaptation, and selection.

selection. If the selection term is negative, a more negative long-difference coefficient may not imply a lack of adaptation. We provide a numerical illustration next.

### 3.2 A numerical illustration

This section extends the stylized example from Section 2.4. In Figure 2, we plot equilibrium crop yields under three scenarios: (a) both adaptation and selection (full effect); (b) adaptation only; and (c) no adaptation or selection. Because climate damage is modeled as a proportional reduction in productivity, scenario (c) corresponds to the 45-degree line: a $d\%$ damage translates directly into a $d\%$ damage directly translates into a $d\%$ decline in yields from the pre-climate change values. We benchmark scenarios (a) and (b) against this no-adjustment case.[6]

The figure reveals several salient patterns. In the absence of selection, the decline in crop yields is consistently smaller than under no adjustment. However, once farmers are allowed to trade land and exit, the full-equilibrium decline in yields (scenario a) is consistently larger than when there is no selection (scenario b). This negative selection implies that the remaining farming population is lower-skilled on average, producing less output from the same inputs. In fact, for damage levels below a certain threshold, the full-equilibrium decline exceeds even the no adjustment case, indicating that the selection effect dominates the adaptation effect. In such cases, if researchers use observed crop yields to assess adaptation while overlooking dynamic selection, they would overstate the negative impact of climate change on agricultural productivity and wrongly conclude that little or no adaptation has occurred.

We use a numerical example to illustrate that the long-difference coefficient depends on both selection and adaptation. We draw 1000 observations from the distribution of the initial damage $d_0 \sim N(0.2, 0.05)$ and the change in damage $\Delta d \sim N(0.1, 0.025)$. We simulate the initial equilibrium when $d = d_0$ and then the post climate change equilibrium when $d = d_0 + \Delta d$. We also simulate three counterfactual equilibria: one with no adjustment, one with adaptation only, and one with adaptation and land reallocation but no exit. We then estimate equation (3) using several dependent variables: changes in crop yields under each scenario, the change in average

[6] To simulate scenario (b), we assume farmers cannot adjust their land input or quit farming. They only decide whether to take up the adaptation technology.

farmer ability (weighted and unweighted by farm size), and the change in median farm size. We repeat this exercise for a higher damage $\Delta d \sim N(0.2, 0.05)$. The results are presented in Table 1.

Column (1) reports the elasticity of crop yields with respect to increased climate damage in the absence of any adjustment—the effect a first-difference estimator would approximate. Column (2) reports the elasticity with both selection and adaptation, which the long-difference estimator approximates. We also report the marginal effect of a percentage point increase in climate damage. In Panel A, the full effect (climate change, adaptation, and selection) is more negative than the pure climate change effect, but in Panel B, it is less negative. This indicates that negative selection dominates adaptation when climate change is less severe (because the return from investing in adaptation is lower). Now we decompose the full effect into adaptation, selection through land reallocation, and selection through exit.

The dependent variables in columns (3) and (4) represent counterfactual changes in crop yields: column (3) excludes both selection mechanisms (no exit and no land reallocation), while column (4) only excludes exit.[7] The difference between columns (1) and (3) captures the adaptation effect, the difference between columns (3) and (4) isolates the land reallocation effect, and the difference between columns (2) and (4) measures the exit effect. In both panels, adaptation has a positive effect, mitigating larger shares of damage when damages are higher (Panel B). In both cases, the land reallocation effect is positive, and the exit effect is negative. Both estimates are larger in magnitude when temperature rises more (Panel B). Without farmer exit, the long difference coefficient would be 40% to 50% less negative, underscoring exit as an important channel of selection in this particular example.

We provide further evidence on farmer exit and farm consolidation in columns (5) to (7). An increase in climate damage is negatively correlated with the average ability of farmers who stay. Since we assume no land exits farming (and equilibrium land prices in our example are positive), greater farmer exit implies larger average farm sizes. Specifically, each percentage point increase in damage is associated with a 0.48% increase in average farm size in Panel A and a 0.37% increase in Panel B. In column (6), the same increase in damage is correlated with a 0.06% decline in median farm size in Panel A and a 1.58% decrease in Panel B. This suggests under more

[7] We can write these dependent variables in the notation of equation (4). The dependent variable in the no adjustment case is $\int_T^{T+\Delta T} \int_0^{m^E} f_T g(m) dm\, dt$. In the adaptation only case it is $\int_T^{T+\Delta T} \int_0^{m^E} (f_T + f_A A_T^*) g(m) dm\, dt$. In the no exit case it is $\int_T^{T+\Delta T} \int_0^{m^E} (f_T + f_A A_T^* + f_L L_T^*) g(m) dm\, dt$.

severe climate change, the farm size distribution becomes significantly more right-skewed, indicating the emergence of large “superstar” farms, consistent with Figure 1.

In column (7), we study the joint effect of the two selection margins, namely how the farmer's ability on an average acreage (i.e., area-weighted ability) changes. This should decrease if top farmers exit, but increase if high-skilled farmers who remain acquire land from low-skilled ones. We find the selection effect is less negative in Panel B. This can explain why the deviation of the adaptation-only from the no-adjustment coefficient is also less negative in Panel B.

The key takeaway is that the long-difference coefficient can be more negative than the benchmark without adjustment. This occurs when the selection effect is negative and dominates the positive effect of adaptation. In practice, decomposing selection from adaptation using county-level aggregate data is challenging. One would need to model the counterfactual crop yields in the absence of any selection.

## 4. Data

We use several datasets to present descriptive evidence of selection effects at three margins in response to rising temperatures: counties halting production, brain drain, and farm consolidation. We compile a county panel dataset by combining multiple datasets on weather, natural disasters, agricultural production, and demographics. We restrict our analysis to counties east of the 100th meridian, which account for 92.5% of U.S. corn production and 99.3% of soy production.

### 4.1 Weather and natural disasters

The first component of our data consists of various climate variables, such as growing degree days (GDDs) and precipitation. We construct these variables using raw GIS files provided by the PRISM Climate Group.[8] These geospatial files offer daily temperature and precipitation data on a 4 km*4 km grid across the United States from 1981 to 2022. Our analysis is restricted to the growing season (April to September). We use temperature data to calculate GDDs, which measure crops’ cumulative heat exposure in a year. Following Burke and Emerick (2016), we construct piecewise linear weather variables using specific thresholds: GDDs below 33°C (beneficial moderate heat), GDDs above 33°C (damaging extreme heat), precipitation below 55

[8] https://prism.oregonstate.edu/

cm (beneficial rainfall), and precipitation above 55 cm (potentially harmful excess rainfall).[9] We create a county-year panel. In our empirical analysis, we use five-year moving averages of GDDs and precipitation.

To complement the heat and drought data, we incorporate additional disaster variables from the Federal Emergency Management Agency (FEMA).[10] The FEMA dataset records all natural disasters since 1950, including their date, type, and location. We focus on hurricanes, storms, and floods. FEMA does not provide data on heat and drought, so we create dummies for these two disasters based on the GDDs and precipitation data.

### 4.2 Agricultural data

We obtain various agricultural variables from the United States Department of Agriculture (USDA) agricultural census.[11] These include yields and farm areas of major crops at the county-year level. The USDA also reports the average farm size and the distribution of farm sizes within each county in every five years. We also obtain census data on farmland prices.

We also calculate the soil quality of each county. The USDA provides satellite-derived data on land use and soil texture. We compute the farmland area-weighted average soil quality for each county. We use available water capacity (AWC), a proxy for a soil's ability to supply moisture for crop growth, as a measure of soil quality.

### 4.3 Demographic data

We merge several variables from the US census database into our county dataset. Our empirical analysis focuses on the adult population and college graduates since 1980. We also obtain county-to-county migration data from the IRS. The dataset reports migration flows between counties in each year from 2011 to 2020, as well as the average income of the migrants.

## 5. Testing for Selection Effects

[9] GDD above/below 33°C is the total time (in days) when the temperature is above/below 33°C. To calculate the precipitation variables, we use the total rainfall during the growing season. Precipitation above/below 55cm is the maximum/minimum of (total rainfall – 55) and zero.

[10] https://www.fema.gov/openfema-data-page/disaster-declarations-summaries-v2

[11] https://www.nass.usda.gov/Data_and_Statistics/

In this section, we test for three margins of selection: counties may exit agriculture entirely and pivot to other sectors; high-skilled farmers may switch occupations or migrate out of the county as the opportunity cost of farming rises; and land may consolidate into larger farms operated by better managers.

### 5.1 County-level agricultural exit

In Figure 3, we plot the number of counties growing corn and soybeans from 1980 to 2020. The counts for both crops have declined steadily over this period. In the early 1980s, more than 2200 counties grew corn and over 1900 grew soybeans, but by 2022 these figures had fallen to approximately 1700 and 1600, respectively. The long-difference regressions (see equation (3)) are based on a balanced panel, so a county is excluded if it stops growing the crop during the sample period. If this selection is negatively correlated with skill levels—that is, if higher-skilled farmers are more likely to exit—then the OLS coefficient will be biased downward, causing researchers to overestimate the negative effects of climate change. We now study the correlates of agricultural exit.

We estimate a probit model where the dependent variable is a dummy $D_{it} = 1$ if county i exits sometime between time $(t-1)$ and $t$. Each period $t$ refers to a decade from 1980 to 2020 in our data. The model is specified as follows:

$$Pr\,(X) \;= \Phi(\beta X_{i,t-1} + \delta_s + \tau_t) \qquad (5)$$

where Φ is the standard normal CDF; X is a vector of climate and demographic attributes in the base year; $\tau_t$ refers to the decade fixed effects; and $\delta_s$ refers to the state fixed effects. We estimate equation (5) using maximum likelihood and cluster standard errors by state. The results are presented in Table 2.

In all specifications, counties are more likely to exit corn or soybean production in periods when extreme heat in the base year is higher, conditional on lagged crop area, which has the expected negative coefficient. We also find that extreme rainfall is positively correlated with exit from corn production but not from soybean production, consistent with soybeans being more tolerant to flooding than corn.

In columns (3) and (4), the coefficient on the proportion of college graduates is significantly positive. Counties with more educated populations are more likely to exit agriculture, providing evidence of negative selection. The opportunity cost of remaining in farming is higher in these counties as individuals can earn higher wages by switching to non-agricultural sectors. We also find that soil quality is negatively correlated with exit. In counties with high soil quality, it is more profitable for the remaining farmers to acquire land from those who exit.

### 5.2 Brain drain in farming counties

The migration literature has shown that people often migrate out of climate-affected areas in search of better job opportunities and a higher quality of life (Cattaneo et al. 2019; Grafton and Neubauer 2025). Using a global panel dataset, Cattaneo and Peri (2016) have shown that people in middle-income countries emigrate or urbanize in response to higher temperatures. This accelerates the transition away from agriculture into more urban areas.

Rural to urban migration has been a classic strategy for coping with sectoral risk. Rural people can always migrate away from places that fail to adapt to climate shocks (Hornbeck 2012; Feng, Oppenheimer, and Schlenker 2015; Hornbeck and Levitt 2024). Better off-farm work opportunities can accelerate agricultural exit, and the literature has documented that the out-migrants have experienced increases in wages (Goetz and Debertin 2001; Grafton and Neubauer 2025).

Rural counties face ongoing population loss, with brain drain hitting educated individuals hardest—women encounter heightened obstacles to staying. Skilled women find few opportunities, stuck in low-pay, male-heavy jobs that ignore their qualifications. Kim and Waldorf (2024) show rural structures amplify gender wage gaps via segregation and limited high-skill advancement, deterring retention and harming dual-career families. Educated rural youth, especially women, rarely return to rural places due to scant professional prospects. Young, educated men will be even less likely to remain in rural areas if they anticipate that their marriage prospects are worse because educated women are leaving these areas.

We test whether educated people disproportionally move out of farming counties affected by climate change. To this end, we estimate the following regression on county i and decade t:

$$\Delta Y_{it} = \beta * Disaster\ counts_{it} + \delta_s + \tau_t + \varepsilon_{it} \qquad (6)$$

where the dependent variable is the change in the demographic variable from decade (t-1) to t, and disaster counts refer to a vector that includes the total number of droughts, heat waves, hurricanes, floods, and storms experienced by county i during the decade. We include state fixed effects and decade fixed effects. The coefficient of interest is $\beta$. We restrict our sample to counties that grow corn. The results are reported in Table 3.

In columns (1) and (3), the dependent variable is the change in adult population. In both periods, the change is more negative in counties that experienced storms. Yet, we do not find evidence that heat and drought are correlated with population loss. In columns (2) and (4), we test whether migrants tend to be more educated than those who remain. In this case, the number of college graduates should decline, conditional on the overall change in population. We find evidence of brain drain in both periods. The coefficients on drought and heat are significantly negative in both periods. They are numerically larger and more statistically significant over the period from 2000 to 2020. These results suggest that the composition of population in farming counties have changed (with larger changes in the more recent decades), although the total population barely shrinks.[12]

The brain drain from rural agricultural counties may have implications for agricultural technological adoption trends. Moscona and Sastry (2023) find that innovation in US agriculture has been redirected since the mid-20th century toward crops increasingly exposed to extreme temperatures, particularly through technologies related to environmental adaptation. Their analysis shows that this induced innovation has offset approximately 20% of potential losses in US agricultural land values since 1960, but it remains incompletely effective in mitigating projected climate damages.

We argue that this can be resolved by recognizing the complementarity between local human capital and the adoption of innovative ideas: skilled farmers provide essential on-the-ground knowledge, demand for new technologies, and ties to research networks that drive directed R&D. If brain drain leaves behind older, deskilled farmers—who are less inclined or able to demand cutting-edge solutions—this diminishes the incentives for innovation, perpetuating slow technological progress and this inhibits adaptation.

---

[12] Urbanization and human capital are complements because cities offer better jobs for high-skilled workers (Glaeser 2012). In Appendix 2, we show that people are migrating from farming counties to urban areas using data from the IRS.
We rely on average education levels in farming counties because we lack data on the farming population's education. In Appendix 6, we show that in counties experiencing more weather shocks, more farmers have shifted to part-time rather than full-time work. These exiting full-time farmers—who are able to secure off-farm jobs—are likely higher-skilled than average.

### 5.3 Farmland consolidation

The motivating example in Section 2 predicts farm consolidation as exiting farmers sell land to those who remain. Land is also reallocated among the latter until marginal returns equalize. Large "superstar" farms emerge if the highest-skilled remaining farmers acquire land from both lower-skilled peers and those who exit. Meanwhile, because adaptation often entails high fixed costs, larger farms gain a comparative advantage through greater human capital and economies of scale (Zivin and Kahn 2016). While we do not have farmland transaction data or farm-level microdata, we use county-level farm size distribution data to provide descriptive evidence on the emergence of larger farms.

In 1980, the average farm size in counties in our sample was 296 acres, and it has grown to 326 acres in 2020. The proportion of large farms has also increased—about 5% of the farms were above 1000 acres in 1980, but 7% were in 2020. Using agricultural census data, we calculate the average farm size and the proportion of farms exceeding 1000 acres in each county in each decade. We plot their distribution in 1980, 2000, and 2020 in Figure 4. It shows that the density of smaller farms sizes has decreased consistently over time, and both distributions have become more right-skewed.

We now test whether the consolidation of farmland primarily takes place in counties affected by climate change. We estimate equation (6) with a different set of dependent variables: the changes in farm counts, average size, and proportions of farms above several area thresholds. We use a county-decade panel dataset. The dependent variables in the regression are the decade-by-decade differences, and the explanatory variables are the total disaster counts in the decade. The estimation results are reported in Table 4.

In columns (1) and (2), we find that heat waves are negatively correlated with the number of farms in a county and positively correlated with the average farm size. This is consistent with our prediction of farmland consolidation.[13] We then study how the distribution of farm sizes has shifted. Specifically, we are interested in testing the emergence of "superstar" farms that operate large acreages. In 2020, approximately 29% of US farms are at least 180 acres, 14% are over 500 acres, and 7% are more than 1000 acres. In columns (3) and (4), we find that heat and hurricane are positively correlated with the proportion of large farms, as benchmarked by 500- and 1000-

[13] When using the change in total farmland area as the dependent variable, we find no evidence that the total area has more negative trends in counties with more heat and drought.

acres cutoffs. The coefficient is numerically larger and more statistically significant for the 1000-acres cutoff. In contrast, the coefficients do not significantly differ from zero when the dependent variable refers to the 180-acres threshold. This suggests that land has primarily consolidated into megafarms, likely commercial farms, rather than medium-sized operations.

In counties with worsening climate, such economies of scale are more important since adaptation often features high upfront costs (Sumner 2014). Large farms also tend to have better managers, and if knowledge spillovers exist, these can improve management practices on smaller farms in the area, thereby boosting local agricultural productivity (Bloom et al. 2019).

## 6. Incorporating Our Selection Findings in Testing for Heterogeneity in Long-Difference Estimates

In Section 3, we have shown that the long difference estimate is a weighted average of each county's marginal effect, and each marginal effect represents the net effect of climate change, adaptation, and selection. The empirical challenge is that researchers do not observe the counterfactual outcomes in the absence of selection. They would underestimate adaptation if there exists negative selection and vice versa. Given the various selection margins documented in Section 5, we illustrate this point using corn yield data from 1980 to 2020.[14]

We partition corn counties based on two characteristics from 1980: educational attainment) and average farm size. If path dependence is present, counties with low educational attainment and small average farm sizes in 1980 should find it harder to attract and retain talented individuals and high-skilled operators capable of managing large farms. For example, these counties may lack established colleges and the infrastructure required for commercial-scale farming. This suggests greater brain drain and less land consolidation, and thus a more negative selection effect. The long-difference coefficient in this low-education/low-farm-size subsample should be more negative than in the full sample. The long-difference coefficient should be more positive in the subsample with more college graduates and larger farms. In practice, we focus on two subsamples: (1) counties where both the proportion of college graduates and the average farm size were below the

[14] In Appendix 3, we use the compiled dataset to replicate Burke and Emerick's 20-year difference results and extend the analysis through 2020. For the period 1980–2000, we successfully replicate their main findings: both the first-difference and long-difference coefficients on GDDs above 33°C are significantly negative, with the long-difference coefficient approximately 30% less negative than the first-difference coefficient. For the period 2000–2020, the long-difference coefficients are statistically insignificant, while the first-difference coefficients remain significantly negative and substantially different from zero.

respective county-level medians in 1980 (computed across all counties in the sample), and (2) counties where both measures were above their respective medians.

In Table 5, we show that the impact of extreme weather on agricultural output dynamics varies as a function of county skill and farm size. We estimate equation (3) using a county-decade panel dataset. We run the long difference specification on the full sample and the two subsamples described above, for both the full period (1980–2020) and the later subperiod (2000–2020). The dependent variable is the decade-by-decade change in corn yields. The key explanatory variables are the corresponding decade-by-decade changes in growing degree days (GDDs) and precipitation, with separate terms for extreme heat (GDD above 33°C) and extreme rainfall (precipitation above 55 cm). We include county fixed effects to control for time-invariant unobserved heterogeneity across counties, as well as decade fixed effects to account for common shocks over time.

As expected, extreme heat and extreme rainfall are negatively correlated with corn yields. These coefficients are significantly negative and numerically large when we estimate the specification with the full sample. We find that the coefficient on extreme heat is less negative in the later period (2000 to 2020) than in the full period (1980 to 2020), but the coefficient on extreme rainfall is more negative. The constant term is significantly positive in all specifications, and most of the decade dummies are also significantly positive. This implies that aggregate crop yields are growing despite climate change, and the growth is faster in later decades than in the base period from 1980 to 1990.

In Tables 2-4, we show that heat is correlated with all three selection mechanisms: agricultural exit, brain drain, and farm consolidation. We thus focus on the coefficient on GDDs above 33°C to illustrate our points on selection.

The results in Table 5 show that in farming counties featuring a relatively more educated populace and larger farms that the weather impact is much smaller. When we estimate the equation over the 1980–2020 period, the below-median subsample estimate is 15% more negative than the full-sample estimate, while the above-median subsample estimate is 22% less negative. This pattern is even more pronounced when we restrict the analysis to the later period (2000–2020): the below-median subsample estimate is more than twice as negative, while the above-median estimate is 48% less negative and not statistically different from zero.[15]

[15] This pattern robust when we use alternative fixed effects or 20-year differences (see Appendix 4).

We construct bootstrap distribution for the differences in coefficients. We implement block resampling using counties as blocks, estimate each specification on every bootstrap sample, and compute the differences in coefficients. We plot the bootstrap distributions of these differences in Figure 5. In all cases, we reject the null hypothesis that the difference is zero at 5% significance level. These results provide statistically significant evidence that the coefficient in the below-median subsample is more negative than in the full sample, and that the coefficient in the above-median subsample is less negative than in the full sample.

## 7. **Conclusion**

Climate change increases the probability of extreme weather events such as heat waves and drought that reduce crop yields. Given that these trends are partially predictable, rational expectations economics predicts that farmers will invest to protect themselves from this threat. We present an optimizing model where farmers make extensive margin decisions, such as selling their land and exiting agriculture. The county/year "macro" agricultural data that is generated by this micro optimization depends on many extensive margin decisions. If skilled farmers quit farming and move to the city, such a "brain drain" shifts the composition of the county. Researchers who estimate county panel regressions to study agricultural productivity dynamics assume a representative agent framework and miss these selection effects.

Due to data limitations, researchers often track how places are coping with shocks. Our analysis helps to clarify how person-based responses to trends influence the generating process of place-based aggregate data. Person-based adaptation responses to weather trends introduce a type of selection bias for place-based empirical studies.

Our empirical work documents three different selection effects. Using U.S. county-level datasets, we document that approximately 20% of counties specializing in corn and soybean production in 1980 had exited agriculture by 2022. These exits are positively correlated with exposure to extreme weather and with higher educational attainment. We provide descriptive evidence of "brain drain" from farming counties, as higher-skilled individuals depart. The rise of large farms partially offsets this negative selection effect. The distribution of farm size has become more right-skewed in counties affected by climate change. These larger farms have the scale and the human capital to adapt. By estimating the canonical long-difference specification across various subsamples, we find substantial heterogeneity in the results, with patterns that align directionally with our selection hypotheses.

The climate change impact literature has focused on the physical output of farmers in different geographic areas at different points in time. An alternative criterion would be total farmer profits. This criterion would incorporate government transfers and other market sources of revenue. In recent years, some farmers have actively engaged in eco-tourism, and some have sold their land to data centers and real estate developers.

Our study highlights the role of resource reallocation at the extensive margin. Future research should examine how government agricultural policies affect this process. We believe many policies have synergistic effects on the pace of reallocation, such as crop insurance subsidies and farmland conservation mandates (Goetz and Debertin 2001). Federal crop insurance subsidies can create moral hazard, as farmers are compensated even when yields are low, which disincentivizes adaptation (Annan and Schlenker 2015). On the other hand, since larger farms receive disproportionately more subsidies, these programs may accelerate the consolidation into large operations with high-skilled managers (Bekkerman, Belasco, and Smith 2019; Berger, Ifft, and Jodlowski 2025). It remains an open question whether these policies facilitate or hinder extensive-margin reallocation of land and labor.

## References


Acemoglu, Daron, and Pascual Restrepo. "Robots and jobs: Evidence from US labor markets." *Journal of political economy* 128, no. 6 (2020): 2188-2244.

Annan, Francis, and Wolfram Schlenker. "Federal crop insurance and the disincentive to adapt to extreme heat." *American Economic Review* 105, no. 5 (2015): 262-266.

Autor, David H., David Dorn, and Gordon H. Hanson. "The China shock: Learning from labor-market adjustment to large changes in trade." *Annual review of economics* 8, no. 1 (2016): 205-240.

Bekkerman, Anton, Eric Belasco, and Vincent H. Smith (2019). "Does Farm Size Matter? Distribution of Crop Insurance Subsidies and Government Program Payments across U.S. Farms." Applied Economics Policy and Perspectives, 41 (3): 498–518.

Berger, Elizabeth, Jennifer Ifft, and Margaret Jodlowski. 2025. "Risk and Capital-Labor Allocation: Evidence from Large-Scale Government Insurance Provision." Working paper. https://ssrn.com/abstract=4034903.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen. "What drives differences in management practices?." *American Economic Review* 109, no. 5 (2019): 1648-1683.

Burke, Marshall, Mustafa Zahid, Mariana CM Martins, Christopher W. Callahan, Richard Lee, Tumenkhusel Avirmed, Sam Heft-Neal, Mathew Kiang, Solomon M. Hsiang, and David Lobell. *Are We Adapting to Climate Change?*. No. w32985. National Bureau of Economic Research, 2024.

Burke, Marshall, and Kyle Emerick. "Adaptation to climate change: Evidence from US agriculture." *American Economic Journal: Economic Policy* 8, no. 3 (2016): 106-140.

Cattaneo, Cristina, and Giovanni Peri. "The migration response to increasing temperatures." *Journal of development economics* 122 (2016): 127-146.

Cattaneo, Cristina, Michel Beine, Christiane J. Fröhlich, Dominic Kniveton, Inmaculada Martinez-Zarzoso, Marina Mastrorillo, Katrin Millock, Etienne Piguet, and Benjamin Schraven. "Human migration in the era of climate change." *Review of Environmental Economics and Policy* (2019).

Costinot, Arnaud, Dave Donaldson, and Cory Smith. "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, no. 1 (2016): 205-248.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. "Temperature shocks and economic growth: Evidence from the last half century." *American Economic Journal: Macroeconomics* 4, no. 3 (2012): 66-95.

Deschênes, Olivier, and Michael Greenstone. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US." *American Economic Journal: Applied Economics* 3, no. 4 (2011): 152-185.

Dustmann, Christian, Sebastian Otten, Uta Schönberg, and Jan Stuhler. "The effects of immigration on places and people–identification and interpretation." (2025).

Federle, Jonathan, André Meier, Gernot J. Müller, Willi Mutschler, and Moritz Schularick. *The price of war*. No. 2262. Kiel Working Paper, 2024.

Feng, Shuaizhang, Michael Oppenheimer, and Wolfram Schlenker. 2015. "Weather Anomalies, Crop Yields, and Migration in the United States." *Nature Climate Change* 5 (6): 579–85. https://doi.org/10.1038/nclimate2631.

Glaeser, Edward. *Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier*. Penguin, 2012.

Goetz, Stephan J., and David L. Debertin. "Why farmers quit: A county-level analysis." *American Journal of Agricultural Economics* 83, no. 4 (2001): 1010-1023.

Goldin, Claudia. "A grand gender convergence: Its last chapter." American economic review 104, no. 4 (2014): 1091-1119.

Grafton, Gabrielle, and Michael Neubauer. 2025. "(The Great) Migration and Those Left Behind." Working paper, Department of Economics, Brown University.

Heckman, James J., and Bo E. Honoré. 1990. "The Empirical Content of the Roy Model." *Econometrica* 58 (5): 1121–49. https://doi.org/10.2307/2938303.

Hornbeck, Richard. "The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe." *American Economic Review* 102, no. 4 (2012): 1477–1507.

Hornbeck, Richard, and Steven T. Levitt. "The Economic Impact of the Eruption of Mount St. Helens." *Journal of Economic Perspectives* 38, no. 4 (2024): 99–122.

Jones, Benjamin, Jacob Moscona, Benjamin A. Olken, and Cristine von Dessauer. *With or Without U? Binning Bias and the Causal Effects of Temperature Extremes*. No. w34671. National Bureau of Economic Research, 2026.

Kim, A. and Waldorf, B.S., 2024. Unpacking the gender wage gap in the US: The impact of rural employment, age, and occupation. *Journal of the Agricultural and Applied Economics Association*, *3*(4), pp.706-723.

Moscona, Jacob, and Karthik A. Sastry. "Does Directed Innovation Mitigate Climate Damage? Evidence from U.S. Agriculture." *Quarterly Journal of Economics* 138, no. 2 (2023): 637-701. https://doi.org/10.1093/qje/qjac057.

Nordhaus, William. "Climate change: The ultimate challenge for economics." *American Economic Review* 109, no. 6 (2019): 1991-2014.

Olley, Steven, and Ariel Pakes. "The dynamics of productivity in the telecommunications equipment industry." (1992).

Ortiz-Bobea, Ariel. "The role of nonfarm influences in Ricardian estimates of climate change impacts on US agriculture." *American Journal of Agricultural Economics* 102, no. 3 (2020): 934-959.

Ortiz-Bobea, Ariel, and Richard E. Just. "Modeling the structure of adaptation in climate change impact assessment." *American Journal of Agricultural Economics* 95, no. 2 (2013): 244-251.

Ortiz-Bobea, Ariel, Toby R. Ault, Carlos M. Carrillo, Robert G. Chambers, and David B. Lobell. "Anthropogenic climate change has slowed global agricultural productivity growth." *Nature Climate Change* 11, no. 4 (2021): 306-312.

Ranson, Matthew. "Crime, weather, and climate change." *Journal of environmental economics and management* 67, no. 3 (2014): 274-302.

Schlenker, Wolfram, and Michael J. Roberts. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of sciences* 106, no. 37 (2009): 15594-15598.

Sumner, Daniel A. "American farms keep growing: Size, productivity, and policy." *Journal of Economic Perspectives* 28, no. 1 (2014): 147-166.

Tack, Jesse, Andrew Barkley, and Lawton Lanier Nalley. "Effect of warming temperatures on US wheat yields." *Proceedings of the National Academy of Sciences* 112, no. 22 (2015): 6931-6936.

Zivin, Joshua Graff, and Matthew E. Kahn. *Industrial productivity in a hotter world: the aggregate implications of heterogeneous firm investment in air conditioning*. No. w22962. National Bureau of Economic Research, 2016.

## Table 1

## Simulated Long Difference Regression Results

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| | $\Delta\log(Yield)$ | | | | | | |
| | No adjustment | Full effect | Adaptation only | No exit | Δlog(Farmer ability) | Δlog(Median size) | log(Acreage ability) |
| | | | Panel A: 0.2 to 0.3 | | | | |
| $\Delta Damage$ | -1.421*** | -1.869*** | -1.341*** | -1.084*** | -0.479*** | -0.0636 | -0.334*** |
| | (0.0132) | (0.0252) | (0.0968) | (0.0232) | (0.00288) | (0.0658) | (0.0183) |
| | | | Panel B: 0.2 to 0.4 | | | | |
| $\Delta Damage$ | -1.664*** | -1.249*** | -1.399*** | -0.599*** | -0.374*** | -1.578*** | -0.0550*** |
| | (0.0177) | (0.0604) | (0.0515) | (0.0370) | (0.0156) | (0.113) | (0.0158) |

These are the estimation results based on equation (3). The changes are taken between the pre climate change equilibrium values and the post climate change values.
Robust standard errors in parentheses. N=1000 in all columns. We have checked that the assumption of complementarity between human capital and non-farming sectors holds in all simulated equilibria.
*** p<0.01, ** p<0.05, * p<0.1

## Table 2

## Weather Dynamics and Rural Counties Exiting Farm Production

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Corn exit | Soy exit | Corn exit | Soy exit |
| GDD above 33C | 0.0185*** | 0.0144* | 0.0190*** | 0.0135* |
| | (0.00679) | (0.00789) | (0.00632) | (0.00816) |
| GDD below 33C | -0.000505* | 0.000256 | -0.000599** | 0.000107 |
| | (0.000272) | (0.000221) | (0.000291) | (0.000248) |
| Precip above 55cm | 0.0145** | 0.0174 | 0.0128** | 0.0126 |
| | (0.00573) | (0.0108) | (0.00585) | (0.0113) |
| Precip below 55cm | -0.00942 | -0.0151 | -0.00882 | -0.0144 |
| | (0.0174) | (0.0211) | (0.0166) | (0.0199) |
| log(Lagged area) | -0.845*** | -0.830*** | -0.826*** | -0.808*** |
| | (0.0601) | (0.0457) | (0.0617) | (0.0463) |
| College | | | 1.356** | 1.755** |
| | | | (0.576) | (0.686) |
| log(Miles to CBD) | | | 0.0428 | 0.128* |
| | | | (0.0743) | (0.0729) |
| log(Soil quality) | | | -0.984** | -0.876*** |
| | | | (0.393) | (0.325) |
| Observations | 7,011 | 6,383 | 6,993 | 6,362 |

This table reports estimates of equation (5). We use a county-decade panel, and the dependent variable is an indicator of whether the county exits corn/soy farming completely during the decade. The explanatory variables are their values at the beginning of the decade,

Standard errors are clustered by state. State FEs and year FEs are included in all columns.

*** p<0.01, ** p<0.05, * p<0.1

**Table 3**

**Demographic Change in Farming Counties**

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | 1980 to 2000 | | 2000 to 2020 | |
| | $\Delta log\ (Population)$ | $\Delta \log\ (College)$ | $\Delta log\ (Population)$ | $\Delta \log\ (College)$ |
| Drought | -0.00267 | -0.00554** | 0.00178 | -0.00733*** |
| | (0.00281) | (0.00206) | (0.00350) | (0.00195) |
| Heat | -0.00274 | -0.00520* | -0.000947 | -0.00699*** |
| | (0.00723) | (0.00274) | (0.00332) | (0.00223) |
| Storm | -0.0389*** | -0.0122 | -0.00365** | 0.00263 |
| | (0.0127) | (0.0146) | (0.00159) | (0.00303) |
| Flood | 0.0481*** | -0.0317*** | -0.000634 | 0.00113 |
| | (0.00616) | (0.00602) | (0.00427) | (0.00435) |
| Hurricane | | | -0.000131 | 0.00139 |
| | | | (0.00645) | (0.00480) |
| Observations | 4,962 | 4,962 | 4,249 | 4,249 |
| R-squared | 0.358 | 0.596 | 0.395 | 0.435 |

This table reports estimates of equation (6). We use a county-decade panel. The sample includes all corn counties.
The dependent variables are changes from the start to the end of the decade, and the explanatory variables are the total disaster counts over the decade. There was no hurricane in our sample between 1980 and 2000. Population refers to the count of people above 25 years old. College refers to the count of college graduates. Standard errors are clustered by states. State FEs, decade FEs, and lagged values of the dependent variable are included in all columns. In columns (2) and (4), we also control for $\Delta log\ (Population)$.
*** p<0.01, ** p<0.05, * p<0.1

**Table 4**

**Farm Size Growth and Climate Change**

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | | | $\Delta\log\ (Proportion\ of\ farms)$ | | |
| | $\Delta\log\ (Count)$ | $\Delta\log\ (Avg\ size)$ | 1000+ acres | 500+ acres | 180+ acres |
| Heat | -0.00642** | 0.00726*** | 0.0223** | 0.0105* | 0.00114 |
| | (0.00284) | (0.00200) | (0.00933) | (0.00536) | (0.00376) |
| Drought | 0.00438 | 0.000924 | 0.0120 | 0.00394 | 0.00584 |
| | (0.00476) | (0.00322) | (0.00999) | (0.00768) | (0.00491) |
| Hurricane | -0.00727 | 0.0191* | 0.0606** | 0.0361** | 0.00927 |
| | (0.0108) | (0.00999) | (0.0238) | (0.0168) | (0.00983) |
| Storm | -0.00408 | 0.000228 | 0.0112 | 0.00411 | -0.00395 |
| | (0.00739) | (0.00667) | (0.0159) | (0.0104) | (0.00557) |
| Flood | -0.00778 | -0.000307 | -0.00921 | 0.00189 | 0.0124 |
| | (0.0120) | (0.0101) | (0.0287) | (0.0183) | (0.00997) |
| Observations | 8,047 | 8,025 | 7,812 | 8,017 | 8,046 |
| R-squared | 0.425 | 0.258 | 0.264 | 0.291 | 0.309 |

These are estimation results of equation (6). We use a county-decade panel from 1980 to 2020. The explanatory variables are the total count of disasters during the decade.
Standard errors are clustered by state. State FEs and decade FEs are included in all columns. Lagged values of the dependent variable are included in the last three columns. Regressions are weighted by the total number of farms
*** p<0.01, ** p<0.05, * p<0.1

# Table 5

## The Yield Damage Function Differs By Baseline Skill and Farm Size

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | | | $\Delta\log(Corn\ yield)$ | | | |
| ΔGDD above 33C | -0.0185*** | -0.0213*** | -0.0144*** | -0.00982*** | -0.0239*** | -0.00515 |
| | (0.00150) | (0.00342) | (0.00213) | (0.00342) | (0.00513) | (0.00481) |
| ΔGDD below 33C | -0.000247*** | -0.000482*** | -0.000289*** | -0.000270** | -0.000474* | -0.000341* |
| | (6.28e-05) | (0.000183) | (0.000109) | (0.000112) | (0.000255) | (0.000178) |
| ΔPrecip above 55cm | -0.00140*** | 0.00146 | -0.00250*** | -0.00459*** | 0.00239 | -0.00531*** |
| | (0.000468) | (0.00101) | (0.000650) | (0.000742) | (0.00252) | (0.000985) |
| ΔPrecip below 55cm | 0.00755*** | 0.00602*** | 0.00654*** | 0.00758*** | 0.00186 | 0.00626*** |
| | (0.000794) | (0.00202) | (0.00113) | (0.00137) | (0.00395) | (0.00197) |
| Decade dummy: 1990 to 2000 | 0.0429*** | 0.0191 | 0.0823*** | | | |
| | (0.00973) | (0.0248) | (0.0146) | | | |
| Decade dummy: 2000 to 2010 | 0.0179*** | -0.00672 | 0.0211* | | | |
| | (0.00665) | (0.0161) | (0.0126) | | | |
| Decade dummy: 2010 to 2020 | 0.0370*** | 0.0369** | 0.0431*** | 0.0324** | 0.0389 | 0.0370 |
| | (0.00529) | (0.0163) | (0.00832) | (0.0155) | (0.0256) | (0.0252) |
| Constant | 0.125*** | 0.147*** | 0.113*** | 0.138*** | 0.136*** | 0.128*** |
| | (0.00370) | (0.0100) | (0.00661) | (0.00654) | (0.0108) | (0.0119) |
| County sample | Full | Below median | Above median | Full | Below median | Above median |
| Average college% in 1980 | 10.5% | 7% | 12.6% | 10.6% | 7% | 12.5% |
| Average farm size in 1980 | 320 | 159 | 523 | 333 | 163 | 522 |
| Decade sample | 1980-2020 | 1980-2020 | 1980-2020 | 2000-2020 | 2000-2020 | 2000-2020 |
| Observations | 7,277 | 1,624 | 2,088 | 3,419 | 728 | 1,024 |
| R-squared | 0.343 | 0.292 | 0.414 | 0.356 | 0.471 | 0.337 |

The table reports the estimates of equation (3). We use a county-decade panel from 1980 to 2020. The differences are taken decade by decade for each county. The sample in the first three columns includes all decades, and the last three columns only include decades 2000 to 2020. Below median counties are those with below-median educational attainment and below-median average farm size in 1980. Above median counties are those with above-median educational attainment and above-median average farm size in 1980. We report the average educational attainment and farm size (in acres) of each sample in 1980.

Standard errors are clustered by county. County FEs are included in all columns. Regressions are weighted by corn area in 1980 in the first three columns and by corn area in 2000 in the last three.

*** p<0.01, ** p<0.05, * p<0.1

**Figure 1**

**Farm Size and Income Dynamics in the Calibration Example**

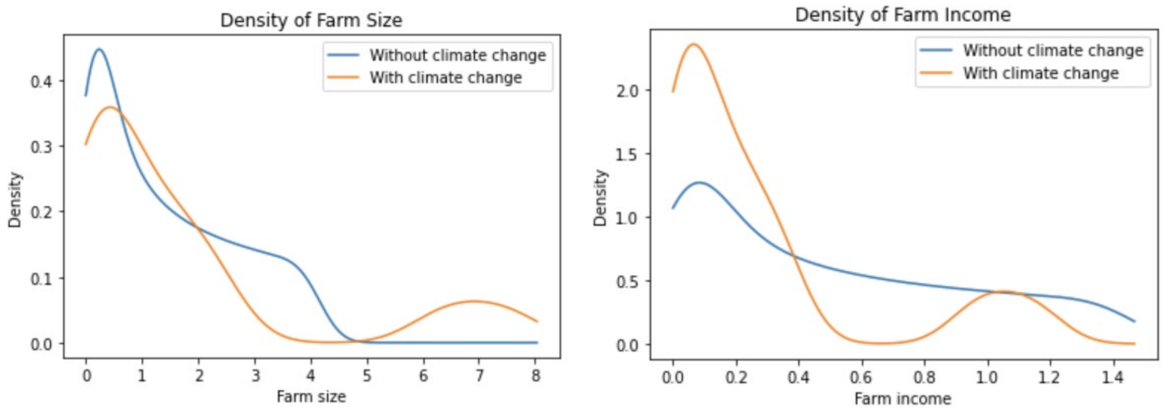

# Figure 2

## Decomposing the Observed Changes in Crop Yields

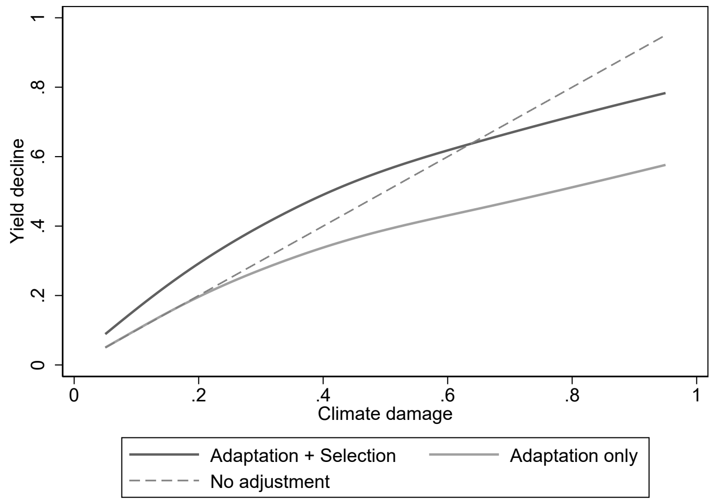

**Figure 3**

**Counts of Farming Counties Over Time**

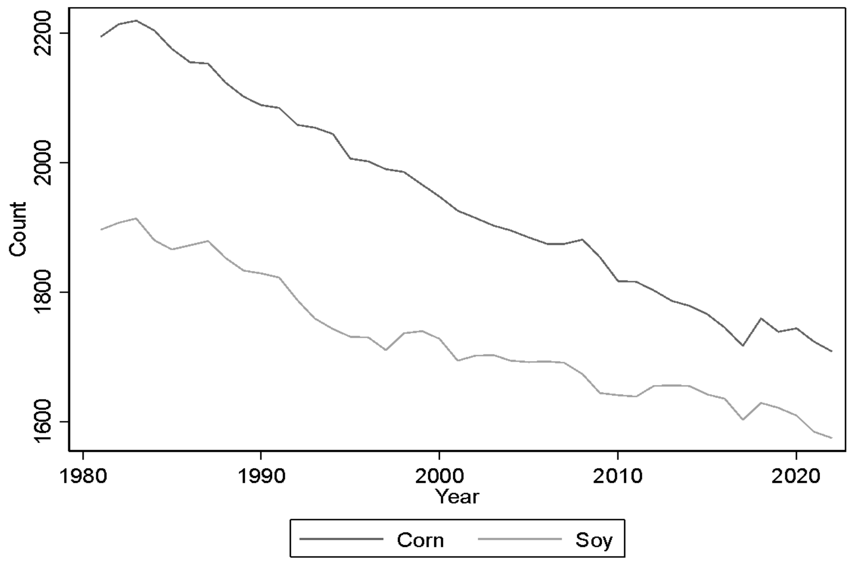

# Figure 4

## The Shifting Distribution of Farm Size

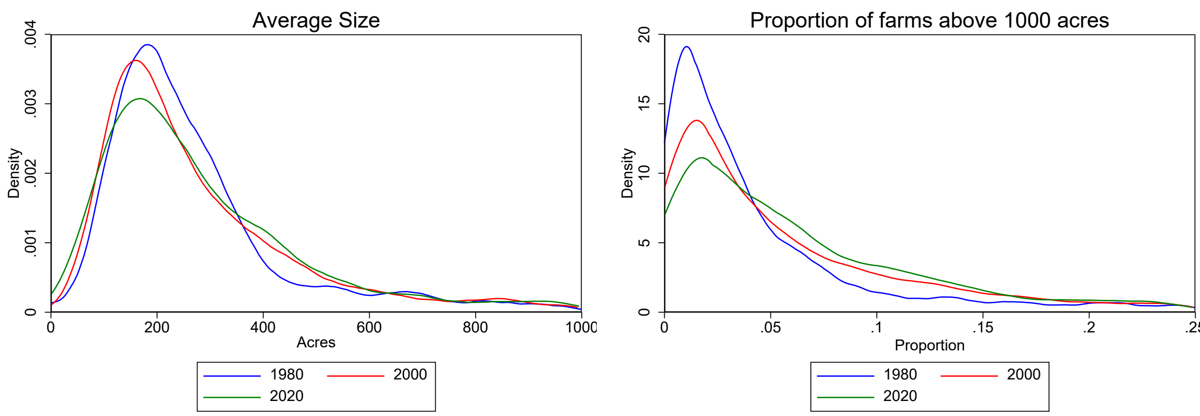

# Figure 5

## Distribution of Coefficient Differences

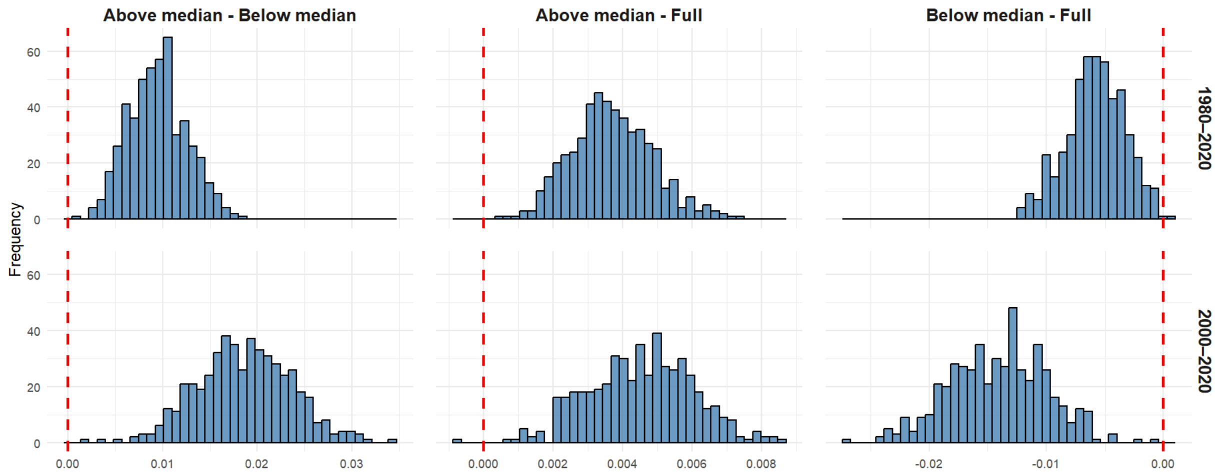

## Appendix 1

## Proposition and Proof

We emphasize that researchers should account for extensive-margin adjustments when benchmarking progress in climate adaptation based on observed crop yields. In what follows, we demonstrate that, in the model presented in Section 2, high-skilled farmers exit farming as land productivity declines due to climate change. Intuitively, this occurs because human capital and the non-farming sector are complements in our framework—meaning that high-skilled farmers face a higher opportunity cost to remain in agriculture. As the remaining farming population becomes less skilled on average, this composition shift correlates with lower observed crop yields. Researchers would therefore understate the positive effects of adaptation if they fail to account for this negative selection effect.

### Proposition

Under the assumption of complementarity between human capital and non-farming sectors, $m^E < m^*$ and $p^C < p^*$. In words, both the exit cutoff and the equilibrium land price are lower in the post-climate change equilibrium.

*Proof*

Consider the following system of equations where k is a constant in (0,1) and C is a nonnegative constant:

$$\frac{p(1-\alpha)}{\alpha}\left(\frac{mk\alpha}{p}\right)^{\frac{1}{1-\alpha}} - C - w(m) = 0 \qquad \text{(A1)}$$

$$\int_0^m \left(\frac{tk\alpha}{p}\right)^{\frac{1}{1-\alpha}} dF(t) = 1 \qquad \text{(A2)}$$

Our goal is to study how the solutions $m^*$ and $p^*$ change as k changes.

By rearranging and simplifying equation (A1), we obtain:

$$w(m) = (1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}} p^{\frac{-\alpha}{1-\alpha}} (km)^{\frac{1}{1-\alpha}} - C \qquad \text{(A3)}$$

By rearranging and simplifying equation (A2), we obtain:

$$p = ak\left(\int_0^m t^{\frac{1}{1-\alpha}} dF(t)\right)^{1-\alpha} \qquad \text{(A4)}$$

Substitute (A4) into (A3), we obtain a single equation that defines the equilibrium exit cutoff:

$$w(m) = (1-\alpha)km^{\frac{1}{1-\alpha}}\left(\int_0^m t^{\frac{1}{1-\alpha}} dF(t)\right)^{-\alpha} - C \qquad \text{(A5)}$$

Now define $\Phi(\mathrm{m}) = (1-\alpha)m^{\frac{1}{1-\alpha}}\left(\int_0^m t^{\frac{1}{1-\alpha}} dF(t)\right)^{-\alpha}$ so that $w(m) = k\Phi(\mathrm{m}) - C$. By our assumption, $w(m) - \Phi(\mathrm{m})$ is increasing in m. Thus, $w(m) - k\Phi(\mathrm{m})$ is also increasing in m since k is in (0,1).

By totally differentiating (A5) and rearranging, we obtain:

$$\frac{dm^*}{dk} = \frac{\Phi(\mathrm{m}^*)}{w\prime(\mathrm{m}^*) - k\Phi'(\mathrm{m}^*)} > 0 \qquad \text{(A6)}$$

Now, the RHS of equation (A4) is an increasing function of both k and $m^*$. This immediately implies $\frac{dp^*}{dk} > 0$. Similarly, we can verify that $\frac{dm^*}{dC} < 0$ and $\frac{dp^*}{dC} < 0$.

When C=0 and k=1, this is the pre-climate change scenario. When C=0 and $k = 1 - d$ where d is the climate damage, this is the case when none of the remaining farmers find it profitable to adapt. When C is the actual cost of adaptation and $k = 1 - (1 - \phi)d$, this is the case when all remaining farmers adapt. In both cases, both the exit cutoff and land price are decreasing in damage d. This implies that the exit cutoff and land price are both lower in the post-climate change scenario.

Heuristically, when only some of the remaining farmers adapt, the solutions should be between these two extreme scenarios and thus lower than the pre-climate change level.

We verify the proposition using the calibration example in section 2.4. Given the calibrated values in this example, the equilibrium exit cutoff and land price are strictly decreasing in damage.

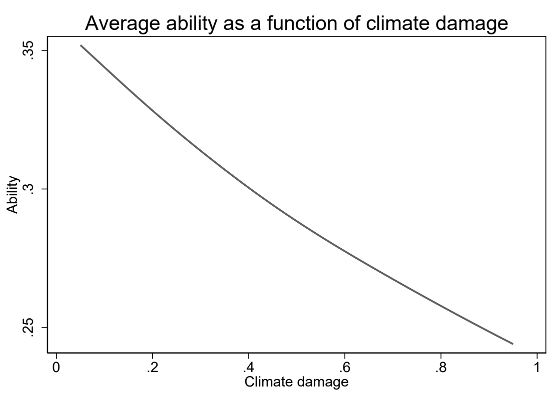


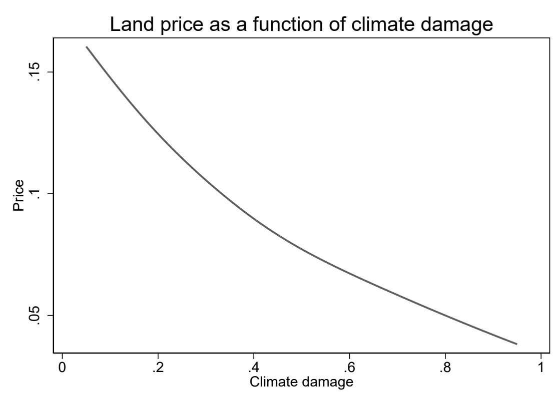

## Appendix 2

## The Urbanization of Farming Population

Using migration data provided by the IRS, we study whether the farming population migrates to more urbanized areas. For each county in our sample, we calculate the average farmland proportion of the destination counties for out-migrants between 2011 and 2021. We benchmark urbanization by the difference between the destination farmland proportion and the origin farmland proportion (measured in 2010). In the following figure, we plot this index of urbanization against the origin county's farmland proportion. We also plot the best-fit lines based on local polynomial regressions. The figure shows that people from farming counties tend to urbanize more. This effect is more pronounced for origin counties that experienced at least five disasters during the period.

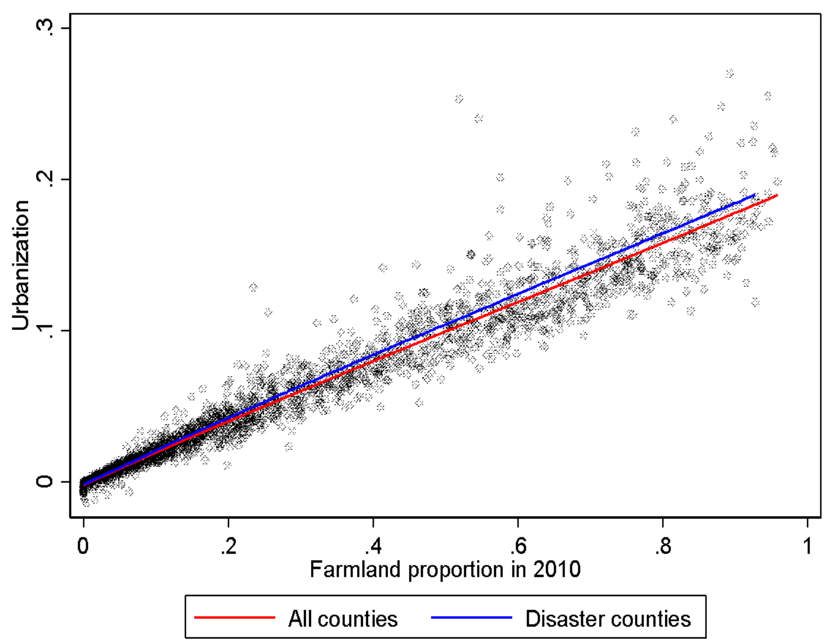

## Appendix 3

## Long Difference Regressions of Crop Yields from 1980 to 2020

2000 to 2020

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Corn | Soy | Corn | Soy |
| | | $\Delta\log(Yield)$ | | |
| ΔGDD above 33C | -0.00886*** | -0.0146*** | -0.00640 | -0.00384 |
| | (0.000849) | (0.000581) | (0.00575) | (0.00747) |
| ΔGDD below 33C | -0.000168*** | 3.48e-05 | -0.000194 | -0.000170 |
| | (2.23e-05) | (2.41e-05) | (0.000201) | (0.000217) |
| ΔPrecip above 55cm | -0.000212* | 0.000445*** | 0.00107 | 0.00218 |
| | (0.000127) | (0.000122) | (0.00130) | (0.00151) |
| ΔPrecip below 55cm | 0.00188*** | 0.00566*** | -0.00140 | 0.00371 |
| | (0.000360) | (0.000310) | (0.00367) | (0.00370) |
| Method | First difference | | Long difference | |
| County FEs | Yes | | No | |
| State FEs | No | | Yes | |
| Year FEs | Yes | | Yes | |
| Observations | 38,489 | 33,691 | 1,692 | 1,528 |
| R-squared | 0.273 | 0.264 | 0.364 | 0.624 |

These are estimation results of equation (3). The unit of analysis is a county. In columns (1) and (2), we use a county-year panel, where the change is taken between year (t-1) and t from 2000 to 2020. In columns (3) and (4), we use a county cross section, where the change is taken between 2000 and 2020.
Standard errors are clustered by states. Regressions are weighted by a county's total corn/soy area in 2000.
*** p<0.01, ** p<0.05, * p<0.1

1980 to 2000

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Corn | Soy | Corn | Soy |
| | | $\Delta\log(Yield)$ | | |
| ΔGDD above 33C | -0.0181*** | -0.0230*** | -0.0130*** | -0.000543 |
| | (0.00102) | (0.000690) | (0.00457) | (0.00309) |
| ΔGDD below 33C | 0.000365*** | 0.000386*** | -0.000106 | -0.000107 |
| | (1.85e-05) | (2.02e-05) | (0.000198) | (0.000207) |
| ΔPrecip above 55cm | -0.000682*** | -0.00114*** | 0.00134 | -0.000796 |
| | (0.000193) | (0.000167) | (0.00125) | (0.00178) |
| ΔPrecip below 55cm | 0.00546*** | 0.00602*** | 0.0106*** | 0.00619*** |
| | (0.000283) | (0.000340) | (0.00289) | (0.00223) |
| Method | First difference | | Long difference | |
| County FEs | Yes | | No | |

| | | | | |
|---|---|---|---|---|
| State FEs | No | | Yes | |
| Year FEs | Yes | | Yes | |
| Observations | 38,550 | 33,067 | 1,868 | 1,602 |
| R-squared | 0.312 | 0.365 | 0.533 | 0.357 |

These are estimation results of equation (3). The unit of analysis is a county. In columns (1) and (2), we use a county-year panel, where the change is taken between year (t-1) and t from 1980 to 2000. In columns (3) and (4), we use a county cross section, where the change is taken between 1980 and 2000.
Standard errors are clustered by states. Regressions are weighted by a county's total corn/soy area in 1980.
*** p<0.01, ** p<0.05, * p<0.1

## Pooled results

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Corn | Soy | Corn | Soy |
| | $\Delta\log(Yield)$ | | | |
| ΔGDD above 33C | -0.0157*** | -0.0178*** | -0.00704** | -0.00975*** |
| | (0.00102) | (0.000570) | (0.00343) | (0.00237) |
| ΔGDD below 33C | 0.000160*** | 0.000239*** | 5.25e-05 | -0.000252 |
| | (1.54e-05) | (1.62e-05) | (0.000220) | (0.000153) |
| ΔPrecip above 55cm | -0.000281** | -0.000123 | 0.00312* | 0.00230* |
| | (0.000120) | (0.000123) | (0.00174) | (0.00118) |
| ΔPrecip below 55cm | 0.00375*** | 0.00616*** | 0.00811** | 0.00631** |
| | (0.000248) | (0.000277) | (0.00344) | (0.00235) |
| Method | First difference | | Long difference | |
| County FEs | Yes | | No | |
| State FEs | No | | Yes | |
| Year FEs | Yes | | Yes | |
| Observations | 77,039 | 66,758 | 3,560 | 3,130 |
| R-squared | 0.269 | 0.291 | 0.263 | 0.496 |

These are estimation results of equation (3). The unit of analysis is a county. In columns (1) and (2), we use a county-year panel, where the change is taken between year (t-1) and t from 1980 to 2020. In columns (3) and (4), we use a county panel with three time periods: 1980, 2000, and 2020. The changes are the 20-year differences.
Standard errors are clustered by states. Regressions are weighted by a county's total corn/soy area in the base year.
*** p<0.01, ** p<0.05, * p<0.1

## Appendix 4

## Robustness Check of Table 5

### State FEs and year FEs

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | | | $\Delta\log\,(Corn\ yield)$ | | | |
| ΔGDD above 33C | -0.0182*** | -0.0214*** | -0.0140*** | -0.00996*** | -0.0258*** | -0.00544 |
| | (0.00490) | (0.00393) | (0.00425) | (0.00347) | (0.00415) | (0.00348) |
| ΔGDD below 33C | -0.000248 | -0.000475* | -0.000305* | -0.000279 | -0.000338* | -0.000400 |
| | (0.000158) | (0.000235) | (0.000179) | (0.000173) | (0.000178) | (0.000248) |
| ΔPrecip above 55cm | -0.00124 | 0.00170 | -0.00238** | -0.00383*** | 0.00340* | -0.00481*** |
| | (0.00131) | (0.00151) | (0.00101) | (0.00101) | (0.00169) | (0.000771) |
| ΔPrecip below 55cm | 0.00759*** | 0.00592** | 0.00655** | 0.00578*** | 0.00210 | 0.00427 |
| | (0.00209) | (0.00227) | (0.00239) | (0.00209) | (0.00354) | (0.00258) |
| County sample | Full | Below median | Above median | Full | Below median | Above median |
| Decades | 1980-2020 | 1980-2020 | 1980-2020 | 2000-2020 | 2000-2020 | 2000-2020 |
| Observations | 7,277 | 1,624 | 2,088 | 3,419 | 728 | 1,024 |
| R-squared | 0.343 | 0.292 | 0.414 | 0.356 | 0.471 | 0.337 |

Below median counties are those with below-median educational attainment and below-median average farm size in 1980. Above median counties are those with above-median educational attainment and above-median average farm size in 1980.
Standard errors are clustered by state. State FEs and decade FEs are included. Regressions are weighted by corn area in 1980 in the first three columns and by corn area in 2000 in the last three.

*** p<0.01, ** p<0.05, * p<0.1

### 20-year differences from 1980 to 2000

| | (1) | (2) | (3) |
|---|---|---|---|
| | | $\Delta\log\,(Corn\ yield)$ | |
| ΔGDD above 33C | -0.0120*** | -0.0193** | -0.00560 |
| | (0.00321) | (0.00847) | (0.00708) |
| ΔGDD below 33C | 0.000281** | 0.000866*** | -0.000231 |
| | (0.000122) | (0.000206) | (0.000229) |
| ΔPrecip above 55cm | 0.00371*** | 0.00609*** | 0.00109 |
| | (0.00102) | (0.00149) | (0.00165) |
| ΔPrecip below 55cm | 0.00865*** | 0.00305 | 0.00775** |
| | (0.00215) | (0.00330) | (0.00328) |
| County sample | Full | Below median | Above median |
| Observations | 3,530 | 770 | 1,032 |
| R-squared | 0.529 | 0.693 | 0.542 |

Standard errors are clustered by county. County FEs and year FEs are included. Regressions are weighted by corn area in 1980.
*** p<0.01, ** p<0.05, * p<0.1

## Appendix 5

## Farmland Price Dynamics

Our model predicts that land prices should decline in areas with rising climate damage. We document that farmland prices (including buildings on the land) have a more negative trend in counties affected by heat between 2000 and 2020.

| | (1)<br>2000 to 2020 | (2)<br>2000 to 2010 | (3)<br>2010 to 2020 |
|---|---|---|---|
| ΔGDD above 33C | -0.00490* | -0.00759*** | 0.000826 |
| | (0.00257) | (0.00150) | (0.00465) |
| ΔGDD below 33C | -0.000234 | 8.87e-05 | -0.000926** |
| | (0.000251) | (0.000209) | (0.000370) |
| ΔPrecip above 55cm | -0.00191 | 0.00148 | -0.00302 |
| | (0.00229) | (0.00197) | (0.00200) |
| ΔPrecip below 55cm | -0.00165 | 0.00489*** | -0.00125 |
| | (0.00334) | (0.00162) | (0.00392) |
| Observations | 2,277 | 2,316 | 2,279 |
| R-squared | 0.756 | 0.289 | 0.726 |

The dependent variable is the change in log farmland price per acre. State FEs and logged base year price are included in all columns.
Standard errors are clustered by state. All regressions are weighted by total farmland area in the base year.
*** p<0.01, ** p<0.05, * p<0.1

**Appendix 6**

**Correlates of Nonfarm Employment and Natural Disasters**

In the model in Section 2, we assume that farmers must exit farming completely to take a nonfarm job. We do not allow a 'partial exit' option for farmers. This is for simplicity to illustrate our points on extensive-margin adjustments. In reality, many farmers, especially those with small farms, hold off-farm jobs. In the table below, we show that in counties affected by climate change, fewer farmers report farming as their primary occupation, and more work off-farm for more than 200 days per year. Because high-skilled farmers are more likely to secure nonfarm employment, this pattern suggests a brain drain from the agricultural sector, as these farmers spend less time managing their farms.

| | (1) | (2) |
|---|---|---|
| | $\Delta\log\,(Farm\ operator\ count)$ | |
| | Primary job: Farming | Days worked off farm: 200+ |
| Drought | -0.00608* | 0.00116 |
| | (0.00356) | (0.00417) |
| Heat | -0.00186 | -0.000329 |
| | (0.00694) | (0.00537) |
| Hurricane | -0.0310*** | 0.0525*** |
| | (0.0109) | (0.00882) |
| Storm | -0.0156** | -0.00538 |
| | (0.00681) | (0.00736) |
| Flood | 0.0122 | -0.0488 |
| | (0.0314) | (0.0300) |
| Observations | 7,200 | 7,200 |
| R-squared | 0.778 | 0.655 |

These are estimates of equation (6). We use a county-decade panel from 1980 to 2010. The dependent variables are the changes in logged farmer count by characteristics.
We control for state FEs, decade FEs, lagged values, and the change in total operator count in the decade.
Standard errors are clustered by state.
*** p<0.01, ** p<0.05, * p<0.1