Supplemental Appendix: "Distortions, Producer Dynamics, and Aggregate Productivity: A General Equilibrium Analysis"*

Stephen Ayerst
International Monetary Fund

Loren Brandt
University of Toronto

Diego Restuccia
University of Toronto and NBER

October 2025

A Data Details

We provide additional details of the data analysis including a comparison of the Red River and Mekong deltas, farm life-cycles, land quality differences between south and north Vietnam, and differences in the composition of crop production.

A.1 Comparing the Red River and Mekong Deltas

Our baseline analysis compares the agricultural sectors in north and south Vietnam. Our theory shows how institutional differences, captured by wedges, distort farm-level decisions to lower productive investments and aggregate productivity. An advantage of our approach is that we exploit historical differences between the north and south to make within-country comparisons. Nevertheless, differences in climate, geography, and crop suitability may still imply gaps in production capabilities across regions. In this Supplemental Appendix, we

*Contact: Ayerst (stephen.b.ayerst@gmail.com); Brandt (loren.brandt@utoronto.ca); Restuccia (diego.restuccia@utoronto.ca).

show that the main empirical differences between the south and north in our analysis hold if we restrict focus to the two rice-growing delta regions: the Red River Delta in the north and Mekong Delta in the south, where technology, geographic, and crop differences are less likely to be a concern. These regions are also economically important to the agricultural sector in Vietnam and capture around 30% and 37% of the observations in our dataset.

Table A.1: Red River and Mekong Delta Comparison

	Red River Delta (North)	Mekong Delta (South)
TFP growth	7.07	8.50
Std log TFP	0.75	1.02
Std log output	0.82	1.63
Std log land	0.74	1.29
Std log employment	0.83	1.15
Measured elasticity	0.94	0.79
Std log wedge	0.73	0.85
Rice farm share	88.1	71.4

Notes: The rice farm share is the share of farms where at least 50% of the crop output value in all periods is from rice.

Table A.1 compares key moments from the two regions. The comparison shows the same patterns emphasized in our main analysis between the north and the south. Growth in the south is higher, accompanied by more dispersion in TFP, in output, and landholdings. The gap in the measured elasticity of distortions between the south and north is also larger than in the baseline analysis. Finally, while these are both primarily rice-growing regions, we find a higher share of perennial farms in the south.

A.2 Farm Life Cycles

Production over the farm life cycle. Section II in the main text reported that farm life cycles in the South relative to the North are characterized by sharper increases in productivity at young ages followed by a similar decline in productivity at older ages. This initial increase is important for explaining why, on average, farms in the South grow faster and are able to achieve higher productivities. Figure A.1 breaks productivity down into output and input dynamics over the life cycle. Dynamics for output and inputs are similar to the composite productivity measure reported in the main text.

Figure A.1: Production over the Farm Life-cycle

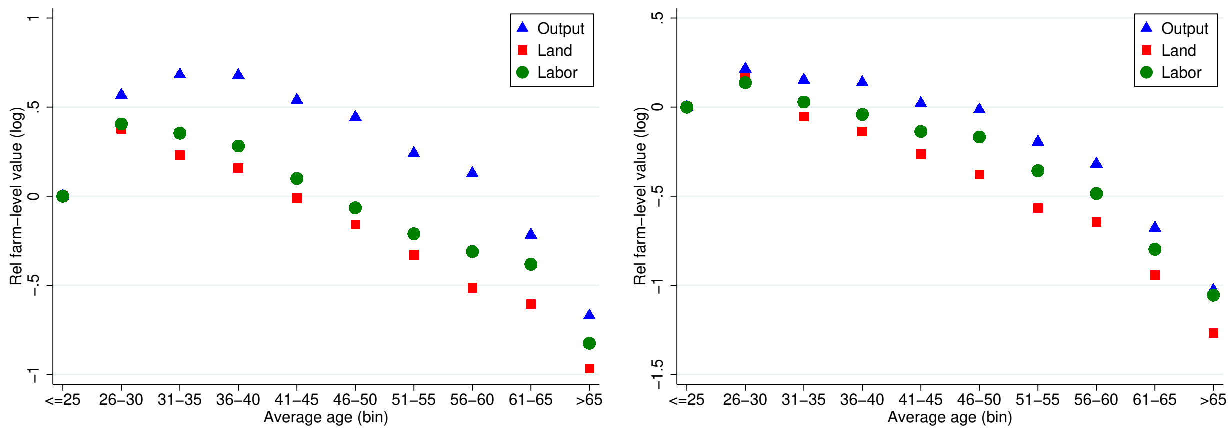

(a) South

(b) North

Notes: The figure reports the estimated of age bin fixed effects c_j^R (for $R \in \{South, North\}$) from the regression $\log Y_{f,t} = \sum_{j\in\mathcal{A}} c_j^R 1_{age_{f,t}\in j} + \Gamma^R + \Gamma_t + \varepsilon_{f,t}$ where Γ^R and Γ_t are region and year fixed effects. The coefficient estimates are normalized such that the youngest bin has value zero.

Alternative age definitions. Table A.2 reports the productivity life cycle of farms in North and South Vietnam using three different measures of age. The baseline measure, discussed in the main text, constructs household age as the average of household members weighted by their time spent working on household crops. The Household Head measure constructs age using the member identified as the household head. The Average measure constructs age as the simple average across household members. The productivity measure is normalized in each region and year such that the regressions do not capture time trends.

We find that in all three cases the two main observations in the main text hold. First, household productivity life cycles in both the North and the South display a hump-shaped pattern where households quickly increase productivity when they are young and then decline at old ages. Second, the dynamics of farms in the South are much sharper than in the North,

Table A.2: Farm Life Cycle

	(1) log TFP	(2) log TFP	(3) log TFP
Age(North)	0.0257***	0.0328***	0.0139***
	(0.00701)	(0.0116)	(0.00537)
Age(South)	0.0480***	0.0433*	0.0235***
	(0.0108)	(0.0223)	(0.00866)
Age^2 (North)	-0.000342***	-0.000368***	-0.000231***
	(0.0000743)	(0.000118)	(0.0000631)
Age^2 (South)	-0.000613***	-0.000527**	-0.000377***
	(0.000106)	(0.000221)	(0.0000939)
Age Definition	Baseline	Household Head	Average
R^2	0.0336	0.0141	0.0261
Observations	10100	9090	10389

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the household level are included in parentheses. All regressions include a region fixed effect. log TFP is normalized at the region-by-year level. Household Head measures age as the age of the household member identified as the head of household. Average measures age as the average age of all household members. Column (2) excludes households where the head of household is older than 70.

where productivity tends to be flatter over the farm's life cycle.

A.3 Differences in Land Quality

Table A.3 compares the quality of land across provinces in Vietnam from Adamopoulos and Restuccia (2022). Land quality is measured as the average potential yield of land (across cells) within the province. We focus on two measures: an average of 27 crops and wet rice, the most prevalent crop in Vietnam. We use the rainfed, low input potential yield which most closely reflects the land quality without human intervention, see Adamopoulos and Restuccia (2022) for details and discussion.

Panel A of Table A.3 describes land quality differences between the North and South for all provinces in the country. Panel B focuses on only the twelve provinces that are included in the VARHS dataset (Central Institute for Economic Management (CIEM) et al., 2018).

Panel C adjusts the mean values of land quality for the relative frequency of observations in our final dataset.

Table A.3: Comparison of Land Quality

A. All Provinces

	Mean Avg.	Std Avg.	R9010 Avg.	Mean Rice	Std Rice	R9010 Rice
North	88.8	0.4	3.2	1.6	0.8	6.6
South	87.3	0.3	2.0	2.1	0.5	3.2
Total	88.0	0.4	2.2	1.9	0.7	5.3

B. In Final Dataset

	Mean Avg.	Std Avg.	R9010 Avg.	Mean Rice	Std Rice	R9010 Rice
North	67.4	0.6	3.7	1.1	1.0	12.4
South	94.9	0.3	2.1	1.7	0.5	4.1
Total	81.2	0.5	3.5	1.4	0.9	11.3

C. In Final Dataset (observation-weighted means)

	Mean Avg.	Std Avg.	R9010 Avg.	Mean Rice	Std Rice	R9010 Rice
North	87.2	0.6	3.7	1.7	1.0	12.4
South	86.3	0.3	2.1	1.8	0.5	4.1
Total	86.8	0.5	3.5	1.8	0.9	11.3

Notes: Values calculated using provinces as unit of observation. “Avg.” refers to statistics calculated on the average potential yield of 27 common crops. “Rice” refers to statistics calculated on the average potential yield of wet rice. “Std” is the standard deviation of the log variable. “R9010” is the ratio between the 90th and 10th percentile observations. Panel C constructs the mean values using the relative frequency of farm-year observations in our data as weights.

Comparing Panel A and Panel C shows that after adjusting the means for the relative frequency of observations there is little difference between our final dataset and the average province in the North and South. The observed differences in land quality are not large enough to explain the productivity gap that we observe between farms in the North and South. Taking the production function in Section III in the main text implies that the impact of land quality on TFP requires differences to be scaled by a factor $\alpha\gamma = 0.35$. This would further reduce the potential impact of any differences between the North and the South.

A.4 Crop Differences

We categorize households as either rice farms, perennial farms, or other (annual) crop farms based on their most valuable crop grown over the survey.[1] We categories farmers as a rice or perennial farmer if more than 50% of their output value, across all years, is in rice or perennial crops. We do not impose strict annual cutoffs because of inter-cropping, crop-rotation, and the fact that farms may devote some of their land to other crops. However, cropping tends to be concentrated in these categories.[2]

Figure A.2: Farm Crop Type Distribution

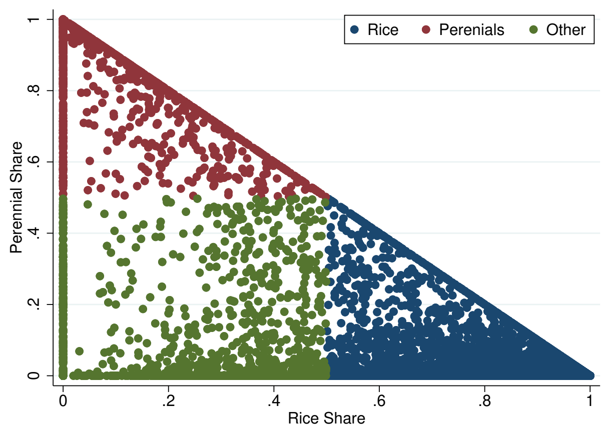

Figure A.2 summarizes the empirical distribution of households across farm types. Unsurprisingly, most households are classified as rice farms as this is the most widely produced

[1]In the survey, perennials include: fruits, coffee, tea, cocoa, cashew nuts, sugarcane, pepper, rubber, medicinal trees and plants, and other perennial crops. Other crops include maize, potatoes, sweet potatoes, cassava, peanuts, soybeans, vegetables, and other annual crops. Farmers for whom more than half of their average yearly crop output is from rice (perennials) are rice (perennial) farmers while the remainder are "other crop" farmers.

[2]For example, over two-thirds of rice farmers have a rice share over 75% and just under half of rice farmers have a rice share over 90%. For perennials, these numbers are slightly higher at 70% and 50% of farmers for the same thresholds. In addition, more than 90% of rice and perennial farm-year observations would have the same classification if classified year by year. Differences are more common in the case of other crop farmers.

crop in Vietnam. The remaining households are split between perennials and other crop farms. Other crop farms grow, on average, around 30% of value in rice, but have a higher production value in other annual crops.

Table A.4: Farm Type Comparison in South Vietnam

	(1) log Output	(2) log Land	(3) log Labor	(4) log TFP	(5) TFP Growth
Perennials	0.685***	0.676***	0.502***	0.185***	5.994***
	(0.104)	(0.0837)	(0.0688)	(0.0608)	(1.732)
Other	-0.267***	-0.0936	0.0698	-0.287***	-2.958
	(0.103)	(0.101)	(0.0759)	(0.0609)	(2.611)
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	4401	4401	4401	4382	3481

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the household level are included in parentheses. All regressions include year fixed effects.

Tables A.4 and A.5 report cross-crop differences for output, land, labor, TFP, and TFP growth by farm type. There are stark differences within regions between crops. In the South, perennials tend to outperform the other crop types in terms of production, productivity, and growth. This is consistent with the fact that perennials are cash crops that incentivize investment. Among perennials, coffee is the most important. Rice and other crops are more likely to be food crops for the household's own consumption, and underperform compared to perennials. In contrast, in the North rice tends to over perform relative to the other crops. Farms in the North also tend to be smaller in terms of land and labor inputs and output, and experience lower growth. These differences motivate our main quantitative experiment.

Figure A.3 reports the TFP distribution by crop for south Vietnam. The figure highlights that while perennial farmers are on average more productive than the other farm types, there is a significant mass of perennial farmers that are less productive than the typical rice or other crop farmer. In contrast, selection based on farmer ability (as in Adamopoulos and Restuccia, 2020) would imply a discrete productivity cutoff in contrast with the data.

Table A.5: Farm Type Comparison in North Vietnam

	(1) log Output	(2) log Land	(3) log Labor	(4) log TFP	(5) TFP Growth
Perennials	-0.578***	-0.471***	-0.408***	-0.376***	-6.037
	(0.148)	(0.114)	(0.109)	(0.0998)	(4.885)
Other	-0.222***	-0.108	-0.114**	-0.223***	-7.053***
	(0.0686)	(0.0772)	(0.0528)	(0.0453)	(2.427)
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	6347	6347	6347	6138	5033

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the household level are included in parentheses. All regressions include year fixed effects.

Figure A.3: Distributions of Farm TFP by Crop

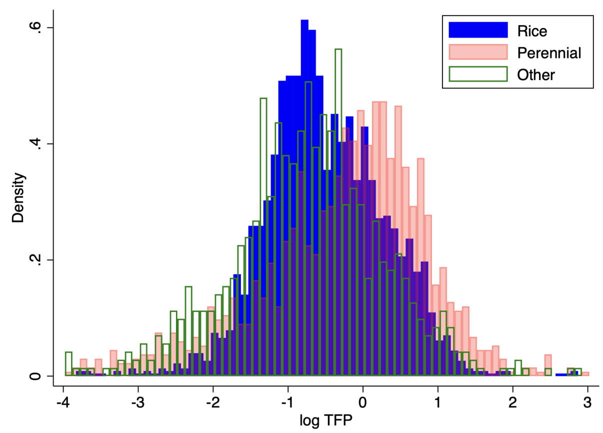

B Other Quantitative Results

We explore other quantitative implications related to measurement error, identification of model parameters, a calibration to north Vietnam, a counterfactual undistorted economy, and the asymmetric effects of correlated distortions.

B.1 Measurement Error

We examine the role of measurement error empirically and in the quantitative model.

Empirical analysis. Bils et al. (2021) develop a methodology to assess the importance of additive measurement error in misallocation models. Following Adamopoulos et al. (2022), we estimate:

$$\Delta \log y_{f,t} = \beta_1 \log Wedge_{f,t} + \beta_2 \Delta \log input_{f,t} + \beta_3 \log Wedge_{f,t} \Delta \log input_{f,t} + F_t + \upsilon_{f,t}, \tag{B.1}$$

where $input_{f,t} = \ell_{f,t}^{\alpha} n_{f,t}^{1-\alpha}$ is the Cobb-Douglas aggregate of farm inputs, F_t is a year fixed effect, and $\upsilon_{f,t}$ is an error term. Following Bils et al. (2021), the estimated value $\hat{\lambda}^{BKR} = 1 + \hat{\beta}_3/\hat{\beta}_2$ captures the ratio of dispersion in the true farm-level distortions (i.e., $\tau_{f,t}$ in the model) to the dispersion in the distortions plus dispersion in the wedge due to measurement error. Estimates of $\hat{\lambda}^{BKR}$ close to one indicate little measurement error, while values close to zero indicate that dispersion in the observed wedge is mostly due to measurement error. Without measurement error, high and low wedge farms adjust output and inputs similarly to changes in productivity and distortions, so the cross-term, β_3, is zero. With measurement error, the cross-term becomes negative as over-reporting output (under-reporting inputs) farms appear to have higher wedges but adjust inputs by less (output by more) than would be implied by the change in output (inputs).

Table B.6 reports the estimated value of $\hat{\lambda}^{BKR}$ from (B.1) and the estimated standard errors. The results indicate a relatively limited role of measurement error in both south and north Vietnam, consistent with estimates for agriculture in other contexts (Adamopoulos et al., 2022; Aragón et al., 2024).

Table B.6: Estimates of Additive Measurement Error

	South	North
$\hat{\lambda}^{BKR}$	0.906	0.987
	(0.030)	(0.022)

Quantitative analysis. A common concern in the misallocation literature is that the estimated wedges, such as those in Section II in the main text, reflect some degree of measurement error. For example, over-reported farm output would result in a higher measured TFP and wedge. We expect measurement error to be similar in both the South and North such that our main experiment—which adjusts distortions in the South to match the North—should not be drastically contaminated by measurement error. Nevertheless, measurement error can impact the level and measurement of distortions and the dynamic productivity implications. As a result, we extend our model to allow explicitly for measurement error in our calibration to limit its impact on our conclusions.

We allow for measurement error on output and the composite input. Measured output is now given by $y_{f,t} = y^*_{f,t} \exp\{\iota^{out}_{f,t}\}$ and measured inputs are given by $\ell_{f,t} = \ell^*_{f,t} \exp\{\iota^{in}_{f,t}\}$ and $n_{f,t} = n^*_{f,t} \exp\{\iota^{in}_{f,t}\}$, where we use stars to indicate unobserved true values. While we interpret $\iota^{out}_{f,t}$ and $\iota^{in}_{f,t}$ as measurement error, these variables also capture any type of shock to production or inputs that farmers are unable to fully adjust to, including, for example, unexpected weather events. In this regard, the important difference between $\iota^{out}_{f,t}$ and $\iota^{in}_{f,t}$ and the random components of productivity $v_{f,t}$ and distortions $\varepsilon_{f,t}$ is that farms cannot adjust production and input decisions based on ι's. Measured farm-level TFP and Wedge are now:

$$\text{TFP}_{f,t} = \frac{y_{f,t}}{(\ell^\alpha n^{1-\alpha}_{f,t})^\gamma} = s^{1-\gamma}_{f,t} e^{\iota^{out}_{f,t} - \gamma \iota^{in}_{f,t}}, \qquad \text{and} \qquad \text{Wedge}_{f,t} = \frac{y_{f,t}}{\ell^\alpha_{f,t} n^{1-\alpha}_{f,t}} = \frac{e^{\iota^{out}_{f,t} - \iota^{in}_{f,t}}}{1 - \tau_{f,t}}.$$

These expressions show that measured TFP and wedges depend on both the model fundamentals, captured by $s_{f,t}$ and $\tau_{f,t}$, but also the measurement errors, $\iota^{out}_{f,t}$ and $\iota^{in}_{f,t}$. Measurement error creates an upward bias in both the standard deviation of wedges and the elasticity of wedges with respect to farm-level productivity.

The measurement errors $\iota^{out}_{f,t}$ and $\iota^{in}_{f,t}$ are drawn from normal distributions with standard deviations $\sigma_{\iota^{out}}$ and $\sigma_{\iota^{in}}$. We jointly target the parameters $\sigma_{\iota^{out}}$ and $\sigma_{\iota^{in}}$ along with other model parameters in the calibration. We include moments on the standard deviation of farm land size and the correlation between within farm changes in output and changes in land to discipline the extent of measurement error. The correlation moment reflects intuition from Bils et al. (2021) that co-movements in producer outcomes can be used to assess the extent of measurement error, which tends to dampen this relationship.

B.2 Identification of Model Parameters

We expand on the discussion in Section IV.A in the main text by deriving the relationship between model parameters, discussing the sensitivity of the moments to changes in model parameters, and examining the robustness of calibration results.

Calibration moments. Our theory describes the evolution and distribution of productivities and how these relate to farm crop decisions and the institutional environment. We leverage the micro data to construct moments that describe the joint distributions of TFP and growth, crop-specific differences across farms, and distortions to discipline the model parameters. We discuss the construction of the moments and provide a qualitative discussion on their connection to parameters below.

- Avg TFP growth. The moment reports the average growth of measured farm-level TFP. In the data, farm-level growth is calculated over a two-year period as $g_{f,t} = (TFP_{f,t} - TFP_{f,t-2})/(0.5 * (TFP_{f,t} + TFP_{f,t-2}))$, averaged over all farm-years. In the simulated data, we similarly construct the growth in TFP from t to $t+2$ and report the average value over all farms that remain active into $t+2$. The moment relates to the size λ and likelihood of improving productivity, which depends on the investment technology (λ, ψ, ζ).

- Std TFP growth. The moment reports the standard deviation of $g_{f,t}$ calculated in the previous moment across all active farms in both the empirical and simulated data. While the moment depends on the size and likelihood of improving productivity, it is moment is closely related to the idiosyncratic component of productivity through σ_v and the measurement error parameters $\sigma_{\iota^{out}}$ and $\sigma_{\iota^{in}}$.

- Std log TFP. The moment reports the standard deviation of measured farm-level $\text{TFP}_{f,t}$ across all active farms in both the empirical and simulated data. The moment acts as a residual measure of farm-level TFP differences to discipline the dispersion in the permanent component of productivity σ_z but also relates to parameters that determine the productivity distribution, such as the investment technology (λ, ψ, ζ), the idiosyncratic component of productivity through σ_v and the measurement error parameters $\sigma_{\iota^{out}}$ and $\sigma_{\iota^{in}}$.

- Std log land. The moment reports standard deviation of measured farm-level land $\ell_{f,t}$ across all active farms in both the empirical and simulated data. The moment is closely linked to parameters that cause variation in productivity, such as σ_z and σ_v, or distortions, such as σ_ε. The moment also helps disciplines measurement error on inputs $\sigma_{\iota^{in}}$ by explaining cross-sectional variation in farm size not linked with productivity or distortions.

- Reg coefficient: TFP growth on log TFP. The moment measures the coefficient from regressing farm-level TFP growth on log TFP. The empirical specification is given by $g_{f,t} = \nu \log \text{TFP}_{f,t-2} + \Gamma_t + \epsilon_{f,t}$ where Γ_t is a year fixed effect and ν is the reported moment. The moment is calculated similarly using the simulated data, without the time fixed effect. The moment helps discipline the investment technology (λ, ψ, ζ) nad other parameters related to the productivity distribution, such as (σ_z, σ_v). In particular, the moment helps discipline curvature of the cost function since higher curvature ζ implies investment is less elastic to incremental profitability. Correlated distortions lead to a flattening of profitability at higher ability levels implying that higher curvature ζ increases relative investment by higher ability farms. Less negative estimates of ν then correspond to higher values of ζ.

- Top 10% land share. The moment is calculated as the share of land held by the 10% largest farmers (by land size) in the empirical and simulated data. A farmer's land size is closely related to both productivity and distortions. The moment helps discipline the size and likelihood of productivity improvements through investment, and is, consequently, closely related to the investment technology (λ, ψ, ζ). Intuitively, larger ability improvements (higher λ) by fewer farmers (higher costs ψ) leads to a more skewed productivity distribution and amore concentrated distribution of land.

- Measured elasticity. The data moment is estimated from regressing the wedge on TFP. We construct the model moment similarly by regressing measured farm-level wedges on measured farm-level TFP where we incorporate measurement error in the construction of both variables. Despite the bias, the data moment remains useful for disciplining the elasticity of distortions ρ in the model and changes in the parameter elasticity ρ closely coincide with the measured elasticity.

- Std log Wedge. The moment reports standard deviation of measured farm-level wedge across all active farms in both the empirical and simulated data. The moment helps discipline the standard deviation of the random component of wedges σ_ε that is independent of the systematic component that relates to dispersion in TFP.

- Corr($\boldsymbol{\Delta}\log \mathrm{TFP}, \boldsymbol{\Delta}\log \ell$). The moment reports the correlation of within-farm changes in measured farm-level TFP and measured farm-level land. In both the empirical and simulated data, the moments are calculated over a two-year period, such that $\Delta \log \mathrm{TFP}_{f,t} = \log \mathrm{TFP}_{f,t} - \log \mathrm{TFP}_{f,t-2}$. The moment is used to discipline the measurement error parameters $\sigma_{\iota^{out}}$ and $\sigma_{\iota^{in}}$ and builds on the intuition from Bils et al. (2021) in using within-farm changes in variables to identify measurement error. Intuitively, farms adjust land holdings with measured TFP when it captures a true change in productivity (e.g., increasing z) but not when it captures a change in measurement error (e.g., changing ι^{out}).

Derivation of model moments. We derive the relationship between the calibration moments and parameters. Most moments do not have explicit closed-form solutions and so we use this to highlight the intuition provided in the main text rather than as a proof on identification. Our quantitative analysis focuses on nine moments.

The measured elasticity of distortions and the standard deviation of measured wedges are given by:

$$elas(Wedge, TFP) = \frac{\rho(1-\gamma)^2\sigma_s^2 + \sigma_{\iota^{out}}^2 + \gamma\sigma_{\iota^{in}}^2}{(1-\gamma)^2\sigma_s^2 + \sigma_{\iota^{out}}^2 + \gamma^2\sigma_{\iota^{in}}^2},$$

$$\sigma_{\log Wedge}^2 = (1-\rho)\sigma_s^2 + (\sigma_{\iota^{out}}^2 + \sigma_{\iota^{in}}^2) + \sigma_\varepsilon^2.$$

The above expression shows that the measured elasticity of distortions depends directly on the elasticity parameter ρ but measurement error terms can bias the estimate towards one. Similarly, the standard deviation of the wedge depends on variation of the fundamental farm productivity s and the wedges σ_ε but also on the measurement error terms.

The variance of measured TFP and land are given by:

$$\sigma_{\log TFP}^2 = (1-\gamma)^2\sigma_s^2 + \sigma_{\iota^{out}}^2 + \gamma^2\sigma_{\iota^{in}}^2 = (1-\gamma)^2[\sigma_a^2 + \sigma_z^2 + \sigma_v^2] + \sigma_{\iota^{out}}^2 + \gamma^2\sigma_{\iota^{in}}^2,$$

$$\sigma_{\log \ell}^2 = (1-\rho)^2\left[\sigma_a^2 + \sigma_z^2 + \sigma_v^2\right] + \sigma_\varepsilon^2 + \sigma_{\iota^{in}}^2.$$

The above expressions show that the variance of both TFP and land depends on the underlying drivers of productivity: the farm's fundamental productivity a and the random component v. Additionally, both the TFP and labor variance depend on the measurement error term for inputs ι^{in}.

We use the change in log TFP to illustrate the identification of the model, rather than the growth rate calculated in the baseline model since it provides simpler expressions. The next three moments are equal to:

$$\mathbb{E}[\Delta \log TFP] = \int_z \sum_{i,h} x_{z,h}((1-\gamma)\log\lambda)\mu_{z,h} d\Phi_z(z),$$

$$\sigma^2_{\Delta \log \text{TFP}} = (1-\gamma)^2 \left(\int_z \sum_h \left((x_{z,h})^2 - \left(\int_z \sum_h x_{z,h} \right)^2 \right) (\log\lambda)^2 \mu_{z,h} + 2\sigma_v^2 \right)$$

$$+ 2(\sigma^2_{\iota^{out}} + \gamma^2 \sigma^2_{\iota^{in}}),$$

$$elas(1+g, TFP) = \frac{(1-\gamma)^2[(Cov(x,\log z) + Cov(x, \log a) - \sigma_v^2] - \sigma^2_{\iota^{out}} - \gamma^2\sigma^2_{\iota^{in}}}{\sigma^2_{\log TFP}},$$

where $1 + g = TFP'/TFP$. Average TFP growth depends on the step size λ and the relative incentives to improve productivity through $x^i_{z,h}$, which itself depends on a collection of parameters. Dispersion in TFP growth depends on farm-level investment choices but also on variation in the random component of productivity v and the measurement error terms. Finally, the elasticity between growth and TFP depends on the covariance between farm investment x and endogenous farm productivity z. This term is more negative if more productive farmers invest less.

We include the correlation between the change in farm-level TFP and the change in farm land to discipline the measurement error terms. The moment is given by:

$$Corr(\Delta \log \text{TFP}, \Delta \log \ell) = \frac{(1-\rho)(1-\gamma)(\sigma^2_{\Delta \log z} + 2\sigma_v^2) - 2\gamma\sigma^2_{\iota^{in}}}{\sigma_{\Delta \log TFP}\sqrt{(\sigma^2_{\Delta \log z} + 2\sigma_v^2 + 2\sigma^2_{\iota^{out}} + 2\sigma^2_{\iota^{in}})}},$$

where $\sigma^2_{\Delta \log z} = \int_z \sum_h \left[(x_{z,h})^2 - \left(\int_z \sum_h x_{z,h} \right)^2 \right] (\log\lambda)^2 \mu_{z,h} d\Phi_z(z)$. The above expression shows that measurement error for inputs enters the expression differently than the random component of productivity and the output measurement error term.

While the expression for land share is difficult to write in the full model, the intuition can be understood through a simpler model in which there is no misallocation or dispersion in the permanent or random components of farm productivity or measurement error. In this context, all farms choose a common investment rate x, then the distribution of farms across

h is approximately $\delta/(x+\delta)[x(1-\delta)/(x+\delta)]^h$. Farms with production technology h have land size λ^h. The land share held by farms with productivity above some node $\bar{h}$ is then given by:

$$\text{Land Share}(\bar{h}) \approx C \sum_{n \geq \bar{n}} \frac{\delta}{x+\delta} \left(\frac{\lambda x(1-\delta)}{x+\delta} \right)^h,$$

for constant $C = (\delta - (\lambda - 1)x)/\delta$. It is straightforward to see from the above expression that the land share by the top farms becomes larger as either farms become more likely to improve productivity (i.e., higher x) or improve productivity by more (i.e., higher λ).

Sensitivity of model moments and targets. Table B.7 summarizes the changes in moments to a 10% change in the model parameters, highlighting that the moments are highly interconnected with the set of parameters. The table also shows that no individual moment identifies an individual parameter. Nevertheless, the table shows that the chosen moments for calibration are informative about the values of parameters in the calibration. The relationship between the moments and parameters is discussed in detail in Section IV.A.

Table B.7: Sensitivity of Moments to Calibrated Parameters (%)

	ψ	ζ	λ	σ_z	σ_v	σ_ε	ρ	$\sigma_{\iota^{out}}$	$\sigma_{\iota^{in}}$
Land share	−0.1	0.6	0.1	1.7	0.9	5.2	−9.5	0.0	1.3
Reg coeff	0.0	−1.8	−0.5	−9.7	6.1	0.1	9.2	1.3	2.1
Avg growth	−2.4	4.7	5.0	−0.6	−1.4	0.9	−14.2	−0.3	−0.4
Corr(ΔTFP, $\Delta\ell$)	−0.1	0.1	0.8	0.3	51.9	−5.8	−113.6	−0.8	−47.6
Meas elas	0.1	−0.1	−0.1	−0.7	−0.3	−0.0	8.8	0.2	1.2
Std TFP	−0.0	0.0	0.1	−0.0	4.9	−0.0	−0.2	1.1	1.7
Std land	−0.1	0.4	0.2	1.7	0.8	6.0	−9.9	0.0	1.2
Std wedge	−0.2	0.8	0.5	3.7	1.7	0.9	2.4	0.6	2.2
Std growth	−0.3	1.1	0.7	5.0	2.3	0.0	−4.9	0.5	0.9

Notes: Percent change in the moments from a 10% change in each parameter relative to the benchmark calibration value. For λ the change is calculated only on the value above one.

As highlighted in Table B.7, the model parameters are jointly chosen to match the calibration moments. There are no parameters that are identified by individual moments. That said, some moments are more useful for identifying specific parameters. To explore this further, we show the sensitivity of the calibrated parameters to changes in the targeted moments, as suggested by Fujimoto et al. (2023) and Andrews et al. (2017). Table B.8 reports the sensitivity of the model parameters.

Table B.8: Sensitivity of Parameters to Moments (%)

	ψ	ζ	λ	σ_z	σ_v	σ_ε	ρ	$\sigma_{\iota^{out}}$	$\sigma_{\iota^{in}}$
Land share $\times 1.25$	38.8	-4.5	6.5	-1.2	-1.0	24.3	0.0	-2.7	-0.3
Reg coeff $+10$	185.9	-10.0	51.8	16.2	-9.9	6.8	3.8	-19.8	-19.0
Avg growth -2%	117.3	-12.3	-2.1	2.8	1.1	1.5	0.3	-2.4	-0.4
Avg growth -4%	186.2	-12.3	-16.4	8.7	-1.8	1.4	-0.2	21.1	-2.9
Corr$(\Delta\text{TFP}, \Delta\ell)$ $\times 1.5$	-10.1	0.4	-0.8	0.0	-5.3	9.9	2.5	43.7	-27.9
Meas elas -0.1	-31.5	-14.0	-8.2	-11.1	-35.1	-57.8	-19.6	111.8	-13.5
Std log land $\times 1.25$	-42.9	-33.6	-4.2	1.6	6.7	17.3	0.5	-84.7	3.9
Std log tfp $\times 1.25$	-8.8	1.4	-1.1	7.1	0.5	-4.8	-0.3	36.3	0.7

Notes: Percent change in the calibrated parameter from alternative calibration in which the indicated data moment is adjusted and all other moments are held at their benchmark values.

The table results highlight the relationship between the calibration moments and parameter values. Increasing the land share of the top 10% of farms increases the size of productivity improvements λ and the cost of investment ψ such that productivity improvements become relatively more infrequent but result in a large shift of resources. Lowering the negative relationship between growth and farm TFP or lowering the growth rate results in shifts in the investment technology parameters (ψ, λ, ζ). The correlation between within-farm changes in TFP and land is mostly absorbed by changes in the measurement error terms. Lowering the measured elasticity of distortions results in a similar, albeit larger, shift in the elasticity parameter ρ as well as compensating shifts in other parameters.

B.3 Calibration to the North

The baseline experiment applies distortions that are set to match north Vietnam to the benchmark economy, which is calibrated to match south Vietnam. We show that the counterfactual economy moves towards the north Vietnam data moments, relative to the benchmark economy. An alternative approach is to re-calibrate the model to match the full set of moments from north Vietnam and then use this to compare with south Vietnam. We explore this approach in this section.

Calibration moments and parameters. The calibration follows the same procedure as in the baseline calibration. We adjust the total stock of land to be $L = 1.10$ to reflect the smaller average farm size in the North. The jointly calibrated parameters are selected to target the same moments as in the baseline calibration, where the values for the North are reported in Table B.9.

Table B.9: Moments Calibrated to North Vietnam

	Model	Data
Avg growth (%)	2.44	2.61
Std growth	80.9	89.2
Std log TFP	0.85	0.84
Std log land	0.99	0.98
Reg coefficient: growth on log TFP	−51.4	−48.2
Top 10% land share (%)	35.7	38.3
Measured elasticity	0.974	0.961
Std log wedge	0.88	0.81
$\mathrm{Corr}(\Delta \log \mathrm{TFP}, \Delta \log \ell)$	−0.048	−0.022

The parameters in the re-calibrated model are summarized in Table B.10. Overall, the parameter values in the re-calibrated model are relatively similar to those in the baseline calibration. This reflects the overall ability of the benchmark economy to match the North data moments when the North distortions were imposed. The main difference between the North and South parameters is in the ability investment function, (λ, ψ, ζ). Relative to the South, investment in the North is substantially cheaper but also has a smaller payoff. The lower return to investment through λ explains the lower farm dynamism in the North compared with the South.

Table B.10: Parameters Calibrated to North Vietnam

Parameter		Value	Parameter		Value
Discount rate	β	0.96	Survival rate	ξ	0.955
Span-of-control	γ	0.7	Land share	α	0.5
Land	L	1.1			
Investment level	ψ	66.89	Investment curvature	ζ	2.21
Ability step size	λ	3.39			
Permanent productivity	σ_z	1.14	Random productivity	σ_v	2.02
Elasticity	ρ	0.90	Random distortion	σ_ε	0.95
Output mismeasurement	$\sigma_{\iota^{out}}$	0.24	Input mismeasurement	$\sigma_{\iota^{in}}$	0.52

Aggregate productivity. The re-calibrated model generates a productivity gap between north and south Vietnam that matches closely the data. Following equation (3) for aggregate output, aggregate total factor productivity in the calibrated economy is calculated as:

$$\frac{A^{North}}{A^{South}} = \frac{Y^{North}/(L^{North})^{\alpha\gamma}}{Y^{South}/(L^{South})^{\alpha\gamma}} = 43.0\%,$$

whereas this ratio is 42% in the data. This implies that the re-calibrated model is able to account for the entirety of productivity differences between north and south Vietnam.

B.4 Undistorted Economy

The undistorted economy represents a hypothetical first-best economy that could be achieved if all institutional distortions were removed. In practice, it is unclear whether this economy is achievable since some baseline distortions may be unavoidable. Additionally, other unmodeled factors (e.g., selection, plot-specific investment) may become more relevant when comparing a highly distorted economy with an undistorted economy. With these caveats noted, we find the undistorted economy useful as a benchmark to understand the full potential gains in productivity.

We calculate the undistorted economy by setting the parameters as in the baseline calibration and setting the elasticity of distortions to $\rho = 0$ and the random component of distortions $\sigma_\varepsilon = 0$.

Table B.11: Comparison with Undistorted Economy

	Benchmark	Undistorted Economy
Productivity	1.00	2.80
Avg growth (%)	6.23	0.55
Std growth (%)	75.6	75.2
Std log TFP	0.99	0.90
Std log land	1.19	2.76
Reg coefficient: growth on log TFP	−34.2	−38.5
Top 10% land share (%)	42.0	86.5
Measured elasticity	0.850	0.213
Std log Wedge	0.90	0.48
$\mathrm{Corr}(\Delta \log \mathrm{TFP}, \Delta \log \ell)$	0.085	0.649

Table B.11 presents the comparison of the undistorted economy with the benchmark economy.The undistorted economy is 2.8 times as productive as the benchmark economy (a productivity gain of 180%). Table 6 in the main text shows that the gains from removing static misallocation in the benchmark economy is around 75% implying that a substantial portion of the productivity gains are coming from improving the productivity distribution through higher investment in ability. Nevertheless, differences in the productivity distribution alone do not account for all the remainder gains because of complementarities between the channels.

Another noticeable difference between the benchmark and undistorted economy is in the average growth rate. This can be understood through two channels. First, removing correlated distortions causes investment in ability to become flat with respect to the farmer's ability because farmers are not disincentivized by larger distortions at higher abilities. All else equal, this causes higher ability farmers to invest more than in the benchmark economy. Second, removing distortions improves productivity and, consequently, increases the wage rate w and cost of land q resulting in lower profits for a given ability level. Lower profits disincentivize investment in ability for all farmers. The net impact is that lower ability farmers invest less in the undistorted economy while higher ability farmers invest more.

This results in both more low ability farmers and more very high ability farmers in the undistorted economy. The productivity gains are then driven by these increases in the top end of the productivity distribution, which is consistent with the concentration of agricultural production in large, highly productive farms in advanced economies.

Finally, despite setting the elasticity of distortions to $\rho = 0$, the measured elasticity remains positive and larger than zero. This is also the case for the standard deviation of wedges, which only falls to around half its initial value. These results are due to the inclusion of the input and output measurement errors, which are held at their benchmark values in the undistorted economy. Measurement error tends to have a larger bias in the undistorted economy, which is also found by Ayerst et al. (2024). We emphasize that while our main results on the difference between the South and North are not substantially affected by the extent of measurement error in the data, the role of measurement error is more substantial in the level effects relative to the undistorted economy.

B.5 Asymmetric Effects of Elasticity of Distortions

The baseline experiment shows that increasing the elasticity of distortions ρ can explain a large share of the productivity gap between the two regions, despite the increase being relatively small. Mechanically, the large productivity cost from increasing ρ is driven by the disincentivizing effect of correlated distortions on investment (Farm Ability). As ρ increases farms invest less because the incremental increase in profits becomes smaller. At the extreme, when $\rho \to 1$ farms have no incentive to invest because the entirety of additional profits is absorbed by higher distortions. This leads to an increasing impact of ρ on productivity that is maximized as ρ gets closer to one.

Table B.12 shows the asymmetric impact of increasing and decreasing ρ on productivity through each channel. The values are reported as log changes in productivity (rather than percent changes) for comparability. The difference in the effects is mainly driven by the farm ability channel due to the disincentivizing effect of ρ on ability investment.

Table B.12: Increasing and Decreasing Elasticity of Distortions

	Increase $\rho = \rho + \Delta\rho$	Decrease $\rho = \rho - \Delta\rho$
Factor Misallocation	-0.19	0.15
Farm Ability	-0.33	0.25
Total	-0.48	0.43

Notes: Values report the log change in output. The elasticity of distortions is set to the South benchmark value, and the change in ρ is set equal to $\Delta\rho = 0.1$, which is the observed difference in the measured elasticity of the North and South in the data. All other parameters are set to the benchmark calibration values.

C Model Extensions

We consider three model extensions to incorporate crop choice, to capture the full farm life-cycle dynamics of productivity, and to capture potential intergenerational transmission of ability.

C.1 Crop Choice

The data indicates that north and south Vietnam differ substantially on the crop composition of production and that the North has much more restrictive land use than the South. We evaluate the importance of differences in crop composition and land use restrictions by extending the baseline model to incorporate crop choices.

Model. We index crops by superscript $i \in \{R, P, O\}$ where farmers can choose between rice ($i = R$), perennials ($i = P$) and other crops ($i = O$). We make the following adjustments:

1. Crop-specific productivity κ^i such that productivity is now $s^i_{f,t} = z_f \times \kappa^i \times a_{f,t} \times e^{v_{f,t}}$.

2. Crop-specific preferences b^i_f for farm managers, such that the utility of farm manager of a crop i is now given by

$$U^i_f = \mathbb{E}\left[\sum_{t=0}^{\infty} (\xi\beta)^t t C^i_t\right] \times b^i_f.$$

where the farmer's idiosyncratic preferences over crop b^i_f is drawn from a Frechet distribution $H(b) = \exp\{-(b/\eta^i)^{-\theta}\}$,

3. Crop-specific distortions φ^i such that the farm's wedge is now given by $1 - \tau^i_{f,t} = [\bar{\tau}\varphi^i(s^i_{f,t})^{-\rho}e^{\epsilon_{f,t}}]^{1-\gamma}$.

4. Government-imposed crop restriction that force farms to grow rice with probability ω. This reflects a direct cropping restriction imposed by the Vietnamese government on individual farms that are quantitatively important for aggregate production (see Le, 2020) and would not be captured by the standard misallocation wedge.

Let $\bar{V}^i_z$ denote the expected value of a new farm with crop i, permanent productivity z, and ability-level $h = 0$ before the shock (v, ε) is realized. Farmers with the government-imposed crop restriction do not choose their crop and are forced to produce rice, $i = R$. For unrestricted farmers, the crop decision is

$$\max_{i \in \mathcal{I}} \bar{V}^i_z \times b^i.$$

The resulting share of farmers that grow crop i is equal to

$$\Omega^i_z = \begin{cases} \omega + (1-\omega)\frac{(\eta^i \bar{V}^i_z)^\theta}{\sum_{i' \in \mathcal{I}}(\eta^{i'}\bar{V}^{i'}_z)^\theta} & \text{for } i = R \\ (1-\omega)\frac{(\eta^i \bar{V}^i_z)^\theta}{\sum_{i' \in \mathcal{I}}(\eta^{i'}\bar{V}^{i'}_z)^\theta} & \text{for } i \neq R \end{cases}. \tag{C.2}$$

The fraction of (unrestricted) farmers that choose a specific crop depends on both the relative expected value of growing that crop $\bar{V}^i_z$ and the relative difficulty of growing that crop, captured by the preference parameter η^i, where θ determines the elasticity of farmers to these factors.

The impact of crops on farm productivity and, consequently, value also implies that the farm value $V^i_{z,h}(v, \epsilon)$ and policy $x^i_{z,h}$ are also crop specific. This also requires that a crop-specific distribution $\mu^i_{z,h}$ be calculated for each crop type. The aggregate distribution of farms across types is then $\mu^i_{z,h}\Omega^i_z\phi_z(z)$. The expressions for $V^i_{z,h}(v,\epsilon)$,$x^i_{z,h}$, and $\mu^i_{z,h}$ are straightforward extensions of the main text.

Production of the agricultural good is given by

$$Y = \left[\frac{\left(\int_{v,\varepsilon} e^{\gamma\varepsilon+(1-\rho\gamma)v} d\Phi_v(v) d\Phi_\varepsilon(\varepsilon)\right) \int_z \sum_i \sum_h (\varphi^i)^\gamma (z\kappa^i\lambda^h)^{1-\rho\gamma} \mu^i_{z,h}\Omega^i_z d\Phi_z(z)}{\left(\int_{v,\varepsilon} e^{\varepsilon+v(1-\rho)} d\Phi_v(v) d\Phi_\varepsilon(\varepsilon)\right)^\gamma \left(\int_z \sum_i \sum_h \varphi^i (z\kappa^i\lambda^h)^{1-\rho} \mu^i_{z,h}\Omega^i_z d\Phi_z(z)\right)^\gamma}\right]$$
$$\times N_F^{1-\gamma}(L^\alpha N_W^{1-\alpha})^\gamma,$$

where it is clear from the above expression that aggregate output and misallocation now depend on the distribution of farms across crops Ω^i_z.

Calibration. We calibrate the model parameter to the moments described in the text. In addition, we need to choose values for the crop-specific parameters related to the distortions φ^i, the preferences η^i, and productivity κ^i as well as the crop restrictions ω and the elasticity of preferences across crops θ. The distortions κ^i and government-imposed restrictions ω are set to directly match the data. The remaining parameters are set in two stages exploiting that the crop-specific preferences η^i can always be chosen to set the crop distribution to match the empirical distribution. The first stage fixes the crop distribution and calibrates the nine parameters in the baseline calibration along with the crop-specific productivities κ^i to match the original nine moments and the relative measured TFP for each of the three crops. The second stage sets the crop-specific preferences η^i to match the crop distribution and sets the the preference elasticity θ to minimize the overall size of the crop-specific preferences (i.e., minimum $\sum_i(\eta_i - 1)^2$). Table C.13 reports the parameters.

The crop-specific distortions are set to match the coefficients from regressing wedges on farm crop types and TFP. We set the government-imposed restriction $\omega = 21\%$ for south Vietnam.[3]

[3] In the data, farmers with multiple plots may report that only some plots face restriction while crop restrictions are a binary variable in the model. To construct the data moment, we take a land-weighted average of crop-restrictions for each farmer and then average this value over all farmers in south Vietnam

Table C.13: Model Parameters

Parameter		Value	Parameter		Value
Discount rate	β	0.96	Survival rate	ξ	0.955
Span-of-control	γ	0.7	Land share	α	0.5
Land	L	2.77			
Crop restriction	ω	0.21	Crop-specific distortion	φ^i	(1.00 , 1.61 , 1.12)
Crop preference elasticity	θ	2.55	Preference shifter	η^i	(1.00 , 0.58 , 0.90)
			Crop-specific productivity	κ^i	(1.00 , 1.36 , 0.41)
Investment level	ψ	1.51	Investment curvature	ζ	1.74
Ability step size	λ	1.44			
Permanent productivity	σ_z	1.51	Random productivity	σ_v	1.81
Elasticity	ρ	0.80	Random distortion	σ_ε	1.00
Output mismeasurement	$\sigma_{\iota^{out}}$	0.25	Input mismeasurement	$\sigma_{\iota^{in}}$	0.46

Counterfactual experiment. We now describe the counterfactual experiment where we adjust distortions in the benchmark economy to match measured distortions in the North. The benchmark economy has four parameters related to distortions: (1) the elasticity of distortions to farm-level productivity ρ; (2) crop-specific distortions φ^i; (3) the government-imposed crop restriction ω; and (4) the random component of distortions σ_ε. Table C.14 summarizes the values estimated for the first three of these for the counterfactual experiment. Other parameters, including the random component of distortions, are held fixed at the benchmark economy values.

Table C.14: Counterfactual Distortions

		Benchmark (South)	Counterfactual (North)
Elasticity	ρ	0.80	0.91
Crop-Specific Distortion	φ^i	(1.00 , 1.61 , 1.12)	(1.00 , 0.68 , 0.92)
Crop Restriction (%)	ω	21	39

Notes: Distortions are ordered for Rice, Perennial, Other crop farm types.

Table C.15 compares the calibration moments and agricultural productivity in the bench-

in the dataset. This implies, for example, a farmer in the data with one-third of their land restricted is captured in the model by three farmers with one farmer having the entirety of their crop choice restricted and the other two being unrestricted. The comparison between the North and the South remains similar with alternative constructions of this moment.

mark and counterfactual economies as well as in the data for north Vietnam. Our main result is the implied productivity gap between the counterfactual and benchmark economy, which is a measure of how much of the observed productivity gap can be explained by differences in the distortions between the North and South. We find that aggregate TFP in the counterfactual economy (North) is 59% of the benchmark economy (South), implying that the model accounts for almost two-thirds ($61\% \approx \log(0.59)/\log(0.42)$) of the productivity gap between the North and the South. The model also accounts almost entirely for differences in the farm share by crop in the data as well as around half of the relative measured TFP of perennial farmers. Other comparisons between the counterfactual and the North remain similar to the baseline experiment.

Table C.15: Results with Crop Choice of Counterfactual Economy Compared with Benchmark and North Data

	Benchmark	Counterfactual	Data
Productivity	1.00	0.59	0.42
Avg growth (%)	6.21	4.62	2.61
Std growth (%)	75.6	75.4	89.2
Std log TFP	0.99	0.91	0.84
Std log land	1.18	1.01	0.98
Reg coefficient: growth on log TFP	−34.2	−39.2	−48.2
Top 10% land share (%)	42.2	36.2	38.3
Measured elasticity	0.850	0.961	0.961
Std log wedge	0.90	0.93	0.81
Corr($\Delta \log \text{TFP}, \Delta \log \ell$)	0.084	−0.058	−0.022
Relative measured TFP	(1.00 , 1.20 , 0.75)	(1.00 , 0.98 , 0.74)	(1.00 , 0.69 , 0.80)
Farm share by crop (%)	(49.1 , 33.1 , 17.8)	(77.6 , 3.6 , 18.8)	(75.1 , 5.0 , 19.9)

Notes: Where applicable, moments are first reported for rice farms, followed by those for perennials and then other crop farms.

We also decompose the relative contributions of factor misallocation, crop choice, and farm ability. Table C.16 summarizes the loss in aggregate productivity from changing channels individually from the benchmark economy to match the counterfactual economy. The sum of the losses does not equal the total gap between the benchmark and counterfactual economy because of interactions between the channels. For example, changes in the ability or crop distributions also affect the potential scope for factor misallocation through their effect on the productivity distribution. We discuss each channel and its calculation below.

Table C.16: Output Loss by Channel

	Change in Output (%)	
Factor Misallocation	-19.0	
Crop Choice	-10.2	
Farm Ability	-31.4	
Sum of Channels	-60.6	
Total		-40.5

Notes: The change in output is equivalent to the change in productivity since aggregate inputs are constant.

We calculate the loss from factor misallocation as the change in aggregate output when distortions, $\tau^i_{f,t}$, are adjusted to match the counterfactual economy but the crop and ability distributions remain fixed at the benchmark distributions. Starting from the distribution of farm-level productivities $s^i_{f,t}$ in the benchmark economy, we recalculate the distortions $\tau^i_{f,t}$ that farmer f would receive with the counterfactual correlation ρ and crop-specific distortions φ^i. We find that factor misallocation lowers agricultural output by 19.0%, accounting for just under half of the productivity gap between the counterfactual and benchmark economies. Partitioning the interaction effects proportionately to each channel, factor misallocation accounts for around one-third ($\approx -19.0/-60.6$) of the resulting productivity loss. Factor misallocation has a negative interaction with the other two channels explaining why the sum of the losses from the individual channels is larger than the total loss in productivity. This is because factor misallocation has a smaller effect on aggregate productivity when productivity is less dispersed, as is the case in the counterfactual economy.

We calculate the loss from the crop distribution as the change in aggregate output when the crop shares Ω^i_z are adjusted to match the counterfactual economy. We fix the within-crop ability distribution $\mu^i_{z,h}$ to that of the benchmark economy. However, the aggregate ability distribution, equal to $\mu^i_{z,h}\Omega^i_z\Phi_z(z)$, changes due to changes in the crop distribution. Average ability falls since perennial farmers are, on average, higher ability than rice farmers and the experiment redistributes around 30% of farmers from perennials to rice. We find that the change in the crop distribution has a relatively small contribution to the overall gap between the counterfactual and benchmark economies compared with the other channels. Nevertheless the output loss from the change in the crop distribution is a non-trivial -10.2%.

We calculate the loss from farmer ability as the change in aggregate output when the ability distribution is adjusted to match the counterfactual economy. We adjust the farmer ability distribution $\mu^i_{z,h}$, conditional on crop i and permanent productivity z, to the counterfactual economy and hold the crop shares Ω^i_z fixed to that of the benchmark economy. The ability distribution in the counterfactual economy results from lower investment by farmers due to more correlated distortions, which makes higher ability levels less profitable. The change in farm ability generates a loss in output of 31.4%, accounting for around three-quarters of the productivity gap between the counterfactual and benchmark economies.

We also examine the role of the individual distortions, rather than channels. We measure the impact on output of individual distortions from unilaterally changing ρ, φ^i, or ω in the benchmark economy to match the estimated distortions parameters for the North. Table C.17 summarizes the results.

Table C.17: Output Loss from Individual Distortions

	ρ	φ	ω	(ρ, φ, ω)
Change in Output (%)	-38.6	-7.4	-1.7	-40.5

Consistent with our baseline model, the main driver of the gap between the benchmark and counterfactual economies is the elasticity of distortions. The elasticity of distortions has a large impact on factor misallocation by reallocating resources from high productivity to low productivity farms. The increase in elasticity also dampens the increase in profits associated with increasing farm productivity, which results in weaker incentives for farmers to invest in ability or select crops based on market factors as opposed to preferences. Our results point to a large productivity effect from seemingly small variation in the elasticity of distortions ρ between the North and the South due to the asymmetric productivity effects from changes in ρ that are magnified as ρ approaches one.

The crop-specific distortions have a more moderate effect on the productivity gap between the benchmark and counterfactual economies. Crop-specific distortions increase factor misallocation by reallocating resources across different farm types. Crop-specific distortions also affect the relative incentives for farmers to invest in improving ability since it changes the relative profitability of crops. Finally, crop-specific distortions affect the crop distribution through changing the relative market value of farm types.

The difference between the North and South in government-imposed crop restriction has the smallest impact on productivity. Part of the reason is that crop restrictions are implemented before farmers make crop choices implying that some farmers would choose to grow rice independent of the restriction. Since around half of farmers grow rice in the benchmark economy, this reduces the impact by a comparable amount. As a back-of-the-envelope calculation, the change in productivity is approximately equal to reducing the productivity of 6% (the change in $\Delta\omega = 0.2$ times the 29.5% share of perennial farmers) of farmers by 20% (the measured productivity of perennials farmers relative to rice farmers). In contrast, Le (2020) finds a larger cost of government-imposed crop restrictions on agricultural productivity due to a negative correlation between crop restrictions and plot suitability for rice, which is absent in our analysis.

C.2 Farm Life Cycle Dynamics

In Section II in the main text, we show that farm productivity is hump shaped over the life cycle with the productivity of young farms increasing quickly and then deteriorating as the farm reaches older ages. Our baseline model focuses on the initial buildup of farm productivity through investments in farm ability but does not account for the decline in productivity of older farmers. We show that the main model results are relatively unchanged if we extend the model to incorporate this feature.

Model. We extend the model to allow for life cycle dynamics following a similar structure of aging as in Acemoglu et al. (2018). Farmers initially enter as young age ($j = Y$) farmers and then with probability ϕ transition to old age ($i = O$) farmers. Old age acts as an absorbing state that all farmers eventually reach (if they do not exit), albeit at different points of time. Young farmers operate as described in the main text while old farmers have ability $a_{f,t} = 1$, regardless of their previous ability h or investment. Farm ability is given by

$$a^j_{f,t} = 1_{j=Y}\lambda^h + 1_{j=O}. \tag{C.3}$$

The age structure allows us to capture the dynamics observed in the data in a reduced form. Intuitively, the transition to old age could capture the deterioration of physical abilities of older farmers.

The model is otherwise as described in the main text. The equilibrium characterization is similar with the exceptions that the value function now accounts for the possibility of transitioning to old age and the type distribution now accounts for farms in old age.

Quantitative analysis. We consider an alternative calibration of the model to focus more on the farm productivity life cycle. The preference shifters η^i, preference curvature θ and all other parameters follow the baseline calibration procedure. We re-calibrate the jointly chosen parameters $\{\lambda, \psi, \zeta, \kappa^i, \sigma_z, \sigma_v\}$ as well as the transition probability ϕ to target a new set of moments. In addition to the baseline moments, we add two moments: (i) the average productivity of 36-40 year old farmers is 0.54 log points higher than the average productivity of 25 and younger farmers and (ii) the average productivity of 65 and older farmers is equal to the average productivity of 25 and younger farmers (both from Figure 3 in the main text). We also remove the moments on the average TFP growth of farms. Table C.18 reports the parameter estimates.

The estimated transition to old age is around 2% indicating that farms spend an average of 50 years at young age. Relative to the baseline parameters, the estimated ability step size

Table C.18: Parameters with Farm Life Cycle Targets

Parameter		Value	Parameter		Value
Discount rate	β	0.96	Survival rate	ξ	0.955
Span-of-control	γ	0.7	Land share	α	0.5
Land	L	2.77			
Transition to old	ϕ	0.02			
Investment level	ψ	1.69	Investment curvature	ζ	1.35
Ability step size	λ	1.52			
Permanent productivity	σ_z	1.49	Random productivity	σ_v	1.58
Elasticity	ρ	0.79	Random distortion	σ_ε	1.08
Output mismeasurement	$\sigma_{\iota^{out}}$	0.40	Input mismeasurement	$\sigma_{\iota^{in}}$	0.41

increases in the extended model, which is necessary to offset some of the negative growth from aging. Figure C.4 reports the relationship between farm TFP and age, highlighting a similar hump-shaped dynamic as the data.

Figure C.4: Farm Productivity Life Cycle

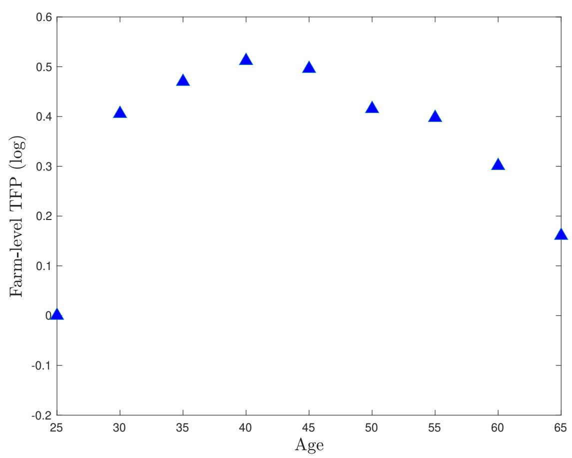

Notes: Age bins are {≤25, 26-30, 31-35, 36-40, 41-45, 46-50, 51-55, 56-60, >60} and plotted according to the oldest age in the group and 65 for the oldest group. The average of farm-level log TFP is calculated using simulated data (as described in Section IV.A in the main text) for 100,000 farms.

The shape of farm dynamics is by construction since key features of the life-cycle productivity profile are targeted in the calibration. The purpose of the recalibration is to assess its impact on the main results. The main experiment adjusts distortions in the benchmark economy, calibrated to south Vietnam, to match distortions in north Vietnam. Table C.19

compares moments in the extended model benchmark economy, calibrated to the South, with the counterfactual economy, which adjusts distortions to match the values in the North. The table highlights that the results in the extended model are in line with the baseline model in the main text. Aggregate productivity drops by 43%, similar to the baseline experiment (41%). We also find similar dynamics when comparing the other moments with the baseline experiment. The counterfactual economy is almost able to replicate entirely the farm crop distribution, in addition to accounting for around half of the change in the standard deviation of TFP, the regression's coefficient of growth on TFP, and the top 10% land share.

Table C.19: Results with Life Cycle Targets of Counterfactual Economy Compared with Benchmark and North Data

	Benchmark Economy	Counterfactual Economy	North Data
Productivity	1.00	0.65	0.42
Avg growth (%)	4.32	3.13	2.61
Std growth (%)	76.4	76.0	89.2
Std log TFP	0.98	0.90	0.84
Std log land	1.18	1.05	0.98
Reg coefficient: growth on log TFP	−34.2	−39.6	−48.2
Top 10% land share (%)	42.3	37.3	38.3
Measured elasticity	0.858	0.960	0.961
Std log wedge	0.90	0.92	0.81
Corr($\Delta \log \text{TFP}, \Delta \log \ell$)	0.083	−0.033	−0.022

The table also shows the average growth rate in the benchmark calibration economy and the counterfactual economy. Unlike the main text, this is no longer a moment that is directly targeted in the calibration. As discussed in Section IV.E in the main text, the growth rate of productivity is potentially related to factors unrelated to ability investment in the model. Table C.19 provides an extreme view on the magnitude of these other factors since it attributes none of the non-life cycle growth to farmer investment. The growth rate lies within the bounds considered in the robustness exercise in Section IV.E in the main text. Overall, the evidence in Table C.19 is reassuring about the robustness of the main results.

C.3 Entrant Ability

In the baseline model, entrants start at the lowest ability node and then progress to higher nodes through investment. In practice, we might expect that some ability is passed on through generational learning, such that some entrants are more productive than others. We show a simple extension of the model that incorporates this feature and find that the quantitative results remain relatively unchanged.

Model. Rather than entering with ability $a = \lambda^0$, we allow entrants to draw ability $a = \lambda^h$ where $h \in \{0, 1, ..., \tilde{h}\}$ is drawn from distribution $m(h, \tilde{h})$ and $\tilde{h}$ is the ability node of the exiting farmer the entrant replaces. We include $\tilde{h}$ as the upper bound to capture the intuition that entrants are learning from the previous generation of (exiting) farmers and note that the distribution would be the same if we instead had entrants learn from active farms, since exit is random. Since our goal is to show the robustness of the baseline results, we set the distribution of entrant productivity to be uniform between 0 and $\tilde{h}$, where we expect that this would tend to overstate the persistence in ability over time.

The model is otherwise as described in the main text. The equilibrium characterization is similar with the exception that the type distribution now accounts for entry into higher nodes. We assume that entrants draw their predecessors' permanent productivity z and preferences η^i and that entrants can only deviate from the crop choice of their predecessor by accepting ability $h = 0$, for tractability. However, this assumption has little impact on the results.

Quantitative analysis. Given that the model parameters are the same as in the main text, the calibration procedure is unchanged. The main experiment adjusts distortions in the benchmark economy, calibrated to south Vietnam, to match distortions in north Vietnam. Table C.20 compares moments in the extended model benchmark economy, calibrated to the South, with the counterfactual economy, adjusting distortions to match values in the North.

The results show a similar productivity loss (57%) as in the benchmark calibration.

Table C.20: Results with Entrant Ability of Counterfactual Economy Compared with Benchmark and North Data

	Benchmark Economy	Counterfactual Economy	North Data
Productivity	1.00	0.63	0.42
Avg growth (%)	6.19	5.17	2.61
Std growth (%)	74.3	74.2	89.2
Std log TFP	0.98	0.91	0.84
Std log land	1.19	1.05	0.98
Reg coefficient: growth on log TFP	-34.5	-39.0	-48.2
Top 10% land share (%)	42.0	37.3	38.3
Measured elasticity	0.854	0.958	0.961
Std log wedge	0.90	0.93	0.81
$\mathrm{Corr}(\Delta \log \mathrm{TFP}, \Delta \log \ell)$	0.086	-0.053	-0.022

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, reallocation and growth. *American Economic Review*, 126(4):1374–1443.

Adamopoulos, T., Brandt, L., Leight, J., and Restuccia, D. (2022). Misallocation, selection, and productivity: A quantitative analysis with panel data from china. *Econometrica*, 90(3):1261–1282.

Adamopoulos, T. and Restuccia, D. (2020). Land reform and productivity: A quantitative analysis with micro data. *American Economic Journal: Macroeconomics*, 12(3):1–39.

Adamopoulos, T. and Restuccia, D. (2022). Geography and agricultural productivity: Cross-country evidence from micro plot-level data. *The Review of Economic Studies*, 89(4):1629–1653.

Andrews, I., Gentkow, M., and Shapiro, J. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *Quarterly Journal of Economics*, 132:1553–1592.

Aragón, F. M., Restuccia, D., and Rud, J. P. (2024). Assessing misallocation in agriculture: plots versus farms. *Food Policy*, 128:102678.

Ayerst, S., Nguyen, D. M., and Restuccia, D. (2024). The micro and macro productivity of nations. Technical report, National Bureau of Economic Research.

Bils, M., Klenow, P., and Ruane, C. (2021). Misallocation or mismeasurement? *Journal of Monetary Economics*, 124:39–56.

Central Institute for Economic Management (CIEM), Institute of Labour Science and Social Affairs (ILSSA), Ministry of Planning and Investment (MPI), and UNU-WIDER (2018). Vietnam access to resources household survey (varhs) 2006-2016. Panel data collected jointly by CIEM, ILSSA, MPI, and UNU-WIDER covering 2006-2016.

Fujimoto, J., Lagakos, D., and Van Vuren, M. (2023). Macroeconomic effects of ’free’ secondary schooling in the developing world. Technical report, National Bureau of Economic Research.

Le, K. (2020). Land use restrictions, misallocation in agriculture, and aggregate productivity in vietnam. *Journal of Development Economics*, 145:102465.