

Risk-Sharing and Land Misallocation*

Davide Pietrobon

Lund University

Alessandro Ruggieri

CUNEF Universidad

March 7, 2025

Abstract

We study the impact of incomplete consumption risk-sharing on land misallocation in rural economies. We develop a general equilibrium model of land cultivation choices, where heterogeneous agricultural households face idiosyncratic output shocks and insure themselves by participating in a risk-sharing arrangement. Incomplete insurance distorts households' choices, leading them away from maximizing expected incomes and resulting in land misallocation. Using the latest ICRISAT panel data from rural India, we quantify the losses attributable to limited risk-sharing. Completing insurance markets leads to output and welfare gains of 19% and 29%, respectively. Improving the functioning of consumption insurance markets within an otherwise undistorted economy can yield gains comparable to those achieved by removing distortions in factor markets.

Keywords: Misallocation, risk-sharing, agriculture, productivity, welfare

JEL Classification: O11, D61, Q12, D52

*Pietrobon: Department of Economics, Lund University, Tycho Brahes väg 1, Lund, Sweden, d.pietrobon91@gmail.com. Ruggieri: Department of Economics, CUNEF Universidad, Calle Leonardo Prieto Castro 2, Madrid, Spain, alessandro.ruggieri@cunef.edu. We thank Giacomo De Giorgi, Andre Groeger, Maia Güell, Nezih Guner, Jonna Olsson, Nicola Pavoni, Josep Pijoan-Mas, Sevi Rodríguez-Mora, and Raul Santaaulalia-Llopis for their useful comments. The usual disclaimers apply.

1 Introduction

Markets in the developing world are plagued by several frictions, and insurance markets are no exception to this general rule. Barriers to risk-sharing affect the allocation of consumption (Townsend, 1994; Udry, 1994; Fafchamps, 2011), production choices (Benjamin, 1992), temporary migration (Morten, 2019), and engagement in non-farm entrepreneurship (De Giorgi et al., 2024). Incomplete insurance also diminishes the incentives for acquiring risky inputs—an argument originally advanced by Arrow (1971)—lowering productivity and increasing consumption inequality among farmers in developing countries (Donovan, 2021).

In this paper, we show that imperfections in insurance markets, taking the form of incomplete consumption risk-sharing, have implications for the allocative efficiency of land among farmers, resulting in large output and welfare losses. A growing literature highlights how farm-specific distortions contribute to factor misallocation in agriculture (Adamopoulos and Restuccia, 2014; Adamopoulos et al., 2022).¹ In the presence of limited consumption risk-sharing, land misallocation can exist *irrespective* of other market frictions. Our quantitative findings for rural India imply that completing insurance markets can substantially reduce land misallocation, leading to output and welfare gains of 19% and 29%, respectively.

In village economies, shocks from harvest failures, illness, and pests leave households vulnerable to severe hardship. Insurance against idiosyncratic income fluctuations often relies on informal arrangements like gift exchanges and personal loans.² The literature indicates that imperfections in these arrangements are pervasive: households are unable to fully insure against idiosyncratic risks (Townsend, 1994; Udry, 1994; Fafchamps, 2011). Lack of insurance not only affects consumption but it can also have a distortionary effect on the allocation of factors of production (Foster and Rosenzweig, 2010; Donovan, 2021). Building on this evidence, we study the impact that limited insurance has on *land* markets, focusing on allocative efficiency—the potential to redistribute cultivated plots among farms thereby increasing overall agricultural yields.

To illustrate our line of reasoning, consider a village economy with a fixed supply of land that is bought and sold (or rented in and out) in an undistorted, competitive market. Under full insurance, household-farms' production decisions are separable from their consumption, ensuring that, in equilibrium, each household chooses how much land to

¹The literature highlights the impact of various factors, including inheritance laws, tax and subsidies, and land and tenancy regulations, as potential sources of distortions to the allocation of land across farms in developing countries. See Subsection 1.1 for a review.

²See Dercon (2002) for a discussion of the sources of idiosyncratic income risks and coping strategies in rural economies.

cultivate to maximize expected profits. Thus, each farmer is driven to the familiar condition of equating the expected marginal product of land to its price, resulting in an efficient allocation of land and maximal aggregate expected output. Under incomplete insurance, the “separation property” breaks apart: households’ equilibrium land choices are not generally characterized by an expected profit maximization condition, which prevents the equalization of the expected marginal products of land across farms. Thus, imperfections in insurance markets lead to land misallocation and lower aggregate output.

Building on the foundational works of Singh et al. (1985) and Singh et al. (1986), we outline a general equilibrium model of risk-sharing in which household-farms with heterogeneous productivity levels insure against idiosyncratic output shocks by sharing the incomes they generate from operating their farms. Each household chooses how much land to buy before the shocks are realized. We characterize the equilibrium land allocation across a range of risk-sharing levels, ranging from full to no insurance. Besides decreasing the expected utility of buying land, lower insurance weakens the link between farm productivity and landholdings. Under full insurance, each household-farm’s equilibrium land choices maximize its expected income, making it impossible to redistribute land from one household to another without lowering aggregate expected income. Incomplete insurance distorts these choices away from expected income maximization by increasing the weights that households attach to states of the world in which income is low. These distortions imply that, in equilibrium, the expected marginal products of land are not equalized across households—i.e., land is misallocated.

We test our model using the latest ICRISAT monthly panel data (2009–2014) from the Indian semi-arid tropics, offering evidence linking risk-sharing to land misallocation. First, we measure risk-sharing by estimating the elasticity of household consumption with respect to idiosyncratic income shocks for each village and year. We find evidence of limited consumption insurance in rural India, consistent with the literature: on average, 22.5% of idiosyncratic income fluctuations are passed through to consumption. Second, we quantify land market misallocation in each village and year using two metrics: the correlation between total land cultivated and household-farm physical productivity, and the variance of the marginal product of land for each village-year pair. The former is a well-known measure of allocative efficiency in the land market, where a higher value indicates that more productive farms cultivate more land (Chen et al., 2023). The latter quantifies the deviation from a benchmark scenario in which land’s marginal products are equalized across farms, indicative of an efficient allocation of land to production units (Restuccia and Rogerson, 2017). Consistent with our theory, we find a significant positive correlation between risk-sharing and the correlation between productivity and land cultivated, and a significant negative correlation between risk-sharing and the variance of the

marginal product of land.

Finally, we leverage the structure of our model to quantify output and welfare gains resulting from improving the functioning of consumption insurance markets in village economies. While parsimonious, our model successfully replicates the negative correlation between risk-sharing and land misallocation across villages as an untargeted moment. Armed with the structural estimates, we conduct a counterfactual analysis to explore the impact of improving consumption insurance markets in Indian villages. We examine how these improvements affect the functioning of land markets and contribute to gains in aggregate output and welfare. Completing insurance markets leads to output and welfare gains of 19%, and 29%, respectively. Output per unit of land increases by 45% under full insurance. This figure is comparable to other estimates in the literature: e.g., Adamopoulos et al. (2022) finds that eliminating farm-specific distortions in rural China increases agricultural TFP by 53%. We check the robustness of our findings to incorporating farm-specific distortions, which capture the impact of land and output market frictions that disproportionately affect high-productivity households. We find that the aggregate output and efficiency gains from completing insurance remain substantial, even in this alternative framework. Thus, our counterfactual exercise allows us to conclude that imperfections in *consumption insurance* markets can be as important as farm-specific distortions in explaining the gains from reallocating inputs across farms in developing countries.

1.1 Related literature

Our paper belongs to the growing literature on misallocation of inputs in agriculture. Gollin et al. (2002) and Restuccia et al. (2008) emphasize the role of the agricultural sector in economic development and its importance in explaining cross-country productivity and income differences. The broad theme of factor misallocation and its influence on cross-country productivity differences is explored in Restuccia and Rogerson (2008), Restuccia and Rogerson (2013), and Restuccia and Rogerson (2017). Chen et al. (2023) find that capital and operational land size are essentially unrelated to farm productivity in Malawi, suggesting the existence of misallocation in the land market. Chen et al. (2022) show how land rental market imperfections in Ethiopia lead to land misallocation, highlighting the output and welfare gains from land certification reforms. Acampora et al. (2022) provide experimental evidence that cultivation rights decrease land misallocation in Kenya. Adamopoulos et al. (2022) argue that within-village frictions in the capital and land markets, linked to land institutions, disproportionately constrain productive farm-

ers in rural China.³ This body of literature generally explains factor misallocation as a consequence of generic distortions in input or output markets (i.e., “wedges”), or institutions that constrain the choices of productive firms. Our research enriches this narrative by introducing a potential micro-foundation for these wedges, highlighting the role of incomplete consumption insurance markets.

Deviations from perfect risk-sharing within village economies are well documented (see Townsend (1994), Udry (1994), and Fafchamps (2011), among others). A body of work has provided several explanations for imperfect risk-sharing, rationalizing them as consequences of primitive frictions such as action unverifiability (Ligon, 1998), limited commitment (Ligon et al., 2002), hidden income (Kinnan, 2021), and localized information constraints (Ambrus et al., 2022). The degree of risk-sharing within a village can affect several aspects of its economy. Morten (2019) studies the interaction between informal insurance and temporary migration as a self-insurance strategy. Mazur (2023) explores the relationship between risk-sharing and irrigation investments. Pietrobon (2024) examines how informal insurance impacts fertilizer use when fertilizer is risky and effort-complementary, and insurance may crowd out effort supply. Donovan (2021) examines the impact of insurance on the use of agricultural intermediates, and suggests that completing financial markets allows farmers to invest in risky inputs, leading to significant increases in labor productivity and input share. We build upon a mechanism similar to that in Donovan (2021) to argue that the lack of insurance might distort farmers’ land cultivation choices. However, our emphasis is distinct: rather than focusing on how imperfect insurance can decrease investments in land, we highlight how these imperfections result in land misallocation.

Finally, our paper contributes to the understanding of how land gets allocated to farmers in developing countries. In the semi-arid tropics of India, land markets exhibit a rich diversity, with many farm households engaging in buying or selling of land, or participating in the land rental market to some extent (see Ray (1998), Chapter 12). The salience of rental markets is emphasized in the literature on sharecropping practices (e.g., Lamb (2003)). Our research intersects with this topic by exploring the interaction between imperfect consumption insurance and input allocation in the land market.

³Misallocation of inputs in agriculture extends beyond the markets for capital and land: for example, Adamopoulos and Restuccia (2022) estimate substantial aggregate productivity gains from the spatial reallocation of crop production.

2 Model

We analyze a static economy in which households with heterogeneous productivity face idiosyncratic output shocks and can insure against these shocks by relying on a risk-sharing arrangement. Our theoretical framework is grounded in the seminal works of Singh et al. (1985) and Singh et al. (1986), which examine the behavior of households acting simultaneously as consumers and producers. To simplify the exposition, we consider an environment where land is the sole factor of production.⁴ Each household operates a farm and decides how much land to purchase for cultivation before its output shock is realized. This choice affects the distribution of the income generated by the farm, which is calculated as the value of agricultural output net of the cost of acquiring land. The risk-sharing pool allows households to share their incomes to hedge against the idiosyncratic output shocks.⁵ We use this model to illustrate how the degree of risk-sharing affects misallocation in the land market. Appendix A.1 contains all the proofs. Appendix A.2 discusses some of our modeling choices. In the discussion that follows, when we refer to households, we specifically mean agricultural households that also operate a farm.

Consider a static economy populated by a unit measure of household *types* indexed by i . For each type i , there is a unit mass of ex-ante identical households. Each household type i is characterized by a productivity level, θ_i , and initial land holdings, $\tilde{\ell}_i$. Let the total quantity of land available in the economy be $L = \int \tilde{\ell}_i di$.^{6,7} Households have identical preferences over consumption, represented by a constant-relative risk averse (CRRA) utility function with a coefficient of relative risk aversion σ . Households of the same type, which are ex-ante identical, differ ex-post only with respect to the realization of an idiosyncratic (household-specific) output shock, ρ , drawn from a cumulative distribution function $Q_\rho(\rho)$ and support on some interval $[\underline{\rho}, \bar{\rho}] \subset \mathbb{R}_+$.

⁴In Appendix A.3 we show that our results extend to economies that include additional inputs, such as capital and labor. Moreover, our model can be extended to include an output shock with both aggregate and idiosyncratic components: since all households receive the same aggregate shock, they can engage in risk-sharing to fully insure against the output variation that arises from the idiosyncratic components of the shocks.

⁵We abstract from modeling borrowing or savings decisions, which are alternative means for households to self-insure. In the context of rural India, households' ability to borrow and save appears to be heavily constrained (Rosenzweig and Wolpin, 1993). However, as demonstrated in Appendix A.4, the effect of incomplete consumption insurance on factor misallocation carries over to environments where agents can borrow and save through a risk-free asset. See also Armangué-Jubert et al. (2025).

⁶With a fixed land supply, risk-sharing does not affect the aggregate size of land cultivated, as in equilibrium the land price adjusts to balance supply and demand. In this way, we can theoretically isolate the effect of risk-sharing on land misallocation.

⁷Our model assumes homogeneous land quality for the sake of simplification, but the main results of the paper would remain valid even if land parcels had heterogeneous productivity. In the empirical analysis, we account for heterogeneous land quality based on observable characteristics.

A household of type i produces output according to the following decreasing-returns-to-scale production function:

$$y_{i\rho} = \theta_i \rho \ell_i^\alpha,$$

where ℓ_i is land cultivated by a household of type i , and $\alpha \in (0,1)$ denotes the land share—the elasticity of agricultural yields with respect to land cultivated.⁸ Let r be the price of land, and let $\pi_{i\rho} = \theta_i \rho \ell_i^\alpha - r(\ell_i - \tilde{\ell}_i)$ denote the income of a household of type i under output shock realization ρ . In the following, we assume that $\underline{\rho}$ is high enough so that π_{it} is bounded away from zero for all possible land cultivation choices.⁹ All the land cultivated by each household is bought and sold in a land market before the output shocks are realized. Unlike most of the misallocation literature, we assume this market to be competitive and frictionless. Additionally, there are no distortions or “taxes” in the output market. These assumptions enable us to distinguish our findings from the more conventional narratives in the misallocation literature, which typically attribute land misallocation to land or goods market distortions.

2.1 Full insurance vs. no sharing

To isolate the impact of insurance on land misallocation, we begin by comparing an economy with complete markets, i.e., with full insurance, to one in which households are hand-to-mouth, i.e., with no risk-sharing. Starting from the former, let $c_i(\rho)$ denote the consumption of a household of type i when the state of the world is ρ , where ρ represents the collection of realizations of the output shock for each household in the economy, drawn from the joint cumulative distribution function $Q_\rho(\rho)$.¹⁰ Moreover, let $\mathbf{c}(\rho) = (c_i(\rho))_i$ and $\boldsymbol{\ell} = (\ell_i)_i$ represent the collections of consumptions (under state of the world ρ) and land cultivation choices of all household types. To characterize an allocation of resources under complete markets, we solve the following planner’s problem

⁸Since households of the same type are ex-ante identical and make land cultivation choices before the output shocks are realized, these choices are identical for all households of the same type. Hence, referring to ℓ_i as the land cultivated by a household of type i is unambiguous.

⁹This assumption guarantees that household consumption is always positive, regardless of the insurance regime, ensuring that utility remains well-defined. Alternatively, positive cash-on-hand can be ensured by introducing an asset and allowing households to borrow and save (see Appendix A.4).

¹⁰The careful reader will observe that our notation implies identical household consumption for households of the same type, conditional on the realization of the output shock. This assumption holds if the equilibrium allocation of resources under full insurance can be computed as the solution to a planner’s problem with a weighted utilitarian social welfare function, using type-specific Pareto weights. We invoke this assumption below.

for a given collection of type-specific Pareto weights $(v_i)_i$:

$$\max_{(c(\rho))_{\rho, \ell}} \int v_i \int \frac{(c_i(\rho))^{1-\sigma}}{1-\sigma} dQ_{\rho}(\rho) di,$$

subject to the land availability constraint

$$\int \ell_i di = \int \tilde{\ell}_i di = L$$

and the feasibility constraint

$$\int \int c_i(\rho) dQ_{\rho}(\rho) di = \int \int y_{i\rho} dQ_{\rho}(\rho) di.$$

Under full insurance, households can completely eliminate the effects of idiosyncratic output shocks. Each household consumes a fixed fraction of the constant aggregate output, where this fraction is proportional to its Pareto weight.¹¹ An optimal consumption allocation satisfies the well-known Borch rule, which states that the ratio of any two households' marginal utilities of consumption is constant across all states of the world.

Claim 1. *Under full insurance, each household consumes a constant fraction of aggregate output, with the fraction being proportional to its Pareto weight.*

An optimal consumption allocation under full insurance ensures that each household's consumption remains constant across all states of the world. Consequently, the planner can disregard how land cultivation decisions affect the distribution of consumption: each household is allocated an amount of land such its expected marginal product equals its shadow price. In a decentralized complete-market economy, this outcome would result from the separation theorem, which states that each household-farm makes production decisions to maximize its expected income. (Bardhan and Udry, 1999). Given that the expected marginal products of land are equalized across households, an allocation of land under full insurance features no *ex-ante* misallocation and maximizes aggregate expected output.

Claim 2. *Under full insurance, the expected marginal products of land are equalized across households and aggregate expected output is maximized.*

Next, we consider the allocation of land that obtains in a competitive equilibrium under no sharing. When risk-sharing is absent, the problem of a household of type i

¹¹Pareto weights are type-specific, meaning that full risk-sharing does not necessarily result in an egalitarian allocation of consumption. In particular, households with larger land endowments or higher productivity may systematically consume more than those with smaller endowments or lower productivity.

reads as follows:

$$\max_{c_i(\rho), \ell_i} \int \frac{(c_i(\rho))^{1-\sigma}}{1-\sigma} dQ_\rho(\rho)$$

subject to the budget constraint

$$c_i(\rho) = y_{i\rho} - r^{IM} (\ell_i - \tilde{\ell}_i),$$

where r^{IM} denotes the equilibrium price of land under no sharing (incomplete markets). Without risk-sharing, household's consumption and its marginal utility depend on the realization of its output shock. This dependency distorts households' land cultivation decisions, implying that an equilibrium land allocation does *not* maximize expected income for each household. In particular, because households internalize how land cultivation choices impact the distribution of consumption across different states of the world, the expected marginal product of land they cultivate will not equate its market price. When this is the case, there exists land misallocation—redistributing land across households can increase aggregate expected output.

Claim 3. *Under no sharing, there is ex-ante land misallocation.*

The key takeaway from Claim 3 is that distortions in *consumption insurance* markets (i.e., lack of insurance) alone are sufficient to cause *ex-ante* misallocation in the *land* market.

2.2 Partial insurance

More broadly, we can investigate the impact of risk-sharing on land misallocation for any degree of risk-sharing. Specifically, we explore the relationship between the elasticity of consumption with respect to own income—used as a measure of lack of insurance—and misallocation in the land market, as measured by the extent to which the marginal returns of land are distorted away from its price across household types and states of the world.

We consider an environment with partial insurance, an intermediate situation between the full insurance and no sharing scenarios discussed in the previous section. To model partial insurance, we define the following consumption function for a household i :

$$c_i(\rho) = \exp \left\{ \beta \log(\pi_{i\rho}) + (1 - \beta) \log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int \pi_{j\rho} dQ_\rho(\rho) dj \right) \right\}. \quad (1)$$

In this formulation, β represents the elasticity of consumption with respect to individual income, while $1 - \beta$ is the elasticity of consumption with respect to aggregate income.

Under full insurance, $\beta = 0$; under no sharing, $\beta = 1$. Any β value between these extremes represents varying degrees of partial insurance, with higher β values indicating worse insurance. Beyond providing a flexible way to model the relationship between household consumption and income, the function in Equation (1) maps into the standard regression equation commonly estimated in the literature to test for efficient risk sharing in village economies (Townsend, 1994; Chiappori et al., 2014). The following theorem shows that misallocation in the land market decreases with the degree of insurance.

Theorem 1. *Land misallocation increases in the elasticity of consumption with respect to own income, β .*

In the full insurance benchmark, where each household’s consumption is constant across all states of the world, the expected marginal products of land are equalized across households. As we deviate from this benchmark, the distribution of each households’ marginal utility of consumption increasingly reflects the impact of its realized output shocks. This is isomorphic to imposing distortions that affect the marginal return on land in each state. Thus, as households’ marginal utilities become increasingly tied to their realized output shocks, land allocation decisions deviate further from the full insurance benchmark.¹²

2.3 Robustness of theoretical results

The model in the previous section focuses on a setting where land is the only factor of production. This assumption is made purely for conceptual clarity: in Appendix A.3, we show that all of our results extend to a model that includes additional factors of production, such as capital and labor. Here, the key point is that as long as *some* factors of production are chosen before the realization of output shocks, incomplete consumption insurance distorts their allocation away from one that maximizes aggregate expected output. This result holds even if other inputs, such as labor, can be adjusted ex post. Moreover, insurance market incompleteness alone is sufficient to generate factor misallocation, even if *all* factor markets operate without frictions. Finally, as shown in Appendix A.4, the effect of incomplete consumption insurance on factor misallocation carries over to environments where households can borrow and save through a risk-free asset.

¹²Our model shows that enhancing risk-sharing reduces land misallocation in equilibrium, but it is silent about the distributional implications of such improvements. While a reduction in risk-sharing could theoretically cause land to shift from less to more productive households, this situation would still lead to greater misallocation. This result follows from the fact that under full insurance, the expected marginal products of land are equalized. Therefore, any distortion in households’ choices due to decreased risk-sharing would lead to deviations from this benchmark, exacerbating misallocation.

3 Risk-sharing and misallocation in Indian villages

3.1 Background and data

We use household panel data collected under the Village Dynamics in South Asia (VDSA) project by the International Crop Research Institute for the Indian Semi-Arid Tropics (ICRISAT). The data are derived from detailed survey interviews conducted between 2009 and 2014 in 18 villages in the Indian semi-arid tropics. Some components of the survey were administered monthly, while others were administered annually. These data allow us to construct *monthly* measures of consumption and income for households in different villages.¹³ This feature offers the advantage of estimating the level of risk-sharing specific to each village and year. We exploit this possibility later in the section, where we relate the level of risk-sharing in each village and year to land misallocation at the village-year level. The data include information from 40 randomly selected households in each village, stratified by landholding size, and including 10 landless laborers, 10 small farmers, 10 medium farmers, and 10 large farmers. The data provide detailed information on the quantity and value of all inputs and outputs in farm activities, as well as expenditures and incomes. We refer to Townsend (1994), Mazzocco and Saini (2012), and Morten (2019) for more detailed descriptions of the data.¹⁴ Appendix B reports a detailed description of the variables used in the analyses.

3.2 Land distribution

Table 1 presents the distribution of cultivated land, measured in hectares, across farms in the Indian semi-arid tropics and compares it to those from Malawi, Belgium, and the United States, as documented in the 1990 World Census of Agriculture (Adamopoulos and Restuccia, 2014).¹⁵

¹³As in Section 2, the term ‘household’ refers to an agricultural household operating a farm.

¹⁴As pointed out by Mazzocco and Saini (2012), it can be difficult to compare some of the information contained in the data (e.g., expenditures) across households and over time, since (1) the frequency of the interviews varies, and (2) the interview dates differ across respondents. Some recall periods can be longer than a month (e.g., a household in Aurepalle reported the amount spent on rice from July 1 to November 8, 2009). Hence, it is impossible to determine how the information provided is distributed over the months that make up recall periods longer than a month. Fortunately, from 2010 onward, the survey gives information on the month to which every piece of information refers. Therefore, we drop the observations that pertain to the year 2009.

¹⁵In the World Census of Agriculture, “[a]n agricultural holding is an economic unit of agricultural production under single management comprising all livestock kept and all land used wholly or partly for agricultural production purposes, without regard to title, legal form or size” (<https://www.fao.org/4/x0187e/x0187e01.htm>). Given this definition, when comparing the ICRISAT data to the World Census of Agriculture, it is more appropriate to consider farm size in terms of cultivated land rather than owned area.

Table 1: Farm size distributions (% of farms by size)

Farm size (hectares)	2009-2014 ICRISAT		1990 World Census of Agriculture		
	India	Malawi	Belgium	United States	
≤ 1	17.93	77.7	14.6	0.0	
1–2	23.88	17.3	8.5	0.0	
2–5	38.26	5.0	15.5	10.6	
5–10	13.45	0.0	14.8	7.5	
≥ 10	6.49	0.0	46.6	81.9	
Average	3.82	0.7	16.1	187.0	

Notes: This table presents the percentage distribution of farm sizes in hectares for India, Malawi, Belgium, and the United States. Data for India are refer to the 2009-2014 ICRISAT panel data. The sample refers to Indian farms with positive amounts of cultivated land. Data for Malawi, Belgium, and the United States are from the 1990 World Census of Agriculture, as documented in Adamopoulos and Restuccia (2014).

Compared to Belgium and the United States, the distribution of land cultivated is more left-skewed. In our sample, around 80% of farms cultivate less than 5 hectares of land, and nearly 95% of the farmlands are under 10 hectares. The average land size in our sample is 3.8 hectares, against a value of 0.7 hectares for Malawian farms, and much larger values of 16.1 and 187.0 hectares for Belgian and American farms, respectively.

3.3 Farmers' physical productivity

For our analysis, it is essential to accurately estimate each household-farm's physical productivity, also known as TFP-Q. We later use these estimates to infer land misallocation. Measuring productivity is a fundamental step in virtually all assessments of factor misallocation, whether direct or indirect (Restuccia and Rogerson, 2017). In turn, nearly all productivity measurements require imposing some structure, specifically by specifying a production function. To assess household-farm productivity, we estimate the agricultural production function outlined in Appendix A.3, which incorporates capital and labor alongside land as factors of production. Specifically, consider the following agricultural production function:

$$y_{i\tau} = A_{v\tau} \theta_i \rho_{i\tau} k_{i\tau}^{\alpha^k} h_{i\tau}^{\alpha^h} \ell_{i\tau}^{\alpha^\ell}, \quad (2)$$

We also maintain that cultivated area is inherently a more accurate measure of farm size because it avoids measurement errors arising from not including plots that are cultivated but not owned including plots that are owned but not cultivated.

where $y_{i\tau}$ denotes the physical quantity of output (measured in kilograms) produced by household i in year τ ; $A_{v\tau}$ captures factors common to village v and year τ ; $k_{i\tau}$ is total capital expenditure; $h_{i\tau}$ denotes total hours of family labor dedicated to farming activities; $\ell_{i\tau}$ represents total land cultivated, adjusted for differences in quality.¹⁶ We decompose the village-level aggregate productivity term as $A_{v\tau} = e^{\kappa\tau} e^{\theta \text{rain}_{v\tau}}$, where $\kappa\tau$ capture common year-specific factors and $\text{rain}_{v\tau}$ denotes the amount of rain (in millimeters) in village v and year τ .¹⁷ We assume that

$$\theta_i = e^{\mu_i}, \quad (3)$$

where μ_i represents permanent unobserved heterogeneity specific to household i . Similarly, we let the output shock, $\rho_{i\tau}$, to be equal to

$$\rho_{i\tau} = e^{\epsilon_{i\tau}}, \quad (4)$$

where $\epsilon_{i\tau}$ is an unobserved error term. Our approach involves treating the output shock, $\rho_{i\tau}$, as a residual after accounting for farm-specific fixed effects and other variation in output originating from observable sources. After taking logs and re-arranging terms, Equation (2) becomes

$$\log(y_{i\tau}) = \alpha^k \log(k_{i\tau}) + \alpha^h \log(h_{i\tau}) + \alpha^\ell \log(\ell_{i\tau}) + \mu_i + \kappa\tau + \theta \text{rain}_{v\tau} + \epsilon_{i\tau}. \quad (5)$$

Finally, we assume $\epsilon_{i\tau}$ to be randomly distributed and estimate Equation (5) using OLS.

In Table 2, we report the dispersion of the estimated farm productivity, $\log \hat{\theta}_i$ among Indian household-farms, and compare it to the same dispersion across farms in Malawi (Chen et al., 2023) and in the US (Adamopoulos and Restuccia, 2014), as well to the dispersion of TFP-Q across manufacturing plants in the US and India, as reported by Hsieh and

¹⁶Specifically, we assume that $\ell_{i\tau} = q_{i\tau} a_{i\tau}$, where $a_{i\tau}$ is the area cultivated by household i in year τ , and

$$\log q_{i\tau} = \delta_1 \text{depth}_{i\tau} + \delta_2 \text{slope}_{i\tau} + \delta_3 \text{fertility}_{i\tau} + \delta_4 \text{degradation}_{i\tau},$$

where $\text{depth}_{i\tau}$, $\text{slope}_{i\tau}$, $\text{fertility}_{i\tau}$, and $\text{degradation}_{i\tau}$ represent measures of the average soil depth, slope, fertility, and degree of degradation, respectively, for the plots cultivated by household i in year τ (see Appendix B for additional details on the construction of these variables). In Appendix C.4, we show that the findings in this section are robust to an alternative definition of land quality. Although measuring land quality is notoriously difficult, and any model relating land quality to observable and unobservable characteristics is susceptible to misspecification errors, any attempt to estimate agricultural production functions cannot sidestep it.

¹⁷Rainfall shocks are major sources of transitory variation in agricultural output in semi-arid tropical India, where the vast majority of land plots are rain-fed. To measure them, we use daily recordings of rainfall levels at the nearest weather station to each village and derive the total annual rainfall for each village by summing these daily measurements over the year.

Table 2: TFP-Q dispersion across farms and manufacturing firms

	Farms			Manufacturing firms		
	India 2010-2014	Malawi 2010-2011	US 1990	India 1987	China 1998	US 1977
St.dev., log	1.08	1.18	0.80	1.16	1.06	0.85
75-25 log ratio	1.02	1.39	1.97	1.55	1.41	1.22
90-10 log ratio	2.50	2.89	2.50	2.77	2.72	2.22

Notes: The first column reports statistics of the estimated farm productivity using the 2009-2014 ICRISAT panel data for the Indian semi-arid tropics. The second column reports statistics of farm productivity in Malawi from Chen et al. (2023). The third column reports statistics of farm productivity in the United States from the calibrated distribution in Adamopoulos and Restuccia (2014) to U.S. farm-size data. The third and fourth columns report statistics of manufacturing plants' productivity in Hsieh and Klenow (2009). St.dev. refers to the standard deviation of log productivity; 75-25 is the log difference between the 75 and 25 percentile; 90-10 is the log difference between the 90 and 10 percentile.

Klenow (2009). TFP-Q is largely dispersed across Indian farms: the standard deviation of log productivity (1.08) is comparable to that of farms in Malawi (1.18) and significantly larger than that of farms in the US (0.80). Moreover, differences in TFP-Q account for differences in agricultural output across Indian farms (see Table C1 in Appendix C.1).

3.4 Land misallocation

We now present direct evidence of land misallocation among households in rural India. To do so, we employ data on operational (cultivated) landholdings, $\ell_{i\tau}$, together with our estimates of physical productivity, $\hat{\theta}_i$, and construct two measures of allocative efficiency in the land market. The first is the village-year unconditional correlation between the log of operational land size and the log of farm productivity:

$$\text{corr}_{v\tau} \left[\log \ell_{i\tau}, \log \hat{\theta}_i \right].$$

Larger misallocation implies a lower correlation between farm size and TFP-Q. The second measure is the dispersion in the marginal product of land in each village-year pair:

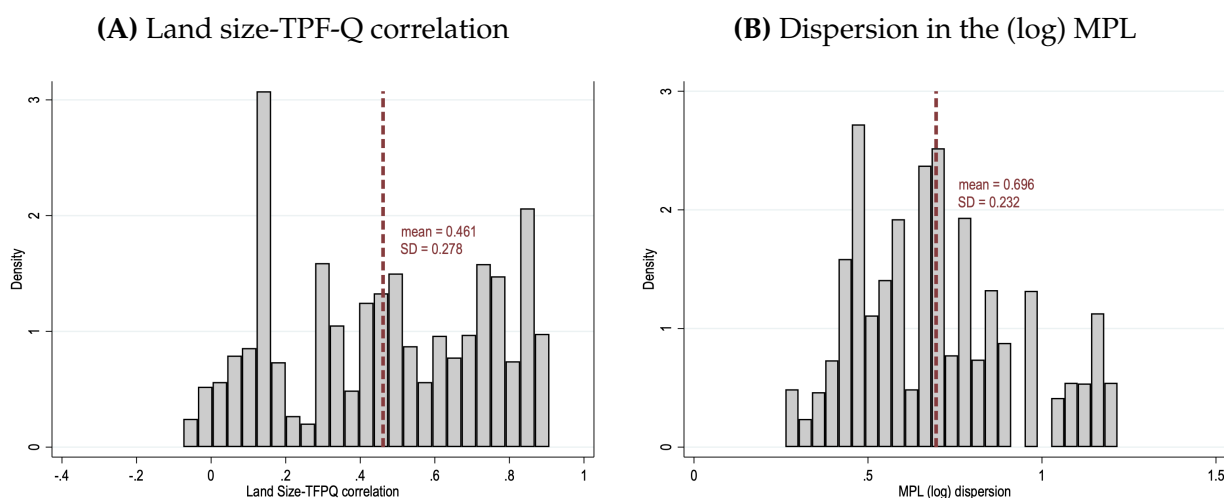
$$\text{st.dev}_{v\tau} \left[\log \text{MPL}_{i\tau} \right],$$

where $\text{MPL}_{i\tau}$ is constructed using the production function specified in Equation (2). Large misallocation implies a higher dispersion of the marginal product of land across farms

(Restuccia and Rogerson, 2017).

There is a large variation in land misallocation across villages and years. Panel A in Figure 1 displays the distribution of $\text{corr}_{v\tau} [\log \ell_{i\tau}, \log \hat{\theta}_i]$ across villages and years. The correlation ranges from negative values in some village-year pairs to nearly 1 in others. That is, while in certain village-year pairs more productive household-farms cultivate, on average, *less* land than less productive households, in other village-year pairs, there is an almost one-to-one relationship between operational landholdings and physical productivity, indicating minimal land misallocation. Panel B in Figure 1 reports the distribution $\text{st.dev}_{v\tau} [\log \text{MPL}_{i\tau}]$ across villages and years. Similarly to Panel A, this dispersion is small in certain village-year pairs, indicating little misallocation, while much greater in others.

Figure 1: Land misallocation in each village



Source: VDSA survey (ICRISAT) and own calculations.

What accounts for the variation in land misallocation across villages and years? The theory developed in Section 2 suggests that misallocation in the land market is negatively correlated with the ability of households to share idiosyncratic risks. In what follows, we present evidence of imperfect risk-sharing among households, and show that land misallocation negatively correlate with the ability to share risks within villages.

3.5 Risk-sharing

It is commonly acknowledged that risk-sharing within villages in developing countries tends to be incomplete. This holds true for the ICRISAT villages as well. To see this, we consider how the per-capita consumption for household i in month t , c_{it} , comoves with its per-capita income, π_{it} . We start by outlining the construction of these two vari-

ables, which are critical to our measure of risk-sharing: the elasticity of consumption with respect to idiosyncratic income shocks. Rather than providing a direct measure of household consumption, the ICRISAT data record detailed individual expenditures on food and non-food items, as well as the monetary value of home-produced commodities that are consumed. This granular level of detail for each consumption item is one of the key advantages of this dataset, as it allows us to compile a comprehensive measure of household consumption by aggregating these recorded values at the household-month level. Notice that, as a result, our measures of household consumption are never imputed from household income, not even for the poorest households. This feature stands in contrast to many household surveys, which often rely on income-based measures of living standards (Deaton and Zaidi, 2002). To derive a measure of monthly household income, we adopt the budget-constraint approach of Mazzocco and Saini (2012). In particular, full income is calculated as total expenditure minus resources borrowed from various sources, plus resources saved in different accounts or lent to others, plus transfers given out, minus transfers received, plus taxes, and minus subsidies. To convert monthly household consumption and income into per-capita terms, we adjust them using the age-sex index proposed by Townsend (1994).¹⁸ Then, we consider the following model:

$$\log c_{it} = \beta \log \pi_{it} + \zeta_i + \zeta_{vt} + \epsilon_{it}. \quad (6)$$

In Equation (6), $\log c_{it}$ and $\log \pi_{it}$ denote the log per-capita consumption and log per-capita income, respectively, for household i in month t ; ζ_i are household fixed effects; and ζ_{vt} represents village-month fixed effects that capture the average resources available to each village in each month. We can interpret $1 - \beta$ as the level of risk-sharing in village economies, where a higher β means a higher elasticity of consumption with respect to idiosyncratic income shocks (indicating a lower degree of risk-sharing). Under perfect risk-sharing, household income should not affect household consumption, conditional on total resources at the village-month level.

Full risk-sharing is rejected.¹⁹ On average, 22.5% of idiosyncratic income fluctuations are passed through to consumption. These values are aligned with what the literature has already documented for Indian villages using alternative empirical specifications (Townsend, 1994; Ravallion and Chaudhuri, 1997; Morduch, 2005).

The average elasticity of consumption with respect to idiosyncratic income shocks in equation (11) hides considerable variation in risk-sharing across villages and years. To uncover this variation, we estimate the degree of risk-sharing across households within

¹⁸See Appendix B for additional details on the construction of the household consumption and income variables.


¹⁹See Table C2 in Appendix C.2.

each village *separately*. More specifically, we estimate the following equation independently for each village v and year τ :

$$\log c_{it} = \beta_{v\tau(t)} \log \pi_{it} + \zeta_i + \zeta_t + \epsilon_{it}. \quad (7)$$

Here, with two slight abuses of notation, index i denotes households within village v , and index t corresponds to the months within year $\tau(t)$, where $\tau(t)$ specifies the year associated with month t .²⁰ As usual, c_{it} and π_{it} indicate the consumption and income of household i during month t , respectively, and ζ_i and ζ_t are household and month fixed effects (where, again, note that ‘months’ here refers to the months within the specific year $\tau(t)$). A high estimate of $\beta_{v(t)}$ indicates that *within village v during year $\tau(t)$* , the elasticity of household consumption in response to idiosyncratic shocks to household income is high—i.e., risk-sharing is low in village v and year $\tau(t)$.

Figure 2: Estimated degrees of risk-sharing in each village



Notes: This figure reports the distribution of risk-sharing parameters estimated for each village and year using equation (7), $\hat{\beta}_{v\tau}$, for each v and τ . Source: VDSA survey (ICRISAT) and own calculations.

Figure 2 plots the distribution of the estimated elasticities of consumption with respect to idiosyncratic income shocks, $\hat{\beta}_{v\tau}$, for each village-year pair. The average estimated $\hat{\beta}_{v\tau}$ across villages and years is 0.223, indicating that, on average, a 1% idiosyncratic increase

²⁰For example, if t corresponds to October of the year 2010, then $\tau(t) = 2010$.

in income results in a 0.223% increase in consumption. As shown in the figure, there is considerable variation in risk-sharing across villages and years, with our estimates ranging from full insurance ($\hat{\beta}_{v\tau} \approx 0$) to several others showing very high elasticities of consumption with respect to idiosyncratic income shocks ($\hat{\beta}_{v\tau} > 0.6$).

3.6 Linking land misallocation to risk-sharing


We are now ready to test whether a higher degree of risk-sharing within villages and years is associated with better allocative efficiency in the land market. To do so, we relate our two measures of misallocation in each village and year, obtained in Section 3.4 and denoted by $\omega_{v\tau}$, to our estimates of the elasticities of consumption with respect to idiosyncratic income shocks at the village-year level, $\hat{\beta}_{v\tau}$, obtained in Section 3.5. In particular, we estimate the following equation:

$$\omega_{v\tau} = \gamma \hat{\beta}_{v\tau} + \varphi_v + \kappa_\tau + \varphi_v \times \tau + \epsilon_{v\tau}, \quad (8)$$

where φ_v and κ_τ are village and year fixed effects, and $\varphi_v \times \tau$ are village-specific linear year trends.

Figure 3: Risk-sharing and land misallocation

(A) Land size-TPF-Q correlation vs. risk-sharing (B) Dispersion in the MPL vs. risk-sharing



Notes: Panel A scatters the unconditional relation between the degree of risk-sharing, β , and the correlation between log farm size and log TFP-Q. Panel B scatters the unconditional relation between the degree of risk-sharing, β , and the standard deviation of the log marginal product of land. In both panels, each dot refers to the average village-year pair in a given 4 percent bin of the estimated risk-sharing. Source: VDSA survey (ICRISAT) and own calculations.

Figure 3 binscatters the unconditional relation between the estimated degrees of risk-sharing, $\hat{\beta}_{v\tau}$, and

1. the correlation between log farm size and log TFP-Q, $\text{corr}_{.v\tau} [\log \ell_{i\tau}, \log \hat{\theta}_i]$, in Panel A;
2. the standard deviation of the log marginal product of lands, $\text{st.dev}_{.v\tau} [\log \text{MPL}_{i\tau}]$, in Panel B.

For each dot, the x -axis shows the average estimated risk-sharing for village-year pairs within a given percentile bin, while the y -axis gives the average measure of land misallocation. As we move from village-year pairs with no risk-sharing ($\hat{\beta}_{v\tau} = 1$) to pairs with full insurance ($\hat{\beta}_{v\tau} = 0$), land misallocation decreases: the correlation between TFP-Q and landholdings rises from 0.3 to around 0.7, while the dispersion of MPL drops from 0.8 to approximately 0.4.

Table 3 reports the estimation outcomes for different specifications of Equation (8). Columns (1) to (4) refer to $\text{corr}_{.v\tau} [\log \ell_{i\tau}, \log \hat{\theta}_i]$, and (5) to (8) to $\text{st.dev}_{.v\tau} [\log \text{MPL}_{i\tau}]$. Standard errors (in parentheses) are computed using village-level clustered bootstrap with 5,000 replications.

Table 3: Risk-sharing and land misallocation

	$\text{corr}_{.v\tau} [\log \ell_{i\tau}, \log \hat{\theta}_i]$				$\text{st.dev}_{.v\tau} [\log \text{MPL}_{i\tau}]$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{\beta}_{v\tau}$	-0.648*** (0.226)	-0.675*** (0.228)	-0.163** (0.080)	-0.229*** (0.109)	0.732*** (0.178)	0.749*** (0.245)	0.161*** (0.063)	0.192*** (0.076)
Observations	90	90	90	90	90	90	90	90
R^2	0.106	0.118	0.897	0.942	0.218	0.223	0.902	0.948
Year FE		✓	✓	✓		✓	✓	✓
Village FE			✓	✓			✓	✓
Village time trends				✓				✓

Notes: The unit of analysis across all columns is a village-year pair. The first four columns present the results of regressing our first measure of land misallocation on the estimated village-and-year-specific consumption elasticities to idiosyncratic income shocks. The following four columns show the results using our second measure of land misallocation. Standard errors in parentheses are computed using village-level clustered bootstrap (5,000 replications) following the procedure in Cameron et al. (2008).

There is a negative correlation between land misallocation and the degree of risk-sharing. For example, column (4) indicates that moving from village-year pairs with full insurance ($\hat{\beta}_{v\tau} = 0$) to no risk-sharing ($\hat{\beta}_{v\tau} = 1$) reduces the correlation between farm size and productivity by approximately 0.229 points, after controlling for village fixed effects, year fixed effects, and a village-specific linear time trend. In terms of magnitudes, the effect is equal to 0.82 times the standard deviation of the correlation between land size and farm productivity across villages and years (Figure 1, Panel A). The dispersion

of the marginal product of land is negatively correlated with the degree of risk-sharing. The estimates in column (8) suggest that moving from full insurance to no risk-sharing is associated with an increase of approximately 0.192 points in the standard deviation of the log of MPL_{it} . Similarly to before, this increase represents 0.83 times the standard deviation in the dispersion of the marginal product of land across villages and years (Figure 1, Panel B).²¹

Our results show a clear negative correlation between the degree of risk-sharing in village economies and the misallocation in the land markets of those villages. How much would land efficiency and aggregate output increase with a better risk-sharing arrangement? In the following section, we leverage the structure of our model to address this question in detail.

4 The gains from full insurance

In this section, we employ our theoretical model to quantify the aggregate gains from completing village consumption insurance markets. To proceed, we need to specify values for the model parameters, such as the land share α , the aggregate (fixed) supply of land L , the level of risk-sharing $1 - \beta$, and the coefficient of relative risk aversion σ . In addition, we also need to define the distributions of household productivity, θ_i , and the output shock, ρ_i . In the following, we describe what we do in detail.

We fit the model to the average village in our data. Some parameters are externally calibrated without solving the model. These parameters are listed in Table 4. We set the output elasticity of land, α , to 0.193, based on estimates obtained from Equation (5). The aggregate land supply, L , is set to the average farm size of 3.819 hectares (Table 1). The elasticity of consumption with respect to idiosyncratic income shocks, β , is set to 0.223, which is the estimate obtained from Equation (6). The estimates from Equation (5) are used to construct household-farms' physical productivity, θ_i , and output shocks, ρ_i . The distributions of these parameters are calibrated to match the empirical frequencies observed across households and years.²²

We are left with only one parameter, the coefficient of relative risk aversion, σ , which

²¹In Appendix C.3, we present estimates from an IV strategy that exploits variation in caste diversity across villages, which may help address potential endogeneity in the relationship between risk-sharing and misallocation. Identifying exogenous variation in risk-sharing across village economies is an extremely challenging task. Even randomized controlled trials of microfinance can generate global spillover effects by altering social networks, which play a crucial role in informal insurance and credit (Banerjee et al., 2024). Therefore, while our IV results are not intended to provide a definitive causal estimate of the impact of risk-sharing on land misallocation, it is reassuring they align with the predictions of our model.

²²To solve the model, we discretize the possible values of physical productivity and output shocks into 100 and 50 bins, respectively, each corresponding to different percentiles within their distributions.

Table 4: Parameters calibrated externally

Parameters	Description	Value	Source
α	Land share	0.193	Equation (5)
L	Aggregate land supply (hectares)	3.819	Table 1
β	Elasticity of consumption with respect to idiosyncratic income shocks	0.223	Equation (6)

Notes: This table reports the parameters that are externally calibrated without solving the model and their sources.

is estimated using the simulated method of moments (SMM). In particular, we choose σ to match the average correlation between log farm size and log productivity, which we denote by

$$\overline{\text{corr.}} \left[\log \ell_{i\tau}, \log \hat{\theta}_i \right] = \frac{\sum_{v\tau} \text{corr.}_{v\tau} \left[\log \ell_{i\tau}, \log \hat{\theta}_i \right]}{\mathcal{VT}},$$

where \mathcal{V} and \mathcal{T} denote the numbers of villages and years in the data. Table 5 reports the point estimate and the model fit. We obtain a value of $\hat{\sigma} = 1.60$, indicating moderate risk aversion. This estimate aligns with the findings in Holden and Quiggin (2017), who estimate a coefficient of relative risk aversion of 1.73 (Table A.4) for a sample of farmers in Malawi.²³

Table 5: Estimated risk aversion


Parameters	Description	Value	Target	Data	Model
σ	Relative risk aversion	1.600	$\overline{\text{corr.}} \left[\log \ell_{i\tau}, \log \hat{\theta}_i \right]$	0.461	0.469

Notes: This table reports the value of the coefficient of relative risk aversion that is estimated to match the average correlation between log farm size and log productivity.

To validate the model, we assess its ability to replicate the observed correlation between the estimated levels of risk-sharing, $\hat{\beta}_{v\tau}$, and the correlation of farm size and productivity across villages and years, $\text{corr.}_{v\tau} \left(\log \ell_{i\tau}, \log \hat{\theta}_i \right)$. To accomplish this, we solve replicas of our model that differ only in the values of β . Figure 4 plots the equilibrium correlation between log farm size and log productivity for different levels of risk-sharing (blue dots) against the corresponding empirical estimates from Figure 3A (red dots). Our model can replicate the negative correlation between risk-sharing and land misallocation observed in the data, even though this correlation was not explicitly targeted in the model's estimation. This suggests that the model effectively captures the relationship between land misallocation and the degree of consumption insurance across villages and years.

²³As shown in Equation (1), solving the model requires us to take a stance on the households' Pareto weights, $(v_i)_i$. In all the exercises performed in this section, we assume that $v_i = 1$, for each i .

Figure 4: Risk-sharing and misallocation: Model vs. data



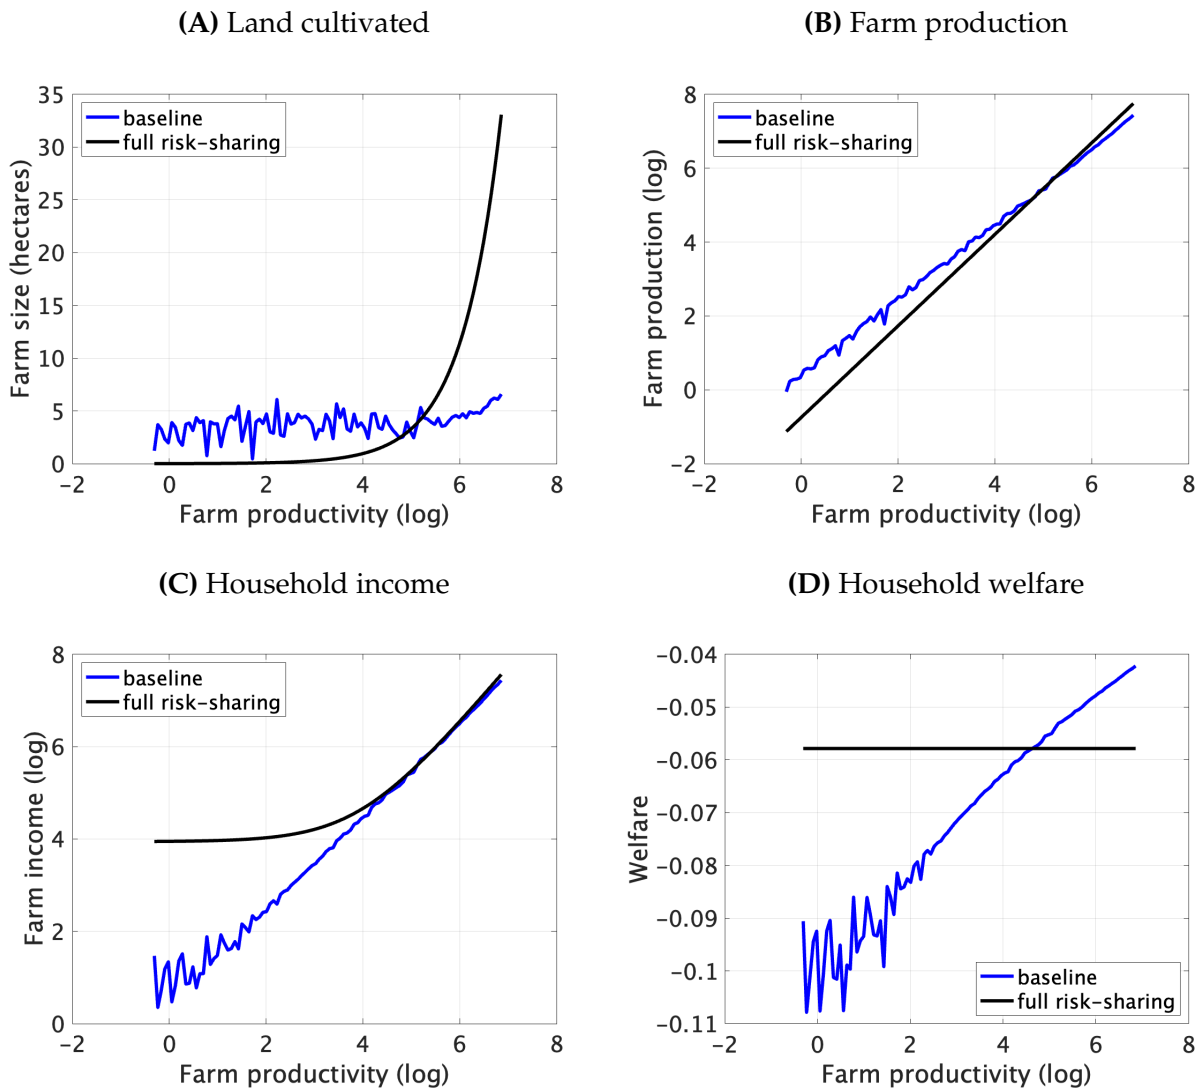
Source: VDSA survey (ICRISAT) and own calculations.

4.1 Counterfactual exercise

How would households’ land cultivation choices differ if village insurance markets were complete? To what extent does full insurance enhance allocative efficiency in the land market? We answer these questions with a counterfactual exercise where we improve the functioning of consumption insurance markets in village economies. Completing the market for insurance against shocks to agricultural output is a natural benchmark of “financial deepening” (see Townsend and Ueda (2006) on this concept). Specifically, we contrast our baseline model to a counterfactual scenario in which $\beta = 0$, corresponding to perfect risk-sharing. We keep all the other parameters at their baseline values, including the overall land supply, and assume identical Pareto weights across households, ensuring that, in the full insurance equilibrium, each household consumes the economy’s average income.²⁴

²⁴If we conceptualize the full insurance equilibrium as the solution to a planner’s problem, as outlined in Section 2, we can imagine the planner having complete information and the ability to maximally punish households that opt out of the risk-sharing arrangement. As a result, she can implement this solution without facing incentive compatibility or participation constraints. The egalitarian allocation, in which

Figure 5: Partial (baseline) vs. full risk-sharing



Source: VDSA survey (ICRISAT) and own calculations.

Figure 5 plots the amount of land cultivated (Panel A), farm output (Panel B), household income (Panel C), and household expected utility (Panel D) on the y -axes against (log) farm productivity (on the x -axis) for both the baseline (blue line) and full insurance economies (black line). As we move from partial to full risk-sharing, land reallocates from low- to high-productivity farms (Panel A). Under full insurance, the most productive household cultivates more than four times as much land compared to what it does under partial insurance, increasing its cultivated area from just over 5 hectares to nearly

each household consumes a constant fraction of aggregate income, is generally not implementable as a decentralized solution in a complete markets economy. Instead, in such a solution, households of the same type (i.e., those who are ex-ante identical) consume the same amount in expectation.

35 hectares. Conversely, those with low productivity cultivate less land under full insurance than under partial insurance. Improved risk-sharing decreases land misallocation, leads to greater output dispersion across farms (Panel B), and simultaneously reduces the dispersion in household income (Panel C). Panel D shows that most households, particularly those with low productivity, experience substantial welfare gains under full sharing compared to partial insurance. Conversely, the most productive households face welfare losses when participating in the full sharing arrangement rather than the partial insurance scheme.

Table 6: Counterfactual exercise

	Baseline (partial insurance) (1)	Counterfactual (full risk-sharing) (2)
β	0.223	0
corr. $[\log \ell_{i\tau}, \log \hat{\theta}_i]$	0.469	0.914
Share of land, top 1% productive farms	0.017	0.086
Share of land, top 10% productive farms	0.155	0.627
Land dispersion (st.dev. $[\log \ell_{i\tau}]$)	0.399	2.589
Aggregate efficiency (output per hectare)	1	1.415
Aggregate output	1	1.186
Aggregate welfare	1	1.286

Source: VDSA survey (ICRISAT) and own calculations.

Table 6 offers a further comparative analysis of baseline and counterfactual economies, evaluated across various dimensions. With full insurance, land misallocation is sizably reduced: the correlation between log farm size productivity and log productivity is nearly twice as large the one in the baseline. Under full insurance, the distribution of cultivated land becomes significantly more unequal: the share of total available land allocated to firms at the top 1% of the productivity distribution increases by approximately six times, while the share going to the top 10% of farms increases by approximately four times. The variance in the distribution of (log) cultivated land increases approximately ninefold. Improved insurance leads to an increase in output of 18.57%, while the overall welfare

gains, measured in consumption-equivalent terms, are equal to 28.75%.^{25,26} These figures are comparable to those quantifying the gains from eliminating distortions in the land markets (e.g., Adamopoulos et al. (2022)). Thus, our counterfactual analysis suggests that inefficiencies in consumption insurance markets may be as significant as land market distortions in explaining the potential gains from improving land allocation across households. This result is confirmed by the increase in output per unit of land of 41.51% with improved risk-sharing.

5 Conclusions

This paper bridges the gap between the literatures on risk-sharing and resource misallocation. We begin with two key observations. First, insurance markets in rural villages are incomplete, leading household income shocks to significantly affect consumption. Second, there is substantial misallocation of factors of production among farmers, resulting in reduced agricultural productivity in village economies. We argue that these two phenomena are deeply interconnected. Specifically, we see the limited functioning of consumption insurance markets as a key factor contributing to land misallocation in rural villages.

We explore how imperfections in insurance markets affect land misallocation. Our theoretical results show that incomplete consumption insurance can increase land misallocation, even when land markets operate without distortions. Empirically, we quantify the losses attributable to limited risk-sharing using the latest ICRISAT data from rural India. Our findings suggest that fully developed insurance markets could significantly enhance the allocation of land, resulting in output and welfare gains of 19% and 29%, respectively. Thus, improving risk-sharing within an otherwise undistorted economy can yield gains comparable in magnitude to those achieved by removing distortions in factor or output markets (Adamopoulos et al., 2022).

The main contribution of this paper is to offer an alternative explanation for land misallocation beyond the well-known land market frictions or farm-specific distortions that disproportionately affect more productive farmers. Our explanation is grounded in the observation that village-level risk-sharing is incomplete, and our analysis quantifies

²⁵In Appendix D.1, we study whether efficiency and output gains from ensuring full risk sharing are robust to incorporating farm-specific distortions in the model. Specifically, we extend our model by introducing farm-specific distortions taking the form of output wedges that are correlated to farm productivity. We then conduct the same comparative analysis as for the baseline model. We find that the efficiency and output gains from achieving perfect risk-sharing remain large and are equal to 19.3% and 13.4%, respectively.

²⁶Results are also robust to an alternative calibration of the land elasticity, α . See Appendix D.2.

how imperfect risk-sharing contributes to land misallocation in rural Indian communities. From a policy perspective, our emphasis on imperfect risk-sharing as a cause of misallocation highlights the potential of financial deepening as a way to enhance the allocation of factors of production.

References

- Acampora, M., Casaburi, L., and Willis, J. (2022). Land rental markets: Experimental evidence from Kenya. Technical report, National Bureau of Economic Research.
- 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. (2014). The size distribution of farms and international productivity differences. *American Economic Review*, 104(6):1667–1697.
- 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.
- Ambrus, A., Gao, W., and Milán, P. (2022). Informal Risk Sharing with Local Information. *The Review of Economic Studies*, 89(5):2329–2380.
- Anderson, T. W. and Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American statistical Association*, 76(375):598–606.
- Armangué-Jubert, T., Pietrobon, D., and Ruggieri, A. (2025). Incomplete Markets as Correlated Distortions. Working paper.
- Arrow, K. J. (1971). The Theory of Risk Aversion. In *Essays in the Theory of Risk-Bearing*, pages 90–120.
- Banerjee, A., Breza, E., Chandrasekhar, A. G., Duflo, E., Jackson, M. O., and Kinnan, C. (2024). Changes in social network structure in response to exposure to formal credit markets. *Review of Economic Studies*, 91(3):1331–1372.
- Bardhan, P. and Udry, C. (1999). *Development Microeconomics*. Oxford University Press.

- Benjamin, D. (1992). Household composition, labor markets, and labor demand: testing for separation in agricultural household models. *Econometrica: Journal of the Econometric Society*, pages 287–322.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.
- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The review of economics and statistics*, 90(3):414–427.
- Chari, A., Liu, E. M., Wang, S.-Y., and Wang, Y. (2021). Property Rights, Land Misallocation, and Agricultural Efficiency in China. *The Review of Economic Studies*, 88(4):1831–1862.
- Chen, C., Restuccia, D., and Santaeuilàlia-Llopis, R. (2022). The Effects of Land Markets on Resource Allocation and Agricultural Productivity. *Review of Economic Dynamics*, 45:41–54.
- Chen, C., Restuccia, D., and Santaeuilàlia-Llopis, R. (2023). Land Misallocation and Productivity. *American Economic Journal: Macroeconomics*, 15(2):441–465.
- Chiappori, P.-A., Samphantharak, K., Schulhofer-Wohl, S., and Townsend, R. M. (2014). Heterogeneity and Risk Sharing in Village Economies. *Quantitative Economics*, 5(1):1–27.
- De Giorgi, G., Di Falco, S., and Pietrobon, D. (2024). Farming, Non-Farm Enterprises, and Migration: Incomplete Markets and Structural Transformation. Working paper.
- Deaton, A. and Zaidi, S. (2002). *Guidelines for constructing consumption aggregates for welfare analysis*, volume 135. World Bank Publications.
- Dercon, S. (2002). Income risk, coping strategies, and safety nets. *The World Bank Research Observer*, 17(2):141–166.
- Donovan, K. (2021). The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity. *The Review of Economic Studies*, 88(5):2275–2307.
- Fafchamps, M. (2011). Risk Sharing Between Households. In *Handbook of Social Economics*, pages 1255–1279. Elsevier.
- Foster, A. D. and Rosenzweig, M. R. (2010). Microeconomics of Technology Adoption. *Annual Review of Economics*, 2(1):395–424.

- Gollin, D., Parente, S., and Rogerson, R. (2002). The role of agriculture in development. *American economic review*, 92(2):160–164.
- Guner, N. and Ruggieri, A. (2022). Misallocation and inequality. *CEPR Discussion Paper No. DP17113*.
- Holden, S. T. and Quiggin, J. (2017). Climate risk and state-contingent technology adoption: shocks, drought tolerance and preferences. *European Review of Agricultural Economics*, 44(2):285–308.
- Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448.
- Kinnan, C. (2021). Distinguishing Barriers to Insurance in Thai Villages. *Journal of Human Resources*. Forthcoming.
- Lamb, R. L. (2003). Inverse Productivity: Land Quality, Labor Markets, and Measurement Error. *Journal of Development Economics*, 71(1):71–95.
- Ligon, E. (1998). Risk Sharing and Information in Village Economies. *Review of Economic Studies*, 67(4):847–864.
- Ligon, E., Thomas, J. P., and Worrall, T. (2002). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies*, 69(1):209–244.
- Mazur, K. (2023). Sharing Risk to Avoid Tragedy: Theory and Application to Farming. *Journal of Development Economics*, 161:103030.
- Mazzocco, M. and Saini, S. (2012). Testing Efficient Risk Sharing with Heterogeneous Risk Preferences. *American Economic Review*, 102(1):428–468.
- Midrigan, V. and Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American economic review*, 104(2):422–458.
- Morduch, J. (2005). *Consumption smoothing across space: Testing theories of risk-sharing in the ICRISAT study region of South India*. Oxford: Oxford University Press.
- Morten, M. (2019). Temporary Migration and Endogenous Risk Sharing in Village India. *Journal of Political Economy*, 127(1):1–46.
- Munshi, K. (2019). Caste and the Indian Economy. *Journal of Economic Literature*, 57(4):781–834.

- Munshi, K. and Rosenzweig, M. (2016). Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap. *American Economic Review*, 106(01):46–98.
- Pietrobon, D. (2024). The Dual Role of Insurance in Input Use: Mitigating Risk versus Curtailing Incentives. *Journal of Development Economics*, 166:103203.
- Ravallion, M. and Chaudhuri, S. (1997). Risk and insurance in village india: Comment. *Econometrica: Journal of the Econometric Society*, pages 171–184.
- Ray, D. (1998). *Development Economics*. Princeton University Press.
- Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720.
- Restuccia, D. and Rogerson, R. (2013). Misallocation and productivity. *Review of Economic dynamics*, 16(1):1–10.
- Restuccia, D. and Rogerson, R. (2017). The causes and costs of misallocation. *Journal of Economic Perspectives*, 31(3):151–174.
- Restuccia, D., Yang, D. T., and Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of monetary economics*, 55(2):234–250.
- Rosenzweig, M. R. and Wolpin, K. I. (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in india. *Journal of political economy*, 101(2):223–244.
- Singh, I., Squire, L., and Strauss, J. (1985). Agricultural household models: A survey of recent findings and their policy implications. Working paper, Yale University, Economic Growth Center.
- Singh, I., Squire, L., and Strauss, J. (1986). *Agricultural household models: Extensions, applications, and policy*. Johns Hopkins University Press, Baltimore.
- Sleet, C. (2006). Endogenously Incomplete Markets: Macroeconomic Implications. *Mimeo. Tepper Business School*.
- Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica*, 62(3):539–591.
- Townsend, R. M. and Ueda, K. (2006). Financial deepening, inequality, and growth: a model-based quantitative evaluation. *The Review of Economic Studies*, 73(1):251–293.
- Udry, C. (1994). Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria. *Review of Economic Studies*, 61(3):495–526.

A Theory appendix

A.1 Proofs

Proof of Claim 1. The first-order conditions for $c_i(\boldsymbol{\rho})$ read as follows:

$$v_i(c_i(\boldsymbol{\rho}))^{-\sigma} - \lambda = 0,$$

or, equivalently,

$$c_i(\boldsymbol{\rho}) = \lambda^{-\frac{1}{\sigma}} v_i^{\frac{1}{\sigma}}, \quad (9)$$

where λ is the Lagrange multiplier attached to the feasibility constraint. Thus, each farmer's consumption is constant across states of the world $\boldsymbol{\rho}$. Integrating the last equation over all farmer types j and states of the world $\boldsymbol{\rho}$, we get that

$$\int \int c_j(\boldsymbol{\rho}) dQ_{\boldsymbol{\rho}}(\boldsymbol{\rho}) dj = \int \int \left(\frac{v_j}{\lambda}\right)^{\frac{1}{\sigma}} dQ_{\boldsymbol{\rho}}(\boldsymbol{\rho}) dj = \lambda^{-\frac{1}{\sigma}} \int v_j^{\frac{1}{\sigma}} dj.$$

Combine this equation with the feasibility constraint to obtain

$$\lambda^{-\frac{1}{\sigma}} = \frac{\int \int y_{j\rho} dQ_{\boldsymbol{\rho}}(\boldsymbol{\rho}) dj}{\int v_j^{\frac{1}{\sigma}} dj}.$$

Substituting this expression back into Equation (9), we get

$$c_i(\boldsymbol{\rho}) = \frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int y_{j\rho} dQ_{\boldsymbol{\rho}}(\boldsymbol{\rho}) dj.$$

□

Proof of Claim 2. The first-order conditions for ℓ_i read as follows:

$$\iota + \lambda \int \frac{\partial y_{i\rho}}{\partial \ell_i} dQ_{\boldsymbol{\rho}}(\boldsymbol{\rho}) = 0, \quad (10)$$

where ι is the Lagrange multiplier attached to the land availability constraint. Thus, the expected marginal products of land are equalized across households. To maximize aggre-

gate expected output, we can solve the following programming problem:

$$\begin{aligned} \max_{\ell} \int \int y_{i\rho} dQ_{\rho}(\rho) di \\ \text{s.t.} \quad \int \ell_i di = L. \end{aligned}$$

The first-order conditions for ℓ_i and ℓ_j imply that:

$$\int \frac{y_{i\rho}}{\partial \ell_i} dQ_{\rho}(\rho) = \int \frac{y_{j\rho}}{\partial \ell_j} dQ_{\rho}(\rho);$$

i.e., an allocation of land that maximizes aggregate expected output is such that the expected marginal products of land are equalized across households. \square

Proof of Claim 3. Under no sharing, the first-order conditions for ℓ_i read as follows:

$$\int (c_i(\rho))^{-\sigma} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^{IM} \right) dQ_{\rho}(\rho) = 0.$$

Thus, unless the households are risk neutral ($\sigma = 0$), the expected marginal products of land are not necessarily equalized across households. \square

Proof of Theorem 1. For each $\beta \in (0, 1]$, the first-order conditions for ℓ_i are

$$\int (c_i(\rho))^{-\sigma} \frac{\partial c_i(\rho)}{\partial \pi_{i\rho}} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^P \right) dQ_{\rho}(\rho) = 0,$$

where

$$\frac{\partial c_i(\rho)}{\partial \pi_{i\rho}} = \exp \left\{ \beta \log(\pi_{i\rho}) + (1 - \beta) \log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int \pi_{j\rho} dQ_{\rho}(\rho) dj \right) \right\} \frac{\beta}{\pi_{i\rho}}.$$

Letting

$$T_{i\rho} = (c_i(\rho))^{-\sigma} \frac{\partial c_i(\rho)}{\partial \pi_{i\rho}}$$

we can rewrite these first-order conditions as

$$\int T_{i\rho} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^P \right) dQ_{\rho}(\rho) = 0.$$

This equation shows that the effect of partial insurance on optimal land cultivation choices can be interpreted as the introduction of type- and state-specific distortions, affecting the

marginal return of land in each state of the world. These distortions imply that the expected marginal returns of land are not equalized to zero across farms. Instead, they vary in proportion to the type- and state-specific distortions. Notice that these distortions become more pronounced the further $T_{i\rho}$ deviates from being constant across states of the world. Since

$$\log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int \pi_{j\rho} dQ_\rho(\rho) dj \right)$$

is a constant, an increase in β amplifies the variance of $T_{i\rho}$ across states of the world. \square

Proof of Claim 4. The proof that each household consumes a constant fraction of aggregate output, proportional to its Pareto weight, follows the same steps as the proof of Claim 2. Therefore, we omit the details to avoid repetition.

The first-order conditions for ℓ_i and k_i read as follows:

$$t^\ell + \lambda \int \frac{\partial y_{i\rho}}{\partial \ell_i} dQ_\rho(\rho) = 0$$

and

$$t^k + \lambda \int \frac{\partial y_{i\rho}}{\partial k_i} dQ_\rho(\rho) = 0,$$

where t^ℓ and t^k are the Lagrange multipliers attached to the land and capital availability constraints, respectively. Thus, the expected marginal products of land are equalized across households, as are the expected marginal products of capital. The first-order conditions for $h_i(\rho)$ read as follows:

$$-\phi + \lambda \frac{\partial y_{i\rho}}{\partial h_i(\rho)} = 0.$$

To maximize aggregate expected output, we can solve the following programming problem:

$$\begin{aligned} & \max_{\ell, k, (h(\rho))_\rho} \int \int y_{i\rho} dQ_\rho(\rho) di \\ \text{s.t.} & \int \ell_i di = L, \quad \int k_i di = K, \quad \int h_i(\rho) di = H(\rho), \quad \forall \rho, \end{aligned}$$

where $H(\rho)$ is the total amount of labor the households are willing to supply under state ρ . The first-order conditions for ℓ_i and ℓ_j imply that

$$\int \frac{y_{i\rho}}{\partial \ell_i} dQ_\rho(\rho) = \int \frac{y_{j\rho}}{\partial \ell_j} dQ_\rho(\rho);$$

i.e., to maximize aggregate expected output, the expected marginal products of land must be equalized across households. Similarly, it follows directly that the equalization of expected marginal products of capital across households is a necessary condition for maximizing aggregate expected output. The first-order conditions for $h_i(\rho)$ and $h_j(\rho)$ imply that

$$\frac{\partial y_{i\rho}}{\partial h_i(\rho)} = \frac{\partial y_{j\rho}}{\partial h_j(\rho)},$$

for each ρ . □

Proof of Claim 5. Under no sharing, the first-order conditions for ℓ_i and k_i read as follows:

$$\int (c_i(\rho))^{-\sigma} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^{IM} \right) dQ_\rho(\rho) = 0$$

and

$$\int (c_i(\rho))^{-\sigma} \left(\frac{\partial y_{i\rho}}{\partial k_i} - q^{IM} \right) dQ_\rho(\rho) = 0.$$

On the other hand, the first-order conditions for $h_i(\rho)$ read as follows:

$$c_i(\rho)^{-\sigma} \frac{\partial y_{i\rho}}{\partial \ell_i} - \phi = 0.$$

Thus, unless the households are risk neutral ($\sigma = 0$), the expected marginal products of land and capital are not necessarily equalized across households. The misallocation of land and capital spills over to labor. That is, even though the marginal products of labor are equalized across households—implying that standard marginal product equalization tests would fail to detect labor misallocation—labor allocation across farms is indirectly distorted by the misallocation of land and capital. □

Proof of Theorem 2. For each $\beta \in (0, 1]$, the first-order conditions for ℓ_i and k_i are

$$\int (c_i(\rho))^{-\sigma} \frac{\partial c_i(\rho)}{\partial \pi_{i\rho}^P} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^P \right) dQ_\rho(\rho) = 0$$

and

$$\int (c_i(\rho))^{-\sigma} \frac{\partial c_i(\rho)}{\partial \pi_{i\rho}^P} \left(\frac{\partial y_{i\rho}}{\partial k_i} - q^P \right) dQ_\rho(\rho) = 0,$$

where

$$\frac{\partial c_i(\rho)}{\partial \pi_{i\rho}^P} = \exp \left\{ \beta \log(\pi_{i\rho}^P) + (1 - \beta) \log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \pi_{j\rho}^P dQ_\rho(\rho) dj \right) \right\} \frac{\beta}{\pi_{i\rho}^P}.$$

Letting

$$T_{i\rho} = (c_i(\rho))^{-\sigma} \frac{\partial c_i(\rho)}{\partial \pi_{i\rho}^P}$$

we can rewrite these first-order conditions as

$$\int T_{i\rho} \left(\frac{\partial y_{i\rho}}{\partial \ell_i} - r^P \right) dQ_\rho(\rho) = 0$$

and

$$\int T_{i\rho} \left(\frac{\partial y_{i\rho}}{\partial k_i} - q^P \right) dQ_\rho(\rho) = 0.$$

These equations show that the effect of partial insurance on optimal land cultivation and capital utilization choices can be interpreted as the introduction of type- and state-specific distortions. These distortions become more pronounced the further $T_{i\rho}$ deviates from being constant across states of the world. Notice that an increase in β amplifies this variance. \square

Proof of Theorem 3. For each $\beta \in (0, 1]$, the first-order conditions for b_{mit} are

$$-\int (c_{mit\rho_t})^{-\sigma} \frac{\partial c_{mit\rho_t}}{\partial \pi_{mit\rho_t}^P} dQ_\rho(\rho_t) + \Xi (1 + R^P) \int (c_{mit+1\rho_{t+1}})^{-\sigma} \frac{\partial c_{mit+1\rho_{t+1}}}{\partial \pi_{mit+1\rho_{t+1}}^P} dQ_\rho(\rho_{t+1}) = 0,$$

which is a standard Euler equation. The first-order conditions for $\ell_{mit}(b_{mit-1})$ and $k_{mit}(b_{mit-1})$ are

$$\int (c_{mit\rho_t})^{-\sigma} \frac{\partial c_{mit\rho_t}}{\partial \pi_{mit\rho_t}^P} \left(\frac{\partial y_{mit\rho_t}}{\partial \ell_{mit}(b_{mit-1})} - r_t^P \right) dQ_\rho(\rho_t) = 0$$

and

$$\int (c_{mit\rho_t})^{-\sigma} \frac{\partial c_{mit\rho_t}}{\partial \pi_{mit\rho_t}^P} \left(\frac{\partial y_{mit\rho_t}}{\partial k_{mit}(b_{mit-1})} - r_t^P \right) dQ_\rho(\rho_t) = 0.$$

The rest of the proof follows the same steps as the proof of Theorem 2. \square

A.2 Discussion of modeling assumptions

Household types. Our model features households that are ex-ante heterogeneous in their “permanent” agricultural productivity. The importance of heterogeneity in farm productivity has been emphasized in various fundamental papers, including Adamopoulos and Restuccia (2014), Chen et al. (2023), and Adamopoulos et al. (2022). By focusing our model on ex-ante heterogeneity in productivity, we can examine the relationship between risk-sharing and the misallocation of inputs chosen ex ante. We adopt this ap-

proach because the production decisions we are interested in are naturally viewed as choices made before any shocks to agricultural yields, such as rainfall or pests, occur.²⁷

Land markets. The model in Section 2 assumes that land markets are undistorted. This modeling choice allows us to distinguish our findings from most of the results in the misallocation literature, where land misallocation typically stems from land market frictions. We relax this assumption in Appendix D.1. The incomplete markets economy introduced in Subsection 2.1 features an environment where households initially possess land endowments and engage in trading these endowments within a competitive land market before production occurs. Alternatively, we may imagine that competitive moneylenders initially own all the plots and sell them to households before farming takes place. Our findings apply in both contexts.

Land ownership and tenancy. In our model, purchasing (respectively, selling) land is essentially equivalent to renting in (respectively, renting out) land; i.e., there is no difference between ownership and tenancy. A dynamic model may feature channels through which imperfect risk-sharing influences the decision to sell versus rent land. In a frictionless environment, a standard arbitrage condition dictates that the selling price of land should equal the net present value of its expected future rental earnings. Missing insurance markets, borrowing constraints, and imperfections in the saving technology might deter farmers from selling land, which serves as a buffer stock. Indeed, the presence of these frictions implies that the cash obtained from selling land cannot be perfectly smoothed over time or states of the world and that a farmer who sells land may be subsequently forced to engage in renting. In this context, the insurance value of owning land may contribute to land misallocation by affecting the relationship between a landowner's productivity and the amount of land owned. Our empirical analysis closely mirrors our theoretical framework by focusing on misallocation in operational landholdings, encompassing both owner-cultivators and renters.

Consumption functions. We model the relationship between consumption and income using consumption functions, as specified in Equation (1). We adopt this approach for two main reasons. First, it directly corresponds to our empirical specification for estimating the elasticity of consumption with respect to idiosyncratic income shocks, as introduced in Equation (6), which is a standard specification in the literature (see, e.g., Morten

²⁷Alternatively, one could model an economy where households are ex-ante homogeneous, receive productivity shocks, and make land cultivation decisions ex-post (after their productivity is revealed). This framework would still allow one to analyze the impact of insurance on misallocation; however, the misallocation would be contingent on the realization of those shocks.

(2019)). This correspondence is crucial because, in our quantitative exercise (Subsection 4), we calibrate the consumption functions so that β matches the estimated elasticity of consumption with respect to idiosyncratic income shocks. Second, this approach falls within the tradition of the literature on exogenously incomplete markets²⁸ and partial insurance (Blundell et al., 2008). In particular, these consumption functions formalize the idea that households participate in a risk-sharing arrangement, allowing them to pool their agricultural incomes to insure against idiosyncratic shocks and are flexible enough to capture a whole range of possible risk-sharing arrangements, from no sharing to full insurance. Note that risk-sharing does not have to be egalitarian: the Pareto weights $(v_i)_i$ allow different types of households to receive different fractions of the constant aggregate output. We deliberately sidestep detailed explanations of the underlying reasons for the specific forms of the consumption functions, which determine the level of insurance in the economy, focusing instead on how different degrees of insurance influence land misallocation. Finally, beyond providing a flexible way to model the relationship between household consumption and income, the consumption function in Equation (1) has the advantage of directly corresponding to our strategy to estimate risk-sharing in Subsection 3.5. Recall that, in our model, $\underline{\rho}$ is set sufficiently high to ensure that household income remains strictly positive for all possible land cultivation choices. Alternatively, positive cash-on-hand can be ensured by introducing an asset that allows households to borrow and save (see Appendix A.4).

A.3 Model with other factors of production

In this appendix, we show that all the results from Section 2 remain valid in a model that incorporates capital and labor as factors of production, where we also allow labor to be chosen after the realization of output shocks.

Consider a version of the model in Section 2 where household types are defined by a productivity level θ_i and an endowment

$$\left(\tilde{\ell}_i, \tilde{k}_i\right),$$

where \tilde{k}_i represents initial capital for a household of type i .

Households decide how much land to cultivate and how much capital to rent before output shocks are realized. Since households of the same type are ex-ante identical and

²⁸Contrast the approach where $(c_i)_i$ are primitives of the model with the perspective taken in the literature on optimal risk-sharing (Townsend, 1994) and endogenously incomplete markets (Sleet, 2006), where agents' consumption functions are derived from optimal consumption allocation problems featuring deeper primitive constraints on monitoring or enforcement technologies.

make these choices prior to the realization of shocks, their land and capital allocations are identical within each type. Thus, we can refer to ℓ_i and k_i as the land cultivated and capital utilized by a household of type i . Besides choosing land and capital, households decide how many hours to work. Labor supply decisions are made after the output shocks are realized, in contrast to capital and land. We assume households have complete information, so they can choose the amount of labor to employ conditional on the realization of all output shocks in the economy, ρ . Since households of the same type are ex-ante identical, we can unambiguously define the function $h_i(\rho)$, which represents the farm labor supply of a household of type i when state ρ is realized. The output produced by a household of type i in state ρ is given by

$$y_{i\rho} = A\theta_i\rho\ell_i^{\alpha^\ell}k_i^{\alpha^k}(h_i(\rho))^{\alpha^h},$$

where A is an aggregate productivity term, $\alpha^\ell, \alpha^k, \alpha^h \in (0, 1)$, and $\alpha^\ell + \alpha^k + \alpha^h < 1$. Let r be the price of land and q the rental price of capital. Given these prices, let

$$\pi_{i\rho} = y_{i\rho} - r(\ell_i - \tilde{\ell}_i) - q(k_i - \tilde{k}_i)$$

denote the income of a household of type i when state ρ is realized.

Let $\mathbf{k} = (k_i)_i$ and $\mathbf{h}(\rho) = (h_i(\rho))_i$. To characterize an allocation of resources under complete markets, we solve the following planner's problem for a given collection of type-specific Pareto weights $(v_i)_i$:

$$\max_{(c(\rho))_\rho, \ell, k, (h(\rho))_\rho} \int v_i \int \left[\frac{(c_i(\rho))^{1-\sigma}}{1-\sigma} - \phi h_i(\rho) \right] dQ_\rho(\rho) di,$$

subject to the land and capital availability constraints

$$\int \ell_i di = \int \tilde{\ell}_i di = L$$

and

$$\int k_i di = \int \tilde{k}_i di = K,$$

and the feasibility constraint

$$\int \int c_i(\rho) dQ_\rho(\rho) di = \int \int y_{i\rho} dQ_\rho(\rho) di.$$

The following claim describes an allocation of resources under complete markets.

Claim 4. *Under full insurance, each household consumes a constant fraction of aggregate output,*

with the fraction being proportional to its Pareto weight. Moreover,

1. the expected marginal products of land are equalized across households,
2. the expected marginal products of capital are equalized across households,
3. the marginal products of labor are equalized across households, in each state of the world.

Finally, aggregate expected output is maximized.

Next, we analyze the allocation of land that arises in a competitive equilibrium under no risk sharing. In this case, the income of a household of type i in state ρ is given by

$$\pi_{i\rho}^{IM} = y_{i\rho} - r^{IM} (\ell_i - \tilde{\ell}_i) - q^{IM} (k_i - \tilde{k}_i),$$

where r^{IM} and q^{IM} denote the equilibrium prices of land, capital, and labor under incomplete markets. The household's optimization problem is then given by:

$$\max_{\ell_i, k_i, (h_i(\rho))_\rho} \int \left[\frac{(\pi_{i\rho}^{IM})^{1-\sigma}}{1-\sigma} - \phi h_i(\rho) \right] dQ_\rho(\rho).$$

The following result demonstrates that incomplete insurance markets lead to factor misallocation, even when *all* factor markets operate without frictions.

Claim 5. *Under no sharing, there is land, capital, and labor misallocation.*

One implication of this result is that the presence of factor markets where inputs are chosen after the realization of output shocks does not undo the misallocation caused by missing insurance markets. The reason is that if some factors of production are chosen before shocks are realized (e.g., land and capital), but insurance markets are missing, households distort their ex-ante input choices to self-insure against risk. After output shocks are realized, households adjust their ex-post input choices (e.g., labor) by equating the realized marginal product of labor to its price, thereby maximizing their income. However, as long as inputs chosen ex post are not perfect substitutes for those chosen ex ante, the distortions in ex-ante input choices caused by missing insurance markets “spill over” to those chosen ex post.

Finally, consider an environment with partial insurance. Let the income of a household of type i in state ρ is given by

$$\pi_{i\rho}^P = y_{i\rho} - r^P (\ell_i - \tilde{\ell}_i) - q^P (k_i - \tilde{k}_i),$$

where r^P and q^P denote the equilibrium prices of land, capital, and labor under partial insurance. We define the following consumption function for a household of type i :

$$c_i(\rho) = \exp \left\{ \beta \log \left(\pi_{i\rho}^P \right) + (1 - \beta) \log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int \pi_{j\rho}^P dQ_\rho(\rho) dj \right) \right\}.$$

The following theorem extends Theorem 1 to settings with multiple factors of production.

Theorem 2. *Land market misallocation increases in the elasticity of consumption with respect to own income, β .*

A.4 Model with borrowing and saving

In this section, we show that Theorem 2 extends to environments where households can borrow and save through a risk-free asset. To do so, we build on the static model in Appendix A.3 by introducing time in a discrete setting. In this framework, households have additively separable preferences over time and discount the future at a constant rate Ξ . Let ρ_t be the state of the world in period t , assumed to be identically and independently distributed over time. Households can borrow and save through a risk-free asset. Let b_{mit} denote the amount of bonds held by household m of type i in period t (where a negative value indicates borrowing). This amount must be repaid in the following period at the interest rate R^P . In each period, households decide how much land to cultivate and how much capital to employ before the realization of ρ_t . After ρ_t is realized, they choose how much labor to supply to their own farm. Let these choices be denoted by

$$(\ell_{mit}(b_{mit-1}), k_{mit}(b_{mit-1}), h_{mit}((b_{mit-1}), \rho_t)).$$

Denote household m 's farm output when state ρ_t realizes by

$$y_{mit\rho_t} = A\theta_i\rho_t (\ell_{mit}(b_{mit-1}))^{\alpha^\ell} (k_{mit}(b_{mit-1}))^{\alpha^k} (h_{mit}((b_{mit-1}), \rho_t))^{\alpha^h}.$$

The income of a household m of type i in period t when state ρ_t realizes is given by

$$\pi_{mit\rho_t}^P = y_{mit\rho_t} - r_t^P (\ell_{mit}(b_{mit-1}) - \tilde{\ell}_i) - q_t^P (k_{mit}(b_{mit-1}) - \tilde{k}_i) + (1 + R^P) b_{mit-1} - b_{mit}.$$

Let consumption for household m of type i in period t when state ρ_t realizes be

$$c_{mit\rho_t} = \exp \left\{ \beta \log \left(\pi_{mit\rho_t}^P \right) + (1 - \beta) \log \left(\frac{v_i^{\frac{1}{\sigma}}}{\int v_j^{\frac{1}{\sigma}} dj} \int \int \int \pi_{njt\rho_t}^P dQ_\rho(\rho_t) dndj \right) \right\}.$$

The following theorem generalizes Theorem 2 to environments where households can borrow and save through a risk-free asset.

Theorem 3. *Land market misallocation increases in the elasticity of consumption with respect to own income, β .*

B Data appendix

We use information from the “Village Dynamics Studies in South Asia” (VDSA) project by ICRISAT, a widely used panel data set (Townsend (1994), Mazzocco and Saini (2012), and Morten (2019), among many others). The data is collected through different modules (general endowment, cultivation schedule, rainfall schedule, among others) containing questions on different topics generally asked to the household head. Questions asked only to the household head generally refer to information about the whole household. Some modules ask questions to a greater subset of the household members (e.g., the questions in the employment schedule are asked to all members who completed 6 years of age).

Most modules are collected at a monthly frequency (e.g., the employment schedule) while others only come at a yearly frequency (for instance, the general endowment schedule, and the questions in the cultivation schedule that refer to agricultural output). We use data from July 2010 to June 2015. We aggregate the individual-level data to the household level. We end up with monthly household-level panel data, containing information on farming, expenditure, and income for families in 18 villages in the Indian semi-arid tropics.

General endowment schedule. This schedule provides annual, individual-level data on various characteristics of household members, including age, sex, education, and primary and secondary occupations. Additionally, it offers household-level details on landholdings, such as ownership status, total and irrigable areas, and various soil characteristics. We leverage the data on these characteristics to build the measures of average soil depth, slope, fertility, and degree of degradation introduced in Subsection 3.3. The schedule also contains yearly household-level data on livestock, farm equipment, buildings, durable consumption goods, stocked items (like crops, cooking fuel, and agricultural inputs), assets, liabilities, gender roles, and coping strategies employed in response to self-reported negative income shocks.

We employ the individual-level demographic data in this schedule to construct an age-sex index at the household-year level, following the methodology described in Townsend (1994). Specifically, we assign individual weights based on age and sex as follows: 1 for males over 18 years, 0.9 for females over 18 years, 0.94 for males aged 13 to 18, 0.83 for females aged 13 to 18, 0.67 for children aged 7 to 12, 0.52 for children aged 4 to 6, 0.32 for toddlers aged 1 to 3, and 0.05 for infants under one year. We then calculate the household-year age-sex index by aggregating these weights for each household annually. In Subsection 3.5, we utilize this index to adjust household-level consumption and income

variables to per capita terms.

The general endowment schedule provides detailed information on each household's landholdings annually, with the plot as the unit of observation. This includes data on ownership status—whether owned, leased, shared, or mortgaged—and identifies the household members associated with each plot. It also details both total and irrigable areas, proximity to the house, irrigation sources and their distances. Additionally, the schedule provides information on plot attributes such as soil type (e.g., red, shallow black, medium black, deep black), fertility (an ordered scale from 0 to 4), slope (an ordered categorical variable indicating the degree of the plot's slope), soil degradation (a categorical variable indicating whether the plot is subject to soil degradation and specifying its type), and average soil depth in centimeters. The schedule also notes the presence of bunds, number of trees, if the plot is owned or leased, potential sale revenue, actual rent paid or received, and an imputed rental value for owned plots. We utilize the detailed information on plot attributes to construct measures of average soil depth, slope, fertility, and degree of degradation, which we use to construct a measure of land quality used in the estimation of household-farms' physical productivity (see SubSection 3.3). Specifically, define $fertility_{pit}$, $slope_{pit}$, $degradation_{pit}$, and $depth_{pit}$ as the fertility, slope, degradation, and soil depth for plot p cultivated by household i in year τ . For each $x \in \{fertility, slope, degradation, depth\}$, we define

$$x_{it} = \sum_{p \in P(it)} \frac{\ell_{pit}}{\sum_{p' \in P(it)} \ell_{p'it}} x_{pit},$$

where $P(it)$ is an index set for the set of plots cultivated by household i in year τ .

Finally, we leverage the farm equipment section of the general endowment schedule to construct a measure of farm capital. Each year, the household head is asked to report the names and values of all farm equipment items owned by the household, including plows, sprayers, dusters, electric motors, diesel pumps, bullock carts, tractors, trucks, threshers, pipelines, rice mills, and flour mills, among others. We aggregate this data at the household-year level to create a measure of farm capital for each household annually. We use this variable in the estimation of household-farms' physical productivity, as detailed in Subsection 3.3.

Cultivation schedule. The cultivation schedule is divided into two main sections: inputs and outputs. We start with an overview of the input section. This part of the schedule gathers detailed monthly data on the inputs utilized by each household for every operation conducted on each plot they farm. Specifically, interviewers asks the household head to detail all operations carried out on each cultivated plot in each month. For

every operation, they collect data on the quantities and costs of the inputs used. Thus, the unit of analysis for this section is the operation conducted on each plot by each household each month.

One fundamental piece of information we can obtain from the input section of the cultivation schedule is the household's labor supply to their farm, which we use in the estimation of household-farms' physical productivity, as detailed in Subsection 3.3. This section meticulously details the total labor hours devoted to each farming operation, categorizing them by family, hired, and exchange labor—specifically distinguishing contributions from females, males, and children—as well as labor provided by bullocks, motors, and other sources. To calculate the total labor supplied by a household to their farm, we aggregate the labor hours contributed by family members to each operation across all plots at the household-month level.

The output section gathers data differently, focusing not on individual operations each month but on crop production each season for each plot. Specifically, interviewers collect information from household heads regarding the quantity (in kilos) and value of each crop harvested during the defined agricultural seasons: Rabi, Kharif, annual, perennial, and summer. A critical measure derived from this section is the total annual output quantity per household, which serves as the dependent variable in Equation (2). To obtain this variable, we compute the total output produced by each household across all cultivated plots and each season throughout the year, aggregating this data at the household-year level. Another crucial variable compiled from the output section of the cultivation schedule is the total size of the plots cultivated by each household annually. We use this variable as a measure of land size in the estimation of household-farms' physical productivity, as explained in Subsection 3.3. An advantage of this variable is that it reflects the total land cultivated by the households, independent of ownership title, legal status, or other formal distinctions (see Subsection 3.2).

Transaction schedule. This schedule meticulously documents every monetary inflow and outflow for each household on a monthly basis, along with the monetary value of all home-produced commodities. Rather than providing a direct measure of household consumption, the expenditure segment records detailed individual expenditures on food and non-food items, as well as the monetary value of home-produced commodities that are consumed. This level of granularity in consumption items is one of the key advantages of the ICRISAT dataset, as it allows us to compile a comprehensive measure of monthly household consumption by aggregating these recorded values at the household-month level. As a result, measures of household consumption are never imputed from household income, not even for the poorest households. To derive a measure of monthly house-

hold income, we adopt the budget-constraint approach of Mazzocco and Saini (2012). Specifically,

- We use the section on financial transactions to track monthly household cash flows from lending and borrowing activities.
- The section on loans allows us to keep track of monthly inflows from loans and repayments by the household.
- From the section on government benefits, we determine monthly receipts of state-provided aid to households.
- The section on product and livestock sales allows us to measure the monthly revenue households earn from agricultural and livestock sales.
- The section on asset sales and purchases allows us to monitor the households' monthly financial activities related to the trading of capital goods.

We construct our measure of total monthly household income, we compute:

$$\begin{aligned} \text{Income}_{it} = & \text{Consumption}_{it} - \text{Cash received}_{it} + \text{Cash lent}_{it} \\ & - \text{Loans received}_{it} + \text{Loans repaid}_{it} - \text{Government benefits received}_{it}. \end{aligned}$$

We use the age-sex index defined above to convert the household-level monthly consumption and income variables to per capita terms. These variables are employed in Subsection 3.5 to estimate the level of risk-sharing in Indian villages.

Rainfall schedule. The rainfall schedule provides detailed information on rainfall levels (measured in millimeters) for each village daily, derived from readings at the nearest weather station. We aggregate these daily measurements over a year, to generate total village-specific annual rainfall. This aggregation yields the total annual rainfall for each village, denoted as rain_{vt} . We utilize this variable to parameterize the impact of observable environmental shocks on output, as specified in Equation (4).

C Further empirical results

C.1 Agricultural output decomposition

In Table C1, we decompose the variance of agricultural output into its different sources. We employ two measures. In column (1) we report the R-squared from regressing household-farms' annual physical output on various production inputs separately, one by one. The R-squared suits our purpose as it indicates the proportion of variance in the dependent variables explained by the regressors.

In column (2) we report the Shapley value (expressed in %) of each production input. The Shapley value quantifies the average marginal contribution of each variable to the explained variance in agricultural output, considering all possible combinations of the explanatory variables. Each Shapley value is computed by averaging the incremental changes in R-squared when an explanatory variable is added to a subset of other variables across all possible subsets.

Table C1: Input contributions to agricultural output

Variable	R-squared (1)	Shapley value (%) (2)
Rainfall shocks: $\text{Var} [\hat{\vartheta}\text{rain}_{v\tau}]$	0.004	0.161
Land quality, $\text{Var} [\log \hat{q}_{i\tau}]$	0.103	3.469
Household-farms' physical productivity: $\text{Var} [\log \hat{\theta}_i]$	0.614	41.57
Family labor: $\text{Var} [\log h_{i\tau}]$	0.382	19.03
Capital: $\text{Var} [\log k_{i\tau}]$	0.292	10.08
Landholdings: $\text{Var} [\log \ell_{i\tau}]$	0.277	10.02

Notes: Column (1) reports the R-squared from regressing household-farms' annual physical output on various production inputs separately, one by one. Column (2) reports the Shapley value (expressed in %) of each production input.

Differences in estimated physical productivity stand as the major sources of variation in production yield across household-farms. Using the R-squared, differences in estimated physical productivity can explain around 60% of the variation in annual yields across households. Using the Shapley value, more than 40% of total output variation across farms can be attributed to differences in estimated physical productivity.

C.2 Imperfect risk-sharing in the ICRISAT villages

Table C2 reports our estimates for equation (6). To complement these results, we also estimate the following regression for household i within village v in month t (following Townsend (1994)):

$$\Delta \log c_{it} = \beta \Delta \log \pi_{it} + \chi_{vt} + \epsilon_{it} \quad (11)$$

Here, $\Delta \log c_{it}$ and $\Delta \log \pi_{it}$ denote changes in log per-capita consumption and log per-capita income, respectively, for household i between two consecutive months. We estimate equation (11) through ordinary least squares and using the first three lags of the change of log income as an instrument to correct for measurement error.²⁹

Table C2: Risk-sharing in the ICRISAT villages

	ln c_{it}		Δ ln c_{it}
	OLS (1)	OLS (2)	IV (3)
ln π_{it}	0.223*** (0.0181)		
Δ ln π_{it}		0.206*** (0.0194)	0.230*** (0.00483)
Household FE	✓		
Village-month FE	✓	✓	✓
Observations	46369	41263	29601
R-squared	0.681	0.319	0.193
First-stage F -statistic			574.54

Notes: The unit of analysis across all columns is the household year. The first column presents the results of regressing monthly consumption on monthly log income while controlling for village-month fixed effects. The second column presents a regression analysis where the monthly change in log consumption is regressed against the monthly change in log income, controlling for village-month fixed effects. The third column shows the second-stage results of an instrumental variable regression, where the first three lags of the monthly change in log income serve as instruments for the contemporaneous change in log income. Standard errors (in parentheses) are clustered at the village-year level across all columns.

²⁹This strategy was inspired by Ravallion and Chaudhuri (1997), who instrument changes in household income using year dummies and the changes in the components of household income from sources other than agricultural cultivation. Our approach is similar to Anderson and Hsiao (1981)'s instrument in panel data econometrics.

C.3 Instrumental variable strategy

In Section 3, we provide evidence that risk-sharing is negatively correlated with land misallocation at the village-year level. In this section, we propose an instrumental variable strategy that may get us closer to identifying the causal effect of risk-sharing on land misallocation. It should be noted that identifying exogenous variation in risk-sharing across village economies is an extremely challenging task. Even randomized controlled trials of microfinance can generate global spillover effects by altering social networks, which play a crucial role in informal insurance and credit (Banerjee et al., 2024). Therefore, while the results in this section are not intended to provide a definitive causal estimate of the impact of risk-sharing on land misallocation, it is reassuring that they align with the predictions of our model.

Our strategy uses village caste composition as an instrument for risk-sharing. Specifically, we propose that villages with a more homogeneous caste structure are better able to provide insurance to their villagers. Considering relevance first, several studies (Munshi and Rosenzweig (2016), Mazzocco and Saini (2012), Munshi (2019)) argue that caste networks play an important role in providing credit and insurance to its members. Secondly, exclusion restriction requires caste composition to be uncorrelated with land misallocation once we control for risk-sharing. To satisfy this criterion, we must ask whether caste networks tend to be more homogeneous in villages that have less land misallocation, irrespective of the capacity of more homogeneous networks to enhance the functioning of credit and insurance markets. A possible concern is that households permanently migrate between villages to capitalize on the benefits of operating in more efficient rural land markets. Even though permanent migration from one rural area to another is not completely uncommon in India, caste composition in our villages is very persistent, with an average coefficient of variation below 0.10. Thus, it seems unlikely that differences in the ability of village factor markets to efficiently allocate production factors among households would affect caste composition through permanent migration. On the other hand, although the caste structure of a village may certainly influence its ability to access resources (e.g., members of upper castes may have better access to financial products, such as loans), we believe there is no clear reason to suggest that it *directly* impacts the extent to which the community manages to efficiently allocate factors of production. Certainly, one could argue that caste-based discrimination could influence how factors of production are allocated among competing uses, thus suggesting a direct role for caste diversity in the extent of land misallocation in village economies. However, the available evidence, particularly in urban labor markets, indicates that caste discrimination is statistical (Munshi (2019)), and there is no compelling evidence to suggest that discrimination patterns would differ significantly in other factor markets, including those in rural economies.

Below we highlight our instrumental variable strategy in detail. The ICRISAT data records each household head's caste (*jati*), sub-caste, and caste group (backward caste, forward caste, nomadic tribe, other backward caste, scheduled tribe, and special backward caste). Although we share Munshi (2019)'s concerns regarding the analysis of caste networks based on broad caste groups rather than *jatis*, our limited sample size necessitates the use of these broader categories. Let \mathcal{J} be an index set for the set of caste groups, where j denotes a typical element in this set. Let n_{jv} be the number of households belonging to a caste group j in village v , with $\sum_{j \in \mathcal{J}} n_{jv} = n_v$. We construct a measure of caste diversity in village v as 1 minus the Simpson's Diversity Index:

$$z_v = 1 - \frac{\sum_{j \in \mathcal{J}} n_{jv} (1 - n_{jv})}{n_v (1 - n_v)}.$$

Because of the very limited variation over time in caste composition, we let our instrumental variable vary across villages only.

Finally, we estimate the following equation independently for each village v and year τ ,

$$\log c_{it} = \beta_{v\tau} \log \pi_{it} + \mu_i + \eta_t + \epsilon_{it}.$$


Then, we relate our two measures of misallocation in each village and year, denoted by $\omega_{v\tau}$, to the elasticities of consumption with respect to idiosyncratic income shocks at the village level, $\hat{\beta}_{v\tau}$:

$$\omega_{v\tau} = \gamma \hat{\beta}_{v\tau} + \mu_\tau + \epsilon_\tau,$$

where we use z_v as an instrumental variable for $\hat{\beta}_{v\tau}$.

Figure C.1 scatters the first-stage regression: caste diversity is significantly and positively correlated with the estimated elasticity of consumption with respect to income at the village and year level. Table C3 reports first and second-stage IV estimates together with F-stat. A one-unit increase in our measure of caste diversity results in a 0.440-point increase in the village-specific elasticity of consumption with respect to idiosyncratic income shocks. Our instrumental-variable strategy suggests that a one-unit increase in this elasticity results in a 1.331-point decrease in the correlation between farm size and productivity and a 2.095-point increase in the dispersion of the marginal product of lands.

Figure C.1: elasticity of consumption with respect to idiosyncratic income shocks and caste diversity



Source: VDSA survey (ICRISAT) and own calculations.

Table C3: Land misallocation and risk-sharing

	corr. $_{v\tau} [\log \ell_{it}, \log \hat{\theta}_i]$	st.dev. $_{v\tau} [\log MPL_{it}]$
	(1)	(2)
$\hat{\beta}_{v\tau}$	-1.331** (0.574)	2.095*** (0.579)
Observations	90	90
R^2	0.783	0.858
First-stage regression		$\hat{\beta}_v$ (1)
z_v		0.415*** (0.126)
Observations		90
F-statistics		10.87

Notes: Standard errors are robust. Source: VDSA survey (ICRISAT) and own calculations.

C.4 Alternative measure of land quality

Here, we perform a robustness check on the results in Section 3 using an alternative measure of land quality. Following De Giorgi et al. (2024), we construct this measure based on plot-level indices of agricultural suitability, which are estimated from a comprehensive regression of plot rental values on land characteristics, while controlling for village-level permanent unobserved heterogeneity and year-specific fixed effects. The idea behind this approach is that the rental value of a plot reflects valuable information about its suitability for agricultural production; i.e., its “quality.” Specifically, we run the following regression:

$$\log(\text{rental}_{pit}) = f(\text{type}_{pit}, \text{fertility}_{pit}, \text{slope}_{pit}, \text{degradation}_{pit}, \text{depth}_{pit}; \psi) + \varphi_v + \kappa_\tau + \epsilon_{pit},$$

where rental_{pit} represents the self-reported rental value of plot p cultivated by household i in year τ ; “type” denotes a categorical variable reflecting the plot’s soil type; “fertility” is an ordered categorical variable measuring the plot’s fertility on a scale from 0 to 4; “slope” is an ordered categorical variable indicating the degree of the plot’s slope; “soil degradation” is a categorical variable indicating whether the plot is subject to soil degradation and specifies its type; “depth” refers to the average depth of the plot measured in centimeters; f is a function that represents linear and higher-order terms for the included variables, along with multiple interaction terms between those variables; ψ is a parameter vector; φ_v are village fixed effects; and κ_τ are year fixed effects. We use the predicted values of this regression, $\widehat{\log(\text{rental}_{pit})}$, as plot-level indices of agricultural suitability. The intuition is that these predicted values contain information about the extent to which the interactions of multiple soil characteristics, village-level permanent unobserved factors, and year-level aggregate shocks influence the suitability of a plot for agricultural production. To generate a household-year level land quality index, we average these suitability indices as follows:

$$q'_{it} = \frac{\sum_{p \in P_{it}} \left(\widehat{\log(\text{rental}_{pit})} \times a_{pit} \right)}{\sum_{p \in P_{it}} a_{pit}},$$

where a_{pit} denotes the area of plot p cultivated by household i in year τ , and P_{it} is an index set for the set of plots cultivated by household i in year τ .

We use this alternative measure of land quality to reproduce two main findings from the main text. First, in Table C4 we report the contribution of each input to the agricultural output. While our alternative measure of land quality explains twice as much of the overall output variation across farms as before (7.5 versus 3.5%, see Table C1), household-

Table C4: Input contributions to agricultural output—alternative land quality measure

Variable	R-squared (1)	Shapley value (%) (2)
Rainfall shocks: Var $[\hat{\vartheta}\text{rain}_{v\tau}]$	0.004	0.191
Land quality, Var $[\log \hat{q}_{i\tau}]$	0.186	7.500
Household-farms' physical productivity: Var $[\log \hat{\theta}_i]$	0.584	36.91
Family labor: Var $[\log h_{i\tau}]$	0.382	19.43
Capital: Var $[\log k_{i\tau}]$	0.292	9.513
Landholdings: Var $[\log \ell_{i\tau}]$	0.277	10.72

Notes: Column (1) reports the R-squared from regressing household-farms' annual physical output on various production inputs separately, one by one. Column (2) reports the Shapley value (expressed in %) of each production input.

farms' physical productivity still represents the biggest source of output variation, with a Shapley value of about 37%.

Second, we construct the two measures of land misallocation using our alternative proxy for land quality and relate them to the degree of risk sharing across villages and years. Specifically, we estimate different specifications of Equation (8) in the main text, and report the outcomes in Table C5. All the new estimates are virtually unchanged.

Table C5: Risk-sharing and land misallocation—alternative land quality measure

	corr. $_{v\tau} [\log \ell_{i\tau}, \log \hat{\theta}_i]$				st.dev. $_{v\tau} [\log \text{MPL}_{i\tau}]$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{\beta}_{v\tau}$	-0.639*** (0.234)	-0.675*** (0.230)	-0.132* (0.078)	-0.187* (0.096)	0.691*** (0.175)	0.715*** (0.230)	0.150*** (0.056)	0.178*** (0.069)
Observations	90	90	90	90	90	90	90	90
R ²	0.100	0.116	0.892	0.942	0.220	0.229	0.901	0.944
Year FE		✓	✓	✓		✓	✓	✓
Village FE			✓	✓			✓	✓
Village time trends				✓				✓

Notes: The unit of analysis across all columns is a village-year pair. The first four columns present the results of regressing our first measure of land misallocation on the estimated village-and-year-specific consumption elasticities to idiosyncratic income shocks. The following four columns show the results using our second measure of land misallocation. Standard errors in parentheses are computed using village-level clustered bootstrap (5,000 replications) following the procedure in Cameron et al. (2008).

D Further counterfactual results

D.1 Incomplete insurance vs. farm-specific distortions

In this appendix, we discuss the robustness of our results to the inclusion of generic farm-specific distortions (wedges). These wedges capture the effects of frictions in output or land markets that may disproportionately affect more productive farmers, such as land ceilings (Adamopoulos and Restuccia, 2020), restrictions on land rental activities (Chari et al., 2021; Acampora et al., 2022), borrowing constraints (Midrigan and Xu, 2014), or size-dependent policies that directly tax/subsidize output (Guner and Ruggieri, 2022). Formally, we assume households of type i are subject to an “output tax,” τ_i . This tax is given by

$$\tau_i = 1 - \theta_i^{-\zeta}, \quad (12)$$

where ζ governs the correlation between output distortions and household productivity. Then, the output of a farmer of type i is

$$y_{i\rho} = (1 - \tau_i) \theta_i \rho \ell_i^\alpha = \theta_i^{1-\zeta} \rho \ell_i^\alpha.$$

Notice that when $\zeta > 0$, distortions are positively correlated with household productivity, meaning high-productivity households face relatively higher distortions. Conversely, when $\zeta < 0$, high-productivity households face relatively lower distortions. When $\zeta = 0$, there are no output distortions in the economy, and output production reverts to the scenario described in the main text.³⁰

To quantify the gains from full insurance using the model outlined in this appendix, we proceed as follows. Unlike what we do in Section 4, where we estimate the coefficient of relative risk aversion, σ , by calibrating it to match the correlation between (log) landholdings and (log) productivity in a model without wedges, here we fix σ at 2. Instead, we target the correlation between (log) landholdings and (log) productivity estimate the correlation in a model with partial insurance to estimate the correlation between wedges and productivity, ζ . All other parameter values are as in Section 4.

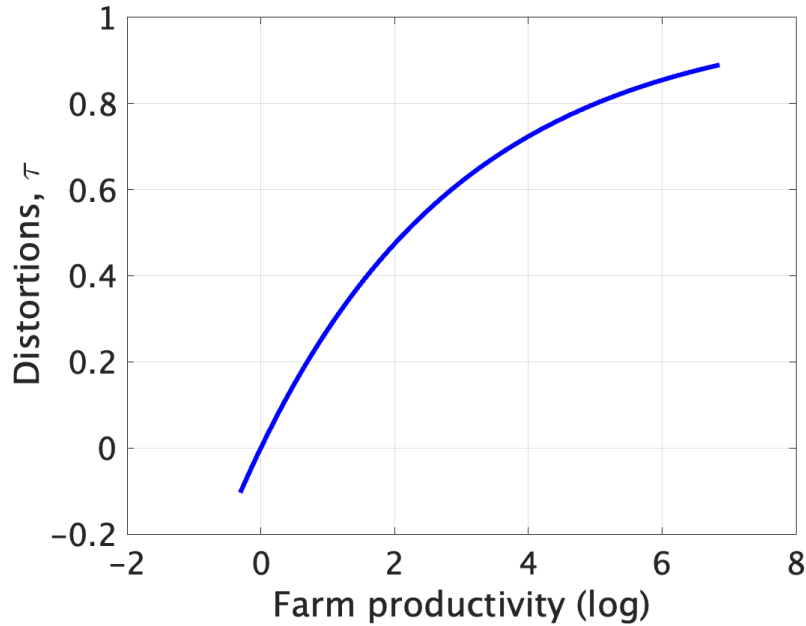
Table D1 reports the estimated value of ζ , along with the empirical and simulated values of the targeted moment. Our estimate of ζ is 0.322, indicating a strong positive correlation of approximately 0.96 between distortions τ_i and (log) productivity θ_i across households. This pattern is illustrated in Figure D.1, which reveals that distortions effec-

³⁰For some combination of ζ and θ_i , τ_i can be negative. In this case, distortions take the form of a subsidy towards household i . For further applications of the functional form in Equation (12) to describe firm-level distortions in developing countries, see Guner and Ruggieri (2022), among others.

Table D1: Estimated parameters

Parameters	Description	Value	Target	Data	Model
ζ	Distortion correlation	0.322	$\overline{\text{corr.}} [\log \ell_{i\tau}, \log \hat{\theta}_i]$	0.461	0.464

Notes: This table reports the value of the correlation between farm-specific distortions and household productivity that is estimated to match the average correlation between log farm size and log productivity.

Figure D.1: Estimated distortions

Source: VDSA survey (ICRISAT) and own calculations.

tively act as an output tax of up to 90% for the most productive households, while the least productive households receive a subsidy of around 10%.

Table D2 presents the results of the same counterfactual exercise described in Section 4, but within a model that incorporates farm-specific distortions. Column 1 corresponds to the baseline scenario, where consumption insurance is partial ($\beta = 0.223$) and households' land cultivation choices are distorted by wedges correlated with their productivity ($\zeta = 0.322$). Column 2 examines a counterfactual scenario in which risk-sharing is perfect ($\beta = 0$), while all other factors remain unchanged.

As shown in Table D2, the aggregate output gains from improving risk-sharing in an economy with farm-specific distortions amount to 14%, closely aligning with the 19% gains observed in an economy without these wedges (Section 4). The welfare and efficiency gains of full insurance, when accounting for output distortions, are 19% and 87%,

Table D2: Counterfactual exercise in a distorted economy

	Baseline (partial insurance) (1)	Counterfactual (full risk-sharing) (2)
β	0.223	0
ζ	0.322	0.322
$\overline{\text{corr.}} [\log \ell_{i\tau}, \log \hat{\theta}_i]$	0.464	0.914
Share of land, top 1% productive farms	0.012	0.059
Share of land, top 10% productive farms	0.131	0.489
Land dispersion (st.dev.[$\log \ell_{i\tau}$])	0.329	1.757
Aggregate efficiency (output per hectare)	1	1.193
Aggregate output	1	1.134
Aggregate welfare	1	1.869

Source: VDSA survey (ICRISAT) and own calculations.

respectively, compared to 29% and 45% in the distortion-free scenario. These results show that assuming the existence of farm-specific distortions only marginally reduces the estimated output, efficiency, and welfare gains from improving risk-sharing.

D.2 Robustness to alternative land elasticity

In this section, we test the robustness of our main counterfactual results to an alternative value of the land elasticity, α . In particular, we fix $\alpha = 0.50$, a value closer to the estimates of the decreasing returns-to-scale parameter in the farming technology found in the literature. We keep the values of all the other parameters unchanged, relative to the main calibration, except for the coefficient of relative risk aversion, σ , which is re-estimated to match the average correlation between \log farm size and \log productivity. We find a value for σ of 1.675, not far from our baseline estimate of 1.60 (Table 5 in the main text).

Table D3: Estimated risk aversion with alternative land elasticity

Parameters	Description	Value	Target	Data	Model
σ	Relative risk aversion	1.675	$\overline{\text{corr.}} [\log \ell_{i\tau}, \log \hat{\theta}_i]$	0.461	0.470

Notes: This table reports the value of the coefficient of relative risk aversion that is estimated to match the average correlation between \log farm size and \log productivity.

Table D4 reports selected outcomes for the baseline and counterfactual economies. Again, Column 1 corresponds to the baseline scenario, where consumption insurance is

Table D4: Counterfactual exercise with alternative land elasticity

	Baseline (partial insurance) (1)	Counterfactual (full risk-sharing) (2)
β	0.223	0
$\overline{\text{corr.}} [\log \ell_{i\tau}, \log \hat{\theta}_i]$	0.470	0.914
Share of land, top 1% productive farms	0.017	0.135
Share of land, top 10% productive farms	0.162	0.797
Land dispersion (st.dev.[$\log \ell_{i\tau}$])	0.320	4.180
Aggregate efficiency (output per hectare)	1	2.131
Aggregate output	1	1.671
Aggregate welfare	1	1.519

Source: VDSA survey (ICRISAT) and own calculations.

partial ($\beta = 0.223$) while Column 2 corresponds to a counterfactual scenario in which risk-sharing is perfect ($\beta = 0$).

All the results are qualitatively the same as those in the main text. On the other hand, a larger calibrated value for α implies larger output and welfare gains from completing insurance markets, equal to 67% and 52%, respectively.