Market Failures and Misallocation

Ajay Shenoy*

University of California, Santa Cruz

March 16, 2017

First Version: April 28, 2011

Abstract

I develop a method to measure and separate the production misallocation caused by failures in factor markets versus financial markets. When I apply the method to rice farming villages in Thailand I find surprisingly little misallocation. Optimal reallocation would increase output by less than 19 percent. By 2007 most misallocation comes from factor market failures. I derive a decomposition of aggregate growth that accounts for misallocation. Declining misallocation contributes little to growth compared to factor accumulation and rising farm productivity. I use a government credit intervention to test my measures. I confirm that credit causes a statistically significant decrease in financial market misallocation, but has no effect on factor market misallocation. (JEL O47, O16, E13)

*Postal Address: Rm. E2-455, University of California, M/S Economics Department, 1156 High Street, Santa Cruz CA, 95064; Email: azshenoy@ucsc.edu

1 Introduction

The average farm worker in the most advanced countries produces 44 times as much output as one in the least (Gollin, Lagakos, and Waugh, 2014a,b). The sheer size of this gap has led some researchers to propose that an underproductive farm sector may explain late industrialization and economic underdevelopment (Gollin, Parente, and Rogerson, 2002, 2007). Farming in poor countries may be unproductive for many reasons, but one of the best documented is that markets for inputs, credit, and insurance are dysfunctional. As a result the land, labor, and capital used in production may be misallocated across farmers. This misallocation may, according to calibrated models, reduce agricultural productivity by as much as one-half (Adamopoulos and Restuccia, 2014). But these models do not directly measure the aggregate misallocation in agriculture, or measure how much is caused by each of many possible market failures. Yet governments cannot design policies to reduce misallocation unless they know its sources.

This paper develops a method to measure and separate the misallocation caused by failures in factor and financial markets, two markets that are critical to production. I show that when markets are perfect the optimal allocation of inputs to a farmer depends only on her productivity, which is idiosyncratic, and the production function, which is common. I use the method of Anderson and Hsiao (1981, 1982) to estimate the parameters and calculate the optimal allocation. I define the increase in output from optimal reallocation to be the total misallocation from factor and financial market failures.

To separate the cost of each market's failure I make a crucial assumption: that in the absence of market imperfections, all inputs may be instantly and costlessly transferred between farmers within a village. As long as all inputs are chosen with equal information about productivity, this assumption lets me model the farmer as though she chooses all of her inputs at the same time before producing any output. Then a perfect factor market lets each farmer optimally divide spending between inputs to achieve the perfect mix.

The intuition of the method is simple. Consider a farmer choosing how many workers to hire and how much land to sow. Even if the credit market is distorted towards one input—if, for example, it is easier to get a loan to buy land than to hire workers because land can serve as collateral—the farmer will still choose the optimal mix of land and labor as long as the factor markets work. The key is that the owner of the land need not be its user. The farmer can buy the land but then immediately rent it out to another farmer who can use it more productively. The key insight behind my method is that rental markets break the link between the ownership of an asset and its use.

This insight implies that with perfect factor markets the farmer can always achieve the perfect mix. By perfecting each farmer's mix of inputs while holding scale constant, I place a lower bound on the aggregate gains from perfecting factor markets. By then perfecting scale I place an upper bound on the gains from subsequently perfecting financial markets. The method always identifies misallocation from mix versus scale, but my key contribution is to show that under the assumptions about costless transfers and timing, perfecting the factor markets would perfect each farmer's mix.

Finally, I decompose aggregate output into three components: an aggregate production function, average farm productivity, and the efficiency of factor allocations. I calculate the counterfactual path of aggregate output if growth in any component were shut down. The counterfactuals show how output would have grown if factor allocations had not improved.

In principle the method could be applied to any production environment, farming or otherwise, that fits the assumptions. But the method is particularly suited to Thailand's rice sector. Since rice production is relatively uniform I can assume a common production function without straying too far from reality. By estimating this production function I ensure its parameters are a valid reflection of how Thai farmers actually produce rice, minimizing the spurious "misallocation" caused by a flawed calibration. Other sources of misallocation like monopoly and taxes are rare in rice farming. Most important, my assumption that inputs may be transferred within a village without cost is not unreasonable. Compared to firms in heavy industry, rice farmers can transfer inputs within a village at relatively little cost.

I find surprisingly little misallocation. In 1996 the overall cost of within-village misallocation is 19 percent of output. It falls to 5 percent by 2008. By then most misallocation comes from factor markets rather than financial markets. Decreases in misallocation contributed little to growth in aggregate rice output relative to growth from factor accumulation and rising average productivity.

Like all structural assumptions, those that underpin my method are simplifications. To confirm that these simplifications do not invalidate the method I study the effects of a government credit program. First studied by Kaboski and Townsend (2011), Thailand's Million Baht Program created exogenous variation across villages in the supply of credit. If valid, the method should show that the program reduces misallocation, and the effect should be mainly through my measure of financial market misallocation. As expected, credit has a statistically significant but small effect. A one percent increase in credit reduces misallocation by 0.1 percentage points, nearly all of which comes from a reduction in financial market misallocation.

This paper's main contribution is methodological. Its approach is most similar to that of Hsieh and Klenow (2009), who also calculating the cost of misallocation caused by distortions to a firm's mix and scale. My key contribution is to show that under some assumptions these two distortions have a direct economic interpretation as distortions in factor versus financial markets. This approach is distinct from that of Adamopoulos and Restuccia (2014) and Midrigan and Xu (2013), who model the market imperfections faced by farmers or firms, then calibrate their models to predict the gains from eliminating these imperfections. Their approach requires that the researcher correctly models and calibrates the parameters of each imperfection.

By contrast, as long as the timing assumptions of the method proposed in this paper are met, it will measure the misallocation caused by each imperfection even if its functional form is unknown. Though well-suited to agriculture the assumptions hold in any sector where the key inputs are all chosen with the same information set before any output is produced, and where moving these inputs between firms has little technical cost. One example might be drug development, where the success of a new drug is not known until the final product is tested, and lab equipment or researchers can move or may outsource their services to other labs.

This paper also complements a vast microeconomic literature that tests for market imperfections in developing countries. Much of this work has focused on imperfections in financial markets, such as how communal divisions inefficiently concentrate capital among incumbent garment manufacturers in India (Banerjee and Munshi, 2004); lending arrangements fail to perfectly insure households in Nigeria (Udry, 1994); and entrepreneurs could reap large returns with small capital investments in Sri Lanka (De Mel, McKenzie, and Woodruff, 2008). But other work has emphasized the importance of imperfections in factor markets, such as markets for labor (e.g. Benjamin, 1992; Petrin and Sivadasan, 2013) and land (e.g. De Janvry et al., 2015). This paper proposes and applies a method for measuring and comparing the aggregate costs of each imperfection.

Finally, this paper extends the macroeconomic literature on agricultural productivity. Much of this literature has focused either on the misallocation of inputs between farming and other sectors (Gollin, Parente, and Rogerson, 2004; Lagakos and Waugh, 2011) or under-investment caused by market distortions (Restuccia, Yang, and Zhu, 2008). The most prominent exception, as mentioned above, is Adamopoulos and Restuccia (2014). Their calibrated model suggests misallocation may account for a large part of cross-country differences in productivity, whereas I measure very little misallocation within Thai villages. Given that their study differs from mine in setting, method, and that they consider reallocation between villages, further research is needed to determine the extent that misallocation explains economic underdevelopment.[1]

To my knowledge this paper is the first to split misallocation into the contributions of factor versus financial markets. By doing so I show that imperfec-

[1] In addition to those I cite in the main text, other recent papers on misallocation include Restuccia and Rogerson (2008); Banerjee and Moll (2010); Peters (2011); Bollard, Klenow, and Sharma (2012); Alfaro, Charlton, and Kanczuk (2008); Bartelsman, Haltiwanger, and Scarpetta (2009); Jones (2011); Osotimehin (2011); Alfaro and Chari (2012); Moll (2010); David, Hopenhayn, and Venkateswaran (2013); Restuccia and Santaeulalia-Llopis (2014); Sandleris and Wright (2014); Keniston (2011).

tions in factor market are no less important than those in financial markets. I also measure misallocation under weaker assumptions than some earlier work. For example, I estimate the production function instead of assuming U.S. parameters apply to Thailand. Finally, compared to papers that calibrate and simulate structural models, I link misallocation to its sources under weaker functional form assumptions about how markets fail. When my assumptions about timing are satisfied the method works under a fairly general set of imperfections.

2 The Method

This section gives the intuition behind the method using a simple static model of farm production. It is similar to the basic model of Singh et al. (1986, p. 17) except that I simplify the market for output while allowing a more complex (and potentially imperfect) market for inputs. I then solve for general equilibrium in the markets for land, labor, and capital to show that when factor markets are perfect, only the scale of input use rather than the mix will be distorted. In the main text I privilege simplicity over generality; in the appendix I show that a more general dynamic model yields identical results.

2.1 Setup

The farmer maximizes her utility from consumption C, which she pays for by producing farm revenue Y using capital K, land T, and labor L. Her output also depends on Hicks-Neutral productivity, part of which (A) she anticipates when choosing factors while the rest (Φ) is random and unanticipated. I normalize

$\mathbb{E}[\Phi] = 1$. The objective is

$$\text{Maximize} \qquad \mathbb{E}\left[\frac{C^{1-\gamma}}{1-\gamma}\right]$$

and her output is

$$Y = A\Phi K^{\theta_K} T^{\theta_T} L^{\theta_L}$$

I assume constant relative risk aversion solely to ease the exposition; the general model in the appendix allows a general utility function that may vary across farmers. I also show that the method works under a more general production function. But unlike the utility function the production function must be estimated, making the Cobb-Douglas assumption useful. I show in Appendix A.5.2 that it is not a bad approximation. Following the literature (for example Restuccia and Rogerson, 2008) I assume $\theta_K + \theta_T + \theta_L < 1$, implying there are decreasing returns to scale. Decreasing returns would arise if the farmer's managerial talent is spread more thinly as her farm grows.

The farmer may either buy or rent land and capital. Total land used in production T is the sum of owned land T^o and rented land $T - T^o$. She buys land at price p^T and rents it at price w^T. Capital is similar. Labor is hired at wage w^L. Total expenditure is

$$Z = w^K(K - K^o) + w^T(T - T^o) + w^L L + p^K K^o + p^T T^o$$

Whether bought or rented, inputs must be paid for at the beginning of the season. Since the farmer earns no revenue until the end of the season she must pay for inputs using her wealth W. If she is too poor to cover the whole cost she must borrow the balance $Z - W$. There are two credit market imperfections.

First, the interest rate she pays may depend on her collateral, which is owned land T^o. Second, there may be an upper bound $\bar{Z}$ on her total borrowing. Then the farmer must satisfy both a budget constraint

$$\lambda : \quad C = Y - R(T^o)(Z - W)$$

and a credit constraint

$$\omega : \quad Z - W \leq \bar{Z}$$

where λ and ω are Lagrange multipliers on these constraints.

Finally, the rental market for land may be imperfect. For simplicity I assume there are upper and lower bounds $\overline{T}$ and $\underline{T}$ on the amount of land that can be rented in or out. If the Lagrange multipliers on these constraints are $\overline{\kappa}$ and $\underline{\kappa}$, the farmer must satisfy

$$\overline{\kappa}, \underline{\kappa} : \quad \underline{T} \leq T - T^o \leq \overline{T}$$

In the extreme case, $\underline{T} = \overline{T} = 0$, preventing poor farmers from renting in land and large farmers from renting out land.

Many of the assumptions made here can be relaxed. The model can be made dynamic, prices and constraints can be allowed to vary across farmers, purchased assets can be collateralized, the borrowing rate may differ from the savings rate, and rental constraints can be applied to labor and capital (see Online Appendix A.3). The timing, however, is crucial. At the beginning of the season the farmer buys and rents all inputs simultaneously knowing A but not Φ. Purchased inputs can immediately be used in production or rented out to other farmers. Then the shock Φ is realized, production occurs, the farmer pays back her loans and consumes whatever income is left. This timing effectively assumes there is no cost in time or money to transfering factors between farmers, and all factors are chosen with equal information about productivity.[2]

Comparison to the textbook model of farm production: This model differs from the basic model of Singh et al. (1986, p. 17) mainly in that I allow for markets in land, labor, and capital that may be imperfect. For example, some farmers may be unable to rent in or rent out as much land as they need. I also allow for an imperfect credit market and an imperfect market for insurance (which is relevant because, unlike in the standard model, I allow for uncertainty in production). Most importantly, I allow constraints to differ across farmers because farmers have different levels of wealth or (in the appendix) pay different interest rates to borrow and different prices for inputs.

At the same time I make two simplifications, both of which are adopted from the literature on misallocation. First, I abstract from the household's decision of how much labor to supply. This assumption is innocuous because, like Singh et al. (1986), I assume hired and family labor are equally productive. (I relax this assumption in Appendix A.5.3 and show that the results are similar.) Given that the two forms of labor are equally productive, the household's labor supply decision affects only the aggregate stock of labor; the individual production outcome depends only on labor employed on the farm, not labor supplied. I define misallocation as the aggregate gain from reallocating factors *conditional on the current stock of factors*—a common definition in the literature (see, for

[2]If farmers are risk neutral the information assumption can be relaxed to require equal information about only the idiosyncratic part of productivity.

example, Hsieh and Klenow, 2009; Adamopoulos and Restuccia, 2014). Given this definition, I can for notational simplicity ignore the household's decision of how much labor to supply.

Second, I simplify the market for output. I assume the household has a single consumption good, and that it buys this good with the revenue it earns from farming rice. This assumption is equivalent to assuming the household cares only about its consumption spending. I also assume that all farmers receive the same price for their output. This allows me to focus on estimating the production function for revenue rather than physical output, as the two quantities differ by only a constant. The latter assumption is the less standard of the two. In fact, all that I require is that prices within a village in a given year are the same, as between-village differences would have no effect on misallocation within the village. As I explain in Section 3, this assumption is not unreasonable.

2.2 Perfect Choices and Distortions

The farmer's optimal choice of land satisfies

$$\big(\mathbb{E}[C^{-\gamma}] + Cov(C^{-\gamma}, \Phi)\big)\theta_T A T^{-(1-\theta_T)} K^{\theta_K} L^{\theta_L} = (R(T^o)\mathbb{E}[C^{-\gamma}] + \omega) w^T + \overline{\kappa} - \underline{\kappa} \tag{1}$$

This expression shows why imperfect markets cause misallocation. The left-hand side is the (shadow) marginal product of land. The term $Cov(C^{-\gamma}, \Phi)$ is the covariance of the marginal product of consumption and the unanticipated shock. The covariance will be more negative for farmers with lower wealth, less insurance, and greater risk aversion. These farmers are driven to farm less land.

The right-hand side is the (shadow) cost of land. It is higher for farmers who buy less land (as opposed to renting land) or are at the borrowing constraint ($\omega > 0$). It will be higher for farmers at the upper bound ($\overline{\kappa} > 0$) and lower for farmers at the lower bound ($\underline{\kappa} < 0$) for renting land. The term $R(T^o)$ will also appear in the optimal choice of capital and labor—this point is crucial to the method, as I show momentarily.

Suppose markets are fully perfect. With perfect insurance, the unanticipated shock does not affect consumption ($Cov[C^{-\gamma}, \Phi] = 0$). With perfect credit markets, the borrowing rate does not depend on collateral ($R(T^o) = R$) and the liquidity constraint does not bind ($\omega = 0$). With perfect factor markets the rental constraints never bind ($\underline{\kappa} = \overline{\kappa} = 0$). Then

$$\theta_T A T^{-(1-\theta_T)} K^{\theta_K} L^{\theta_L} = \tilde{w}^T \tag{2}$$

where $\tilde{w}^T = Rw^T$ is the effective rental price of land. This is the standard neoclassical result that firms set the marginal product equal to the price. Since all farmers set their marginal products equal to the same price, all farmers have equal marginal products, implying the allocation of land is efficient. Likewise, the allocation of labor and capital are efficient.

Now suppose factor markets are perfect ($\underline{\kappa} = \overline{\kappa} = 0$) but financial markets are not. Then the optimal choices of land and labor satisfy

$$\big(\mathbb{E}[C^{-\gamma}] + Cov(C^{-\gamma}, \Phi)\big)\theta_T A T^{-(1-\theta_T)} K^{\theta_K} L^{\theta_L} = (R(T^o)\mathbb{E}[C^{-\gamma}] + \omega)w^T$$
$$\big(\mathbb{E}[C^{-\gamma}] + Cov(C^{-\gamma}, \Phi)\big)\theta_L A L^{-(1-\theta_L)} K^{\theta_K} T^{\theta_T} = (R(T^o)\mathbb{E}[C^{-\gamma}] + \omega)w^L$$

Divide the first-order condition for land by the condition for labor:

$$\frac{L}{T} = \frac{\theta_L}{\theta_T}\frac{w^T}{w^L} \tag{3}$$

Assuming all farmers have the same production function, all farmers will satisfy this condition. Since under perfect factor markets all farmers pay the same prices w^T and w^L, all farmers in the village will employ the same labor per unit of land, and by similar logic they will employ the same capital per unit of land. Regardless of how badly the financial market fails, as long as factor markets are perfect the farmer can choose the right mix of inputs.

Why does $R(T^o)$ not appear in this expression? A naive intuition might suggest that when financial markets take land but not labor as collateral, they distort the mix of land to labor. This intuition does not hold because a farmer who buys land need not farm it. With perfect factor markets a farmer can buy land, take advantage of its collateral value, but rent it out to someone who could more profitably farm it.

To link this result to the prior literature, consider two farmers i and j. The land-labor ratio of i can be written

$$\frac{L_i}{T_i} = \frac{1+\tau_j}{1+\tau_i}\frac{L_j}{T_j} \tag{4}$$

for some τ_i and τ_j such that $\tau_i = \tau_j = 0$ if factor markets are perfect.[3] The τ term is what Hsieh and Klenow (2009) call the "capital distortion" because in their model it distorts the capital-labor ratio. Since it vanishes when factor markets are perfect, in this model the distortion has an economic interpretation: the

[3] $1+\tau_i = \frac{(R(T_i^o)\mathbb{E}[C_i^{-\gamma}]+\omega_i)+\frac{\overline{\kappa}_i-\kappa_i}{w^T}}{(R(T_i^o)\mathbb{E}[C_i^{-\gamma}]+\omega_i)}$.

factor market distortion.

2.3 Optimal Allocations

Suppose i is a farmer in village I observed to use $\bar{K}_i, \bar{T}_i, \bar{L}_i$. How would perfect markets allocate the aggregate stocks of land, labor, and capital across farmers? The observed aggregate stocks are $T_I = \sum_{i \in I} \bar{T}_i$ and so on. They do not change because I only reallocate the village's existing resources. (In Section 8 I consider reallocating resources between villages as well.) With aggregate stocks pinned down, I can ignore the supply side of the market and normalize $R = 1$.

Use (3) to eliminate K_i and L_i from (2) and define the production returns to scale $\sigma = \theta_K + \theta_T + \theta_L$, which I assume is less than one. Combine this expression with the market clearing condition $T_I = \sum_{j \in I} T_j^*$ to solve for the optimal allocations with fully perfect markets:

$$T_i^* = \frac{A_i^{\frac{1}{1-\sigma}}}{\sum_{j \in I} A_j^{\frac{1}{1-\sigma}}} T_I \tag{5}$$

Optimal capital and labor are similar. Call farmer i's output with perfect allocations $Y_i^* = A_i \Phi_i (K_i^*)^{\theta_K} (T_i^*)^{\theta_T} (L_i^*)^{\theta_L}$.

Now suppose factor markets are perfected but financial markets left untouched. Farmers choose the optimal mix of factors, as given by Equation 3, but their overall scale is unknown. Any assumption about scale would define an allocation. I consider a hypothetical case in which the farmer takes her original choices—those I observe in the data—and trades them in the perfected factor markets as though they were endowments. Let K_i^+ be the farmer's new choice of capital while $\bar{K}_i$ is still her original choice. Then the value of her new choices must add up to the value of her endowment, as determined endogenously in the market under the new prices w^{K+}, w^{T+}, w^{L+}:

$$w^{K+}K_i^+ + w^{T+}T_i^+ + w^{L+}L_i^+ = w^{K+}\bar{K}_i + w^{T+}\bar{T}_i + w^{L+}\bar{L}_i$$

I effectively drop the farmers into an Edgeworth economy. The farmer choosing a profit-maximizing mix of factors behaves like a consumer choosing a utility-maximizing bundle of goods. The resulting allocation is easy to compute and perfects each farmer's mix while leaving her scale unchanged. Each farmer's allocation may differ from what she would have choosen under perfect factor markets without the constraint that her scale is unchanged. But under an assumption I explain in Section 2.4, gains from moving to the computed allocation are a lower bound on true misallocation from imperfect factor markets.

Again taking $T_I = \sum_{j\in I} T_j^+$ as the market-clearing condition, the allocations under perfect factor markets are

$$T_i^+ = \frac{1}{\theta_K + \theta_T + \theta_L}\left[\theta_K\frac{\bar{K}_i}{K_I} + \theta_T\frac{\bar{T}_i}{T_I} + \theta_L\frac{\bar{L}_i}{L_I}\right] T_I \tag{6}$$

Optimal capital and labor are similar. Call farmer i's output with perfect factor markets $Y_i^+ = A_i\phi_i(K_i^+)^{\theta_K}(T_i^+)^{\theta_T}(L_i^+)^{\theta_L}$.

2.4 Costs of Misallocation

For each of the three scenarios, aggregate output in village I is the sum of each farmer's output under that scenario. Call actual aggregate output Y_I, counterfactual output with fully perfect markets Y_I^*, and counterfactual output with only perfect factor markets Y_I^+. I use two measures of misallocation: the gains from making markets efficient, and the fraction of efficient output achieved. The gains from reallocation (or simply "misallocation") measure how much output a village loses from misallocation. The fraction of efficient output achieved (or "efficiency") compares the real world to the world with perfect markets and appears naturally in the aggregate production function I derive in Section 2.5. Define

$$G_I = \frac{Y_I^* - Y_I}{Y_I} \qquad G_I^{FACT} = \frac{Y_I^+ - Y_I}{Y_I} \qquad G_I^{FIN} = \frac{Y_I^* - Y_I^+}{Y_I}$$
$$E_I = \frac{Y_I}{Y_I^*} \qquad E_I^{FACT} = \frac{Y_I}{Y_I^+} \qquad E_I^{FIN} = \frac{Y_I^+}{Y_I^*}.$$

The gains from perfecting each market add up to the overall gains ($G_I = G_I^{FACT} + G_I^{FIN}$), and overall efficiency is the product of factor and financial market efficiency ($E_I = E_I^{FACT} \cdot E_I^{FIN}$). The overall gains are a decreasing function of efficiency ($G_I = \frac{1}{E_I} - 1$).

My measure of factor market misallocation (G_I^{FACT}) may not equal the true gains from perfecting factor markets. I compute factor market misallocation by holding each farmer's scale of production fixed, but if factor markets actually became perfect a productive farmer would probably increase her scale. Factor market failures might directly distort a farmer's scale—for example, if she had to pay more for all inputs. Alternatively, a perfect mix might make farming more profitable (and a larger scale more attractive) because the farmer can allocate each dollar to the factor she needs most. Proposition 1, which I prove in Appendix A.1, formalizes this argument:

Proposition 1 *Let $\tilde{K}_i^+$ be the level of capital farmer i would choose if factor markets were perfected, financial markets left untouched, and the endowment constraint were not imposed. Assume $\mathbb{E}[A_i((\tilde{K}_i^+)^\sigma - (K_i^+)^\sigma)] > 0$. Then in expectation G_I^{FACT} is a lower bound on the true gains from perfecting factor markets and G_I^{FIN} an upper bound on the true gains from subsequently perfecting financial markets.*

The assumption states that with perfect factor markets the most productive farmers will increase their scale relative to the actual outcome.[4] The assumption would fail only if factor market failures somehow compensate for financial market failures—for example, if the farmers who cannot get bank loans can rent land more cheaply, and they are also the most productive farmers. The scenario is implausible in a poor rural village, where those shut out of financial markets are usually shut out of factor markets as well.[5]

2.5 Decomposing Aggregate Output and Growth

Growth accounting traditionally measures changes in per capita output rather than per firm output. To match the literature I decompose the growth in the village's rice output per household instead of per farmer. Suppose $\mathcal{I}$ is the set of all households (rice-farming or otherwise) in village I, and let $\mathcal{Y} = \frac{Y}{|\mathcal{I}|}$ be per household rice output. Let $Z_{it} = A_{it}\Phi_{it}$ denote overall productivity, Z_{It} its mean and $\tilde{Z}_{it}$ deviations from the mean. I use overall productivity to be consistent with traditional growth accounting (which computes an overall Solow residual)

[4]Capital simply stands in for scale of production. I could have phrased the proposition in terms of land or labor just as easily because the ratios of all factors are fixed by (3).

[5]The proof is an equivalency result. If the assumption failed and $\mathbb{E}[A_i((\tilde{K}_i^+)^\sigma - (K_i^+)^\sigma)] < 0$, the computed gains would be an upper bound. In the knife-edge case where $\mathbb{E}[A_i((\tilde{K}_i^+)^\sigma - (K_i^+)^\sigma)] = 0$ the computed gains equal the actual gains.

and because it is unclear how to split village-level shocks into anticipated and unanticipated parts.[6]

Then

$$\begin{aligned}\mathcal{Y}_{It} &= Z_{It}E_{It}\cdot\frac{1}{|\mathcal{I}|}\sum_{i\in I}\tilde{Z}_{it}(K_{it}^*)^{\theta_K}(T_{it}^*)^{\theta_T}(L_{it}^*)^{\theta_L}\\ &= Z_{It}E_{It}F(K_{It},T_{It},L_{It};\{\tilde{A}_{it}\}\{\tilde{\Phi}_{it}\})\end{aligned}$$

Recall from (5) the optimal factor allocations $K_{it}^*, T_{it}^*, L_{it}^*$ are only functions of the aggregate stocks and relative productivity. Taking the relative productivity distribution $\{\tilde{A}_{it}\},\{\tilde{\Phi}_{it}\}$ as a parameter, F is a function of only aggregate capital, land, and labor—the aggregate production function.

There is a similar decomposition for sample-wide output $\mathcal{Y}_t$, but since households were sampled into the survey in multiple stages the decomposition must weight villages by their size. Let φ_{It} be the population of village I as a fraction of the total population of all villages surveyed, and let $\chi_{It} = \frac{\varphi_{It}\mathcal{Y}_{It}^*}{\sum_I\varphi_{It}\mathcal{Y}_{It}^*}$ be the share of sample-wide output it produces under optimal within-village allocations. Let $Z_t=\sum_I\chi_{It}Z_{It}$ denote the output-weighted mean and $\tilde{Z}_{It}$ deviations from the mean of village productivity. Let E_t be sample-wide allocative efficiency with reallocation still within villages. One can show that $E_t=\sum_I\chi_{It}E_{It}$, which means

$$\begin{aligned}\mathcal{Y}_t &= Z_t\cdot E_t\cdot\sum_I\tilde{Z}_{It}\kappa_{It}F(K_{It},T_{It},L_{It})\\ &= Z_tE_tF(\{K_{It},T_{It},L_{It}\}) \qquad (7)\end{aligned}$$

[6]For example, how much of a district-year dummy is anticipated? The distinction does not matter for within-village reallocation but will affect how much growth is assigned to anticipated versus unanticipated aggregate productivity. Rather than make an arbitrary and misleading distinction I combine the two into overall productivity.

Since reallocation is within a village, the sample-wide aggregate production function depends on the aggregate factor stocks of each village.

Define g_t^V as the log change of any variable V over baseline. Since $g_t^{\mathcal{Y}} = g_t^Z + g_t^E + g_t^F$, growth in per household rice output decomposes into the contributions of improvements in productivity, improvements in factor allocations, and aggregate factor accumulation. Setting g_t^Z and g_t^E to zero, for example, shows how output would have grown if the rice sector had made no improvements to productivity or efficiency.

3 Context and Assumptions

Thailand is the world's largest exporter of rice. Yet rice farming absorbs a much larger share of the workforce than it contributes towards output. In 2013 nearly a quarter of Thailand's workforce labored to grow a harvest that accounted for only 1 percent of GDP. This disparity between labor spent and output gained is a hallmark of agriculture in developing countries. But aside from being similar to other farm sectors in other countries, Thai rice farming is a convenient sector to study because it is well-suited to the assumptions of my method.

My first assumption, which lets me measure overall misallocation from factor and financial market failures, is that all misallocation comes from these two sources. Comparing the original allocation of land, labor, and capital to the allocation that equalizes marginal products does not work if measured marginal products differ for other reasons. For example, the econometrician might miscalculate marginal products and find misallocation where there is none. He may assume the wrong technology or incorrectly assume firms use the same technology. Unanticipated productivity shocks might change firms' marginal products after they choose their factors, making the allocation look inefficient even when markets are perfect. Real forces other than weak factor and financial markets might also drive marginal products apart. Firms may pay different taxes, have adjustment costs, or be monopolists.

Such issues cause fewer problems in rice farming than in manufacturing. Rice production is relatively uniform. Though not all farmers in Thailand grow the same type of rice, they grow each variety following a similar technique.[7] And unlike in many developing countries nearly all farmers use modern pesticides and fertilizers. Figure 1 shows that nearly 100 percent of my sample uses modern farming technology (fertilizers or pesticides) throughout the entire sample period.[8] This is not to say everyone farms rice identically, but rather that assuming a common production function for rice is safer than assuming one for manufacturing. Since identifying the sources of anticipated versus unanticipated productivity is easier in rice production—a farmer knows his own talent but does not know whether rats will eat his harvest—it is easier to model productivity as described in Section 5.

Monopoly, taxation, and adjustment costs are less likely to distort the rice sector. Rice is a commodity and Thai farmers are all price-takers who sell their output to mills and merchants at market prices. In Shenoy (2014) I show that farmers' selling prices move with the international rice price. Though the gov-

[7]The farmer seeds a nursery plot and transplants the seedlings to a flooded paddy where they grow to adulthood. Farmers fertilize and apply pesticides until the rice matures and they harvest, thresh, dry, and sell the grains. According to the International Rice Research Institute, most of Thailand's farmers use this lowland rain-fed method to grow their rice.

[8]Thailand was an early ally of the U.S. during the Cold War and received American aid to modernize its rice sector in the 1960s. Despite fears to the contrary, both small and large farmers adopted the new seeds and fertilizers. The adoption of fertilizer was so rapid that, according to Baker and Phongpaichit (2009), a Japanese anthropologist visiting in 1970 found "Villagers who had described the local rituals to him only a decade ago now exclaimed 'the rice spirit is no match for chemical fertilizer'."

Figure 1
Characteristics of Thai Rice Farmers

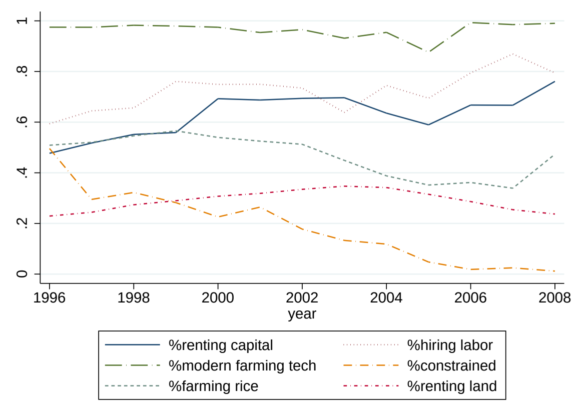

Note: Descriptive statistics of the sample. I describe the sample in more depth in Section 4.

ernment often supports prices, price subsidies will affect all farmers equally and leave allocations unchanged. As a result the prices received by farmers within a village show little variation. According to the monthly survey, in 75 percent of village-years the coefficient of variation in prices is less than 10 percent.

Though Thailand's farmers grow rice more commercially than their Indian or Chinese counterparts, they enjoy a similar freedom from taxes. Of the roughly 1500 survey households who reported any agricultural activity in 1996, only eight reported paying land taxes. Less than two percent of rice farmers in the monthly survey report paying any income tax.

My assumption that farmers in a village can exchange land, labor, and capital without adjustment costs is also plausible. Tractors and bullocks can be driven across the village, most on-farm machinery can be moved, and one farmer can store his crops in another's granary. There is no cost to hiring or firing a casual farm worker. Exchanging land is not difficult, as over three quarters of the rice paddies cultivated in 1996-1997 are no more than two kilometers from the village.[9]

The assumptions I make about functional forms are also more plausible among farmers. Hicks-Neutral productivity, for example, may not hold in manufacturing and services. But since a rice farmer's inputs work to make a single product, year-to-year productivity shocks like poor rainfall will damage the end crop rather than the contribution of the workers versus the tractors. Decreasing returns to scale, which I assume in my derivations, is plausible because rice farming in Thailand is still labor-intensive. A large farm with a large workforce is harder to manage than a small one. It is harder to reason whether or not farm production is Cobb-Douglas. Though this assumption is standard in the literature on misallocation, I show that it approximates reality in Appendix A.5.[10]

By using a household survey I avoid the problem of selective attrition. Even if a household stops farming it remains in the panel and re-enters my sample if it starts farming again. Thus the farmers and factors I observe give an accurate reflection of the current state of the rice sector.[11]

Finally, my method assumes I can model the farmer as though she chooses

[9] Clearly building a tractor or raising a bullock takes time. But with perfect factor markets, whatever capital exists in the village can flow to the farmers who can make best use of it as long as there are no costs to exchanging factors. This is what matters for measuring misallocation.

[10] The general model and procedure I present in Appendix A.8 does not require any assumptions beyond concavity, decreasing returns, and twice-differentiability. But estimating a more complicated production function requires stronger assumptions about productivity and factor input choices. Moreover, it is unlikely a more complicated production function would produce a convenient closed-form expression for the optimal allocation.

[11] The main caveat is that I must exclude from the sample farmers who farm only once because I cannot calculate a fixed effect for them. If they are severely over- or under-allocated, villages will falsely appear efficient. A one-time farmer, however, can only enter and exit if the rental or purchase markets work well. As villages with efficient markets have little misallocation, if anything the bias works towards finding too much misallocation because the efficient villages receive less weight in the sector-level calculation of efficiency.

and pays for her inputs at the same time. This assumption lets me interpret distortions to mix as factor market misallocation. This assumption is reasonable if the cost of transferring land and capital between farmers is small and the farmer has equal information about productivity when choosing all inputs. I have already argued adjustment costs are small. The information assumption is clearly a simplification. The farmer will know more about rainfall and other random events later in the season when hiring workers for harvest. The question is whether the variation in productivity caused by this new information is large compared to what the farmer knows at planting and what is known only after harvest. If the assumption fails, the measures measures of misallocation would respond incorrectly to an exogenous injection of credit. But I show in Section 7 that the measures do respond as they should, suggesting the simplification is not too extreme.

4 Data

I construct my sample from the Townsend Thai Annual Household Survey (1997). The Townsend Thai Project collected a baseline survey of households to be representative of four provinces. Two are from Thailand's underdeveloped northeast, and two from its more developed center. The Project chose these provinces to capture the country's regional disparities. Within each province twelve sub-districts were chosen by stratified random sampling to ensure the province's environmental diversity was represented. From each sub-district four villages were chosen at random, and within each village 15 households were chosen at random. Despite the stratification, I show in Appendix A.7 that the distribution of agricultural landholding in the baseline data is broadly representative.

The Project subsampled one third of the survey villages and resurveyed the sampled households every year to construct a panel. It later added two more provinces and sampled new households to counter attrition. I use the rounds collected from 1997 through 2009. Since the survey recall period is from June of the previous year through May of the survey year, I label the period covered by the 1997 survey as 1996 and so on. The Project followed sixteen of the villages excluded from the annual survey to collect the Townsend Thai Monthly Household Survey (2012). I use the first two years of the monthly survey throughout the paper to compute stylized facts not found in the annual survey. In the appendix I use district-level precipitation data computed from the University of Delaware Climactic Project and NASA's Tropical Rainfall Measuring Mission to validate the self-reported household measures of rainfall shocks.

Land is the number of rai (6.25 rai = 1 acre) of paddy the household cultivated (whether owned or otherwise). Labor is the sum of hired and family labor in days worked. Hired labor is the household's expenditures on farm workers divided by the median daily wage in the village. Using the median wage is not ideal, but the survey does not ask directly about the amount of labor hired and the within-village variation in unskilled wages is relatively low (the coefficient of variation is less than 0.19 for most village-years). I count the number of household members who report being unpaid family laborers with primary occupations in farming of any sort (or who mention "FIELDS" in a freeform response). The annual survey gives no information on the days each member worked. Instead I use the more detailed labor data in the monthly survey to calculate the median days any individual works on his family's fields (conditional on working at all), and multiply the median—60 days—by the number of family laborers counted in the annual data.

Capital is the sum of the value of owned mechanical capital, the value of owned buffalo, and the value of rented capital and expenses (including intermediate inputs). I do not compute the value of owned capital using perpetual inventory because households do not report the value of assets they sell, meaning I cannot measure disinvestment. Instead I assign a purchase value to each asset the household owns. I deflate and depreciate the purchase value of assets owned at baseline. For assets acquired afterwards I use the purchase price. The survey only reports assets in classes, so if the household has multiple assets of the same type I must treat them as if they have identical value and use the most recent purchase price (most households own one or fewer assets of any type). If I cannot identify a price I drop the asset from the calculation (I can identify a price for the vast majority). I then depreciate the purchase price to get the value in a given year assuming 2 percent depreciation for structures (House and Shapiro, 2008), 10 percent depreciation for machines, and (I treat them as vehicles) 20 percent depreciation for tractors (Levinsohn and Petrin, 2003). Owned mechanical capital in a year is the total value of the assets. I treat intermediate inputs—seeds, fertilizers, pesticides, and fuel—as capital with a 100 percent deprecation rate. I add maintenance, which I treat as investment that takes immediate effect. I then add the purchase price of rented capital, which I approximate with total rental expenses divided by an interest rate of .04 plus the average rate of deprecation for all types of capital (a user cost).[12] Fi-

[12]Given the presence of misallocation, how can I assume a common interest rate and a user cost? Recall my objective is to construct a consistent measure of the value of capital. Reweighting using household-specific interest rates would be equivalent to calling a tractor more valuable because the household renting it pays a higher mortgage. This is not to say farmers do pay the same user cost for capital, only that the productive value of each asset is independent of financial market imperfections. It is a bigger problem if there is variation in the prices farmers actually pay for capital they rent, as I would measure a farmer as having more capital when he

nally, I add the value the household reports for its buffalo. Since households do not report whether they rented out their capital I cannot lower it to reflect how much they actually use. The error might inflate estimated misallocation because unproductive farmers who rent out their machinery will appear to have too much capital.[13]

To borrow the expression of Hsieh and Klenow (2009), my procedure "heroically makes no allowance" for measurement error. Yet my measures of land, labor, and capital are noisy. I show in Section 5.3 that my estimates of productivity are not driven by measurement error in land, and show in Appendix A.5.4 that measurement error would not qualitatively change the results. Nevertheless, measurement error is a limitation of any study that takes a model seriously. This limitation makes it critical that I validate my measures of misallocation in Section 7.

In Section 5.2 I model productivity using several catastrophes the household reports about its income. I use indicators for illness, death in the family, flooding, problems with crop-eating pests, poor rainfall, low yield for other reasons, and a low price for output. Malnourishment might also lower the farmer's productivity. To proxy for it I use the share of the household's consumption budget devoted to rice, the staple food, including the value of home-produced rice. Jensen and Miller (2010) argue as households become less hungry they substitute away from the staple, so a larger share implies more hunger. All monetary variables are deflated to 2005 Thai baht. I describe all the variables in more detail in Appendix A.7.

---

only pays more for it. This is a limitation of any study that measures capital with its cost.

[13]The survey does ask households about "Payments for other rentals," but the rented goods might not be capital and almost no one reports receiving any (only three do in the latest year when we would expect the best rental markets).

Table 1
Sample Descriptives

Revenue from rice	47911.7	Fraction who report. ..	
	(77335.2)		
Capital	74541.4	Illness	0.01
	(93835.1)	Death in Family	0.01
Land	19.4	Flood	0.05
	(15.9)	Crop-Eating Pests	0.04
Labor	181.6	Bad Rainfall	0.22
	(191.9)	Low Yield for Other Reason	0.21
Rice Budget Share	0.4	Low Price	0.15
	(0.2)		
Households	775	Observations	6228
Villages	69		

Note: I restrict the sample to households with at least two years of positive revenue and positive capital, land, and labor. All variables are annual. Revenue and capital are in 2005 baht, land in rai, and labor in human-days. The share of the household's consumption budget spent on rice is my measure of hunger.

Table 1 reports household-year averages of each variable for the sample. I restrict my analysis to the households I observe with positive rice revenue and levels of all factors for at least two years, as I cannot calculate a household fixed-effect for anyone else. At 2005 exchange rates the average annual revenue from rice was roughly 1200 dollars. Farms are small and most farmers plant only 19.4 rai (3.1 acres) of paddy.

5 Estimating the Production Function

To apply the expressions for optimal land, labor, and capital derived in Section 2, I must estimate the production function. Estimating a production function is never easy because the choices of inputs depend on unobserved productivity. Misallocation complicates the task because I cannot assume firms choose their inputs optimally. As I explain, however, misallocation lets me identify the production function using the Anderson-Hsiao estimator.

5.1 Challenges to Estimating the Production Function under Imperfect Markets

A common approach in macroeconomics is to calibrate the production elasticities $\theta_K, \theta_T, \theta_L$ to the cost of each input as a share of revenue. Define the calibration estimate of the (inverse) elasticity of land as

$$\widehat{\left(\frac{1}{\theta_T}\right)} = \sum_i \frac{Y_i}{\tilde{w}^T T_i}$$
$$= \frac{1}{\theta_T} \sum_i \frac{\Phi_i(R(T_i^o)\mathbb{E}[C_i^{-\gamma}] + \omega_i) + \frac{\overline{\kappa}_i - \underline{\kappa}_i}{w^T}}{R\big(\mathbb{E}[C_i^{-\gamma}] + Cov(C_i^{-\gamma}, \Phi_i)\big)}$$

where the last equality follows from Equation 1.[14] This estimate is consistent if $\omega_i = \overline{\kappa}_i = \underline{\kappa}_i = Cov(C_i^{-\gamma}, \Phi_i) = 0$ and $R(T_i^o) = R$—that is, if markets are perfect. In other words, in my model this calibration succeeds only if there is no misallocation.[15]

For similar reasons, misallocation rules out a class of methods Ackerberg, Caves, and Frazer (2006) call "structural techniques." These methods assume

[14] By Jensen's Inequality the more intuitive estimator $\hat{\theta}_T = \sum_i \frac{\tilde{w}^T T_i}{Y_i}$ is inconsistent even when markets are perfect if Φ_i varies by farmer.

[15] Some prior work (e.g. Hsieh and Klenow, 2009; Adamopoulos and Restuccia, 2014) calibrates these parameters using data from the U.S., which is valid if the U.S. uses a similar production technology to developing countries but has relatively well-functioning markets. In my context, the assumption would imply that farmers in California use the same technique as farmers in Thailand. Farmers in the U.S. rely on everything from mechanized combines to aerial drones to grow their crops, whereas those in Thailand may still farm with bullocks and elephants. This makes it unappealing to assume the two countries use the same production technology.

the firm chooses intermediate inputs to match its anticipated productivity. All else equal, more intermediate inputs imply higher anticipated productivity, allowing the econometrician to control for productivity. But if markets for intermediates are imperfect—for example, if some farmers lack the credit to buy fertilizer while others can borrow from their family—these methods may also be inconsistent.[16]

To test whether markets are perfect, suppose each observation varies by household i and year t. Substitute

$$A_{it}T_{it}^{-(1-\theta_T)}K_{it}^{\theta_K}L_{it}^{\theta_L} = \mathbb{E}[\frac{y_{it}}{K_{it}}].$$

into Equation 2, the first-order condition *under perfect markets*:

$$\theta_T \mathbb{E}[\frac{y_{it}}{T_{it}}] = \tilde{w}_t^T$$
$$\Rightarrow \frac{y_{it}}{K_{it}} = \upsilon_{It} + \varepsilon_{it}$$

where υ_{It} is a village-year dummy and ε_{it} is a rational expectations error. In other words, if markets are perfect the output-capital ratio should not be correlated with anything—for example, the log of land or labor—after controlling for village-year fixed-effects.

Table 2 regresses the ratio of output to land, labor, and capital on the log of each factor. I exclude the log of land from the regression for the ratio of output to land (and so on) because measurement error in land might create a correlation even if markets are perfect. I report the p-value on the F-test that none

[16]Ackerberg et al. would say the "scalar unobservable" assumption fails. In principle, one can deal with this problem by adding controls for unequal access to factors. But the case described here is an example where few datasets would contain the necessary variable (having wealthy family members).

Table 2
Test of Imperfect Markets

	(1) MP(K)	(2) MP(T)	(3) MP(L)
Log(Land)	0.12**		96.23***
	(0.06)		(11.45)
Log(Labor)	-0.08	205.29***	
	(0.06)	(64.02)	
Log(Capital)		21.79	7.51
		(62.72)	(8.14)
Observations	6228	6228	6228
Households	775	775	775
Pval on F-Test	0.09	0.00	0.00

Note: Under perfect markets the ratio of output to any factor should not be correlated with the logs of factors after controlling for village-year fixed-effects. I report the results of such regressions. Standard errors are clustered within household. I exclude the log of capital from the output-capital ratio (and so on for the other factors) because measurement error in capital might create a correlation even if markets are perfect.

of the regressors are significant. In all three cases I reject no correlation at the 10 percent level, and in two cases I reject at the 1 percent level. The regressions suggest farmers who have a lot of labor lack land and vice-versa, consistent with failures in the markets for land and labor. Farmers with large families should be sending some relatives to work on the farms of their landed neighbors but instead employ them at home.

5.2 The Anderson-Hsiao Estimator

Since the assumptions required for calibration and structural methods do not hold, I take a different approach. The first step is to model what makes one farmer more productive than another. Much of what determines a firm's revenue productivity in manufacturing or services—a successful marketing campaign, a new product line, the monopoly power born of a competitor's demise—

are absent in agriculture. Many of the most obvious determinants of a farmer's productivity—rainfall, crop-eating pests, illness, accidental misapplication of fertilizer—either affect everyone in the village or are unanticipated. As I argued in Section 3, Thai farmers have used modern seeds, pesticides, and fertilizers for decades, making it unlikely any farmer has a technological edge.

Since malnourishment might lower the farmer's productivity I construct a measure of hunger. What remains is the farmer's own managerial talent: his knowledge of how to eke the most output from his inputs of land, labor, and capital. Managerial talent is a fixed characteristic of the farmer I can capture in a household fixed-effect. I model the log of anticipated productivity a_{it} and the log of unanticipated productivity ϕ_{it} as follow:

$$\begin{aligned} a_{it} &= [Household\ Fixed\ Effect]_i + a^H[Hunger]_{it} \\ &\quad + \sum_k a_k^D[\textit{District-Year Dummies}]_{k,it} \end{aligned} \tag{8}$$

$$\phi_{it} = \sum_j a_j^S[Dummy\ Shocks]_{j,it} + [Overall\ Error]_{it} \tag{9}$$

$$\tag{10}$$

In the dummy shocks I include indicators for illness, death in the family, retirement, flooding, problems with crop-eating pests, poor rainfall, low yield for other reasons, and a low price for output. The overall error includes both measurement error and unanticipated idiosyncratic shocks not covered in the dummies.

The district-year dummies are assigned to anticipated productivity here, but actually will not affect my measures of misallocation. Any shock that affects every farmer in the village equally will simply divide out of the expressions for optimal allocations in Section 2.3. In the aggregate output decomposition of Section 2.5 I combine both anticipated and unanticipated productivity into overall productivity, making the distinction unimportant. Finally, in Section 8 I assess how much the main results change when farmers anticipate all productivity.

Let lowercase letters denote logs (i.e. $y_{it} = \log Y_{it}$), let $\tilde{\phi}_{it} = \phi_{it} - [Overall\ Error]_{it}$, and let $\tilde{a}_{it} = a_{it} - [Household\ Fixed\ Effect]_i$. Since the bulk of a farmer's anticipated productivity is fixed, it seems natural to estimate

$$y_{it} = [Household\ Fixed\ Effect]_i + \tilde{a}_{it} + \tilde{\phi}_{it} + \theta_K k_{it} + \theta_T t_{it} + \theta_L \ell_{it} + [Overall\ Error]_{it}$$

with the within-household estimator or equivalently OLS with household dummies. But the key assumption for its consistency—what Wooldridge (2002) calls strict exogeneity—fails. Strict exogeneity requires that unexpectedly high or low output in either the past or future will not affect a farmer's input decisions today. But a credit-constrained farmer that suffered a bad harvest last year may be unable to rent land this year. Aside from potentially causing misallocation, the situation violates strict exogeneity.

The Anderson-Hsiao estimator (1981; 1982) estimates the production function under a weaker assumption called sequential exogeneity. Sequential exogeneity assumes a farmer will not base her input decisions on unexpectedly high or low *future* output, but makes no assumptions about past output. In other words, current and future error terms are unanticipated shocks to productivity. I implement the estimator by taking first-differences to eliminate the fixed-effect and instrumenting the differenced factors with their lagged levels.

The identification assumptions are

$$\mathbb{E}\big(k_{i,t-1}[Overall\ Error]_{i,t-1}\big) = \mathbb{E}\big(k_{i,t-1}[Overall\ Error]_{it}\big) = 0$$
$$\mathbb{E}\big(k_{i,t-1}\Delta k_{it}\big) \neq 0$$

with similar assumptions about land and labor. The first assumption—that input choices are uncorrelated with the residual error term—is implied by the assumption that ϕ_{it} (and each of its components) is unanticipated.[17] I already make this assumption to compute misallocation.[18]

The second assumption is that lagged capital is informative about changes in capital (that is, the instrument is relevant). Since there are no adjustment costs, if markets were perfect this assumption would not hold. A farmer would choose the optimal level of each input without regard to its level in the past. But ironically the imperfect markets that cause misallocation also create a correlation between levels and changes of inputs. A farmer who lacks credit might fire workers after a bad harvest and slowly rebuild his labor force. The resulting mean reversion creates variation that can be used by the Anderson-Hsiao estimator to identify the production function.

Table 3 reports the estimates of the production function. As expected, rice farming is relatively labor- and land-intensive, and each shock to productivity has the expected sign (except death in family, which is insignificant and close to zero). The first-stage regressions of factor changes on their lags easily satisfy the usual standards for strength (Stock, Wright, and Yogo, 2002). The produc-

[17] That is, if ϕ_{it} were correlated with k_{it} it would imply that the farmer's choice of inputs is a function of ϕ_{it}, which is only possible if the farmer anticipates ϕ_{it}.

[18] Although other dynamic panel estimators—Arellano-Bond or Blundell-Bond—make even weaker assumptions, these estimators also have higher variance. I choose Anderson-Hsiao as a compromise between bias and variance.

Table 3
Production Function Estimates

Production Elasticities		Productivity Modifiers	
Capital		Hunger	-0.03
-Share (θ_K)	0.11		(0.08)
	(0.04)	Illness	-0.15
-1st Stage F-Stat	128.78		(0.09)
		Death in Family	0.03
Land			(0.12)
-Share (θ_T)	0.25	Flood	-0.10
	(0.05)		(0.04)
-1st Stage F-Stat	112.67	Pests	-0.06
			(0.04)
Labor		Bad Rain	-0.13
-Share (θ_L)	0.31		(0.03)
	(0.04)	Low Yield	-0.15
-1st Stage F-Stat	134.37		(0.02)
		Low Price	-0.06
Returns to Scale:			(0.02)
-Estimate (σ)	0.67		
	(0.07)		
Households:	734	Observations:	4856
R-Squared:	0.76	Kleibergen-Paap Stat.	114.85

Note: The table reports the Anderson-Hsiao estimates of the production elasticities and the effects of each component of productivity (see Section 5.2 for details). The variable "Death" refers to a death in the (extended) family. I cluster standard errors by household.

tion function has decreasing returns, justifying the span-of-control assumption made in Section 2. Decreasing returns is not surprising; unlike the highly mechanized farming of the U.S., rice farming in Thailand still requires the farmer to manage workers and animals with her limited span-of-control.

Table 4
Sample Sizes and Productivity in Rice-Farming Villages

Farmers Per Village		Productivity Dispersion	
25th Pctl.	6	75/25	1.73
50th Pctl.	10	90/10	2.56
75th Pctl.	12	95/5	3.09

5.3 Characteristics of Productivity

Table 4 reports sample sizes and the median dispersion of anticipated productivity $\hat{A}$ among the villages of my sample. The median 90/10 ratio for productivity within a village is 3.09, a number close to the range of 1 to 3 that Gandhi, Navarro, and Rivers (2013) find for the gross production functions of several manufacturing industries in Colombia and Chile. Productivity in a rice-farming village is distributed much like in a typical manufacturing industry. Hsieh and Klenow (2009) find much larger 90/10 ratios in their sample, possibly because they assume a value-added production function. Gandhi et al. show that, relative to the gross production function assumed here, value-added production functions tend tend to inflate the dispersion of productivity.

Though the distribution of productivity seems reasonable, where does it come from? The level of misallocation depends crucially on whether the measure of productivity is meaningful. Section 5.2 interpreted anticipated productivity as managerial talent. That suggests it should be correlated with measures of human capital. It would be a problem if estimated productivity were instead absorbing unmeasured land quality. If wealthy farmers were also able to accumulate the best land in the village, they would appear productive even if their impoverished neighbors could be equally productive if given the same land. This might bias my measures of misallocation towards zero.

The Townsend Thai Annual survey has only a few measures of land quality. One is the price of land. Though the price may not perfectly reflect the underlying quality of land (especially if land markets are dysfunctional), it would nevertheless be a problem if price were correlated with productivity. The other two measures are indicators for whether a plot of land has access to water in the hot or cold season. I compute the fraction of a farmer's land with access in each season (in practice, nearly all farmers have a fraction of either 1 or 0). These two measures are only available at baseline. However, since anticipated productivity consists largely of the household fixed effect, the cross-sectional correlation using only data from 1996 is no less informative than a panel regression. I regress the log of anticipated productivity on these three measures of land quality and two measures of human capital: the years of primary schooling of the head of household, and the decades the head has spent farming rice.

Column 1 of Table 5 shows that access to water in the hot season and the price of land are correlated with productivity, but Column 2 shows that that these correlation vanish after I control for village fixed-effects. This suggests the biggest differences in land quality are between villages rather than within villages—for example, between the fertile central provinces and the poorer northeast. This is not to say there are no differences in land quality within a village, but rather that the differences are small relative to the differences between villages and to differences in other sources of productivity. Since I estimate the gains from reallocating resources within villages, differences in productivity between villages will have no effect on measured misallocation.

By contrast, the measures of human capital are highly significant predictors of productivity. As expected, farmers with more primary schooling are also

Table 5
Correlates of Anticipated Productivity (1996)

	(1)	(2)
Value of Land:		
Water in Hot Season	0.20**	0.07
	(0.10)	(0.05)
Water in Cold Season	0.13	-0.06
	(0.08)	(0.05)
Land Price (Log)	0.17***	-0.01
	(0.04)	(0.02)
Human Capital:		
Primary Schooling	0.04**	0.04***
	(0.02)	(0.01)
Decades Farming Rice	-0.07***	-0.05***
	(0.03)	(0.02)
Households	356	356
Village FEs		X

more productive. This is not necessarily the return to schooling, as farmers with higher ability may stay in school longer. Nevertheless it suggests that farmers with more human capital are more productive. More surprising is the robust negative correlation between the time spent farming and productivity. It suggests that older farmers may be using outdated methods of planting. Taken together, the results support the interpretation of productivity as a measure of the farmer's managerial talent.

6 Results

I plug the estimates of production elasticities and anticipated productivity into the expressions for the allocations under fully perfect markets (5) and perfect

Table 6
Correlates of Under-Allocated Farmers

Correlates of Being Under-Allocated in 1996			
Age of Head (Decades)	Rents Land?	Decades Farmed Rice	Illness
-4.9*	-9.4	-5.8**	-7.3
High Risk Aversion	Cash Savings	Factor Constrained	
9.6*	0.0	14.5**	

I define "under-allocated" to mean the household produces more after reallocation than before. I regress a dummy for being under-allocated in the first year (1996) on province fixed effects and a set of variables that might cause a household to be allocated too much or too little. The table reports the coefficients and significance levels from the regression. The sample size is 350. Most of the loss in sample is because the risk aversion question was not asked until 2003, meaning I use only farmers who were surveyed in both 1996 and 2003.

factor markets (6).[19]

6.1 Which Households Are Under-Allocated?

Table 6 reports the correlation between the probability of being under-allocated in 1996 and several characteristics of the household. I call a farmer under-allocated if he produces more after reallocation than before. I predict under-allocation with the farmer's age and years farming rice, his cash savings, whether he rents land, whether illness lowered his income, whether he reports being unable to acquire inputs needed to profitably expand his business, and whether he chose the most risk-averse options in two questions that measure risk preferences.[20]

Ten more years of age or experience reduce by 5 and 6 percentage points the chance of being under-allocated, suggesting farmers accumulate factors as they

[19] I drop all observations from village-years with only a single farmer because they by construction have no within-village misallocation. The sample loses 13 households, 8 villages, and 47 observations.

[20] Each question asked whether he would accept a gamble that could with equal probability double or reduce to two-thirds his current income. If he refused he was offered a similar gamble where the worse outcome would give 80 percent of his current income. I mark the farmer as highly risk averse if he refused both offers. The question was first asked in the 2003 survey, so I linked the 2003 response to the 1996 status of under-allocation.

age. As expected, a risk-averse farmer is 9.6 percentage points more likely to be under-allocated. A farmer who fears risk is less likely to gamble on a large farm even if he is more talented than his neighbors. Savings do not predict under-allocation, perhaps because they may equally be a sign of wealth or unwillingness to invest. Finally, farmers who report being "factor constrained" (unable to acquire inputs) are, not surprisingly, more likely to be under-allocated.

6.2 Village-Level Misallocation

I then estimate the misallocation in each village using the expression for G_I from Section 2.4 and plot the distribution in Figure 2.A. Even in the earliest year of my sample (1996) the total cost of factor and financial market failures was less than 18 percent in most villages. Over time the distribution shifts downward, and misallocation falls below 10 percent for most villages by 2008. Many villages appear to have negative misallocation because estimated misallocation is a random variable. When true misallocation is low the probability a normally distributed estimator falls below zero is high.[21]

Table 7 shows, after controlling for province fixed-effects, the correlation between the cost of village-level misallocation and socioeconomic features of the village. As these characteristics are only available at baseline, these estimates can be interpreted as no more than small sample correlations. Nevertheless these correlations are useful in gauging what aspects of the village predict imperfect markets. In addition to studying overall misallocation G_I I calculate G_I^{FACT} and G_I^{FIN} as defined in Section 2.4, which separate the misallocation

[21] The estimates are random for two reasons: sampling error, as arises with any estimate made with a finite sample, and the unanticipated shocks. The optimal allocations I compute are ex ante perfect, but their ex post efficiency will depend on the realization of shocks. If the most efficient farmers get unlucky draws the optimal allocation will look less efficient.

caused by imperfect factor markets from that caused by imperfect financial markets.

Table 7 shows that misallocation is roughly 1 percentage point higher for every 10 kilometers of distance between a village and the nearest two-lane paved road. Since these roads typically connect villages to the outside world, it is not surprising that proximity to one predits lower financial market misallocation. More accessible villages are likely better integrated into markets, granting them better access to credit and through trade providing some insurance against income shocks. Accessibility does not predict lower factor market misallocation, suggesting linkages between villages do not improve the markets for land, labor, or capital within a village.

Villages where a larger fraction of households have access to electricity have lower factor market misallocation. This could be because electricity facilitates rental markets. It is also possible that, by providing better outside options, it allows large households to move their excess labor into more productive activities (for example, by making it easier for children to attend school).

The most surprising result is that the fraction of the village that speaks a minority language is strongly negatively correlated with misallocation. This effect may be causal—for example, if minority ethnic networks are better able to enforce contracts and thus provide credit or exchange inputs (see, for example, Greif, 1993). It may also be driven by selection bias—for example, if minorities are more likely to settle in wealthy villages with well-functioning markets. One might expect a village with many minorities to suffer from a fragmented market. Table 7 suggests a village with a high degree of ethnolinguistic fractionalization has more misallocation, though the effect is not significant.[22] It bears repeating

[22]For each of several minority languages the village survey asks how many households speak

Table 7
Correlates of Village-Level Misallocation

	(1) Overall Cost	(2) Cost, Fact.	(3) Cost, Fin.
Distance to Road (km)	0.105***	-0.009	0.114***
	(0.020)	(0.006)	(0.019)
Electrification	-1.882	-6.814**	4.932
	(9.691)	(2.585)	(8.358)
Fraction Minority	-37.236***	-8.137***	-29.099***
	(9.007)	(2.995)	(8.433)
Fractionalization	13.747	0.272	13.475
	(9.819)	(1.824)	(9.243)
Province FEs	X	X	X
Observations	54	54	54

These regressions show several correlates of each type of misallocation. “Distance to Road” is the distance to the nearest paved, two-lane road. “Electrification” is the fraction of households that have electricity. “Fraction Minority” is the fraction of households that speak a minority language at home. “Fractionalization” is the linguistic fractionalization (the probability two randomly chosen households do not speak the same language). The predictors are only available for 1996. Standard errors are robust to heteroskedasticity.

that these correlations may not be causal. In Section 7 I validate the measures of misallocation using exogenous variation.

6.3 Aggregate Misallocation and the Growth Decomposition

To calculate the aggregate cost of within-village misallocation, Figure 2.B depicts misallocation across all villages from the original four provinces surveyed at baseline.[23] I estimate the gains in sample-wide output from reallocating fac-

the language at home. A household may speak more than one language at home, meaning in some cases the sum of households speaking minority languages is greater than the number of households in the village. Lacking a better option I define the "gross number of households" as the maximum of the sum and the total households in the village. I define the number of Thai-only speakers as the difference between gross households and the sum of minority-speaking households. I define fractionalization as $1-\sum^k f_k^2$ where f_k is the fraction of (gross) households speaking language k (where Thai-only households are among the set of k).

[23] I restrict the sample to households from the original four provinces surveyed at baseline to avoid the artificial jump that comes from adding a new province partway through.

Figure 2
Panel A: Density of Within-Village Misallocation; **Panel B:** Sample-Level Overall, Factor, and Financial Market Misallocation

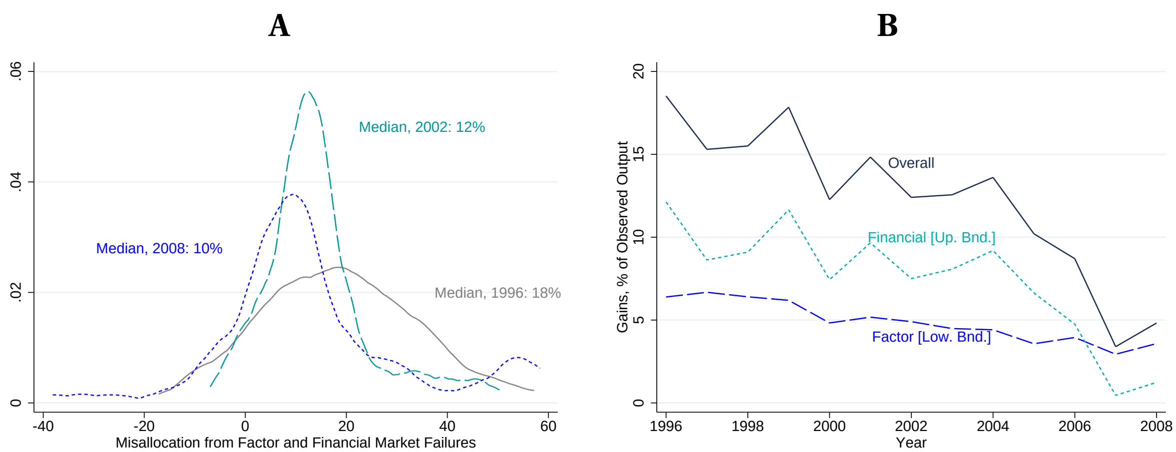

Note: Panel A plots the shifting distribution of misallocation within each village . I report misallocation as the fraction of observed output foregone because factors are misallocated. Panel B calculates the overall cost to the rice sector from misallocation within villages for every year of my sample. It also splits overall misallocation into misallocation from factor versus financial markets, where my measures bound the gains from perfecting first the factor markets and then the financial markets.

tors within each village after weighting by population (see Section 2.5). Sample-wide misallocation is never more than 19 percent and falls to below 5 percent by 2008. The results suggest factor and financial market imperfections never cause much misallocation, and what little they cause falls over time.

I likewise aggregate my decomposed measures of factor versus financial market misallocation to the sample-level. Both types of misallocation fall from 1996 to 2008. Since the factor market measure is a lower bound while the financial market measure is an upper bound, neither market unambiguously causes more misallocation until 2007 when financial market misallocation drops to nearly zero. Policymakers and donors often target interventions at financial markets, but the graph suggests factor markets cause as much or even more misallocation.

Table 8
Estimated Misallocation

	Overall	Factor Market	Financial Market
1996	18.5	6.4	12.1
	(8.3, 29.8)	(3.9, 7.5)	(3.6, 23.0)
2002	12.4	4.9	7.5
	(7.6, 16.9)	(3.3, 5.6)	(4.2, 11.8)
2008	4.8	3.6	1.2
	(1.1, 8.2)	(2.3, 4.0)	(-1.9, 4.6)

Bootstrapped 90 percent confidence intervals are given in parentheses.

Table 8 gives the point estimates of misallocation in 1996, 2002, and 2008. The 90 percent bootstrapped confidence intervals are given in parentheses below each point estimate. Though the estimates are noisy, they confirm that misallocation has declined. For all three measures the level of misallocation in 2008 lies below the confidence interval of the estimate in 1996. Though I can reject that overall misallocation and factor market misallocation are zero in 2008, I cannot reject that financial market misallocation is zero. As suggested by Figure 2.B, financial market misallocation has almost vanished by the end of the sample, whereas factor market misallocation remains present.

Figure 3 decomposes growth into changes in aggregate factor stocks $F(\cdot)$, revenue productivity Z, and the efficiency of factor allocations E. Each line shows how log output would have grown since 1996 if some parts of growth had been shut down. The solid line shows output if productivity and efficiency were fixed at their 1996 level and only aggregate factor stocks changed. Without growth in productivity and efficiency, rice output would have fallen since 1996 as factors flowed out of rice farming. Since Thailand has rapidly industrialized over the past two decades, agriculture's decline is not surprising.

The middle dashed line shows growth if changes in productivity are turned

Figure 3
Decomposition of Growth in Aggregate Rice Output in the Sample

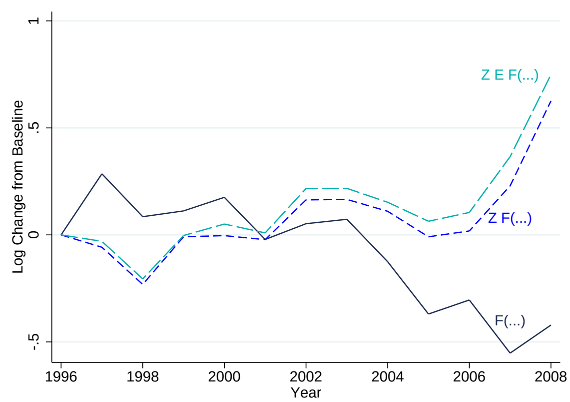

Note: I decompose aggregate output and compute changes in the log of the aggregate factor stocks $F(\cdot)$, revenue productivity Z, and the efficiency of factor allocations E. Each line plots the counterfactual change in output holding all components except the indicated component fixed (so the lowest line holds average productivity and allocative efficiency fixed while letting aggregate factor stocks change).

back on, and comparing it to the solid line shows the contribution of productivity to growth. Rising productivity since 1998 overwhelmed the outflow of factors and produced net gains in rice revenue. Revenue productivity rose for two reasons: better yields and higher prices. Average yields might have improved as less productive farmers left farming and those who stayed became more skilled, but the spike in productivity after 2006 comes entirely from rising food prices. As I show in Appendix A.6.1 physical productivity and prices often move in opposite directions. Prices went into a long decline from 1998 to 2002, which may have driven the least productive farmers into other sectors. As prices recovered from 2002 to 2008, these unproductive farmers were likely coaxed back into rice farming. But the decline in physical productivity was outweighed by the rise in prices, causing revenue productivity to rise.

But the most important feature of Figure 3 is the highest line, which shows the change in output when changes in efficiency are turned back on. Comparing it to the line below shows the contribution of reductions in misallocation. It is trivial. Compared to the other two sources of growth, efficiency barely changed the trajectory of rice output.

6.4 Comparison to Prior Literature

The level of misallocation reported in Table 8 is much lower than what has been found in earlier work. Hsieh and Klenow (2009), for example, find that India could raise manufacturing output by 40 to 60 percent if its labor and capital were allocated as efficiently as that of the U.S., and by over 100 percent if the allocations were perfect. The difference in results may suggest there is less misallocation in agriculture than in manufacturing. Early work in development (e.g. Townsend, 1994; Benjamin, 1992) has found that insurance and labor markets within villages in developing countries are surprisingly efficient.[24]

In contrast to this older literature, Adamopoulos and Restuccia (2014) find that developing countries could double labor productivity in agriculture by reallocating land. There are several possible explanations for the difference. One is that Thai farmers suffer fewer distortions than those in most countries. Thailand is the largest exporter of rice in the world. Though its success may in part rely on its geographic advantages or the technical sophistication of its farmers, it may also have unusually efficient markets. Average landholdings, which Adamopoulos and Restuccia (2014) take as a measure of greater consolidation and thus greater efficiency, are higher in Thailand (3.2 hectares) than the median Asian (1.7 hectares) or African (2.7 hectares) country. This may in part

[24]However, Midrigan and Xu (2013) find levels of misallocation in manufacturing similar to what I find in rural Thailand.

be because what makes Thailand an ideal setting for accurately estimating the production function—the homogeneity of its methods for producing rice—also makes it easier for markets to allocate inputs efficiently.

Another possible explanation is that whereas Adamopoulos and Restuccia (2014) measure misallocation assuming inputs may be reallocated across the entire country, I allow reallocation only within a village. Doing so minimizes the risk of measuring misallocation caused by adjustment costs as opposed to market failures. But if there are large potential gains from reallocating inputs across villages, these gains will be missed. Though I show in Section 8.2 that allowing reallocation across villages within a sub-district does not much increase the gains from reallocation, it is possible that much of the potential gain lies in reallocation across provinces.

Finally, Adamopoulos and Restuccia (2014) use a different method for measuring misallocation. For example, they measure the misallocation predicted by a model calibrated with aggregate statistics, and they abstract from modeling labor as an input.[25] Without further research it is difficult to assess whether differences in context, differences in the scope of reallocation, or differences in methods can best reconcile the difference in results.

[25] On the first point: they model distortions not as explicit market imperfections but as an implied "tax" on farm revenue that is an increasing function of the farmer's productivity. They calibrate this function using aggregate cross-country data on crop-level distortions rather than farm-level data. On the second point: Labor markets may help alleviate misallocation in other markets. If labor markets are more functional than markets for land and capital, productive farmers who are unable to get enough land or capital might compensate by hiring more labor. There are other differences in their assumptions about the production function, but these are unlikely to explain the difference in results. They assume the production function is a CES combination of capital and land-augmenting productivity. However, they calibrate the elasticity of substitution to be 1.36, which is not far from Cobb-Douglas. I also give evidence in Appendix A.5.2 suggesting that the Cobb-Douglas assumption is a reasonable approximation of reality. The other difference is that they calibrate the parameters of the production function to match farming in the U.S. If farming in the U.S. had higher returns to scale it might cause similar market imperfections to have larger aggregate costs. But they calibrate the returns to scale at 0.54, which is somewhat lower than the returns to scale I find in Table 3.

7 The Effect of Credit on Misallocation: The Million Baht Program

Between May 2001 and May 2002 the Thai government gave one million baht of capital to the public lending fund of every village in my sample. The aptly named Million Baht Program in effect gave smaller villages more credit per-household. Kaboski and Townsend (2011) explain that village boundaries have little economic meaning and come from a bureaucratic tangle with statistically random outcomes. Since village sizes are random, the per-household rise in credit is also random. Kaboski and Townsend verify there are no differential trends between the villages that received more or less credit (see Table I on p. 1369 of their paper).

I exploit the program to test my measures of misallocation. By increasing the supply of credit the program improved financial markets, which should decrease misallocation from financial market imperfections.[26] I regress a village's misallocation in each year on year dummies, village fixed-effects, the log of the per-household credit injection (one million divided by the number of households), and the interaction between the log credit injection and 2001, the year of implementation, and 2002, the year after. The standard errors are twoway-clustered within villages and across province-years. The coefficients on the interactions measure the semi-elasticities of misallocation with respect to credit.

Table 9 reports the results, which are rescaled to show the change in the dependent variable caused by a 1 percent increase in per household credit. In

[26]The program might increase misallocation if the village funds lent out credit unfairly. But as Kaboski and Townsend (2011) explain, villagers elected panels of managers to administer the funds. The decisions were transparent and the main criterion was whether the managers thought the borrower could repay the loan.

Table 9
Effects of the Million Baht Credit Intervention

	(1) Overall Cost	(2) Cost, Fact.	(3) Cost, Fin.	(4) Agg. K	(5) Agg. T	(6) Agg. L
Log of Credit	0.047	-0.001	0.048	2.043	0.190***	0.528
	(0.039)	(0.008)	(0.036)	(474.195)	(0.058)	(0.503)
2001 X Credit	-0.090**	-0.003	-0.088**	-389.571	0.117	-0.334
	(0.039)	(0.006)	(0.038)	(541.579)	(0.099)	(1.389)
2002 X Credit	-0.015	0.006	-0.021	-215.139	-0.029	-0.140
	(0.044)	(0.012)	(0.036)	(426.952)	(0.098)	(1.073)
Year FEs	X	X	X	X	X	X
Village FEs	X	X	X	X	X	X
Villages	64	64	64	64	64	64
Province-Years	57	57	57	57	57	57
Observations	734	734	734	734	734	734

Note: The table shows the causal effect of the credit intervention on misallocation. Log of Credit is the log of the per-household credit injection ($\frac{one\ million}{\#\ of\ households}$), and the program was implemented in 2001 and 2002. The coefficient on the interactions gives the percentage point decrease in misallocation caused by a 1 percent increase in per household credit availability. A 1 percent increase in credit reduces overall misallocation by .09 percentage points, and almost all of the effect came from a reduction in misallocation from financial market failures (third column) rather than factor market failures (second column). The program had no effect on (per household) aggregate factor stocks.

each column the coefficient on 2001 X Credit gives the effect of the intervention in its first year. Column 1 shows that a 1 percent increase in credit decreases overall misallocation by 0.1 percent of observed output. Recall that overall misallocation is the sum of factor market and financial market misallocation. By measuring the effect of the intervention on each of these two components, I can decompose the reduction in overall misallocation into how much is caused by reductions in factor market versus financial market misallocation. Columns 2 and 3 show that there is little or no reduction in factor market misallocation; nearly all of the reduction in overall misallocation is caused by a reduction in financial market misallocation. Since a credit intervention should primarily affect financial markets, the results suggest the decomposition of factor versus financial market misallocation is behaving as it should.

The program had no significant effect on aggregate land, labor, or capital, confirming Townsend and Kaboski's (2009; 2011) finding that the program did not affect average investment. Since aggregate stocks did not change but misallocation fell, it suggests most farmers cut back their use of inputs while the most productive farmers scaled up. Kaboski and Townsend's structural model similarly showed the program did not affect all households equally.

Though statistically significant, the effect is small. Even a fifty percent increase in credit would only reduce misallocation by five percentage points. With so little misallocation at baseline, the result is not surprising. Since financial markets cause little misallocation, improving them will not produce spectacular results. But whatever the program's effects, they seem to have faded by its second year. The interaction of average credit injected and the year after implementation (2002) is a third the size and insignificant. One possibility is that households changed how they used their credit in the second year of the program. But since the estimated program effect is noisy, this interpretation should be taken with caution. I cannot reject that the impact in 2002 is identical to the impact in 2001.

8 Robustness and Alternative Specifications

8.1 Perfect Foresight

Even under perfect markets there will be misallocation because farmers make decisions based on anticipated productivity rather than realized productivity. I avoid assigning this ex post misallocation to factor and financial market imperfections by making assumptions in Section 5.2 about how much productivity the farmer anticipates. But if I underestimate how much the farmer knows about productivity I will also underestimate the amount of misallocation.

I bound the resulting bias by recalculating misallocation assuming, as is common in the literature, that farmers anticipate all productivity. I calculate this allocation by subbing $A_{it}\phi_{it}$ in for A_{it} in expression (5). This allocation is unrealistically perfect because farmers do not have perfect foresight. Thus it serves as an upper-bound on true misallocation.

Figure 4.A compares the upper bound to my preferred specification. Misallocation is higher, but in all years after 1996 it still falls short of the numbers reported by Hsieh and Klenow (2009) for Indian manufacturing, much less those that Adamopoulos and Restuccia (2014) report for developing country agriculture.

Figure 4

Panel A: Sample-Wide Misallocation When All Productivity is Anticipated;
Panel B: Sample-Wide Misallocation with Reallocation Within Sub-districts

Note: Panel A plots misallocation using the breakdown between anticipated and unanticipated productivity in the main text next to misallocation assuming all productivity is anticipated. Panel B plots misallocation assuming reallocation within villages (as in the main text) next to misallocation with reallocation within sub-districts (tambons), most of which contain four villages.

8.2 Small Samples and Between-Village Misallocation

Misallocation often happens because the most talented producers do not get enough inputs. If talented farmers are rare they might not appear in my small per-village samples, biasing down my estimates of misallocation. Even if the social planner favors reallocating everything to a small productive elite, I show in Section 5.3 that the dispersion of productivity in most villages is similar to what Gandhi, Navarro, and Rivers (2013) find in manufacturing industries. This suggests the dispersion of productivity in my sample is no lower than in other contexts.

Nevertheless, I test for whether small sample sizes are driving the results by regressing village-level misallocation on the number of farmers sampled from the village. Table 10 shows that estimated efficiency E_I is no lower in villages

Table 10
Correlation Between Village Misallocation and Sample Size

	(1) Efficiency b/se	(2) Gains b/se
Number of Farmers	0.021	-0.445*
	(0.17)	(0.23)
Villages	65	65
Observations	735	735

Note: I regress efficiency and gains from reallocation with the village on the number of farmers in the village in my sample. Efficiency is reported as a percentage of efficient output. Gains are reported as a percentage of observed output. I exclude villages with only one farmer (they have zero misallocation by construction). I find no evidence that small samples bias me towards finding too little misallocation. I cluster standard errors at the village-level.

with larger samples, and equivalently the gains from reallocation G_I are no higher in villages with larger samples (if anything, they are lower). The results give no reason to suggest I would find more misallocation if my sample were larger.

In Figure 4.B I recalculate sample-wide misallocation allowing reallocation between villages within a sub-district. The procedure does not differ much from within-village reallocation except I must account for the sample design (see Appendix A.2). Since sub-districts contain several villages they have larger samples of farmers. Reallocating within sub-district also increases the potential gains because the social planner can reallocate between as well as within villages. Figure 4.B shows the combined effect of both forms of reallocation is small. Reallocating between villages does not much increase output. In summary, sample size does not seem to be a serious source of bias, and between-village misallocation is small relative to within-village misallocation.

9 Conclusion

I derive a method to measure and separate the misallocation caused by factor and financial market failures. I find that in rural Thailand neither market failure causes much misallocation.

These resuls should not be taken to mean market failures are generally unimportant or can never cause misallocation. Misallocation may be a bigger problem in the farm sectors of countries outside of Thailand, or in sectors outside of farming. However, the results suggest imperfect markets need not cause severe misallocation. The method proposed in this paper is a tool that may allow future work to directly measure the level of misallocation. Only by measuring the misallocation caused by each market failure can we know when and where misallocation really matters.

Acknowledgements: Special thanks to Raj Arunachalam, Chris House, David Lam, Dan Ackerberg, Jagadeesh Sivadasan, Rebecca Thornton, Jeff Smith, Dmitriy Stolyarov, Matthew Shapiro, Miles Kimball, Penelope Goldberg, Isaac Sorkin, Eric Lewis, Ryotaro Tashiro, Jochen Gunter, Lasse Bruno, Adam Larson, and seminar participants for valuable feedback. I thank Doug Gollin and two anonymous referees for helpful suggestions. I also acknowledge generous financial support from the Thai Studies Grant Program at the Center for Southeast Asian Studies at the University of Michgian. Thanks as well to Robert Townsend and the staff of the Townsend Thai Project for data access.

References

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer. 2006. "Structural Identification of Production Functions."

Adamopoulos, Tasso and Diego Restuccia. 2014. "The Size Distribution of Farms and International Productivity Differences." *The American Economic Review* 104 (6):1667–1697.

Alfaro, Laura and Anusha Chari. 2012. "Deregulation, Misallocation, and Size: Evidence from India."

Alfaro, Laura, Andrew Charlton, and Fabio Kanczuk. 2008. "Plant-Size Distribution and Cross-Country Income Differences."

Anderson, T.W. and C. Hsiao. 1981. “Estimation of Dynamic Models with Error Components.” *Journal of the American Statistical Association* :598–606.

———. 1982. “Formulation and Estimation of Dynamic Models Using Panel Data.” *Journal of econometrics* 18 (1):47–82.

Baker, Chris and Pasuk Phongpaichit. 2009. *A History of Thailand*. Cambridge University Press.

Banerjee, A. and K. Munshi. 2004. “How Efficiently Is Capital Allocated? Evidence from the Knitted Garment Industry in Tirupur.” *Review of Economic Studies* 71 (1):19–42.

Banerjee, A.V. and B. Moll. 2010. “Why Does Misallocation Persist?” *American Economic Journal: Macroeconomics* 2 (1):189–206.

Bartelsman, E.J., J.C. Haltiwanger, and S. Scarpetta. 2009. “Cross-Country Differences in Productivity: The Role of Allocation and Selection.”

Benjamin, D. 1992. “Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models.” *Econometrica* 60 (2):287–322.

Bollard, Albert, Peter J Klenow, and Gunjan Sharma. 2012. “India’s Mysterious Manufacturing Miracle.” *Review of Economic Dynamics* .

David, Joel M, Hugo A Hopenhayn, and Venky Venkateswaran. 2013. “Information, Misallocation and Aggregate Productivity.”

De Janvry, Alain, Kyle Emerick, Marco Gonzalez-Navarro, and Elisabeth Sadoulet. 2015. “Delinking land rights from land use: Certification and migration in Mexico.” *The American Economic Review* 105 (10):3125–3149.

De Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. “Returns to Capital in Microenterprises: Evidence From a Field Experiment.” *The Quarterly Journal of Economics* 123 (4):1329–1372.

Gandhi, Amit, Salvador Navarro, and David Rivers. 2013. “On the Identification of Production Functions: How Heterogeneous Is Productivity?”

Gollin, Douglas, David Lagakos, and Michael E Waugh. 2014a. “Agricultural Productivity Differences Across Countries.” *The American Economic Review* 104 (5):165–170.

———. 2014b. “The Agricultural Productivity Gap.” *Quarterly Journal of Economics* 129 (2).

Gollin, Douglas, Stephen Parente, and Richard Rogerson. 2002. “The Role of Agriculture in Development.” *The American Economic Review* 92 (2):160–164.

Gollin, Douglas, Stephen L Parente, and Richard Rogerson. 2004. “Farm Work, Home Work and International Productivity Differences.” *Review of Economic Dynamics* 7 (4):827–850.

———. 2007. “The Food Problem and the Evolution of International Income Levels.” *Journal of Monetary Economics* 54 (4):1230–1255.

Greif, Avner. 1993. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition.” *The American Economic Review* :525–548.

House, Christopher L and Matthew D Shapiro. 2008. “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation.” *The American Economic Review* :737–768.

Hsieh, C.T. and P.J. Klenow. 2009. “Misallocation and Manufacturing TFP in China and India.” *Quarterly Journal of Economics* 124 (4):1403–1448.

Jensen, R.T. and N.H. Miller. 2010. “A Revealed Preference Approach to Measuring Hunger and Undernutrition.” Tech. rep., National Bureau of Economic Research.

Jones, Charles I. 2011. “Intermediate Goods and Weak Links in the Theory of Economic Development.” *American Economic Journal: Macroeconomics* 3 (2):1–28.

Kaboski, J. and R. Townsend. 2009. “The Impacts of Credit on Village Economies.” *MIT Department of Economics Working Paper No. 09-13* .

Kaboski, J.P. and R.M. Townsend. 2011. “A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative.” *Econometrica* 79 (5):1357–1406.

Keniston, D. 2011. “Experimental vs. Structural Estimates of the Return to Capital in Microenterprises.” Tech. rep., mimeo Yale.

Lagakos, David and Michael E Waugh. 2011. “Selection, Agriculture and Cross-Country Productivity Differences.” *Agriculture and Cross-Country Productivity Differences (November 1, 2011)* .

Levinsohn, J. and A. Petrin. 2003. “Estimating Production Functions Using Inputs to Control for Unobservables.” *Review of economic studies* 70 (2):317–341.

Midrigan, V. and D. Xu. 2013. “Finance and Misallocation: Evidence from Plant-Level Data.” *American Economic Review* .

Moll, B. 2010. “Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?” *University of Chicago mimeo* .

Osotimehin, Sophie. 2011. “Aggregate Productivity and the Allocation of Resources over the Business Cycle.”

Peters, Michael. 2011. “Heterogeneous Mark-Ups and Endogenous Misallocation.” *Manuscript, Massachusetts Institute of Technology* .

Petrin, Amil and Jagadeesh Sivadasan. 2013. “Estimating Lost Output from Allocative Inefficiency, with an Application to Chile and Firing Costs.” *Review of Economics and Statistics* 95 (1):286–301.

Restuccia, D. and R. Rogerson. 2008. “Policy Distortions and Aggregate Productivity with Heterogeneous Establishments.” *Review of Economic Dynamics* 11 (4):707–720.

Restuccia, Diego and Raul Santaeulalia-Llopis. 2014. “Land Misallocation and Productivity.”

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu. 2008. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics* 55 (2):234–250.

Romer, Christina D. 1986. "Is the Stabilization of the Postwar Economy a Figment of the Data?" *The American Economic Review* 76 (3):314–334.

Sandleris, Guido and Mark LJ Wright. 2014. "The Costs of Financial Crises: Resource Misallocation, Productivity, and Welfare in the 2001 Argentine Crisis." *The Scandinavian Journal of Economics* 116 (1):87–127.

Shenoy, Ajay. 2014. "Risky Income or Lumpy Investments? Testing Two Theories of Under-Specialization."

Singh, Inderjit, Lyn Squire, John Strauss et al. 1986. *Agricultural Household Models: Extensions, Applications, and Policy.* Johns Hopkins University Press.

Stock, James H, Jonathan H Wright, and Motohiro Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business & Economic Statistics* 20 (4).

Sundaram, R.K. 1996. *A First Course in Optimization Theory.* Cambridge Univ Pr.

Townsend, R., A. Paulson, S. Sakuntasathien, T. J. Lee, and M. Binford. 1997. "Questionnaire Design and Data Collection for NICHD Grant *Risk, Insurance and the Family* and NSF Grants." URL http://hdl.handle.net/1902.1/10672.

Townsend, R.M. 1994. "Risk and Insurance in Village India." *Econometrica: Journal of the Econometric Society* :539–591.

Townsend, Robert M. 2012. "Townsend Thai Project Monthly Survey (1-60) Initial Release." URL http://hdl.handle.net/1902.1/14795. Murray Research Archive [Distributor] V1 [Version].

Udry, C. 1994. "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *The Review of Economic Studies* 61 (3):495–526.

———. 1996. "Gender, Agricultural Production, and the Theory of the Household." *Journal of Political Economy* 104 (5):1010–1046.

Wooldridge, J.M. 2002. *Econometric Analysis of Cross Section and Panel Data.* The MIT press.

A.1 Proof of the Bounding Condition

I prove that my measure of factor market misallocation is a lower bound on the truth. Call the land-capital and labor-capital ratios derived with the endowment assumption $\tau = \frac{T_i}{K_i}, \Lambda = \frac{L_i}{K_i}$, and define the ratios without the assumption $\tilde{\tau}, \tilde{\Lambda}$ similarly. According to (3) they must be identical for all farmers. Consider the market-clearing condition for land with the endowment assumption:

$$\begin{aligned} \sum T_i^+ &= T_I \\ \Rightarrow \sum K_i^+ \tau &= T_I \\ \Rightarrow \tau K_I &= T_I \\ \Rightarrow \tau &= T_I / K_I \end{aligned}$$

Identical reasoning shows $\tilde{\tau} = T_I/K_I$ as well, so $\tau = \tilde{\tau}$ and similarly $\Lambda = \tilde{\Lambda}$

The difference between aggregate output with and without the endowment condition is

$$\begin{aligned}
\tilde{Y}_I^+ - Y_I^+ &= \sum A_i\phi_i(\tilde{K}_i^+)^{\theta_K}(\tilde{T}_i^+)^{\theta_T}(\tilde{L}_i^+)^{\theta_L} - \sum A_i\phi_i(K_i^+)^{\theta_K}(T_i^+)^{\theta_T}(L_i^+)^{\theta_L} \\
&= \sum A_i\phi_i(\tilde{K}_i^+)^{\sigma}\tau^{\theta_T}\Lambda^{\theta_L} - \sum A_i\phi_i(K_i^+)^{\sigma}\tau^{\theta_T}\Lambda^{\theta_L} \\
&= \tau^{\theta_T}\Lambda^{\theta_L}\sum A_i\phi_i\left[(\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma}\right] \\
\Rightarrow \frac{1}{Y_I}\left[(\tilde{Y}_I^+ - Y_I) - (Y_I^+ - Y_I)\right] &= \frac{1}{Y_I}\tau^{\theta_T}\Lambda^{\theta_L}\sum A_i\phi_i\left[(\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma}\right] \\
\Rightarrow \tilde{G}_I^{FACT} - G_I^{FACT} &= \frac{\tau^{\theta_T}\Lambda^{\theta_L}}{Y_I}\sum A_i\phi_i\left[(\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma}\right] \\
\Rightarrow \mathbb{E}_i[\mathbb{E}_\phi[\tilde{G}_I^{FACT} - G_I^{FACT}]] &= \frac{\tau^{\theta_T}\Lambda^{\theta_L}|I|}{Y_I}\mathbb{E}_i\left(A_i\left[(\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma}\right]\right)
\end{aligned}$$

where the last step follows because unanticipated productivity ϕ is independent and mean 1, and the size of village I is $|I|$. Since $\frac{\tau^{\theta_T}\Lambda^{\theta_L}|I|}{Y_I} > 0$, $\mathbb{E}_i[\mathbb{E}_\phi[\tilde{G}_I^{FACT} - G_I^{FACT}]] > 0$ if and only if $\mathbb{E}_i\left(A_i\left[(\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma}\right]\right) > 0$. In words, G_I^{FACT} underestimates the gains from perfect factor markets. Since $G_I = G_I^{FACT} + G_I^{FIN}$ it must be that G_I^{FIN} overestimates the subsequent gains from financial market perfection. ■

A.2 Robustness: Subdistrict-Level Reallocation

This appendix explains how to calculate misallocation at the sub-district level. The firm's problem is the same as before, and its optimal capital choice is $K_i^* = \eta A_i^{\frac{1}{1-\sigma}}$ where $\eta = \left[\left(\frac{\theta_K}{w^K}\right)^{1-\theta_T-\theta_L}\left(\frac{\theta_T}{w^T}\right)^{\theta_T}\left(\frac{\theta_L}{w^L}\right)^{\theta_L}\right]^{\frac{1}{1-\sigma}}$. Replace the village-level market-clearing condition with village and subdistrict-level conditions:

$$\sum_{i\in I} K_i^* = K_I^* \qquad \forall I$$
$$\sum_I W_I K_I^* = \sum_I W_I \bar{K}_I$$

where K_I^* is the amount of capital the village is optimally allocated, $\bar{K}_I$ is the village's initial allocation, and W_I is an inverse-probability weight (total number of households in the village divided by the number of households sampled).

Then

$$\eta \sum_I W_I \sum_{i \in I} A_i^{\frac{1}{1-\sigma}} = \sum_I W_I \bar{K}_I$$
$$\Rightarrow \eta = \frac{\sum_I W_I \sum_I w_I \bar{K}_I}{\sum_I W_I \sum_{i \in I} A_i^{\frac{1}{1-\sigma}}}$$

Sub this back into the individual demand:

$$K_i^* = \frac{A_i^{\frac{1}{1-\sigma}}}{\sum_I W_I \sum_{i \in I} A_i^{\frac{1}{1-\sigma}}} \sum_I W_I \bar{K}_I$$

Let y_i^* be output with the optimal capital, land, and labor. Optimal aggregate output is

$$Y^* = \sum_I W_I \sum_{i \in I} y_i^*$$

A.3 More General Model (For Online Publication)

In this appendix I show that many of the simlifying assumptions made in Section 2 can be dropped without changing the result.

Models of farm production must deal with the problem of non-separability, meaning the household as a consumer cannot be treated separately from the household as a producer. The constraints and imperfections that cause misallocation also break separability (Singh et al., 1986; Benjamin, 1992). Thus the household's optimization is a nearly intractable dynamic problem.

But I show that this dynamic problem need not be solved. Assuming the farmer's observed choices do solve the problem, I need only derive whatever choices the farmer would have made under perfect factor and financial markets. I show that with perfect markets the dynamic problem collapses to a simple static problem. From this I derive the optimal allocations. Most importantly I show that when my assumptions are met, perfect factor markets eliminate all distortions to the farmer's mix.

A.3.1 Environment

In every year t the farmer aims to maximize her discounted lifetime utility from consumption, given a per-period utility function u that depends on a vector $\boldsymbol{\gamma_i}$ that captures her individual preferences (most importantly, risk-aversion). She solves

$$\text{Maximize} \qquad \mathbb{E}[\sum_{j=0}^{\infty} \rho^j u(c_{i,t+j}; \boldsymbol{\gamma_i})]$$

To earn income she uses capital K, land T, and labor L to produce farm revenue y. Her output also depends on Hicks-Neutral productivity, part of which (A) she anticipates when choosing factors while the rest (ϕ) is random and unanticipated. I normalize $\mathbb{E}[\phi] = 1$. Her revenue is

$$y_{it} = A_{it}\phi_{it}K_{it}^{\theta_K}T_{it}^{\theta_T}L_{it}^{\theta_L}$$

Though assuming Cobb-Douglas production is a simplification, I show in Appendix A.5.2 that it is not a bad approximation. Let $\boldsymbol{X}$ be a vector that contains the land, labor, and capital used in production. They may come from her stock of owned factors $\boldsymbol{X}^o$ or those she rents from factor markets $\boldsymbol{X} - \boldsymbol{X}^o$. She may buy factors $\boldsymbol{I}$, and owned factors depreciate at rate $\boldsymbol{\delta}$. The law of motion for owned assets is

$$\boldsymbol{X}_{it}^o = \boldsymbol{I}_{it} + (1 - \boldsymbol{\delta}) \otimes \boldsymbol{X}_{i,t-1}^o$$

where $\otimes$ is the element-wise product. To make the notation simple I assume no time to build, meaning investment is immediately productive, but all I require is that as soon as the investment does become productive it can instantly and costlessly be transfered to a different farmer.

The vector of "owned" factors $\boldsymbol{X}^o$ includes family labor, which I assume exogenous for notational simplicity (making labor decisions endogenous changes nothing, as the production outcome depends only on labor employed on the farm, not labor supplied). I assume family labor and hired labor are perfect substitutes, though relaxing the assumption does not change the main results (see Appendix A.5.3).

Thus far the problem is entirely standard and leads to an efficient production outcome. But now I introduce market imperfections and constraints that distort the outcome. Farmer i faces a set of rental prices $\boldsymbol{w}$ that may differ from those others pay. This is the first factor market imperfection. Together with the assets she buys at prices $\boldsymbol{p}$, her total farm expenditure is

$$z_{it} = \boldsymbol{w}_{it} \cdot (\boldsymbol{X}_{it} - \boldsymbol{X}_{it}^o) + \boldsymbol{p} \cdot \boldsymbol{I}_{it}$$

In every period the farmer must meet her budget constraint. Any farm expenditure beyond gross interest on her savings from last year $R^b b_{it}$ is borrowed at gross rate R^z and repaid after the harvest. I allow the interest rate paid to vary with the amount of land owned, as the bank might offer lower rates to those who can offer collateral. (It changes nothing to let the farmer offer capital or other assets as collateral.) The constraint is

$$\lambda : \quad c_{it} + b_{i,t+1} = y_{it} - R_{it}^z(T_{it}^o)(z_{it} - R_{it}^b b_{it})$$

where λ is the Lagrange multiplier. The gap between the interest rate on borrowing versus savings is the first financial market imperfection. A second imperfection, which is implicit, is that since the farmer lacks perfect insurance her consumption may be correlated with the unanticipated shock ϕ. The effect on her decisions depends on her risk aversion.

The third financial market imperfection is a flow-of-funds constraint. The farmer must buy and pay rents to all factors before planting. To finance these payments she may borrow beyond her savings, but only up to a limit $\bar{z}$. It too may depend on how much land she owns and may differ across farmers, perhaps because some farmers have more cosigners or rich relatives. Also, some fraction of her purchases of factors $\boldsymbol{I}$ may not count one-for-one towards the constraint because the bank is more willing to finance hard assets it can seize if she defaults. Let the vector $0 < \boldsymbol{\zeta} \leq 1$ capture what fraction of each dollar is discounted for each asset. The constraint is

$$\omega : \quad z_{it} - R^b_{it} b_{it} - \boldsymbol{p} \cdot (\boldsymbol{\zeta} \otimes \boldsymbol{I}_{it}) \leq \bar{z}_{it}(T^o_{it})$$

where ω is the Lagrange multiplier.

Finally, there may be missing or limited rental markets for land, labor, and capital. For example, if property rights are weak a farmer may refuse to rent out land for fear that the renter will squat. The farmer cannot rent in more factors than $\overline{\boldsymbol{X}}$ or rent out more factors than $\underline{\boldsymbol{X}}$. The constraint, together with Lagrange multipliers, is

$$\underline{\kappa}, \overline{\kappa} : \quad \underline{\boldsymbol{X}}_{it} \leq \boldsymbol{X}_{it} - \boldsymbol{X}^o_{it} \leq \overline{\boldsymbol{X}}_{it}$$

The timing is as follows:

1. The farmer learns anticipated productivity A_{it}
2. The farmer buys factors $\boldsymbol{I}$ and rents factors $\boldsymbol{X}_{it} - \boldsymbol{X}^o_{it}$, borrowing if necessary. Any purchased assets can immediately be used in production or rented out[27]
3. Uncertainty is resolved and production completed

[27] In fact, all I require is that as soon as the asset is usable it is costless and takes no adjustment period to allow someone else to use it. Clearly building a tractor or a granary takes time, but since my optimal allocations simply move factors between farmers the real assumption is that moving a tractor or renting space in a granary takes little time.

4. The farmer pays off her loans and makes consumption and savings decisions

This timing effectively assumes there is no cost in time or money to transferring factors between farmers, and all factors are chosen with equal information about productivity.[28]

A.3.2 Perfect Choices and Distortions

The farmer's optimal choice of capital satisfies

$$\mathbb{E}[\lambda_{it}\phi_{it}]\theta_K A_{it} K_{it}^{-(1-\theta_K)} T_{it}^{\theta_T} L_{it}^{\theta_L} - (R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] + \omega_{it}) w_{it}^K + \underline{\kappa} - \overline{\kappa} = 0$$

Optimal land and labor choices satisfy similar conditions. Note that ζ, which captures how some assets are easier to collateralize than others, does not appear. This is because buying a tractor and using that tractor are not the same. What matters for her production decision is not the price she paid for the tractor or the ease with which she borrowed funds to buy it. All that matters is the opportunity cost of farming it herself, which is simply the rent w_{it}^K she earns by letting someone else farm it.

Suppose markets are fully perfect. With perfect insurance, the unanticipated shock does not affect the shadow value of consumption ($\mathbb{E}[\lambda_{it}\phi_{it}] = \mathbb{E}[\lambda_{it}]\mathbb{E}[\phi_{it}]$). With perfect credit markets, farmers pay the same borrowing rate ($R_{it}^z[T_{it}^o] = R_t^z$) and the liquidity constraint does not bind ($\omega_{it} = 0$). With perfect factor markets, farmers pay the same rental prices ($w_{it}^K = w_t^K$) and factor market constraints do not bind ($\underline{\kappa}_{it} = \overline{\kappa}_{it} = 0$). Then

$$\begin{aligned} \theta_K A_{it} K_{it}^{-(1-\theta_K)} T_{it}^{\theta_T} L_{it}^{\theta_L} &= R_t^z w_t^K \\ &= \theta_K A_{jt} K_{jt}^{-(1-\theta_K)} T_{jt}^{\theta_T} L_{jt}^{\theta_L} \quad \forall j \end{aligned} \tag{11}$$

The expression implies marginal products are equalized across all farmers

[28] If farmers are risk neutral the information assumption can be relaxed to require equal information about only idiosyncratic productivity.

in the village. Moreover, the optimal choice depends only on current variables, meaning the choice that solves the static problem solves the dynamic problem.

Now suppose factor markets are perfect, so $w_{it}^K = w_t^K, \underline{\kappa}_{it} = \overline{\kappa}_{it} = 0$, but financial markets are not. Then the optimal choices of capital and land satisfy

$$\begin{aligned} \mathbb{E}[\lambda_{it}\phi_{it}]\frac{\mathbb{E}[y_{it}]}{K_{it}} &= \frac{1}{\theta_K}(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] + \omega_{it})w_t^K \\ \mathbb{E}[\lambda_{it}\phi_{it}]\frac{\mathbb{E}[y_{it}]}{T_{it}} &= \frac{1}{\theta_T}(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] + \omega_{it})w_t^T. \end{aligned}$$

Divide the capital condition by the land condition:

$$\begin{aligned} \frac{T_{it}}{K_{it}} &= \frac{\theta_T}{\theta_K}\frac{w_t^K}{w_t^T} \\ &= \frac{T_{jt}}{K_{jt}} \quad \forall j \end{aligned} \tag{12}$$

The condition implies capital-land ratios are equalized across farmers throughout the village, and by similar logic the land-labor ratios are as well. In short, if factor markets are perfect the farmer can choose the right mix of inputs even if financial markets are imperfect, and the right mix again is the static optimum.

One can rewrite (12) as
$$\frac{T_{it}}{K_{it}} = \frac{1+\tau_j}{1+\tau_i}\frac{T_{jt}}{K_{jt}}$$
where $\tau_i = \tau_j = 0$ if factor markets are perfect.[29] The τ term matches up to the "capital distortion" of Hsieh and Klenow (2009), but in my model it vanishes when factor markets are perfect. Thus, I have given the distortion an economic interpretation.

A.3.3 Optimal Allocations

How would perfect markets allocate the aggregate factor stock? Since perfect markets make the solution to the farmer's dynamic optimization equal the period-by-period static optimum, I can suppress time subscripts.

Suppose i is a farmer in village I observed to use $\bar{K}_i, \bar{T}_i, \bar{L}_i$. The observed

[29] $1+\tau_i = \frac{(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] - \omega_{it})w_{it}^K + \underline{\kappa} - \overline{\kappa}}{(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] - \omega_{it})w_{it}^T + \underline{\kappa} - \overline{\kappa}}$.

factor stocks are $K_I = \sum_{i \in I} \bar{K}_i$ and so on, and they do not change because I only reallocate the village's existing resources. (In Section 8 I consider reallocating resources between villages as well.) With aggregate stocks pinned down, I can ignore the supply side of the market and normalize $R^z = 1$. Use (12) to eliminate T_i and L_i from (11) and define the production returns to scale $\sigma = \theta_K + \theta_T + \theta_L$, which I assume is less than one. Combine with the market clearing condition $K_I = \sum_{j \in I} K_j^*$ and solve for the optimal allocations with fully perfect markets:

$$K_i^* = \frac{A_i^{\frac{1}{1-\sigma}}}{\sum_{j \in I} A_j^{\frac{1}{1-\sigma}}} K_I \tag{13}$$

Optimal land and labor are similar. Call farmer i's output with perfect allocations $y_i^* = A_i \phi_i (K_i^*)^{\theta_K} (T_i^*)^{\theta_T} (L_i^*)^{\theta_L}$.

Now suppose factor markets are perfected but financial markets left untouched, which means farmers choose the optimal mix of factors. Equation 12 gives the optimal mix but not the overall scale for each farmer. Any assumption about scale would define the allocation, so I consider a hypothetical case in which the farmer takes her original choices—those I observe in the data—and trades them in the perfected factor markets as though they were endowments. Let K_i^+ be the farmer's new choice of capital while $\bar{K}$ is still her original choice. Then the value of her new choices must add up to the value of her endowment, as determined under the new prices w^{K+}, w^{T+}, w^{L+}:

$$w^{K+} K_i^+ + w^{T+} T_i^+ + w^{L+} L_i^+ = w^{K+} \bar{K}_i + w^{T+} \bar{T}_i + w^{L+} \bar{L}_i$$

I effectively drop the farmers into an Edgeworth economy where their original input choices are like endowments. The farmer choosing a profit-maximizing mix of factors behaves like a consumer choosing a utility-maximizing bundle of goods. The resulting allocation is easy to compute and perfects each farmer's mix while leaving her scale untouched. Each farmer's allocation may differ from what she would choose given perfect factor markets and no extra constraints on scale. But under an assumption I explain in Section 2.4, gains from moving to the computed allocation are a lower bound on true misallocation from imperfect factor markets.

Again taking $K_I = \sum_{j \in I} K_j^+$ as the market-clearing condition, the allocations under perfect factor markets are

$$K_i^+ = \frac{1}{\theta_K + \theta_T + \theta_L}\left[\theta_K \frac{\bar{K}_i}{K_I} + \theta_T \frac{\bar{T}_i}{T_I} + \theta_L \frac{\bar{L}_i}{L_I}\right] K_I \tag{14}$$

Optimal land and labor are similar. Call farmer i's output with perfect factor markets $y_i^+ = A_i \phi_i (K_i^+)^{\theta_K}(T_i^+)^{\theta_T}(L_i^+)^{\theta_L}$.

A.4 Unequal Information about Village-Level Shocks (For Online Publication)

This appendix shows that under additional assumptions the method still works when farmers have more information about aggregate productivity when choosing some inputs than others. I work from the more general model of Appendix A.3 rather than the simplified model in the main text.

Consider the choice of labor and capital. There is an aggregate shock ϕ_t^A that is unknown when the farmer chooses capital but known when the farmer chooses labor. The aggregate shock has mean 1 and is uncorrelated with idiosyncratic productivity ϕ_{it}, which remains unknown until all choices are made. Then the conditions for an optimal choice under perfect factor markets are

$$\mathbb{E}[\lambda_{it}\phi_{it}\phi_t^A]\frac{\mathbb{E}[y_{it}]}{K_{it}} = \frac{1}{\theta_K}(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] - \omega_{it})w_t^K$$
$$\mathbb{E}[\lambda_{it}\phi_{it}]\phi_t^A\frac{\mathbb{E}[y_{it}]}{L_{it}} = \frac{1}{\theta_L}(R_{it}^z(T_{it}^o)\mathbb{E}[\lambda_{it}] - \omega_{it})w_t^L.$$

Suppose first that the variance of ϕ_t^A is small. Then

$$\mathbb{E}[\lambda_{it}\phi_{it}\phi_t^A] \approx \mathbb{E}[\lambda_{it}\phi_{it}].$$

Alternatively, assume the variance of ϕ_t^A is arbitrarily large, but the farmer is risk neutral. Then $\lambda_{it} = u'(C_t; \boldsymbol{\gamma}_i) = 1$, and

$$\mathbb{E}[\lambda_{it}\phi_{it}\phi_t^A] = \mathbb{E}[\phi_{it}\phi_t^A] = 1.$$
$$\mathbb{E}[\lambda_{it}\phi_{it}]\phi_t^A = \mathbb{E}[\phi_{it}]\phi_t^A = \phi_t^A.$$

In either case, divide the condition for optimal land by the condition for optimal labor:

$$\begin{aligned}\frac{L_{it}}{K_{it}} &= \frac{\theta_L}{\theta_K}\frac{w_t^K}{\phi_t^A \cdot w_t^L} \\ &= \frac{\theta_L}{\theta_K}\frac{w_t^K}{\tilde{w}_t^L} \\ &= \frac{L_{jt}}{K_{jt}} \qquad \forall j\end{aligned}$$

where $\tilde{w}_t^L$ is a normalized price. Since ϕ_t^A does not vary within the village farmers still have equal capital-labor ratios. The effect of a positive aggregate shock is to proportionally raise or lower every farmer's demand for labor. The effect is absorbed into the price ratio; the relative price of labor rises such that the optimal capital-labor ratio is the same regardless of the shock. By a similar argument the optimal scale of each farmer also remains unchanged.

A.5 Empirical Appendix (For Online Publication)

This appendix runs tests whether rainfall shocks are well-specified, whether the production function is approximately isoelastic (Cobb-Douglas) and whether dropping the assumption of perfect substitution between family and hired labor changes the results.

A.5.1 Are Rainfall Shocks Well-Specified?

Given the limitations of the data I am unable to observe any continuous measure of rainfall shocks at the level of the household. Instead I measure rainfall shocks with a dummy for whether the household reports poor rainfall this year (see Appendix A.7). How accurate is this measure, and how well does it correspond with an objective, remotely sensed measure of aggregate rainfall?

I first construct a measure of aggregate rainfall. The Townsend Thai project has given me district identifiers, which allow me to link average rainfall in the district to survey reports of revenue from rice farming. I construct three measures of rainfall shocks. The first is simply total rainfall during the year ("Total Rain"). I then calculate mean rainfall across years for each month (for example, average rainfall in March across all years). Using this mean I compute the deviation from the mean for each month of each year as a fraction of the mean. The second measure of rainfall is the sum of these monthly deviations within a year ("Sum of Deviations"). The third measure computes total yearly rainfall and takes the deviation in yearly rainfall from the average ("Deviation of Sums"). This third measure is most similar to those used in the literature.

All of these measures would be perfectly colinear with the district-year dummies I used to estimate the production function. Moreover, district-level shocks would not affect within-village misallocation. However, the three measures are useful in validating the self-reported shock.

I first pick out the most meaningful of the three measures. I regress household-level revenue from rice farming on each measure. Since the measures do not vary within districts, I cluster standard errors by district. As the number of districts is small I bootstrap rather than relying on asymptotic approximations. Column 1 of Table A.I shows that total rainfall is not a useful predictor. Columns 2 and 3 show that both the sum of deviations and the deviation of sums have positive and statistically significant effects on revenue, but the deviation of sums is the better of the two. Column 4 shows that the measure becomes even more informative when I drop district-years that had extreme values (deviations greater than 0.2). Since the Deviation of Sums measure seems to be the most informative, I restrict my attention to it in the subsequent analysis.

The top panel of Figure A.I shows the correlation by dividing the range of the rainfall shock into bins and plotting the average revenue of households within each bin. Higher levels of rainfall cause higher revenue, and the relationship appears to be linear. How does the district-level measure match up with the household-level measure? The bottom panel of Figure A.I plots a similar graph where the dependent variable is now the self-reported indicator for poor rain-

Table A.I
Measures of Rainfall

	(1) Revenue	(2) Revenue	(3) Revenue	(4) Revenue
Total Rain	-0.00			
	(0.00)			
Sum of Deviations		0.05^{**}		
		(0.02)		
Deviation of Sum			0.47^{***}	0.71^{**}
			(0.15)	(0.32)
Observations	6228	6228	6228	5687
Districts	15	15	15	15
R-Squared	0.006	0.001	0.003	0.004

Note: Column 4 drops outliers in the distribution of rainfall shocks. Standard errors are computed by bootstrapping districts.

fall. The relationship here is even clearer. Fewer households report poor rainfall in districts that received more rainfall.

Columns 1 and 2 of Table A.II report the results of a regression of the self-reported shock on the district-level rainfall shock. The results suggest a near one-for-one relationship. A 10 percent increase above the mean in annual rainfall predicts a 10 percent decrease in the fraction of households reporting poor rainfall.

Does the household's self-reported measure capture all the information in the district-level shock? To answer this question I run a "horse race" in which I regress revenue on both the self-reported shock and the district dummy. Columns 3 and 4 show that the coefficient on the district-level shock becomes smaller and statistically insignificant when I control for the self-reported shock. By contrast, the self-reported shock is highly significant. Moreover, the R-squared in regressions controlling for the household-level dummy are an order of magnitude greater than those in Table A.I that control only for the district-level shock. Not only does the household dummy capture all of the information, but it seems to add quite a bit more. This is not surprising, as a household-level measure will inevitably be more accurate than a district-level measure.

Though the self-reported measure may be spatially precise, is it too coarse in measuring the size of a shock? The household shock is simply a dummy; it

Figure A.I

District-Level Rainfall Predicts Revenue and Self-Reported Shock

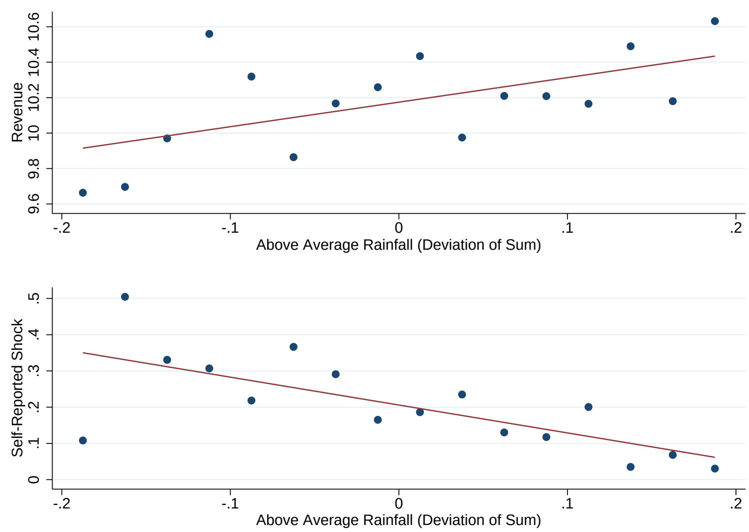

Note: The range of the rainfall measure is divided into equally sized bins. Each dot gives the mean of rainfall within a bin.

Table A.II

Does the Household-Level Dummy Capture the Content of the Continuous Measure?

	(1) Self-Reported Shock	(2) Self-Reported Shock	(3) Revenue	(4) Revenue
Deviation of Sum	-0.86***	-0.93***	0.10	0.30
	(0.12)	(0.19)	(0.17)	(0.26)
Self-Reported Shock			-0.43***	-0.44***
			(0.12)	(0.13)
Observations	6228	5687	6228	5687
Districts	15	15	15	15
R-Squared	0.059	0.045	0.031	0.032

Note: Column 4 drops outliers in the distribution of rainfall shocks. Standard errors are computed by bootstrapping districts.

cannot provide any information about whether a shock was merely bad versus catastrophic. One approach to handle this problem is to combine the household measure, which is precise in measuring which household was affected, with the district shock, which is precise in measuring how bad was the shock. I construct three measures that aim to combine the benefits of each. In all three cases I interact the self-reported shock with some function of the district-level deviation of sums. This interaction converts the purely extensive household shock into something intensive.

Since a household only reports a shock when rainfall is low, the first measure is $[Self]X[-Dev.]$, or the product of the dummy and the negative of the deviation. This measure is large and positive when the household reports a bad shock and district rainfall is low. The second measure combines the first measure with another dummy for whether the district-level deviation is negative: $[Self]X[-Dev.]X[I(Dev. < 0)]$. This measure is positive only when the household reports low rainfall and the level of rainfall in the district is below average. The third measure is simply $[Self]X[abs(Dev.)]$, the product of the self-reported shock and the absolute value of the district level measure.

I add each of these three measures to the regressions of revenue on household and district rainfall shocks. Table A.III shows that all of these measures are statistically significant, though they lose their significance in Columns 4, 6, and 8, where I drop outliers in the district-level shock. Moreover, given the way the measures are defined, their coefficients should be negative. Instead, all are positive, suggesting a problem of colinearity. These measures may not be adding useful information beyond that contained in the self-reported shock.

Nevertheless, I test whether my measures of misallocation change when I add the third measure (which seems the strongest of the three) to the set unanticipated shocks. Figure A.II is the analog of Figure 2.B in the main text. I recalculate misallocation and plot the gains from reallocation across all years. The figures show that the results are almost unchanged. Again, this is not surprising; since rainfall shocks are unanticipated, failing to control for them should not bias the estimates of the production function or change the optimal allocation.

Table A.III
Does Interacting District and Household-level Variables Improve the Measure?

	(1) Revenue	(2) Revenue	(3) Revenue	(4) Revenue	(5) Revenue	(6) Revenue	(7) Revenue	(8) Revenue
Self-Reported Shock	-0.44***	-0.43***	-0.46***	-0.45***	-0.51***	-0.47***	-0.55***	-0.49**
	(0.11)	(0.12)	(0.11)	(0.14)	(0.13)	(0.14)	(0.14)	(0.20)
Deviation of Sum		0.10	0.22	0.34	0.26	0.35	0.21	0.32
		(0.19)	(0.19)	(0.33)	(0.19)	(0.30)	(0.18)	(0.31)
$[Self]X[-Dev.]$			0.57*	0.23				
			(0.31)	(0.44)				
$[Self]X[-Dev.]X[I(Dev. < 0)]$					1.07**	0.49		
					(0.47)	(0.76)		
$[Self]X[abs(Dev.)]$							1.28***	0.61
							(0.47)	(1.03)
Observations	6228	6228	6228	5687	6228	5687	6228	5687
Districts	15	15	15	15	15	15	15	15
R-Squared	0.030	0.031	0.031	0.032	0.032	0.032	0.032	0.032

Note: Columns 4, 6, and 8 drop outliers in the distribution of rainfall shocks. Standard errors are computed by bootstrapping districts.

Figure A.II
Misallocation Using Interaction Measure 3

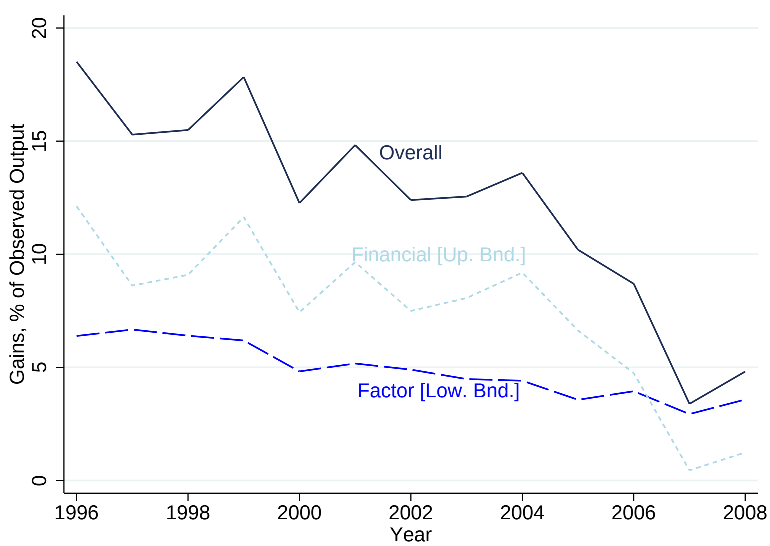

Note: I re-estimate misallocation assuming the third interaction measure is part of the unanticipated shock.

Table A.IV
CES Production Function Estimates

	(1)
	NLS
ϵ	1.01
	(0.01)
σ	0.75
	(0.04)
α	0.31
	(0.01)
β	0.38
	(0.01)
ϑ	0.00
	(0.00)
Observations	6230
Households	775
Pval: $\epsilon = 1$	0.40

Note: Estimated using fixed-effects nonlinear least-squares. Standard errors are bootstrapped with resampling at the household-level.

A.5.2 Is the Production Function Isoelastic?

The Cobb-Douglas production function is a special case of the class of constant-elasticity production functions where the elasticity of substitution ϵ between factors is 1 (hence alternative label "isoelastic"). I follow the procedure in Udry (1996) and suppose

$$y_{it} = A_{it}\phi_{it}\left[\alpha K_{it}^{\frac{\epsilon-1}{\epsilon}} + \beta T_{it}^{\frac{\epsilon-1}{\epsilon}} + (1-\alpha-\beta)L_{it}^{\frac{\epsilon-1}{\epsilon}}\right]^{\sigma\frac{\epsilon}{\epsilon-1}} \tag{15}$$

where σ denotes the returns to scale. For computational simplicity I assume $A_{it} = A_i e^{\vartheta t}$, the product of a fixed effect and a time trend. Take logs of both sides and subtract away the within household mean to eliminate the fixed-effect.

Column 1 of Table A.IV reports the results of estimating the transformed equation with nonlinear least squares. The test of interest is whether ϵ differs substantially from 1. As (1) of Table A.IV indicates, this null is actually rejected. However, the point estimate is almost identical to one ($\hat{\epsilon} = 1.01$) and rejection occurs mainly because the variance of these estimates is very small. The envelope theorem guarantees misallocation does not change much with small changes in the elasticity of substitution, so a tiny deviation from Cobb-Douglas production should not change the results much.

Of course, the ideal estimator is not fixed-effects but the nonlinear equivalent of Anderson-Hsiao: applying GMM to the first-differenced form of (15) using legged factors as instruments. Unfortunately this estimator does not converge, which may itself be a sign that the data reject the extra parameters.

A.5.3 Does Dropping the Assumption of Labor Substitutability Change the Results?

Suppose some imperfection makes hired labor less productive than family labor. For example, hired farm hands might shirk when the farmer is not watching. We can treat this monitoring problem as a factor market imperfection that would not exist with perfect factor markets. To be precise, suppose the production elasticity θ_L is the elasticity of family labor, meaning labor with no monitoring problem. Each unit of hired labor L^H is worth only $f(L^H) \leq 1$ units of family labor. I choose these functional forms to be concrete; they are not crucial to the argument. Then observed output is

$$y_{it} = K_{it}^{\theta_K} T_{it}^{\theta_T} [L_{it}^F + f(L_{it}^H)]^{\theta_L}.$$

Perfecting the market would raise output through two channels: the gains from making hired workers more productive, and the gains from reallocation. Since this paper aims to measure misallocation I need to isolate the second channel. I define the gains from reallocation as

$$G = \frac{Y^* - Y^H}{Y^H}$$

where Y^H is output with the original allocations but all workers are as productive as family workers. Then Y^* is output with both perfect allocations and fully productive workers.

With perfect markets labor substitutability holds and the optimal allocations derived in Section 2.3 are still valid. But since I assume hired labor is as productive as family labor— $f(L_{it}^H) = L_{it}^H$—I have created measurement error that will

Figure A.III
Misallocation with a Higher Production Elasticity for Labor

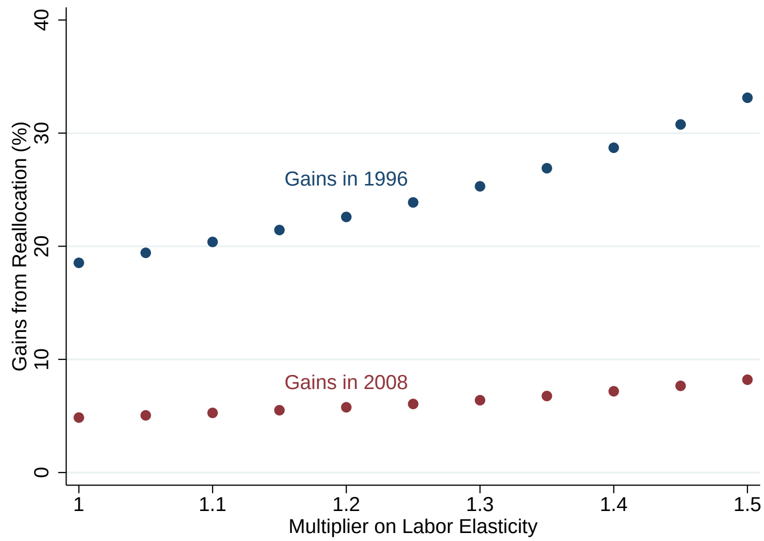

bias the production elasticity of labor downward. Rather than estimating the gains from effective family labor θ_L I estimate some combination of the gains from family labor and hired labor. If I knew the true elasticity I could define

$$Y_t^H = \sum_i A_{it}\phi_{it}\bar{K}_{it}^{\theta_K}\bar{T}_{it}^{\theta_T}\bar{L}_{it}^{\theta_L}$$
$$Y_t^* = \sum_i A_{it}\phi_{it}(K_{it}^*)^{\theta_K}(T_{it}^*)^{\theta_T}(L_{it}^*)^{\theta_L}$$

where factor choices with bars are those I observe the farmer choose and factor choices with stars are the optimal allocations. Since I do not know the true elasticity I assume my biased estimate of the elasticity $\tilde{\theta}_L = (1+\gamma)\theta_L$ for some multiplier γ. I then recalculate G for different values of $(1+\gamma)$ and graph the results in Figure A.III.

Raising the elasticity raises the gains from reallocation. This does not mean the original estimates were wrong. The gains would rise even if the original estimates were right because changing the estimates makes marginal products look unequal and thus gives the illusion of misallocation. But changing the estimates does bound how much bias might be caused by falsely assuming hired labor is as efficient as household labor.

Figure A.III shows that even if the true θ_L is 25 percent higher than my estimate, misallocation in 1996 only rises from 18 percent to 24 percent. Misallocation in 2008 changes even less, from about 5 percent to about 6 percent. In short, although dropping the assumption of labor substitutability does cause me to find more misallocation the difference is relatively small.

A.5.4 How Would Measurement Error in Land Change the Results?

If there is error in measured land—because the survey respondent gave an inaccurate response or because land differs in its quality—the estimates of misallocation may be biased. I cannot estimate misallocation in the counterfactual world where land is more accurately measured. But in the spirit of Romer (1986), I can estimate misallocation in the counterfactual world where land is *less* accurately measured. That is, I can adulterate the existing measure of land and redo the entire procedure for estimating misallocation. Comparing the artifically coarsened estimate to the actual estimate gives some sense of the direction and magnitude of the bias.

I consider two forms of measurement error in land. The first is a completely random shock, as would happen if the respondent gave the interviewer a rough guess of the area of paddy farmed by the household. I construct measured land T as true land T^* plus a normally distributed shock, where T^* is the measure of land used in the main text. I assume the shock has a standard deviation of 1, 5, or 10 percent of the standard deviation of land. I re-estimate the production function and the sample-wide gains from reallocation for each of 100 draws of measurement error.

Panel A of Table A.V reports the mean of all simulation estimates, with the 5th and 95th percentiles reported in parentheses. The random error has almost no effect on the gains from reallocation. Even with an error of 10 percent the estimated gains are barely changed, and the original estimate lies within the 90 percent confidence interval.

Table A.V
Do the Results Change When Measurement Error is Added?

A. Random Error

	No Error	Error: 1 percent	Error: 5 percent	Error: 10 percent
1996	18.5	18.5	18.1	17.6
		(18.1, 19.6)	(16.9, 20.0)	(15.9, 19.8)
2002	12.4	12.4	12.3	12.0
		(12.2, 13.0)	(11.7, 13.1)	(11.1, 13.2)
2008	4.8	4.8	4.9	5.0
		(4.8, 5.0)	(4.7, 5.1)	(4.7, 5.5)

B. Systematic Error

	No Error	Error: 1 percent	Error: 5 percent	Error: 10 percent
1996	18.5	18.2	16.8	15.3
		(18.1, 18.2)	(16.6, 17.0)	(15.0, 15.5)
2002	12.4	12.2	11.6	11.0
		(12.2, 12.3)	(11.6, 11.7)	(10.8, 11.1)
2008	4.8	4.8	4.9	4.9
		(4.8, 4.9)	(4.8, 4.9)	(4.9, 5.0)

Note: Each entry gives the gains from reallocation in the given year (indicated in the row header) and assuming the given level of error (indicated in the column header). I compute the mean across all simulations and report the 5th and 95th percentiles in parentheses.

I then consider measurement error that actually raises output, which works like unobserved land quality. Suppose each unit of land has quality $e^{f \cdot q}$, where q is correlated with output and f is the importance of quality in true (quality-adjusted) land. Then $T^* = Te^{fq}$, which implies I can construct "observed" land by dividing my measure of land by e^{fq}. (I also rescale the resulting measure to have the same standard deviation as the original measure of land.) Let

$$q = \frac{1}{2}y + \frac{1}{2}\nu$$

where ν is normally distributed noise with standard deviation equal to that of log output y. Note that q now inherits some of the persistence of y, meaning it is neither fixed nor completely independent between years.

As before, I run 100 simulations; the mean and percentiles of the simulations are given in Panel B of A.V. Unlike before, now there is some reduction in estimated misallocation in 1996. As expected, unmeasured quality causes me to find less misallocation than exists in truth. However, the reduction is not large. Even when f equals 10 percent misallocation falls by roughly 3 percentage points. The reduction is even smaller in 2002, with the new estimate just 1.5 percentage points lower. The estimate in 2008 is essentially unchanged.

Though this is only a rough check, it is informative. Unmeasured land quality may cause me to underestimate misallocation at baseline, though the difference is unlikely to be large. It is unlikely that even with substantial error I would find estimates of misallocation as large as found by, say, Adamopoulos and Restuccia (2014). More importantly, the qualitative result—that misallocation falls by 2008 to a level that is trivial—seems fairly robust.

A.6 Additional Analysis (For Online Publication)

A.6.1 Revenue Versus Physical Productivity

This appendix shows the relationship between the rice price and physical productivity, the two components of revenue productivity. Suppose physical output is given by

$$Q = Z^Q EF(K, L, T)$$

where Z^Q is physical productivity, and the terms E and $F(K, L, T)$ are aggregate allocative efficiency and the aggregate production function. Assuming there is a single rice price—the international price of rice—then aggregate revenue Y is

$$\begin{aligned} Y &= PQ \\ &= PZ^Q EF(K, L, T) \\ \Rightarrow Z^Q &= \frac{Y}{PEF(K, L, T)} \\ &= \frac{Z}{P} \end{aligned}$$

where, as in the main text, Z is revenue productivity. In other words, under the assumption that there is a single rice price I can define physical productivity as revenue productivity divided by the price.

The top panel of Figure A.IV shows the log change over baseline of both numerator and denominator. Though revenue productivity often rises with the price—as one would expect—the relationship is not one-for-one. Clearly physical productivity is also changing over the sample. The bottom panel shows the log change in physical productivity over baseline. The graph suggests physical productivity fell in the wake of the Asian financial crisis. This could be because the crisis caused workers in manufacturing to lose their jobs. These workers may have turned to farming to survive. People who take up farming only as a last resort are likely less productive than those who choose farming over other options. The entry of such farmers would likely cause physical productivity to fall.

Physical productivity then rose from 1998 to 2001, likely for two reasons. First, Thailand recovered from the crisis, creating new jobs in manufacturing. Second, the price of rice fell. As the returns outside farming rose and those within farming fell, the least productive farmers may have returned to manufacturing.

Finally, from 2002 to 2008 the price of rise again rose. This may have gradually tempted less productive farmers back to producing rice, which gradually lowered physical productivity.

Figure A.IV
Physical Productivity Rises when Prices Fall

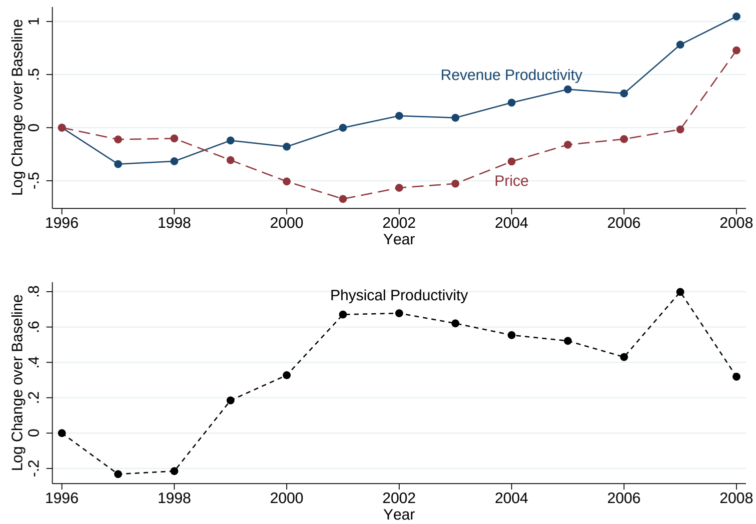

A.7 Data Appendix (For Online Publication)

A.7.1 How Representative is the Sample?

Table A.VI compares the distribution of agricultural landholdings in the baseline Townsend Thai sample to the population distribution.[30] The population distribution is taken from the Food and Agriculture Organization's World Census of Agriculture, which reports data from the 2003 Thai Census of Agriculture. As the census does not separate agricultural land by crop, I report total agricultural landholdings for both the sample and the census. I report numbers from the baseline sample rather than the year 2003 because, by design, a panel survey is only representative of the total population at baseline. The table shows that the distribution in the sample and in the census are broadly similar.

[30] I define agricultural landholdings as all land that is not used exclusively for housing or for aquaculture.

Table A.VI
Distribution of Land

	Distribution According to...	
Land (Ha)	Baseline Data	Agricultural Census
less than 1	18.93%	22.69%
1–2	20.50%	22.81%
2–5	39.24%	35.85%
5–10	15.66%	14.32%
10–20	4.10%	3.58%
20–50	1.43%	0.70%

Note: This table shows the percentage of agricultural landholdings that fall into each of six bins. The agricultural census was conducted in 2003.

A.7.2 Definitions of Village (or higher) Level Variables

International Rice Prices From the IMF's commodity price data. I took the yearly average.

Village Wage Rates From Section V of the annual household survey.

For 1996: For each household, find any worker in the "other" category who lists their occupation as related to "labor" or "labour" and compute their daily wage. Construct medians by village, subdistrict, etc.

For 1997-2008: For each household, find any worker listed as general agricultural laborer of any sort or in the "other" category reporting an occupation related to "labor" or "labour" and compute their daily wage. Construct medians by village, subdistrict, etc.

Village Population From Section iii of the annual key informant survey. Survey records both number of households and population of the village.

Precipitation I obtained gridded monthly rainfall estimates to cover Thailand from 1996 through 2008. The estimates for 1996 and 1997 were .5 x .5 lat-long degree grids from the University of Delaware Climate Project's Terrestrial Precipitation Gridded Monthly Time Series. Those for the rest of the year were .25 x .25 lat-long degree grids from NASA's Tropical Rainfall Measuring Mission (Product 3B-43). I used ArcGIS's Topo-to-Raster tool to create an interpolated raster for the rainfall data. I then used district-level boundaries from the Global Administrative Areas Project (GADM) to construct district-level monthly averages. I converted the levels to fractional deviations from the mean for each month (where the monthly mean was computed over the sample period). The rainfall shocks relevant for a particular year match the survey response period (so rainfall in 1996 is rainfall from May 1996 through April 1997).

Rice Suitability Index I obtained the Food and Agriculture Organization's Global Agro-Ecological Zone (GAEZ) data for the climactic suitability of rice and maize. The data for rice suitability came from Plate 38: Suitability for rain-fed and irrigated Rice (high input). The data for maize suitability came from Plate 30: Suitability for rain-fed Grain Maize (intermediate inputs). I computed a zonal mean for each district: an average value indexing the climate suitability of the district for each crop. I inverted the index so higher values correspond to greater suitability.

A.7.3 Definitions of Household-Level Variables

7.3.1 Factors of Production

Land I use the land cultivation data from Section XIV (Landholdings) of the Annual Household Survey. Households report the quantity and value of land they cultivate (regardless of ownership) by use; that is, they separately report land for rice, field crops, orchards, and vegetables. I total the area of the plots for each use and mark this as the land cultivated for each crop. I also deflate the reported value of the plots and total for each crop to form the value of the land owned.

Capital *Owned Mechanical Capital:* I use Section XII (Agricultural Assets) of the Annual Household Survey. The survey reports the number of assets of each type, where I group the following assets into broad categories by depreciation rates: tractors (walking tractors, large and small four-wheel tractors), machines (sets, sprinklers, and threshing machines) and structures (crop storage buildings). I depreciate tractors like vehicles, so the depreciation rates I use are 2 percent for structures, 10 percent for machines, and 20 percent for tractors. I correct clear errors in the series of asset classes where an asset disappears and reappears without any record of a sale or appears and disappears without any record of a purchase. I then construct the value of assets owned at the beginning of the first survey round by deflating and depreciating the purchase price by the year of purchase. I then attempt to follow each asset over time, where I label a piece of equipment a separate asset if the quantity of a certain type of equipment rises from zero to some positive number. I unfortunately must treat the addition of new pieces of equipment to an existing stock as identical to the existing assets of that class; but it is fairly rare that a household has more than one piece of equipment of a certain class. I then assign a "price" to each asset with the sale value at the very latest transaction date I can find for it (where the initial value in the first survey round is also considered a transaction). I adjust that price for deprecation in preceding and following years and compute the asset value in a given year by multiplying the price by the quantity held. [Recall the quantity is almost always one if the household owns any.] If I cannot identify a price, I am forced to drop the asset from my calculations. [In rare cases where I can identify a year of acquisition but not a price, I use the intertemporal median of village, subdistrict, or district medians for the equipment type.] I then aggregate the value of all assets for each household in each year to construct the value of owned mechanical capital.

Buffalos: I assume buffalo are the only animal used to harvest rice and compute the value of buffalo using the appropriate responses from Section XII (Agricultural Assets) of the Annual Household Survey. The household reports the total current value of all buffalos owned, which I deflate. Missing values for this variable generally mean the question does not apply (e.g. the household owns no buffalo), so I treat missing values as zero.

Capital Expenses: For rented capital, maintenance expenses, and intermediate inputs (which I treat as capital) I use the portion on farm expenses in Section XVI (Income) of the Annual Household Survey. After deflating all currency, I compute intermediate inputs as the sum of expenses on seeds, fertilizers, pesticides/herbicides, and fuel. I then rescale the value of rented capital by a user cost: the depreciation rate plus an interest rate, which I set as 4 percent in line with the literature. It may seem strange to assume a common interest rate given the possibility of financial market failures; but recall my objective is to create a consistent measure of the productive value of the capital owned. Allowing the productive value of a tractor to vary based on the household's borrowing cost makes no sense. I do not know how much of the rental cost goes to machines versus vehicles, so I take the depreciation rate as the average of the rates for each type of asset. Finally, I add the value of maintenance expenses, which is investment (recall I assume investment is immediately productive).

Total capital is the sum of the value of owned mechanical capital, buffalos, and capital expenses.

Labor *Family Labor:* I first construct the quantity of family labor using Section V (Occupation) of the Annual Household Survey. In each year I count the number of household members who report being unpaid family laborers in their primary occupation and report farming of any sort as their primary or secondary occupation (or report working in the "FIELDS" if their occupation is not categorized). (Some farmers grow several types of crops, but the survey only allows two responses for occupation. To deal with this problem, I reason that a household growing rice will use its working family members on all of its fields, so any family member who works in the fields necessarily works in rice.) I define the number of family workers as the number of household members who satisfy this criterion. I have no intensive margin information on how much the household works, so I assume all members work sixty days of the year in the fields (the median number of days worked from the two years of the Monthly Household Survey available at the time of writing). I aggregate the per-member days worked for each household-year to compute the quantity of family labor. (In other words, I multiply the number of family workers by the median days an individual works on their fields conditional on working at all.)

Hired Labor: The only measure of hired labor is the expenditure on wages recorded among the farm expenses in Section XVI (Income) of the Annual Household Survey. I divide the total expenses on wages by the village-level median daily wage (see above) to construct a measure of days worked by hired labor. Total labor is simply the sum of family and hired labor.

7.3.2 Productivity Modifiers

Catastrophes/Bad Income Shocks I use the questions about bad income years from Section II (Risk Response) of the Annual Household Survey. The household reports the worst of the last several years for income (including the response year), and the reason for it being atypically bad. If a household chooses the response year as the worst, I mark it as suffering one of the following catastrophes based on the reason it gives:

- Reports low income this year due to illness
- Reports low income this year due to death in family
- Reports low income this year due to retirement
- Reports low income this year due to flooding
- Reports low income this year due to crop-eating pests
- Reports low income this year due to poor rainfall
- Reports low income this year due to low yield for other reasons
- Reports low income this year due to low price for output

Hunger/Undernourishment I have no direct measure of calories and instead adapt the notion of the staple budget share (SBS) introduced by Jensen and Miller (2010). I use consumption expenditure data from Section XV (Expenditure) to compute the fraction of the household's budget spent on the staple food in Thailand: rice. This measure includes the value of rice the household grew itself. The intuition behind this measure is that as a household becomes wealthier (and less hungry) it substitutes away from the staple crop towards other foods (which are superior goods). The higher the SBS, the more likely it is the household is hungry.

7.3.3 Other Variables

Revenue I use the questions about gross income from Section XVI (Income) of the Annual Household Survey. Households report their revenue from each of several sources, including rice farming and other agricultural activities. Enumerators explicitly reminded households to include the value of crops they produced and then consumed. I deflated and constructed income variables for each of the following sources: Rice Farming, Corn Farming, Vegetable Farming, Orchard Farming, Other Farming.

Education I use the questions about age and education from Section IV (Household Composition) of the Annual Household Survey. I keep information about the age, highest grade completed, and school system of the household head. I defined separate variables for number of years spent in primary school (generally from 1-6 for P1-P6), number of years spent in secondary school (generally 1-6 for M1-M6, unless the individual chose the vocational rather than academic track, in which case I set years of secondary school to 3), years of vocational school (from 1-3 for PWC1-PWC3, PWS1-PWS3, or PWT1-PWT3), and years of university (from 1-4).

Rice Farming Experience Households report the number of years spent at their primary occupation in Section V (Occupation) of the Annual Household Survey. I record the years spent for individuals who report rice farming as their primary occupation and categorize themselves as "owners" of the business. I take this as a measure of the rice-farming experience of the household (head). In the rare cases where multiple household members claim to be the owner of a rice-farming business, I take the median as the household-level experience.

Constraints I use the questions on farm expansion from Section XII (Agricultural Assets) of the Annual Household Survey. A household is labeled as "constrained" if it reports there is room for profitable expansion in its business. I label it credit-constrained if it reports insufficient money for labor, land, or equipment among the reasons for not expanding. I label it factor-constrained if it reports not enough land or labor (a distinct response from insufficient money) among the reasons for not expanding. A household can be both credit- and factor-constrained. I further label households as exclusively credit- or factor-constrained if they report a constraint in one but not the other.

Risk-Aversion In 2003 the survey started posing to households a hypothetical choice between staying at their current income forever and taking a job that with 50-50 chance pays either double or two-thirds their current income. If they choose their current job the interviewer gives them the same choice except the alternate job now has a 50-50 chance of paying either double or 80 percent of their current income. If the respondent chose his current job for both questions I marked it as having "high risk aversion." Since the question was not asked in 1996, I use the 2003 question in Section 6.

Savings I use Section XIX (Savings) of the annual household survey. I take the total savings each household has deposited with commercial banks, agricultural cooperatives, the Bank for Agriculture and Agricultural Cooperatives, PCG village funds, and a rice bank.

A.8 Asymptotics (For Online Publication)

This appendix presents a fully general form of the model I use in Section 2 to derive my measures of misallocation. I use this fully general model to derive the asymptotic behavior of the measure of misallocation. Suppose households are the sole economic actors, and as ever they live to maximize their lifetime utility from consumption. They earn income by selling or renting out the factors they own (including labor) and by operating firms. Aside from the usual budget constraint, they face potentially binding constraints on their choices of factors. For example, if no labor market exists they are constrained to use exactly their labor endowment. They may also be constrained in their period-to-period liquidity. They save and borrow at interest rates that need not be common across households, and may also have to pay an external finance premium to borrow. A household's access to insurance may be imperfect, which means its consumption depends on the profits of its firm. Finally, households differ in their preferences (notably their risk tolerance) and the productivity of the firms they operate.

Suppose household i owns and operates firm $i \in I_t$, where I_t is some group of firms (a village or a sector). The household maximizes present discounted utility from consumption over an infinite horizon:

$$\max_{(c_{i,t+j},\boldsymbol{X}_{i,t+j},\boldsymbol{I}_{i,t+j})_{j=0}^{\infty}} \mathbb{E}[\sum_{j=0}^{\infty} \rho^j u_i(c_{i,t+j}) \mid \mathcal{I}_{it}].$$

Subject to:

$$\begin{aligned}
c_{it} + b_{i,t+1} &= y_{it} + [1 + r_{it} + \zeta_{it}(z_{it} - b_{it}; b_{it}, \boldsymbol{X}^o_{i,t-1})](z_{it} - b_{it}) && (\mathit{Budget\ Const.})\\
y_{it} &= f(\boldsymbol{A}_{it}, \boldsymbol{\phi}_{it}, \boldsymbol{X}_{it}; \boldsymbol{\theta}_i) && (\mathit{Production})\\
z_{it} &= \boldsymbol{w}^T_{it}(\mathbf{X}_{it} - \boldsymbol{X}^o_{it}) + \boldsymbol{p}^T_{it}\boldsymbol{I}_{it} && (\mathit{Expenditure})\\
X^o_{k,it} &= (1 - \delta^k) X^o_{k,i,t-1} + I_{k,it} && \forall k = 1, \dots, K\\
z_{it} - b_{it} &\le \omega_{it}(b_{it}, \boldsymbol{X}^o_{i,t-1}) && (\mathit{Liquidity\ Const.})\\
\underline{\mathbb{X}}_{it} &\le \mathbb{X}_{it}(\boldsymbol{X}_{it} - \boldsymbol{X}^o_{it}, \boldsymbol{I}_{it}) \le \overline{\mathbb{X}}_{it} && (\mathit{Factor\ Choice\ Const.})
\end{aligned}$$

where $\mathcal{I}$ is the information set, f a strictly concave decreasing returns revenue production function, $\boldsymbol{A}$ anticipated revenue productivity, $\boldsymbol{\phi}$ unanticipated revenue productivity, $\boldsymbol{\theta}$ a vector of production parameters, $\boldsymbol{X}$ factor levels used in production, c is consumption, b borrowing, r the borrowing rate, $\zeta(\cdot)$ an external finance premium, y_{it} revenue, $\boldsymbol{X}^o$ owned factors, $\boldsymbol{I}$ purchase of factors, $\boldsymbol{p}$ a vector of factor purchase prices, z input expenditure, and $\omega(\cdot)$ a liquidity constraint. $\mathbb{X}_{it}$ is a continuously twice-differentiable factor choice transformation function, and $\underline{\mathbb{X}}_{it}, \overline{\mathbb{X}}_{it}$ upper and lower bounds on (transformed) factor choice; they bound a household's access to factors beyond those it owns (rented factors). Assume all past and currently dated variables are elements of $\mathcal{I}$ except ϕ. Rental prices $\boldsymbol{w}$ and all other prices can vary by household/firm i.

For notational simplicity I model insurance markets implicitly as the correlation between a household's unanticipated productivity and its consumption (perfect insurance ensures zero correlation). I have assumed away output and asset taxes because they are not important in the empirical application; accounting for them is straightforward if the tax schedule is known.[31]

Let X_{k,I_t} be the aggregate stock of factor k among the unit measure of firms

[31] For output taxes, for example, one would simply modify the budget constraint to be $c_{it} + b_{i,t+1} = (1-\tau_{it})y_{it} + \cdots$ and then perform all subsequent operations conditional on the presence of the taxes to account for the fact that they will continue to distort even the counterfactual optimal scenarios where market failures are eliminated.

in I_t. Define an *allocation vector* as a set of K-dimensional vectors $\{\mathbf{X}'_{it}\}_{i\in I_t}$ such that $\int_{i\in I_t} X'_{k,it}\, di = X_{k,I_t} \forall k$.

Varying factor prices and savings rates, external finance premiums, liquidity constraints, and factor choice constraints can all distort realized allocations away from the frictionless benchmark. Eliminating them separates the household problem from the firm problem and produces the production allocations of the frictionless neoclassical world. Denote outcomes in the world with no constraints or market imperfections by asterisks, and characterize it with these conditions:

$$\begin{aligned}
\text{(Law of One Price) } & \boldsymbol{w}_{it} = \boldsymbol{w}_{I_t}, \boldsymbol{p}_{it} = \boldsymbol{p}_{I_t} \forall i \in I_t m_k \forall k \\
\text{(Unconstrained Factor Choices) } & \underline{\mathbb{X}}_{it} = \mathbb{X}(-\boldsymbol{X}^o_{it}, -\boldsymbol{X}^o_{i,t-1}), \overline{\mathbb{X}}_{it} = \mathbb{X}(\infty, \infty) \ \ \forall i \in I_t \\
\text{(Perfect Credit Markets) } & \omega_{it} = \infty, \zeta_{it}(\cdot) = 0, r_{it} = r_{I_t} \ \ \forall i \in I_t \\
\text{(Perfect Insurance Markets) } & c_{it} \perp \boldsymbol{\phi}_{it} \ \ \forall i \in I_t.
\end{aligned}$$

Under these assumptions the firm maximizes per-period expected profit independently of the household's dynamic consumption problem:

$$\max_{\boldsymbol{X}_{it}} \qquad \mathbb{E}[f(\boldsymbol{A}_{it}, \boldsymbol{\phi}_{it}, \boldsymbol{X}_{it}; \boldsymbol{\theta}_i) - \boldsymbol{w}^T_{I_t}\boldsymbol{X}_{it} \quad | \, \mathcal{I}_{it}]$$

Then the following first-order conditions and market-clearing conditions characterize the unique general equilibrium outcome:

$$\mathbb{E}[f_{X_{k,it}}(\boldsymbol{A}_{it}, \boldsymbol{\phi}_{it}, \boldsymbol{X}^*_{it}; \boldsymbol{\theta}_i)] = w_{k,I_t} \ \ \forall k \tag{16}$$

$$\int_{i\in I_t} X^*_{k,it}(w_{k,I_t}, \boldsymbol{A}_{it}, X^*_{-k,it}; \boldsymbol{\theta}_i)\, di = X_{k,I_t} \ \ \forall k \tag{17}$$

Solving these equations solves for the optimal allocation vector $\{\boldsymbol{X}^*_{it}\}$. The optimal allocations solve a system of equations that contain only observables and production parameters estimable from observables. This makes calculating the counterfactual scenario with production and factor data possible.

To solve for the outcome where factor markets are perfect I must assume unanticipated shocks are Hicks-Neutral. That is, $y_{it} = \phi_{it} f(\boldsymbol{A}_{it}, \boldsymbol{X}_{it}; \boldsymbol{\theta}_i)$. Since all factors are equally risky in production, imperfect insurance only affects overall expenditure and not expenditure on capital versus labor. Consider the following hypothetical: Firm $i \in I_t$, which uses $\bar{\boldsymbol{X}}_{it}$ in production, now has those factors as "endowments." Each firm can then trade factors with $-i \in I$ subject to its expenditures being equal to the value of its endowment until its factor mix is optimal.

Since firms cannot change their total expenditure for a period, they again optimize period-by-period:

$$\max_{\boldsymbol{X}_{it}} \qquad \mathbb{E}[\phi_{it} f(\boldsymbol{A}_{it}, \boldsymbol{X}_{it}; \boldsymbol{\theta}_i) - \boldsymbol{w}_{I_t}^T \boldsymbol{X}_{it}]$$

Subject to:

$$\boldsymbol{w}_{I_t}^T(\boldsymbol{X}_{it} - \bar{\boldsymbol{X}}_{it}) = 0$$

Essentially, I have dropped the firms into an Edgeworth economy where their profit function plays the role of a utility function. The following equations characterize the unique outcome:

$$\frac{\mathbb{E}[f_{X_k}(\boldsymbol{A}_{it}, \boldsymbol{X}_{it}^+; \boldsymbol{\theta}_i)]}{\mathbb{E}[f_{X_j}(\boldsymbol{A}_{it}, \boldsymbol{X}_{it}^+; \boldsymbol{\theta}_i)]} = \frac{w_{k,I_t}}{w_{j,I_t}} \quad \forall k, j\, \forall i \tag{18}$$

$$\boldsymbol{w}_{I_t}^T(\boldsymbol{X}_{it} - \bar{\boldsymbol{X}}_{it}) = 0 \quad \forall i \tag{19}$$

$$\int_{i \in I_t} X_{k,it}^+(w_{k,I_t}, \boldsymbol{A}_{it}, X_{-k,it}^+; \boldsymbol{\theta}_i)\, di = X_{k,I_t} \quad \forall k \tag{20}$$

Define efficiency and the gains from reallocation as in the main text. Assume anticipated productivity is also Hicks-Neutral, the production function is homogeneous, and common production parameters. I can prove this theorem about my measure of the costs of factor versus financial market failures:

Proposition 2 (Bounding) *Assume $y_{it} = A_{it}\phi_{it} f(\boldsymbol{X}_{it}; \boldsymbol{\theta}_i)$, f is homogeneous of degree σ, $\boldsymbol{\theta}_i = \boldsymbol{\theta}_{I_t}\ \forall i \in I_t$, and $\mathbb{E}[A_{it}\{(\tilde{X}_{1,it}^+)^\sigma - (X_{1,it}^+)^\sigma\}] > 0$, where $\tilde{X}_{1,it}$ is the quantity of the first factor i would choose at time t under Unconstrained Factor Choices and the Law of One Price without being constrained by endowments. Then G_I^{FACT} is a lower-bound on the true gains from moving to Unconstrained Factor Choices and the Law of One Price. Likewise, G_I^{FIN} is an upper-bound on the true gains from subsequently creating Perfect Credit and Insurance Markets.*

Proof: Consider the (unobservable and incalculable) outcome where the Law of One Price and Unconstrained Factor Choices hold without the extra restriction of endowment conditionality. Call this the True Perfect Factor Market outcome. Denote with superscript $+$ a variable specific with perfect factor markets and the endowment condition: the Endowment-Conditional Perfect Factor Market (ECPFM) outcome. Let overset tilde and superscript $+$ variables come from the True Perfect Factor Market (TPFM) outcome. An over-bar denotes variables from to the observed/realized outcome. For notational ease, suppress the common production parameter $\boldsymbol{\theta}_I$. I first prove two lemmas useful to the main result.

Lemma 1 *For any level or vector of factor choices X_{it}, let $\ddot{X}_{it} = X_{it}/X_{1,it}$. Then $\ddot{\tilde{\boldsymbol{X}}}_{it}^{+} = \ddot{\tilde{\boldsymbol{X}}}_{I_t}^{+}$ and $\ddot{\boldsymbol{X}}_{it}^{+} = \ddot{\boldsymbol{X}}_{I_t}^{+}$ for all $i \in I$ (that is, all firms in both outcomes will employ factors in exactly the same proportions).*

The optimality condition for ECPFM is

$$\frac{f_{X_k}(\boldsymbol{X}_{it}^{+})}{f_{X_1}(\boldsymbol{X}_{it}^{+})} = \frac{w_k^{+}}{w_1^{+}}.$$

By homogeneity,

$$\begin{aligned} \frac{X_{1,it}^{\sigma-1} f_{X_k}(\boldsymbol{X}_{it}^{+}/X_{1,it})}{X_{1,it}^{\sigma-1} f_{X_1}(\boldsymbol{X}_{it}^{+}/X_{1,it})} &= \frac{w_k^{+}}{w_1^{+}} \;\; \forall k \\ \Rightarrow \frac{f_{X_k}(\ddot{\boldsymbol{X}}_{it}^{+})}{f_{X_1}(\ddot{\boldsymbol{X}}_{it}^{+})} &= \frac{w_k^{+}}{w_1^{+}} \;\; \forall k \end{aligned}$$

Since f satisfies strictly decreasing returns, $\boldsymbol{X}_{it}$ is unique and thus the $\ddot{\boldsymbol{X}}_{it}$ that satisfies the above conditions is also unique for each i, t. But the above conditions are not functions of any variables unique to i (e.g. A_{it}), and thus $\ddot{\boldsymbol{X}}_{it} = \ddot{\boldsymbol{X}}_{I_t}$ for all $i \in I_t$. A similar argument shows that $\ddot{\tilde{\boldsymbol{X}}}_{it} = \ddot{\tilde{\boldsymbol{X}}}_{I_t}$ for all $i \in I_t$. ■

Lemma 2 $\ddot{\tilde{\boldsymbol{X}}}_{I_t}^{+} = \ddot{\boldsymbol{X}}_{I_t}^{+}$ *for all $i \in I_t$, that is the mixes of factors will be identical with TPFM and ECPFM.*

Since $\{\boldsymbol{X}_{it}^{+}\}$ and $\{\tilde{\boldsymbol{X}}_{it}^{+}\}$ are both allocation vectors, each factor must aggregate to the total observed stock. For the latter, for example, for all k

$$\int_{i \in I_t} X_{k,it}^{+} di = X_{k,It}$$

$$
\begin{aligned}
\Rightarrow \qquad & \int_{i\in I_t} X_{1,it}\ddot{X}^+_{k,I_t} di = X_{k,I_t} \\
\Rightarrow \qquad & \ddot{X}^+_{k,I_t}\int_{i\in I_t} X_{1,it} di = X_{k,I_t} \\
\Rightarrow \qquad & \ddot{X}^+_{k,I_t} X_{1,It} = X_{k,I_t} \\
\Rightarrow \qquad & \ddot{X}^+_{k,I_t} = \frac{X_{k,I_t}}{X_{1,I_t}}
\end{aligned}
$$

where the second line follows from Lemma 1. Parallel arguments show that $\ddot{\tilde{X}}^+_{k,It} = \frac{X_{k,It}}{X_{1,It}}$. Then

$$\ddot{X}^+_{k,It} = \ddot{\tilde{X}}^+_{k,It} \quad \forall k$$

implying $\ddot{\boldsymbol{X}}^+_{It} = \ddot{\tilde{\boldsymbol{X}}}^+_{It}$. ■

To prove the main result, write the difference between the TPFM and ECPFM outcome as

$$
\begin{aligned}
\tilde{Y}^+_{I_t} - Y^+_{I_t} &= \int_{i\in I_t} A_{it}\phi_{it} f(\tilde{\boldsymbol{X}}^+_{it})\, di - \int_{i\in I_t} A_{it}\phi_{it} f(\boldsymbol{X}^+_{it})\, di \\
&= \mathbb{E}[A_{it}\phi_{it} f(\tilde{\boldsymbol{X}}^+_{it})] - \mathbb{E}[A_{it}\phi_{it} f(\boldsymbol{X}^+_{it})] \\
&= \mathbb{E}[A_{it}\phi_{it}\{f(\tilde{\boldsymbol{X}}^+_{it}) - f(\boldsymbol{X}^+_{it})\}] \\
&= \mathbb{E}[\phi_{it}]\mathbb{E}[A_{it}\{f(\tilde{\boldsymbol{X}}^+_{it}) - f(\boldsymbol{X}^+_{it})\}] \\
&= \mathbb{E}[A_{it}\{f(\tilde{\boldsymbol{X}}^+_{it}) - f(\boldsymbol{X}^+_{it})\}]
\end{aligned}
$$

where the second equality comes from the measure 1 normalization, the fourth the independence of anticipated and unanticipated variables, and the fifth the unit mean normalization of ϕ_{it}.

By the homogeneity of f and Lemma 1,

$$
\begin{aligned}
f(\tilde{\boldsymbol{X}}^+_{it}) &= (\tilde{X}^+_{1,it})^\sigma f(\ddot{\tilde{\boldsymbol{X}}}^+_{I_t}) \\
f(\boldsymbol{X}^+_{it}) &= (X^+_{1,it})^\sigma f(\ddot{\boldsymbol{X}}^+_{I_t t})
\end{aligned}
$$

and by Lemma 2 $f(\ddot{\boldsymbol{X}}^+_{it}) = f(\ddot{\tilde{\boldsymbol{X}}}^+_{it}) = a$. Applying these results to the above expressions, we have that

$$\begin{aligned}\tilde{Y}_{I_t}^+ - Y_{I_t}^+ &= \mathbb{E}[A_{it}\{f(\tilde{\boldsymbol{X}}_{it}^+) - f(\boldsymbol{X}_{it}^+)\}] \\ &= \mathbb{E}[A_{it}\{(\tilde{X}_{1,it}^+)^\sigma a - (X_{1,it}^+)^\sigma a\}] \\ &= a\mathbb{E}[A_{it}\{(\tilde{X}_{1,it}^+)^\sigma - (X_{1,it}^+)^\sigma\}]\end{aligned}$$

Since $a > 0$ and $\mathbb{E}[A_{it}\{(\tilde{X}_{1,it}^+)^\sigma - (X_{1,it}^+)^\sigma\}] > 0$ by assumption, TPFM output is greater than in the ECPFM outcome, and so the calculated gains will be as well.

QED

Now suppose that consistent estimators of $\{\boldsymbol{A}_{it}, \boldsymbol{\phi}_{it}, \boldsymbol{\theta}_i\}$ are available. Then it is a numerical exercise to solve the sample analogs of (16) and (17) for estimates of the CCM allocations and (18), (19), and (20) for sample analogs of the PFM allocations. Plug the computed allocations into the expressions for efficiency and the gains. The following proposition summarizes the asymptotic properties of these estimators:

Proposition 3 *Suppose $\hat{I}$ is a random sample of I, and define $\hat{E}_{\hat{I}}, \hat{E}_{\hat{I}}^{FACT}, \hat{E}_{\hat{I}}^{FIN}$, and the estimators of the other components as described above. Finally, assume the expectations and variances of $\boldsymbol{A}_{it}, \boldsymbol{\phi}_{it}, \boldsymbol{\theta}_i, y_{it}, \boldsymbol{X}_{it}$ are finite. Then the estimators are all consistent and asymptotically normal.*

Proof:

Consistency: I will prove the consistency of $\hat{E}_{\hat{I}_t}$; demonstrating the consistency of the other estimators is similar. Suppress time subscripts for notational simplicity. I will first identify the population parameters in terms of their moments, and then demonstrate that the sample analogs are consistent estimators.

Recall that (16) characterizes any interior solution to the population optimization - in other words, if an interior solution exists, the function $\boldsymbol{X}^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \boldsymbol{w}_I)$ characterizes firm i's optimal allocation as a function of i-specific parameters and the prices. Since f is concave and satisfies DRS, the solution is not only interior but also unique. Moreover, the Maximum Theorem Under Convexity (see Sundaram, 1996, p. 237) guarantees that $\boldsymbol{X}^*$ is a continuous function of the population prices $\boldsymbol{w}_I$ (see the Lemma below). Inspection demonstrates that the optimality conditions are identical for the sample optimization, and thus the derived $\boldsymbol{X}^*$ is as well.

Applying the measure-one normalization, (17) in population reduces to

$$\mathbb{E}[X^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \boldsymbol{w}_I)] = \mathbb{E}[X_i].$$

Note that the sample analog of (17) in random sample $\hat{I}$ reduces to the sample analog of this moment condition trivially:

$$\sum_{i \in \hat{I}} X^*(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\theta}}_i, \boldsymbol{w}_I) = \sum_{i \in \hat{I}} \boldsymbol{X}_i$$
$$\frac{\sum_{i \in \hat{I}} X^*(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\theta}}_i, \boldsymbol{w}_I)}{|\hat{I}|} = \frac{\sum_{i \in \hat{I}} \boldsymbol{X}_i}{|\hat{I}|}$$

Since $\{\hat{\boldsymbol{A}}_{it}\}, \{\hat{\boldsymbol{\theta}}_i\}$ are consistent estimators for the population technology parameters, then together with the Lemma below this implies that $\hat{\boldsymbol{w}}_I$ is a consistent GMM estimator for $\boldsymbol{w}_I$.

Applying the measure 1 normalization to the definition of E, we have

$$E_I = \frac{\mathbb{E}[y_i]}{\mathbb{E}[f(\boldsymbol{A}_i, \boldsymbol{\phi}_i, \boldsymbol{X}^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \boldsymbol{w}_I); \boldsymbol{\theta}_i)]}$$

while the estimator of E is

$$\begin{aligned} \hat{E}_{\hat{I}} &= \frac{\sum_{i \in \hat{I}} y_i}{\sum_{i \in \hat{I}} f(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\phi}}_i, \boldsymbol{X}^*(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\theta}}_i, \hat{\boldsymbol{w}}_I); \hat{\boldsymbol{\theta}}_i)} \\ &= \frac{\sum_{i \in \hat{I}} y_i}{|\hat{I}|} \Big/ \frac{\sum_{i \in \hat{I}} f(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\phi}}_i, \boldsymbol{X}^*(\hat{\boldsymbol{A}}_i, \hat{\boldsymbol{\theta}}_i, \hat{\boldsymbol{w}}_I); \hat{\boldsymbol{\theta}}_i)}{|\hat{I}|}. \end{aligned}$$

Since $\hat{I}$ is a random sample and y_i has finite expectation, the numerator of $\hat{E}$ is consistent for $\mathbb{E}[y_i]$ by Kolmogorov's Law of Large Numbers. And since f and $\boldsymbol{X}^*$ are continuous in all their arguments and $\{\hat{\boldsymbol{A}}_{it}\}, \{\hat{\boldsymbol{\phi}}_{it}\}, \{\hat{\boldsymbol{\theta}}_i\}, \hat{\boldsymbol{w}}_{\hat{I}}$ are all consistent for their respective population parameters, the denominator is consistent for $\mathbb{E}[f(\boldsymbol{A}_i, \boldsymbol{\phi}_i, \boldsymbol{X}^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \boldsymbol{w}_I); \boldsymbol{\theta}_i)]$ by Kolmogorov's Law of Large Numbers and the Continuous Mapping Theorem. Then applying the Continuous Mapping

Theorem again, the ratio of the two consistent estimators is consistent for the population ratio. Thus, $\hat{E}_{\hat{I}}$ is consistent for E_I.

Lemma (Price Estimator and GMM Consistency): Consider each of the requirements for consistency in turn.

Consistent Weighting Matrix: There are exactly as many market-clearing conditions as prices, implying the estimator is just-identified and thus the weighting matrix irrelevant.

Global Identification: Recall that the optimal outcome will be identical to the solution to the planner's problem. By the assumption that f satisfies decreasing returns, this will be a strictly concave optimization with a unique maximizer $\{\boldsymbol{X}_i\}$. By (16), the condition $f_{\boldsymbol{X}}(\boldsymbol{X}_i^*) = \boldsymbol{w}$ is satisfied for all i. Since $f_{\boldsymbol{X}}$ is a function, the uniqueness of $\boldsymbol{X}_i^*$ implies the uniqueness of $\boldsymbol{w}$. Thus, $\boldsymbol{w}$ uniquely satisfies the market-clearing conditions.

$\boldsymbol{X}^*(\boldsymbol{w})$ Is Continuous at all w: Observe that $\boldsymbol{X}^*$ is the solution to i's optimization problem, and thus it suffices to show the conditions of the Maximum Theorem hold (see Sundaram, 1996, p. 237). Observe that since f is assumed continuous in $\boldsymbol{X}_i$, the continuity condition is satisfied, so one need only show that the constraint set is a compact-valued continuous function of $\boldsymbol{w}$. The firm is implicitly constrained to choose positive values of all factors, so $\mathbf{0}$ is a lower bound. Meanwhile, since f satisfies decreasing returns, for each w_k there exists some $\tilde{X}_k(w_k)$ such that $f(0, \ldots, \tilde{X}_k(w_k), \ldots, 0) - w_k\tilde{X}_k(w_k) = -100$. Define $\mathcal{W}(\boldsymbol{w}) = \sum_k w_k\tilde{X}_k(w_k)$. Since the firm always can choose zero of all factors and earn profit zero, we can impose that $\boldsymbol{w}^T\boldsymbol{X}_i \leq \mathcal{W}(\boldsymbol{w})$ and the outcome will be identical to that of the unconstrained problem. This "budget constraint" is like any other from consumer theory and thus continuous, and is closed and bounded (thus compact) for all $\boldsymbol{w}$. Thus, $\boldsymbol{X}^*$ is continuous by the Maximum Theorem.

$\boldsymbol{w} \in \boldsymbol{\Theta}$,Which Is Compact: Since $f_{\boldsymbol{X}} > 0$ by assumption, $\boldsymbol{w} >> 0$. Then some $\varepsilon > 0$ exists such that $\boldsymbol{w} >> (\varepsilon, \ldots, \varepsilon)$. Meanwhile, aggregate demand $\mathbb{E}[X_k^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \boldsymbol{w}_I)]$ for any factor k is continuous and strictly decreasing in w_k. Then some $\tilde{\boldsymbol{w}}$ exists such that $\mathbb{E}[X^*(\boldsymbol{A}_i, \boldsymbol{\theta}_i, \tilde{\boldsymbol{w}}_I)] = \mathbb{E}[X_i]/2$, and $\tilde{w}_k > w_k$ for all k. Then $\boldsymbol{w} \in [\varepsilon, w_1] \times \cdots \times [\varepsilon, w_K]$, a closed and bounded subset of $\mathbb{R}^K$, which is thus compact.

$\mathbb{E}[\sup_{w\in\Theta}||\boldsymbol{X}^*(\boldsymbol{w})||] < \infty$ **:** Note that $\boldsymbol{X}^*$ is determined by the satisfaction of (16) (and the non-negativity constraint). Since f_{X_k} is strictly decreasing (by strict concavity) and strictly positive (by assumption), for any finite $\boldsymbol{w}$, either the condition will be satisfied by some finite positive $\boldsymbol{X}^*$ or the non-negativity constraint will bind and $\boldsymbol{X}^*$ will have one or more zero elements. Observe that Θ as defined above is closed and bounded, so any $\boldsymbol{w} \in \Theta$ is finite.

Asymptotic Normality: I again prove the result only for $\hat{E}_{\hat{I}}$; similar algebra and applications of limiting statistics prove the result for the other estimators.[32] Suppress time subscripts for notational simplicity.

$$
\begin{aligned}
\sqrt{|\hat{I}|}(\hat{E}_{\hat{I}} - E_I) &= \sqrt{|\hat{I}|}\left(\frac{\sum_{i\in\hat{I}} y_i}{\sum_{i\in\hat{I}} y_i^*} - \frac{\mathbb{E}[y_i]}{\mathbb{E}[y_i^*]}\right) \\
&= \sqrt{|\hat{I}|}\left(\frac{\mathbb{E}[y_i^*]\sum_{i\in\hat{I}} y_i - \mathbb{E}[y_i]\sum_{i\in\hat{I}} y_i^*}{\mathbb{E}[y_i^*]\sum_{i\in\hat{I}} y_i^*}\right) \\
&= \sqrt{|\hat{I}|}\left(\frac{\mathbb{E}[y_i^*]\bar{y}_i - \mathbb{E}[y_i]\bar{y}_i^*}{\mathbb{E}[y_i^*]\bar{y}_i^*}\right) \\
&= \sqrt{|\hat{I}|}(\mathbb{E}[y_i^*]\bar{y}_i - \mathbb{E}[y_i^*]\mathbb{E}[y_i] + \mathbb{E}[y_i^*]\mathbb{E}[y_i] - \mathbb{E}[y_i]\bar{y}_i^*)(\mathbb{E}[y_i^*]\bar{y}_i^*)^{-1} \\
&= \sqrt{|\hat{I}|}\Big[\mathbb{E}[y_i^*](\bar{y}_i - \mathbb{E}[y_i]) - \mathbb{E}[y_i](\bar{y}_i^* - \mathbb{E}[y_i^*])\Big](\mathbb{E}[y_i^*]\bar{y}_i^*)^{-1} \\
&= \Big[\mathbb{E}[y_i^*]\cdot\underbrace{\sqrt{|\hat{I}|}(\bar{y}_i - \mathbb{E}[y_i])}_{\mathbb{A}} - \mathbb{E}[y_i]\cdot\underbrace{\sqrt{|\hat{I}|}(\bar{y}_i^* - \mathbb{E}[y_i^*])}_{\mathbb{B}}\Big]\underbrace{(\mathbb{E}[y_i^*]\bar{y}_i^*)^{-1}}_{\mathbb{C}}
\end{aligned}
$$

By Kolmogorov's Law of Large Numbers and the Continuous mapping theorem, $\mathbb{C} \xrightarrow{p} \mathbb{E}[y_i^*]^{-2}$. By the Lindeberg-Lévy Central Limit Theorem, $\mathbb{A} \xrightarrow{d} N(0, \sigma_y)$ and $\mathbb{B} \xrightarrow{d} N(0, \sigma_{y^*})$ for some finite σ_y, σ_{y^*}. Then by the Mann-Wald Continuous Mapping Theorem and the replication property of the normal distribution, $(\mathbb{E}[y_i^*]\mathcal{A} - \mathbb{E}[y_i]\mathcal{B})$ is asymptotically normal. Finally, by the Slutsky Transformation Theorem, the product of this term and $\mathbb{C}$ is also asymptotically normal.

[32] For example, the gains from reallocation are actually a continuous function of E: $\hat{\mathbb{G}}_{\hat{I}_t} = \frac{1}{\hat{E}_{\hat{I}}} - 1$. Simply apply the Delta method to prove the asymptotic normality of $\hat{\mathbb{G}}_{\hat{I}_t}$.