

Market Failures and Misallocation: The Costs of Factor and Financial Market Failures in Rural Thailand

Ajay Shenoy*

University of Michigan, Ann Arbor

January 8, 2014

First Version: April 28, 2011

[Job Market Paper]

Abstract

I estimate the contributions of factor and financial market failures to production misallocation among Thai rice farmers. I find surprisingly little misallocation. Optimal reallocation would increase output in most villages by less than 15 percent. Financial market imperfections cause no more misallocation than factor market imperfections, and neither causes much. I derive an aggregate production function that accounts for misallocation and use it to decompose output growth. Declining misallocation contributes little to growth compared to factor accumulation and rising farm productivity. A government credit intervention made statistically significant but small improvements to efficiency, acting almost entirely through financial markets. (JEL O16, O13, O47, O43, D24, E13)

*Address: Lorch 351A, 611 Tappan St., Ann Arbor, MI 48109; Email: shenoy@umich.edu; Special thanks to Raj Arunachalam, Chris House, David Lam, Dan Akerberg, 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, four anonymous referees, and many seminar participants for valuable feedback. I also acknowledge generous financial support from the Thai Studies Grant Program at the Center for Southeast Asian Studies at the University of Michigan. Thanks as well to Robert Townsend and the staff of the Townsend Thai Project for data access.

1 Introduction

Two decades of research in growth have concluded that differences in aggregate resources cannot explain differences in aggregate output, and economists have now asked whether the allocation of aggregate resources can. Hsieh and Klenow (2009), for example, argue India could raise manufacturing output by 40 to 60 percent if it allocated its labor and capital among firms as efficiently as the U.S., and by 100 percent if the allocations were perfect. Yet their study cannot tell where the observed misallocation comes from. Theorists cannot build accurate models and policymakers cannot design useful interventions unless they know what frictions cause misallocation.

I develop a method to measure and separate the misallocation caused by factor and financial market failures, two frictions that set poor countries apart from rich. With perfect markets each firm's optimal allocation depends only on two sets of parameters: its productivity and the common production function. I use a dynamic panel approach 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 exploit how factor markets govern the firm's mix of inputs. When factor markets are perfect, a firm can optimally divide its spending between inputs. By perfecting each firm's mix of inputs while holding its scale constant, I place a lower bound on the aggregate gains from perfecting factor markets, and by then perfecting scale I place an upper bound on the gains from subsequently perfecting financial markets. Finally, I decompose aggregate output into three components: an aggregate production function, average firm-level 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.

My method only works in a setting where factor and financial market failures cause all misallocation, and I study Thailand's rice sector because it fits the assumption well. Since rice production is relatively uniform I can estimate a common production function and correctly calculate each farmer's marginal product, minimizing the spurious misallocation caused by measurement error.

Other sources of misallocation like monopoly, taxes, and adjustment costs are rare in rice farming. It is also plausible that a farmer chooses land, labor, and capital with equal information about her productivity, letting me calculate her ideal mix of inputs, which I use to separate factor and financial market failures.

I find surprisingly little misallocation. The overall cost is 15 percent of output in 1996 and falls to 4 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. As Thailand industrializes, labor and capital have left the rice sector. I find that their departure would have caused rice output to fall if farms had not also grown more productive. Finally, I assess the impact on misallocation of the quasi-experimental Million Baht credit program first studied by Kaboski and Townsend (2011). Additional credit has a statistically significant but small effect. A one percent increase in credit reduced misallocation by .1 percentage points, and as expected the program worked almost entirely by reducing misallocation from financial market failures.

Economists started measuring misallocation in response to evidence that financial markets in poor countries do not allocate capital efficiently.¹ 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 et al., 2008). Until recently, however, few studies tried to quantify the aggregate costs of misallocation, prompting Banerjee and Duflo (2005) to emphasize the question and its implications. Jeong and Townsend (2007) explored the idea with a structural model where credit constraints prevent households from buying capital and switching sectors. Their model reproduces much of the change in Thailand's Solow residual even with zero technological progress. With few exceptions (e.g. Benjamin, 1992), the literature has focused primarily on financial markets—markets for credit and insurance—and paid little attention to factor

¹In 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 et al. (2012); Alfaro et al. (2008); Bartelsman et al. (2009); Jones (2011); Osoimehin (2011); Alfaro and Chari (2012); Moll (2010).

markets, where households hire labor and rent land and capital. But if factor markets cause most misallocation, governments ought to worry less about getting credit to the poor and focus on, for example, eliminating regulations that distort the market for land.

Hsieh and Klenow's study (2009), on the other hand, calculates the cost of all misallocation regardless of the source. They build a model of monopolistic competition where firms pay "taxes" to sell output and rent capital. The authors derive the aggregate gains to reallocating factors within manufacturing industries and plug in production functions calibrated with U.S. data. They argue India and China's distorted factor and financial markets are why they have more misallocation than the U.S., but their approach cannot actually link misallocation to its cause. I too compare the actual allocation to more optimal allocations, but my method and context let me attribute misallocation to factor and financial market failures. Midrigan and Xu (2013) approach a question similar to mine with a different method and find a similar answer. Instead of measuring misallocation caused by financial market failures, they calibrate a model of credit-constrained firms and find it cannot predict much of the variation in the marginal product of capital.

To my knowledge this paper is the first to split misallocation into the contributions of factor versus financial markets, and I find earlier work may have been wrong to ignore factor markets. I also measure misallocation under more plausible assumptions than earlier work. For example, since rice production is relatively uniform I can estimate the production function instead of assuming U.S. parameters describe Thailand's production. In addition, I measure misallocation at each point in time using a household survey, making sample attrition less likely to bias the result. Finally, I link misallocation to its sources under weaker assumptions than earlier work, which uses structural models to simulate how much misallocation would occur if each constraint firms face were eased or tightened. Only if the constraints are modeled correctly can they accurately link misallocation to its causes. In contrast, I do not have to assume I perfectly model constraints to measure misallocation from factor versus financial markets.

2 Misallocation among Thai Rice Farmers

My crucial assumption is that all misallocation in the Thai rice sector comes from factor or financial market failures. Comparing the original allocation of land, labor, and capital to the allocation that equalizes marginal products does not work if marginal products differ for reasons other than factor and financial market failures. 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.

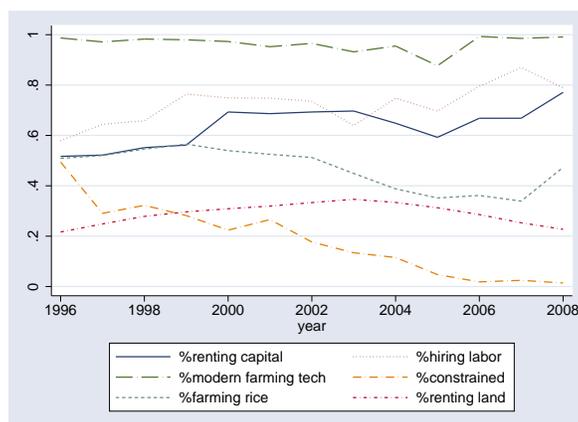
Such issues cause fewer problems in rice farming than in manufacturing. Rice production is relatively uniform and transparent. Though not all farmers in Thailand grow the same type of rice, they grow each variety following a similar technique.² And unlike in many developing countries nearly all farmers 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.³ I can estimate a common production function for all farmers without departing too far from the truth. 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 crop after the harvest—I can model productivity as described in Section 5.2.

Monopoly, taxation, and adjustment costs are not big problems in the rice sector, either. Rice is a commodity and Thai farmers are all price-takers who sell

²The 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.

³Thailand 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.

their output to mills and merchants at market prices. In Shenoy (2013) I show that farmers' selling prices move with the international rice price. Though the government often supports prices, price subsidies will affect all farmers equally and leave allocations unchanged. Thailand's farmers grow rice more commercially than their Indian or Chinese counterparts, but they enjoy a similar lack of 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 factors without adjustment costs is also plausible. Tractors have wheels and bullocks have legs, most on-farm machinery is not too large to move, 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.

Not only are other sources of misallocation absent, but also the assumptions I make to calculate the optimal allocation are more plausible among Thai rice farmers. Two of those assumptions are that productivity is Hicks-Neutral and technology is Cobb-Douglas. Productivity need not be Hicks-Neutral—that is, affect all factors of production equally—in manufacturing or services.

If demand for sneakers rises, Nike's marketing team does more to increase its sales than does the quality of its machinery. Different factors at Nike essentially make different products. Machines and unskilled labor make shoes while skilled workers make a marketing campaign. In rice farming, however, all factors work to make a single product. The farmer's year-to-year productivity shocks—rainfall and pests—will damage the end crop, not the individual contribution of the workers or the tractors. If the farmer falls ill the lost productivity comes from a reduction in her managerial ability, as she can hire others to supply labor. Major advances like the invention of the tractor are unlikely in my thirteen-year sample. The Cobb-Douglas assumption—that the optimal share of expenditure on labor versus capital does not change with scale—is also more plausible because the farmer will farm each additional plot of land much like the previous. I also verify the assumption approximates reality in Appendix A.5.⁴

To split factor market and financial market misallocation I approximate what would happen if each village perfected its factor markets but left its financial markets untouched. Assuming (as I have argued) adjustment costs are small, with perfect factor markets a farmer will always choose the optimal mix of factors regardless of overall expenditure. To calculate the optimal mix I assume the farmer chooses all factors with the same information about within-village productivity shocks. The assumption, which implies all farmers choose the same mix, would be too strong if it applied to all shocks. A farmer does not know how much rain will fall when he chooses how much land to sow, but he does when he hires workers at harvest. Rainfall and other village-level shocks, however, affect all farmers in the village and have no allocative effects. Most within-village shocks, such as when pests eat freshly harvested rice, happen after all factors are chosen or show no effects until harvest. Shocks like illness, however, can strike after renting land but before hiring labor or renting tractors. If such shocks are common I may incorrectly measure the misallocation from factor market imperfections. But illness is not common (see Section 4), and I show in Section 7 that my measures of factor and financial market misallocation re-

⁴The general model and procedure I present in Appendix A.3 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.

spond as expected to a credit intervention.⁵

A common problem in panels of manufacturing firms is that unproductive firms leave the industry and equally unproductive ones enter, but the entrants do not appear in the panel. Over time the sample becomes more productive than the population. Since households still farm most of Thailand's rice, however, I can measure production using a household survey, which follows both farmers and non-farmers. New entrants to farming—households who did not farm at baseline but start farming later on—will enter my sample of farmers at the same rate as the population because the survey was representative at baseline. Migration, which entirely removes households from the survey, is relatively rare. Since the Project sampled new households to replace the ones who left, the sample does not become too unrepresentative. I do not need to assume a stable sample of farmers because exit does not bias my estimate of misallocation. When a farmer exits, he takes his factors out of rice production. The farmers and factors that remain represent the sector's current state, and misallocation is still the output gained from redistributing the factors more efficiently.⁶

3 Measuring Misallocation

I now present a model of household rice production and derive from it my measures of misallocation. In building the model I sacrifice generality in favor of clarity and consistency. The model reflects rice production in Thailand as described in Section 2 and imposes the functional forms I use to compute my measures in Section 6. I describe a more general approach in Appendix A.3.

⁵What if farmers can get credit to purchase some factors but not others? The most obvious example is store credit for seeds and fertilizer. With perfect factor markets a farmer can arbitrage borrowed seeds and fertilizer to acquire land or capital, so my argument about perfect mix still holds. Regardless, few households report borrowing anything on store credit at baseline.

⁶I 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.

3.1 Environment

The farmer maximizes her discounted lifetime utility from consumption. She uses capital K , land T , and labor L to produce farm revenue y . Let \mathbf{X} denote the vector of productive factors chosen. The factors may come from her stock of owned factors \mathbf{X}^o or the factors she rents from factor markets $\mathbf{X} - \mathbf{X}^o$. She buys I_{it} factors, and owned factors depreciate at rates recorded in the matrix δ . The vector of “owned” factors \mathbf{X}^o includes family labor, which I assume exogenous for notational simplicity (making labor decisions endogenous changes nothing because I only care about labor allocations between farms). I implicitly assume family labor and hired labor are perfect substitutes, though relaxing the assumption does not change the result much.⁷ The farmer buys factors (except labor) at a vector of prices p and rents them at prices w . Her output also depends on Hicks-Neutral productivity, part of which she anticipates when choosing factors (A) while the rest (ϕ) is random and unanticipated.

Farmer i faces her own set of prices, a budget constraint, a liquidity constraint, and a factor market constraint. At the beginning of each year she spends z on renting and buying factors. Any spending beyond gross interest on her savings from last year R^b is borrowed at gross rate R^z and repaid after the harvest. She faces a liquidity constraint \bar{z} , which is a function of the owned land she offers as collateral and differs across farmers (perhaps because they differ in how many cosigners or rich relatives they have). She cannot rent in more factors than $\bar{\mathbf{X}}$ or rent out more factors than $\underline{\mathbf{X}}$. Since she lacks perfect insurance her consumption may be correlated with the unanticipated shock ϕ , and the effect on her utility depends on the parameters γ_i . The Lagrange multipliers λ , ω , and $\underline{\kappa}, \bar{\kappa}$ capture the shadow costs of the budget, liquidity, and factor market constraints.

Formally, the farmer’s problem is

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

Subject to:

⁷The sensitivity analysis is available upon request.

$$\begin{aligned}
y_{it} &= A_{it}\phi_{it}K_{it}^{\theta_K}T_{it}^{\theta_T}L_{it}^{\theta_L} && (\text{Production Function}) \\
z_{it} &= \mathbf{w}_{it}(\mathbf{X}_{it} - \mathbf{X}_{it}^o) + \mathbf{p}_{it}\mathbf{I}_{it} && (\text{Expenditure}) \\
\mathbf{X}_{it}^o &= \mathbf{I}_{it} + (1 - \delta)\mathbf{X}_{i,t-1}^o && (\text{Owned Assets}) \\
\lambda : \quad c_{it} + b_{i,t+1} &= y_{it} - R_{it}^z(T_{it}^o)(z_{it} - R_{it}^b b_{it}) && (\text{Budget Const.}) \\
\omega : \quad z_{it} - R_{it}^b b_{it} &\leq \bar{z}_{it} && (\text{Liquidity Const.}) \\
\underline{\kappa}, \bar{\kappa} : \quad \underline{\mathbf{X}}_{it} &\leq \mathbf{X}_{it} - \mathbf{X}_{it}^o \leq \bar{\mathbf{X}}_{it} && (\text{Factor Market Const.})
\end{aligned}$$

3.2 Perfect Choices and Distortions

The farmer's optimal choice of capital satisfies

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

Optimal land and labor choices satisfy similar conditions. When markets are fully perfect the unanticipated shock does not affect consumption ($\mathbb{E}[\lambda_{it}\phi_{it}] = \mathbb{E}[\lambda_{it}]\mathbb{E}[\phi_{it}]$), farmers pay the same borrowing rate ($R_{it}^z(T_{it}^o) = R_t^z$), the liquidity constraint does not bind ($\omega_{it} = 0$), farmers pay the same rental prices ($w_{it}^K = w_t^K$), and factor market constraints do not bind ($\underline{\kappa}_{it} = \bar{\kappa}_{it} = 0$). Then

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

The expression implies marginal products are equalized across all farmers in the village.

Now suppose factor markets are perfect, so $w_{it}^K = w_t^K$, $\underline{\kappa}_{it} = \bar{\kappa}_{it} = 0$, 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:

$$\begin{aligned}\frac{T_{it}}{K_{it}} &= \frac{\theta_T w_t^K}{\theta_K w_t^T} \\ &= \frac{T_{jt}}{K_{jt}} \quad \forall j\end{aligned}\tag{2}$$

The condition implies capital-land ratios are equalized across farmers throughout the village, and by similar logic the land-labor ratios are as well.

3.3 Optimal Allocations

How would perfect markets allocate the aggregate factor stock? Suppress time subscripts for notational simplicity and suppose i is a farmer in village I observed to use $\bar{K}_i, \bar{T}_i, \bar{L}_i$.⁸ The observed 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. With aggregate stocks pinned down, I can ignore the supply side of the market and normalize $R^z = 1$. Use (2) to eliminate T_i and L_i from (1):

$$K_i^* = \left[A_i \left(\frac{\theta_K}{w^K} \right)^{1-\theta_T-\theta_L} \left(\frac{\theta_T}{w^T} \right)^{\theta_T} \left(\frac{\theta_L}{w^L} \right)^{\theta_L} \right]^{\frac{1}{1-\sigma}}$$

Since the aggregate stock does not change, the market-clearing condition is

⁸Since factor markets must clear at any point in time and the conditions for perfect choices (1 and 2) do not involve the past or future, dynamics are irrelevant from here onwards.

$$\begin{aligned}
K_I &= \sum_{j \in I} K_j^* \\
&= \left[\left(\frac{\theta_K}{w^K} \right)^{1-\theta_T-\theta_L} \left(\frac{\theta_T}{w^T} \right)^{\theta_T} \left(\frac{\theta_L}{w^L} \right)^{\theta_L} \right]^{\frac{1}{1-\sigma}} \sum_{j \in I} A_j^{\frac{1}{1-\sigma}} \\
\Rightarrow \frac{K_I}{\sum_{j \in I} A_j^{\frac{1}{1-\sigma}}} &= \left[\left(\frac{\theta_K}{w^K} \right)^{1-\theta_T-\theta_L} \left(\frac{\theta_T}{w^T} \right)^{\theta_T} \left(\frac{\theta_L}{w^L} \right)^{\theta_L} \right]^{\frac{1}{1-\sigma}}
\end{aligned}$$

Plug this back into the individual factor choice to derive the optimal allocations with fully perfect markets:

$$K_i^* = \frac{A_i^{\frac{1}{1-\sigma}}}{\sum_{j \in I} A_j^{\frac{1}{1-\sigma}}} K_I \quad (3)$$

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. Equation 2 gives the optimal mix but not the overall scale of factors for each farmer, so I impose that farmers' purchases in the perfected factor markets add up to the value of their endowments:

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

where K_i^+ is the farmer's new choice of capital while \bar{K} is still her original choice, now treated like an endowment. Any assumption about scale would define the allocation, but this one is appealing because it casts the farmers into an Edgeworth economy where their original factor 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 of factors while leaving her scale fixed. 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 3.4, gains from moving to the computed allocation are a lower bound on true misallocation from imperfect factor markets.

The farmer spends a constant fraction of her wealth on each factor, so the condition for capital is

$$K_i^+ = \frac{\theta_K}{\theta_K + \theta_T + \theta_L} \cdot \frac{w^K \bar{K}_i + w^T \bar{T}_i + w^L \bar{L}_i}{w^K}.$$

The aggregate factor stock does not change, so markets clear as follows:

$$\begin{aligned} K_I &= \sum_{j \in I} K_j^+ \\ &= \frac{\theta_K}{\theta_K + \theta_T + \theta_L} \frac{1}{w^K} \sum_{j \in I} (w^K \bar{K}_j + w^T \bar{T}_j + w^L \bar{L}_j) \\ &= \frac{\theta_K}{\theta_K + \theta_T + \theta_L} \frac{1}{w^K} (w^K K_I + w^T T_I + w^L L_I). \end{aligned}$$

Divide by the analogous condition for land to derive an aggregate version of (2):

$$\frac{K_I}{T_I} = \frac{\theta_K}{\theta_T} \frac{w^T}{w^K}.$$

Substitute this back into the farmer's optimal choice to derive the allocations under perfect factor markets:

$$\begin{aligned} K_i^+ &= \frac{\theta_K}{\theta_K + \theta_T + \theta_L} \cdot \left[\frac{w^K}{w^K} \bar{K}_i + \frac{w^T}{w^K} \bar{T}_i + \frac{w^L}{w^K} \bar{L}_i \right] \\ &= \frac{1}{\theta_K + \theta_T + \theta_L} \left[\theta_K \frac{\bar{K}_i}{K_I} + \theta_T \frac{\bar{T}_i}{T_I} + \theta_L \frac{\bar{L}_i}{L_I} \right] K_I \end{aligned} \quad (4)$$

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

3.4 Costs of Misallocation

For all three scenarios, aggregate output in village I is the sum of each farmer's output: actual output Y_I , output with fully perfect markets Y_I^* , and output with only perfect factor markets Y_I^+ . I use two measures of misallocation: the gains

from perfecting one or more markets, and the fraction of efficient output achieved. The gains from reallocation (or simply “misallocation”) measure how much output a village loses from misallocations. 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 3.5. Define

$$\begin{aligned} G_I &= \frac{Y_I^* - Y_I}{Y_I} & G_I^{FACT} &= \frac{Y_I^+ - Y_I}{Y_I} & G_I^{FIN} &= \frac{Y_I^* - Y_I^+}{Y_I} \\ E_I &= \frac{Y_I}{Y_I^*} & E_I^{FACT} &= \frac{Y_I}{Y_I^+} & E_I^{FIN} &= \frac{Y_I^+}{Y_I^*}. \end{aligned}$$

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 ($G_I = \frac{1}{E_I} - 1$).

The factor market misallocation I compute (G_I^{FACT}) need 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. Since she can allocate each dollar to the factor she needs most, every dollar spent on rice farming yields more revenue than it would with imperfect factor markets. 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.⁹ The only rea-

⁹Capital simply stands in for scale of production. I could have phrased the proposition in

son the assumption might fail is if factor market failures somehow compensate for financial market failures, such as if the farmers who cannot get bank loans can rent land more cheaply than everyone else, 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.¹⁰

3.5 Decomposing Aggregate Output and Growth

Growth accounting traditionally measures changes in per capita output and not per firm output, so 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) and because it is difficult to split aggregate shocks into anticipated and unanticipated parts.¹¹

Then

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

Recall from (3) the optimal factor allocations $K_{it}^*, T_{it}^*, L_{it}^*$ are only functions of the aggregate stocks and relative productivity. So 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.

I can derive 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 κ_{It} be the population of village I as a

terms of land or labor just as easily because the ratios of all factors are fixed by (2).

¹⁰The proof is an equivalency result, so 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 $\mathbb{E}[A_i((\tilde{K}_i^+)^{\sigma} - (K_i^+)^{\sigma})] = 0$, computed gains equal actual gains.

¹¹For 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 and call the result overall productivity.

fraction of the total population of all villages surveyed, and let $\chi_{It} = \frac{\kappa_{It}\mathcal{Y}_{It}^*}{\sum_I \kappa_{It}\mathcal{Y}_{It}^*}$ 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. Trivial algebra shows $E_t = \sum_I \chi_{It}E_{It}$, which means

$$\begin{aligned} \mathcal{Y}_t &= Z_t \cdot E_t \cdot \sum_I \tilde{Z}_{It}\kappa_{It}F(K_{It}, T_{It}, L_{It}) \\ &= Z_t E_t F(\{K_{It}, T_{It}, L_{It}\}) \end{aligned} \quad (5)$$

Since reallocation is within-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 $g_t^{\mathcal{Y}} = g_t^Z + g_t^E + g_t^F$, I can decompose growth in per household rice output into the contributions of improvements in productivity, improvements in factor allocations, and aggregate factor accumulation. By setting g_t^Z and g_t^E to zero, for example, I can examine how output would have grown if the rice sector had made no improvements to productivity or efficiency.

4 Data

I construct my sample from the Townsend Thai Annual Household Survey (1997). The Townsend Thai Project collected a baseline of households from four rural provinces using a stratified design. The Project subsampled one third of the survey villages and resurveyed the sampled households every following year to construct a panel. It later added two more provinces and sampled new households to counter attrition. I use the 1997 through 2009 rounds. The survey response period is June of the previous year through May, so I label the period covered by the 1997 survey as 1996. 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 confirm facts not found in the annual survey. I use district-

level precipitation data computed from the University of Delaware Climactic Project and NASA's Tropical Rainfall Measuring Mission. I also test my micro measures of productivity in the Online Appendix using agro-climactic suitability data from the FAO.

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, and multiply the overall 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 imputed value of rented capital and expenses (including intermediate inputs). I do not compute the value of owned capital using perpetual inventory because households frequently report changes in capital ownership without reporting investment, either because they forget to report the investment or because they receive the new machines as gifts or inheritance. 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 large 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 de-

preciation 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 depreciation 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 depreciation for all types of capital (a user cost).¹² Finally, 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

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. To proxy for malnourishment 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.6.

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. The average household spent only 40 percent of its annual budget on rice, a sign that hunger is not widespread in my sample. Very

¹²Given 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 capital in terms of value. Reweighting using household-specific interest rates would be equivalent to calling a tractor more valuable because the household renting it pays a higher mortgage. Variation in the actual rental price is more problematic, but a relatively small fraction of total capital is rented, making this less of a problem.

¹³The 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	47918.6 (77327.9)	Fraction who report. ..	
Capital	74560.7 (93826.4)	Illness	0.01
Land	19.4 (15.9)	Death in Family	0.01
Labor	181.6 (191.9)	Flood	0.06
Rice Budget Share	0.4 (0.2)	Crop-Eating Pests	0.04
		Bad Rainfall	0.22
		Low Yield for Other Reason	0.21
		Low Price	0.15
Households	775	Observations	6230
Villages	69		

Note: I construct the sample in two steps: I first restrict it to household-years with positive revenue and positive capital, land, and labor. I further restrict the sample to household with at least two years of positive revenue and factors. 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.

few households report illness or deaths in the family relative to the number who report aggregate (rainfall, floods, prices) or completely unanticipated (pests) production shocks. It is reassuring that illness, the type of shock that might undermine my decomposition (see Section 2), is rare.

5 Estimating the Production Function

I cannot use the expressions for optimal land, labor, and capital derived in Section 3 without an estimate of the production function. Estimating a production function is never easy, and misallocations complicate the task because I cannot assume firms choose their inputs optimally. As I explain, however, misallocation enables me to separately identify each factor's production share using a dynamic panel estimator.

5.1 Estimation with Misallocation

I assume in Section 2 that farmers make all their input choices with the same information about within-village productivity. Akerberg et al. (2006) explain why the assumption might make identifying each factor's production elasticity difficult. If a firm using Cobb-Douglas technology selects all its inputs freely and simultaneously they will all be functions of a single variable: anticipated productivity. Equation 3 shows that the farmer's optimal choices of (log) capital, land, and labor are all perfectly correlated. Akerberg et al. solve the problem by assuming the firm chooses factors optimally but at different times.

With imperfect markets the solution is simpler: the firm does not choose factors optimally. If one farmer can rent tractors more cheaply than his neighbor they will use land and capital in different proportions. Even relatively little variation (or misallocation) breaks the perfect correlation between factors. To test whether markets are perfect, take Expression 2 and add a term for measurement error:

$$\begin{aligned}\frac{T_{it}}{K_{it}} &= \frac{\theta_T w_t^K}{\theta_K w_t^T} + \varepsilon_{it} \\ &= v_{It} + \varepsilon_{it}\end{aligned}$$

The land-capital ratio is a function of invariant parameters, noise, and prices that do not vary within a village with perfect factor markets. If markets are perfect it should not be correlated with anything—in particular, the levels of any factor—after controlling for village-year fixed-effects. Table 2 shows that the capital-labor, capital-land, and land-labor ratios are each strongly correlated with the level of at least one factor. I reject the null in five of nine coefficients, which means I can reject perfect markets. Although the markets cause little misallocation (see Section 6), they are imperfect enough to let me identify the production function.

Imperfect markets do, however, rule out a class of methods Akerberg et al. (2006) call "structural techniques." The econometrician assumes the firm chooses intermediate inputs to match its anticipated productivity. All else equal, more intermediate inputs imply higher anticipated productivity. The assumption

Table 2
Test of Imperfect Markets

	(1) $\frac{Capital}{Labor}$	(2) $\frac{Capital}{Land}$	(3) $\frac{Land}{Labor}$
Capital	0.006*** (0.00)	0.089*** (0.02)	-0.000 (0.00)
Land	-5.504*** (1.60)	-299.477*** (45.94)	0.004*** (0.00)
Labor	-1.411** (0.44)	-1.080 (2.21)	-0.000 (0.00)
Observations	6230	6230	6230
Households	775	775	775

Note: Under perfect markets the ratios of factor choices should not be correlated with anything after controlling for village-year fixed-effects. I report the results of such regressions. Standard errors are clustered at the household level.

yields powerful results when satisfied: the econometrician can non-parametrically estimate productivity. It is also why structural techniques fail when factor markets fail. Take a concrete example and suppose two farmers have the same amount of land, workers, machinery, and productivity. One is married to the son of the local fertilizer merchant. She gets a discount on fertilizer and naturally buys more. The econometrician wrongly concludes she is more productive than her clone. The structural techniques fail because they assume the farmer's choices are only a function of her productivity and not her constraints or privileges.¹⁴

5.2 Modeling Productivity

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 2, Thai farmers have used modern seeds, pes-

¹⁴Akerberg et al. would say the “scalar unobservable” assumption fails.

ticides, and fertilizers for decades, so nobody has a technological edge. Land quality does not appear to vary much within a village.¹⁵ Malnourishment might lower the farmer's productivity, and he certainly knows when he is hungry, so 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 anticipated and unanticipated productivity as follow:

$$\log A_{it} = [Household\ Fixed\ Effect]_i + a^H [Hunger]_{it} \quad (6)$$

$$\begin{aligned} \log \phi_{it} = & \sum_j a_j^S [Dummy\ Shocks]_{j,it} + \sum_k a_k^D [District-Year\ Dummies]_{k,it} \quad (7) \\ & + \sum_m a_m^R [Monthly\ Precipitation]_{m,it} + [Overall\ Error]_{it} \end{aligned}$$

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.¹⁶ Why do I call all aggregate shocks and (district-level) precipitation unanticipated? For village-level reallocation the distinction does not matter: a village-level shock will not affect the optimal allocation. But the distinction would change how much growth I credit to anticipated versus unanticipated productivity in the aggregate output decomposition of Section 3.5. I therefore combine both types of productivity into an overall term in the decomposition.

The reader may doubt that I can ever know as much about the farmer's productivity as the farmer himself. In Section 8 I assess how much the main results change when farmers anticipate all productivity.

¹⁵I took the average price per rai as proxy for land quality and included it in unreported production function estimates. The coefficient was small and insignificant, which I take to mean quality does not vary much and price reflects location rather than fertility.

¹⁶Self-reports of bad rainfall and prices probably do not add much information to the district-level variables for precipitation and shocks, but I include them anyways to ensure I correctly estimate the production function.

5.3 Dynamic Panel Estimation

If the bulk of a farmer's anticipated productivity is fixed, it seems natural to estimate

$$\log y_{it} = \log A_{it} + \log \tilde{\phi}_{it} + \theta_K \log K_{it} + \theta_T \log T_{it} + \theta_L \log L_{it} + [Overall\ Error]_{it}$$

with the within-household estimator, where $\log \tilde{\phi}_{it} = \log \phi_{it} - [Overall\ Error]_{it}$. 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 suppose a credit-constrained farmer suffered a bad harvest last year and spent her savings on food, leaving less money to rent land this year. Aside from potentially causing misallocations the situation also violates strict exogeneity.

The Anderson-Hsiao estimator (Anderson and Hsiao, 1981, 1982) can estimate 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. Lagged levels are uncorrelated with the combined error term by sequential exogeneity, so the instruments are valid.

Table 3 reports the Anderson-Hsiao estimates of the production function. As expected, rice farming is relatively labor- and land-intensive. The production function has decreasing returns, and each shock to productivity has the expected sign. The first-stage regressions of factor changes on their lags easily satisfy the usual standards for strength (Stock et al., 2002).

5.4 Sample Characteristics

Table 4 reports sample sizes and median productivity dispersion 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 et al. (2011) find for

Table 3
Production Function Estimates

Production Elasticities		Productivity Modifiers	
Capital		Hunger	-0.032
-Share (θ_K)	0.110		(0.08)
	(0.04)	Illness	-0.147*
-1st Stage F-Stat	111.720		(0.08)
Land		Death	0.029
-Share (θ_T)	0.244		(0.12)
	(0.05)	Flood	-0.101***
-1st Stage F-Stat	128.387		(0.04)
Labor		Pests	-0.057
-Share (θ_L)	0.310		(0.04)
	(0.04)	Bad Rain	-0.131***
-1st Stage F-Stat	133.424		(0.03)
Returns to Scale:		Low Yield	-0.152***
-Estimate (σ)	0.664		(0.02)
	(0.07)	Low Price	-0.065***
			(0.02)
Households:	734	Observations:	4856
R-Squared:	0.760	Cragg-Donald Stat.	204.435

Note: The table reports the Anderson-Hsiao estimates of the production elasticities and the effects of each component of productivity (see Section 5.3 for details). I cluster standard errors by household.

Table 4
Sample Sizes and Productivity in Rice-Farming Villages

Farmers Per Village		Productivity Dispersion	
25th Pctl.	5	75/25	1.81
50th Pctl.	9	90/10	3.09
75th Pctl.	12	95/5	3.64

Table 5
Correlates of Under-Allocated Farmers

Correlates of Being Under-Allocated in 1996						
Age of Head	Rents Land?	Years Farmed Rice	Illness	High Risk Aversion	Cash Savings	
-0.6**	-2.3	-0.5**	-4.6	11.2**	0.0	

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 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. After dropping households for whom all variables are not defined the sample size is 350.

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, most likely because they use value-added rather than gross production functions. Gandhi et al. show that value-added representations require strong assumptions about gross production functions and tend to inflate the dispersion of productivity by an order of magnitude.

6 Results

I plug the estimates of production elasticities and anticipated productivity into the expressions for fully perfect and perfect factor market allocations: (3) and (4) for capital and similar expressions for land and labor. I drop all observations from village-years with only a single farmer because they by construction have no within-village misallocation.¹⁷

Table 5 reports predictors of being under-allocated in the earliest year as per a linear probability regression. I call a farmer under-allocated if he pro-

¹⁷The sample loses 13 households, 8 villages, and 47 observations.

duces more after reallocation than before, and predict under-allocation with the farmer's age and years farming rice, whether he rents land, whether illness lowered his income, his cash savings, and whether he chose the most risk-averse options in two questions that measure risk preferences.¹⁸ Ten more years of age or experience reduce by 6 and 5 percentage points the chance of being under-allocated, suggesting farmers accumulate factors as they age. As expected, a risk-averse farmer is 11.2 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.

I then estimate the misallocation in each village using the expression for G_I in Section 3.4 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 15 percent in most villages. Over time the distribution shifts downward, and misallocation falls below 8 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.

To reduce the noise and represent the whole rice sector, Figure 2.B depicts misallocation across all villages from the original four provinces surveyed at baseline.¹⁹ I estimate the gains in sample-wide output from reallocating factors within each village after weighting by population (see Section 3.5). Sample-wide misallocation is never more than 17 percent and falls to below 4 percent by 2008. The results suggest factor and financial market failures do not produce much costly misallocation, and what little they do produce falls over time.

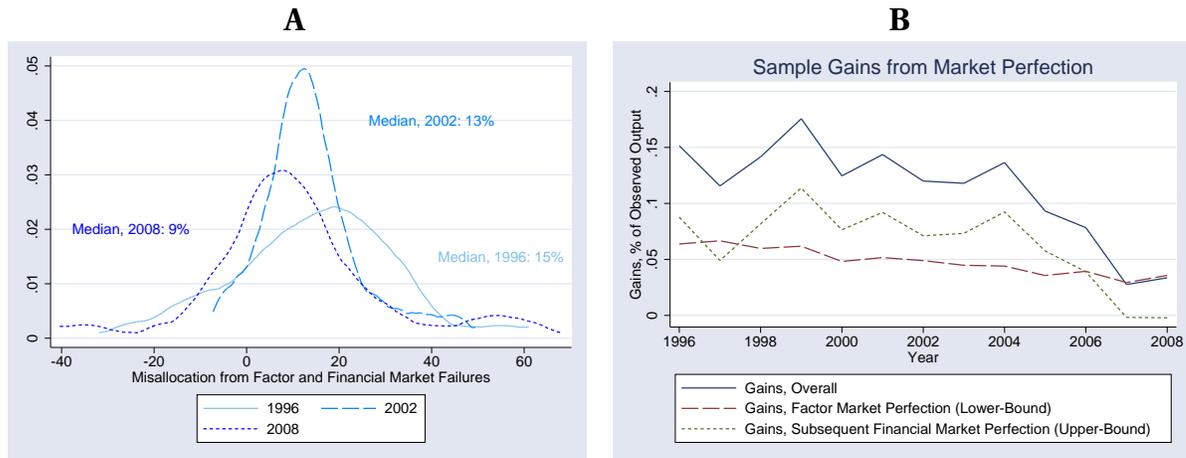
Figure 2.B also separates the misallocation caused by imperfect factor markets from that caused by imperfect financial markets. I calculate G_I^{FACT} and G_I^{FIN} as defined in Section 3.5 and reweight them to the sample-level. Both types of misallocation fall from 1996 to 2008. Since the factor market measure

¹⁸Each 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.

¹⁹I 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



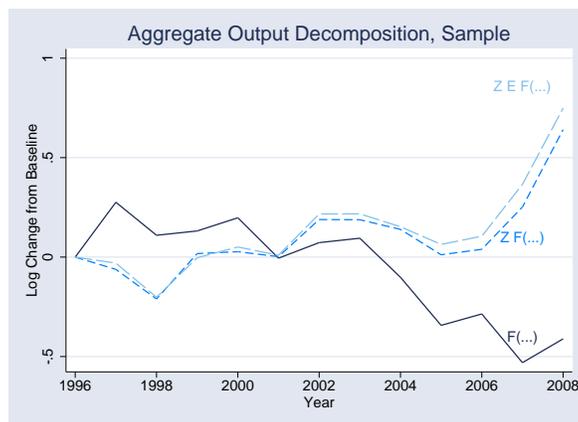
Note: Panel A plots the shifting distribution of misallocation within each village. I report misallocation as the fraction of observed output foregone because factors are misallocated. Panel B calculates the overall cost to the rice sector from misallocation within villages for every year of my sample. 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.

is a lower bound while the financial market measure is an upper bound, neither market unambiguously causes more misallocation until 2006 when financial market misallocation drops to nearly zero. Development economists often blame the financial markets for underdevelopment, but the graph suggests factor markets cause as much or even more misallocation. Like in Figure 2.A, the apparently negative financial market misallocation in 2008 is an artifact of sampling error. If true financial market misallocation is almost zero the probability of estimating it to be negative is large. The estimator for financial market misallocation is also more volatile than the estimator for factor market misallocation, so the apparent contrast between the slow march of factor markets versus the drunken stumble of financial markets may be an artifact.

Figure 3 decomposes growth in the four provinces 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 fac-

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).

tor 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 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. It 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. The final line in Figure 3 shows output when changes in efficiency are turned back on, and comparing it to the middle dashed line shows the contribution of improving efficiency to output growth. It is trivial. Compared to the other two sources of growth, efficiency barely changed the trajectory of rice output.

Why do I find so much less misallocation than earlier work (e.g. Hsieh and Klenow, 2009)? One possibility is that markets in agriculture really are more efficient than in manufacturing. The possibility contradicts popular perception that rural villages are backwards, although Townsend (1994) and Benjamin

(1992) find surprising efficiency in rural insurance and labor markets. Another possibility is that equally bad markets cause more misallocation in manufacturing than in farming because some feature of manufacturing production makes misallocations more costly. One possibility is that productivity is more dispersed in manufacturing, but I show in Section 5.4 that productivity is no less dispersed in a Thai village than a typical manufacturing industry.²⁰ More generally, the distribution of village-level misallocation in Figure 2 shows that some villages have lots of misallocation (nearly 60 percent), demonstrating that costly misallocation is possible in my context; it simply does not happen often. A third possibility is that most misallocation in earlier studies did not come from factor and financial market failures. For example, studies like Hsieh and Klenow's import to poor countries parameters calibrated with U.S. data. If rich country production differs from poor country production and the difference grows with the gap in development, then poorer countries might seem to have more misallocation. That may be why Midrigan and Xu (2013) could not explain much misallocation with financial market constraints. Perhaps little of the misallocation Hsieh and Klenow find in India and China comes from factor and financial market failures, an unexpected result given how much attention they receive in the literature on development.

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

Between May 2001 and May 2002 the Thai government gave the public lending funds of all the villages in my sample one million baht. 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 eco-

²⁰Another possibility is that manufacturing has higher returns-to-scale that make misallocation more costly. Consider a simple example: two firms that produce only with capital. Suppose K_1^* and K_2^* are the optimal levels of capital but x units of Firm 2's capital have been inefficiently allocated to Firm 1. Then efficiency is $E = \frac{A_1(K_1^*+x)^\sigma + A_2(K_2^*-x)^\sigma}{A_1(K_1^*)^\sigma + A_2(K_2^*)^\sigma}$ and $\frac{dE}{dx} = -\sigma \frac{\frac{A_2}{(K_2^*-x)^{1-\sigma}} - \frac{A_1}{(K_1^*+x)^{1-\sigma}}}{A_1(K_1^*)^\sigma + A_2(K_2^*)^\sigma} < 0$ which is negative because the marginal products are equal at the optimum, and Firm 1 has more capital than is optimal while Firm 2 has less. So efficiency falls faster for higher returns-to-scale, and equivalently misallocation rises faster.

conomic 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, and 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 efficiency. By increasing the supply of credit the program improved financial markets and should decrease misallocation from financial market imperfections.²¹ 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 coefficients on the interactions measure the semi-elasticities of misallocation with respect to credit.

Table 6 reports the results, which are rescaled to show the change in the dependent variable due to a one percent increase in per household credit. A one percent increase in credit decreases misallocation by .1 percent of observed output, nearly all of which comes from decreases in financial market misallocation. Since a credit intervention should affect financial markets, the results validate my measures of factor versus financial market misallocation. 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. If the average did not change but misallocation fell, most households must have cut back their scale 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. Improving financial markets that do not cause much misallocation 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 im-

²¹The 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 6
Effects of the Million Baht Credit Intervention

	Overall Cost b/se	Cost, Fact. b/se	Cost, Fin. b/se	Agg. K b/se	Agg. T b/se	Agg. L b/se
Log of Credit	0.002 (0.05)	-0.002 (0.01)	0.003 (0.04)	9.766 (494.52)	0.191*** (0.06)	0.542 (0.72)
2001 X Credit	-0.095** (0.05)	-0.002 (0.01)	-0.093** (0.04)	-396.228 (432.99)	0.114 (0.08)	-0.343 (1.34)
2002 X Credit	-0.028 (0.04)	0.006 (0.01)	-0.034 (0.04)	-221.249 (583.20)	-0.031 (0.09)	-0.147 (1.58)
Year FEs	X	X	X	X	X	X
Village FEs	X	X	X	X	X	X
Villages	65.000	65.000	65.000	65.000	65.000	65.000
Observations	735.000	735.000	735.000	735.000	735.000	735.000

All standard errors clustered by village

Note: The table captures the causal effects of the credit intervention on misallocation. Log of Credit is the log of the per-household credit injection ($\frac{\text{one million}}{\# \text{ of households}}$), 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.

plementation (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 recall that measured misallocation is noisy, making the estimated program effect noisy. The impact could be identical in both years, but the variance of the estimators might make it seem different.

8 Robustness

8.1 Do Farmers Anticipate More Productivity than Assumed?

I must make assumptions about how much productivity the farmer anticipates to avoid blaming factor and financial markets for random events. If I underestimate how much the farmer knows about his productivity, however, I will underestimate the amount of misallocation. I can bound the bias by recalculating misallocation under an extreme assumption: farmers anticipate all productivity. The allocation I calculate by substituting $A_{it}\phi_{it}$ in for A_{it} in expression (3) will produce an unrealistically perfect outcome because it assumes farmers know everything: not just their own managerial talent, but also how much of the harvest rats will eat in three months. Given even the extreme upper-bound on misallocation, do I find as much misallocation as Hsieh and Klenow (2009)?

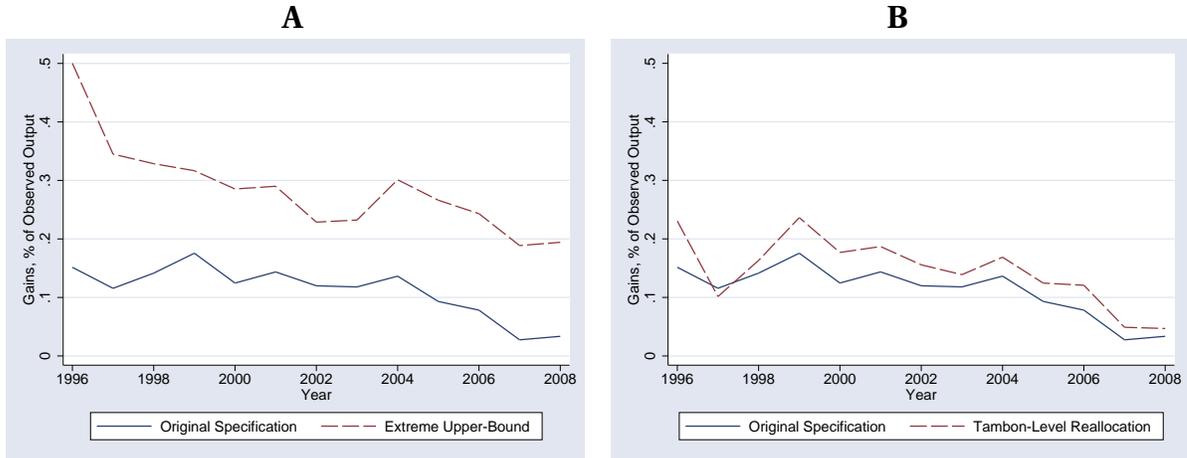
Figure 4.A, which compares the upper bound to my preferred specification, suggests otherwise. Misallocation is higher in all years but still falls short in most years of the 40 to 60 percent Hsieh and Klenow report India could gain from reallocating to U.S. efficiency, much less the 100 percent it could get with perfect efficiency. By the end of the sample misallocation costs less than 20 percent of observed output—half of what Hsieh and Klenow find for the U.S.

8.2 Does a Small Sample Cause Underestimates of Misallocation?

Misallocation often happens because the most talented producers do not get enough factors. If very talented farmers are rare they might not appear in my small per-village samples and I might underestimate misallocation. Even if the social planner favors reallocating everything to a small productive elite, I show

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.

Table 7

Correlation Between Village Misallocation and Sample Size

	(1)	(2)
	Efficiency	Gains
	b/se	b/se
Number of Farmers	0.000 (0.00)	-0.004* (0.00)
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. 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.

in Section 5.4 that the dispersion of productivity in most villages is similar to what Gandhi et al. (2011) find in manufacturing industries. Table 7 reports regressions of efficiency on sample size. Estimated efficiency is no lower in villages with larger samples, and equivalently the gains from reallocation are no higher in villages with larger samples (if anything, they are lower). The results give no reason to fear I would find more misallocation if my sample were larger. In Figure 4.B I recalculate sample-wide misallocation assuming I can reallocate factors 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 forces is small. Reallocating between villages does not much increase misallocation. Sample size is not a problem, and between-village misallocation is small relative even to within-village misallocation.

9 Conclusion

Imperfect factor and financial markets do not cause much misallocation among Thai rice farmers. Of the little they do cause, factor markets cause as much or more than financial markets. Declining misallocation makes positive but trivial contributions to growth in aggregate rice output, whereas most growth is caused by the overall flow of factors out of rice production and an overall increase in average farm productivity. The Thai government's Million Baht credit program did raise efficiency through the expected channel: a statistically significant decrease in misallocation from poor financial markets. But the effects were small, most likely because there was so little misallocation to begin with.

The results do not mean misallocation never matters or market failures have no costs. The assumptions behind my method might hold for other crops and other countries, but that does not mean applying it to India or Africa would reveal as little misallocation as I find in Thailand. It is also possible, as I discussed in Section 6, that factor and financial market failures cause more misallocation

outside agriculture. Finding little misallocation within rice farming is not the same as finding little misallocation between farming and industry, and Jeong and Townsend (2007), Buera et al. (2011), and Midrigan and Xu (2013) seem to find between-sector misallocations do matter. Even among Thai farmers, I find in Shenoy (2013) that risk causes costly under-specialization in economic activities.

What the results do suggest is that evidence of a market failure is not always evidence of costly misallocation. Only by measuring the costs of each type of misallocation caused by each market failure can we know when and where misallocation really matters.

References

- Akerberg, D. A., Caves, K., and Frazer, G. (2006). Structural Identification of Production Functions.
- Alfaro, L. and Chari, A. (2012). Deregulation, Misallocation, and Size: Evidence from India.
- Alfaro, L., Charlton, A., and Kanczuk, F. (2008). Plant-Size Distribution and Cross-Country Income Differences.
- Anderson, T. and Hsiao, C. (1981). Estimation of Dynamic Models with Error Components. *Journal of the American Statistical Association*, pages 598–606.
- Anderson, T. and Hsiao, C. (1982). Formulation and Estimation of Dynamic Models Using Panel Data. *Journal of econometrics*, 18(1):47–82.
- Baker, C. and Phongpaichit, P. (2009). *A History of Thailand*. Cambridge University Press.
- Banerjee, A. and Duflo, E. (2005). Growth Theory Through the Lens of Development Economics. *Handbook of Economic Growth*, 1:473–552.
- Banerjee, A. and Moll, B. (2010). Why Does Misallocation Persist? *American Economic Journal: Macroeconomics*, 2(1):189–206.

- Banerjee, A. and Munshi, K. (2004). How Efficiently Is Capital Allocated? Evidence from the Knitted Garment Industry in Tirupur. *Review of Economic Studies*, 71(1):19–42.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (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, A., Klenow, P. J., and Sharma, G. (2012). India's Mysterious Manufacturing Miracle. *Review of Economic Dynamics*.
- Buera, F., Kaboski, J., and Shin, Y. (2011). Finance and Development: A Tale of Two Sectors. *American Economic Review*, 101(5):1964–2002.
- De Mel, S., McKenzie, D., and Woodruff, C. (2008). Returns to Capital in Microenterprises: Evidence From a Field Experiment. *The Quarterly Journal of Economics*, 123(4):1329–1372.
- Gandhi, A., Navarro, S., and Rivers, D. (2011). On the Identification of Production Functions: How Heterogeneous Is Productivity?
- House, C. L. and Shapiro, M. D. (2008). Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation. *The American Economic Review*, pages 737–768.
- Hsieh, C. and Klenow, P. (2009). Misallocation and Manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4):1403–1448.
- Jensen, R. and Miller, N. (2010). A Revealed Preference Approach to Measuring Hunger and Undernutrition. Technical report, National Bureau of Economic Research.
- Jeong, H. and Townsend, R. (2007). Sources of TFP Growth: Occupational Choice and Financial Deepening. *Economic Theory*, 32(1):179–221.

- Jones, C. I. (2011). Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics*, 3(2):1–28.
- Kaboski, J. and Townsend, R. (2009). The Impacts of Credit on Village Economies. *MIT Department of Economics Working Paper No. 09-13*.
- Kaboski, J. and Townsend, R. (2011). A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative. *Econometrica*, 79(5):1357–1406.
- Levinsohn, J. and Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *Review of economic studies*, 70(2):317–341.
- Midrigan, V. and Xu, D. (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*.
- Osootimehin, S. (2011). Aggregate Productivity and the Allocation of Resources over the Business Cycle.
- Peters, M. (2011). Heterogeneous Mark-Ups and Endogenous Misallocation. *Manuscript, Massachusetts Institute of Technology*.
- Restuccia, D. and Rogerson, R. (2008). Policy Distortions and Aggregate Productivity with Heterogeneous Establishments. *Review of Economic Dynamics*, 11(4):707–720.
- Shenoy, A. (2013). Risk and Economic Under-Specialization: Why the Pin-Maker Grows Cassava on the Side.
- Stock, J. H., Wright, J. H., and Yogo, M. (2002). A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments. *Journal of Business & Economic Statistics*, 20(4).
- Sundaram, R. (1996). *A First Course in Optimization Theory*. Cambridge Univ Pr.

- Townsend, R. (1994). Risk and Insurance in Village India. *Econometrica: Journal of the Econometric Society*, pages 539–591.
- Townsend, R., Paulson, A., Sakuntasathien, S., Lee, T. J., and Binford, M. (1997). Questionnaire Design and Data Collection for NICHD Grant *Risk, Insurance and the Family* and NSF Grants.
- Townsend, R. M. (2012). Townsend Thai Project Monthly Survey (1-60) Initial Release. 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.
- Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. *Journal of Political Economy*, 104(5):1010–1046.
- Wooldridge, J. (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 (2) 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

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

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 $\mathbb{E}_i \left(A_i \left[(K_i^+)^{\sigma} - \tilde{K}_i^+{}^{\sigma} \right] \right) > 0$. 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:

$$\begin{aligned}
\sum_{i \in I} K_i^* &= K_I^* \quad \forall I \\
\sum_I w_I K_I^* &= \sum_I w_I \bar{K}_I
\end{aligned}$$

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

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

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 Technical Appendix: General Model (For Online Publication)

This appendix presents a more general form of the model I use in Section 3 to derive my measures of misallocation, and also derives their asymptotic behavior. 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:

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

Subject to:

$$c_{it} + b_{i,t+1} = y_{it} + [1 + r_{it} + \zeta_{it}(z_{it} - b_{it}; b_{it}, \mathbf{X}_{i,t-1}^o)](z_{it} - b_{it}) \quad (\text{Budget Const.})$$

$$y_{it} = f(\mathbf{A}_{it}, \phi_{it}, \mathbf{X}_{it}; \theta_i) \quad (\text{Production})$$

$$z_{it} = \mathbf{w}_{it}^T (\mathbf{X}_{it} - \mathbf{X}_{it}^o) + \mathbf{p}_{it}^T \mathbf{I}_{it} \quad (\text{Expenditure})$$

$$\mathbf{X}_{k,it}^o = (1 - \delta^k) \mathbf{X}_{k,i,t-1}^o + \mathbf{I}_{k,it} \quad \forall k = 1, \dots, K$$

$$z_{it} - b_{it} \leq \omega_{it}(b_{it}, \mathbf{X}_{i,t-1}^o) \quad (\text{Liquidity Const.})$$

$$\underline{\mathbf{X}}_{it} \leq \mathbf{X}_{it}(\mathbf{X}_{it} - \mathbf{X}_{it}^o, \mathbf{I}_{it}) \leq \overline{\mathbf{X}}_{it} \quad (\text{Factor Choice Const.})$$

where \mathcal{I} is the information set, f a strictly concave decreasing returns revenue

production function, \mathbf{A} anticipated revenue productivity, ϕ unanticipated revenue productivity, θ a vector of production parameters, \mathbf{X} factor levels used in production, c is consumption, b borrowing, r the borrowing rate, $\zeta(\cdot)$ an external finance premium, y_{it} revenue, \mathbf{X}^o owned factors, \mathbf{I} purchase of factors, \mathbf{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 \mathbf{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.²²

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

$$\text{(Law of One Price)} \quad \mathbf{w}_{it} = \mathbf{w}_{I_t}, \mathbf{p}_{it} = \mathbf{p}_{I_t} \forall i \in I_t, \forall k$$

$$\text{(Unconstrained Factor Choices)} \quad \underline{\mathbb{X}}_{it} = \mathbb{X}(-\mathbf{X}_{it}^o, -\mathbf{X}_{i,t-1}^o), \overline{\mathbb{X}}_{it} = \mathbb{X}(\infty, \infty) \quad \forall i \in I_t$$

$$\text{(Perfect Credit Markets)} \quad \omega_{it} = \infty, \zeta_{it}(\cdot) = 0, r_{it} = r_{I_t} \quad \forall i \in I_t$$

$$\text{(Perfect Insurance Markets)} \quad c_{it} \perp \phi_{it} \quad \forall i \in I_t.$$

²²For output taxes, for example, one would simply modify the budget constraint to be $c_{it} + b_{i,t+1} = (1 - \tau_{it})y_{it} + \dots$ 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.

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

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

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

$$\mathbb{E}[f_{X_{k,it}}(\mathbf{A}_{it}, \phi_{it}, \mathbf{X}_{it}^*; \boldsymbol{\theta}_i)] = w_{k,I_t} \quad \forall k \quad (8)$$

$$\int_{i \in I_t} X_{k,it}^*(w_{k,I_t}, \mathbf{A}_{it}, X_{-k,it}^*; \boldsymbol{\theta}_i) di = X_{k,I_t} \quad \forall k \quad (9)$$

Solving these equations solves for the optimal allocation vector $\{\mathbf{X}_{it}^*\}$. 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, $y_{it} = \phi_{it} f(\mathbf{A}_{it}, \mathbf{X}_{it}; \boldsymbol{\theta}_i)$. 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 $\bar{\mathbf{X}}_{it}$ 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:

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

Subject to:

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

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

acterize the unique outcome:

$$\frac{\mathbb{E}[f_{X_k}(\mathbf{A}_{it}, \mathbf{X}_{it}^+; \boldsymbol{\theta}_i)]}{\mathbb{E}[f_{X_j}(\mathbf{A}_{it}, \mathbf{X}_{it}^+; \boldsymbol{\theta}_i)]} = \frac{w_{k,I_t}}{w_{j,I_t}} \quad \forall k, j \quad \forall i \quad (10)$$

$$\mathbf{w}_{I_t}^T (\mathbf{X}_{it} - \bar{\mathbf{X}}_{it}) = 0 \quad \forall i \quad (11)$$

$$\int_{i \in I_t} X_{k,it}^+(w_{k,I_t}, \mathbf{A}_{it}, X_{-k,it}^+; \boldsymbol{\theta}_i) di = X_{k,I_t} \quad \forall k \quad (12)$$

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 $y_{it} = A_{it}\phi_{it}f(\mathbf{X}_{it}; \boldsymbol{\theta}_i)$, f is homogeneous of degree σ , $\boldsymbol{\theta}_i = \boldsymbol{\theta}_{I_t} \quad \forall i \in I_t$, 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 $\ddot{\mathbf{X}}_{it}^+ = \ddot{\mathbf{X}}_{I_t}^+$ and $\bar{\mathbf{X}}_{it}^+ = \bar{\mathbf{X}}_{I_t}^+$ 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

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

By homogeneity,

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

Since f satisfies strictly decreasing returns, \mathbf{X}_{it} is unique and thus the $\ddot{\mathbf{X}}_{it}$ 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 $\ddot{\mathbf{X}}_{it} = \ddot{\mathbf{X}}_{I_t}$ for all $i \in I_t$. A similar argument shows that $\tilde{\mathbf{X}}_{it} = \tilde{\mathbf{X}}_{I_t}$ for all $i \in I_t$. ■

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

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

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

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

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

implying $\ddot{\mathbf{X}}_{I_t}^+ = \tilde{\mathbf{X}}_{I_t}^+$. ■

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

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

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,

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

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

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

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 $\{A_{it}, \phi_{it}, \theta_i\}$ are available. Then

it is a numerical exercise to solve the sample analogs of (8) and (9) for estimates of the CCM allocations and (10), (11), and (12) 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 \mathbf{A}_{it} , ϕ_{it} , θ_i , y_{it} , \mathbf{X}_{it} are finite. Then the estimators are all consistent and asymptotically normal.*

Proof:

Consistency: I will prove the consistency of $\hat{E}_{\hat{I}}$; 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 (8) characterizes any interior solution to the population optimization - in other words, if an interior solution exists, the function $\mathbf{X}^*(\mathbf{A}_i, \theta_i, \mathbf{w}_I)$ 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 \mathbf{X}^* is a continuous function of the population prices \mathbf{w}_I (see the Lemma below). Inspection demonstrates that the optimality conditions are identical for the sample optimization, and thus the derived \mathbf{X}^* is as well.

Applying the measure-one normalization, (9) in population reduces to

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

Note that the sample analog of (9) in random sample \hat{I} reduces to the sample analog of this moment condition trivially:

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

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

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

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

while the estimator of E is

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

Since \hat{I} is a random sample and y_i has finite expectation, the numerator of $\hat{E}_{\hat{I}}$ is consistent for $\mathbb{E}[y_i]$ by Kolmogorov's Law of Large Numbers. And since f and \mathbf{X}^* are continuous in all their arguments and $\{\hat{\mathbf{A}}_{it}\}, \{\hat{\boldsymbol{\phi}}_{it}\}, \{\hat{\boldsymbol{\theta}}_i\}, \hat{\mathbf{w}}_{\hat{I}}$ are all consistent for their respective population parameters, the denominator is consistent for $\mathbb{E}[f(\mathbf{A}_i, \boldsymbol{\phi}_i, \mathbf{X}^*(\mathbf{A}_i, \boldsymbol{\theta}_i, \mathbf{w}_I); \boldsymbol{\theta}_i)]$ 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 $\{\mathbf{X}_i\}$. By (8), the condition $f_{\mathbf{X}}(\mathbf{X}_i^*) = \mathbf{w}$ is satisfied for all i . Since $f_{\mathbf{X}}$ is a function, the uniqueness of \mathbf{X}_i^* implies the uniqueness of \mathbf{w} . Thus, \mathbf{w} uniquely satisfies the market-clearing conditions.

$\mathbf{X}^*(\mathbf{w})$ Is Continuous at all \mathbf{w} : Observe that \mathbf{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 \mathbf{X}_i , the continuity condition is satisfied, so one need only show that the constraint set is a compact-valued continuous function of \mathbf{w} . The firm is implicitly constrained to choose positive values of all factors, so 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 $f(0, \dots, \tilde{X}_k(w_k), \dots, 0) - w_k \tilde{X}_k(w_k) = -100$. Define $\mathcal{W}(\mathbf{w}) = \sum_k w_k \tilde{X}_k(w_k)$. Since the firm always can choose zero of all factors and earn profit zero, we can impose that $\mathbf{w}^T \mathbf{X}_i \leq \mathcal{W}(\mathbf{w})$ 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 \mathbf{w} . Thus, \mathbf{X}^* is continuous by the Maximum Theorem.

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

$\mathbb{E}[\sup_{\mathbf{w} \in \Theta} \|\mathbf{X}^*(\mathbf{w})\|] < \infty$: Note that \mathbf{X}^* is determined by the satisfaction of (8) (and the non-negativity constraint). Since f_{X_k} is strictly decreasing (by strict concavity) and strictly positive (by assumption), for any finite \mathbf{w} , either the condition will be satisfied by some finite positive \mathbf{X}^* or the non-negativity constraint will bind and \mathbf{X}^* will have one or more zero elements. Observe that Θ as defined above is closed and bounded, so any $\mathbf{w} \in \Theta$ is finite.

Asymptotic Normality: I again prove the result only for \hat{E}_j ; similar algebra

and applications of limiting statistics prove the result for the other estimators.²³ Suppress time subscripts for notational simplicity.

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

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, $(\mathbb{E}[y_i^*] \mathbb{A} - \mathbb{E}[y_i] \mathbb{B})$ is asymptotically normal. Finally, by the Slutsky Transformation Theorem, the product of this term and \mathbb{C} is also asymptotically normal.

A.4 More Details on Estimating the Production Function (For Online Publication)

This appendix explains why I use the Anderson-Hsiao estimator instead of fixed-effects and shows measurement error in the factors of production might not bias my estimates of the production function too badly under plausible assumptions.

²³For example, the gains from reallocation are actually a continuous function of E : $\hat{G}_{\hat{I}_t} = \frac{1}{\hat{E}_{\hat{I}_t}} - 1$. Simply apply the Delta method to prove the asymptotic normality of $\hat{G}_{\hat{I}_t}$.

4.1 Fixed-Effects Versus Anderson-Hsiao

Farmers have individual invariant productivity terms, so the standard approach is to estimate the production function

$$y_{it} = A_{it}\phi_{it}K_{it}^{\alpha}T_{it}^{\beta}L_{it}^{\lambda}$$

in logs using fixed-effects (the within estimator). In other words, sub (6) and (7) into the logged production function, and estimate by OLS

$$\begin{aligned} \log y_{it} = & [Household\ Fixed\ Effect]_i + a^H[Hunger]_{it} + \sum_j a_j^S[Dummy\ Shocks]_{j,it} \\ & + \sum_k a_k^D[District - Year\ Dummies]_{k,it} + \sum_m a_m^R[Monthly\ Precipitation]_{m,it} \\ & + \theta_K \log K_{it} + \theta_T \log T_{it} + \theta_L \log L_{it} + [Overall\ Error]_{it}. \end{aligned} \quad (13)$$

However, the consistency of the estimates requires what Wooldridge (2002) calls the strict exogeneity assumption:

$$\mathbb{E}[[Overall\ Error]_{it} | K_{it}, T_{it}, L_{it}] = 0 \quad \forall s, t = 1, \dots, T$$

The unexplained productivity component $[Overall\ Error]_{it}$ must be completely cross-temporally uncorrelated with the regressors of interest, in particular the productive factors.

To see why this might be problematic, consider a farmer who gets a bad unobserved shock because he accidentally used too much fertilizer. Then although the accident will not affect contemporaneous or past factor input choices (because it is unpredictable), it might affect future input choices if the resulting low yield drives him to sell land or capital for food. A drop in the possession of these factors would be informative about an earlier bad productivity shock. If the unobserved shock is big relative to household income then OLS will estimate production shares inconsistently.

Instead I make a weaker assumption of sequential exogeneity Wooldridge (2002). Sequential exogeneity only requires that factor levels be uncorrelated

with future values of unexplained productivity:

$$\mathbb{E}[[\text{Overall Error}]_{is} \mid K_{it}, T_{it}, L_{it}] = 0 \quad \forall t = 1, \dots, T \text{ and } s \leq t$$

A bad shock last year can affect capital stock this year as in the previous example without damaging consistency. Problems only arise if farmers anticipate some part of the error term, like a technological innovation they acquire but I do not observe. If farmers buy more land or capital to exploit the innovation the estimator will not be consistent. Large technological innovations are unlikely in Thailand and affect everyone, so the time and time-district dummies should capture them. Assuming sequential exogeneity is the same as assuming \hat{A}_{it} completely captures the anticipated component of TFP. I must assume as much, anyways, for the optimal allocation I compute to be accurate.

To implement the assumption I first-difference away the fixed component and instrument the differenced factors with their lagged levels. The first stage of the regression for capital is

$$\begin{aligned} \Delta \log K_{it} = & +b^H \Delta[\text{Hunger}]_{it} + \sum_j b_j^S \Delta[\text{Dummy Shocks}]_{j,it} \\ & + \sum_k b_k^D [\text{District} - \text{Year Dummies}]_{k,it} + \sum_m b_m^R \Delta[\text{Monthly Precipitation}]_{m,it} \\ & + \nu_K \log K_{it} + \nu_T \log T_{it} + \nu_L \log L_{it} + \Delta[\text{Overall Error}]_{it}. \end{aligned} \quad (14)$$

and the second stage is

$$\begin{aligned} \Delta \log y_{it} = & a^H \Delta[\text{Hunger}]_{it} + \sum_j a_j^S \Delta[\text{Dummy Shocks}]_{j,it} \\ & + \sum_k a_k^D [\text{District} - \text{Year Dummies}]_{k,it} + \sum_m a_m^R \Delta[\text{Monthly Precipitation}]_{m,it} \\ & + \theta_K \Delta \log K_{it} + \theta_T \Delta \log T_{it} + \theta_L \Delta \log L_{it} + \Delta[\text{Overall Error}]_{it}. \end{aligned} \quad (15)$$

The identification assumption is

$$\mathbb{E}[(\log K_{i,t-1})[Overall Error]_{i,t-1}] = \mathbb{E}[(\log K_{i,t-1})[Overall Error]_{i,t}] = 0$$

and so on for land and labor, as well.

4.2 Measurement Error in Factor Choices

Given how I construct the measures of capital and labor, measurement error is likely. But under some assumptions about the nature of the error it should not bias estimates too badly.

Suppose the error ς follows an autoregressive process, relates the true value to the observed value multiplicatively, and is independent of the true value. For simplicity, focus on capital and ignore TFP modifiers as well as labor and land; this simplification should not qualitatively affect the argument. Formally, let K_{it}^o be the observed value of capital and suppose

$$\begin{aligned}\log y_{it} &= c_i + \alpha \log K_{it} + \varepsilon_{it} \\ K_{it} &= \exp(\varsigma_{it}) K_{it}^o \\ \varsigma_{it} &= \rho \varsigma_{i,t-1} + \xi_{it}\end{aligned}$$

Assume that ρ is close to 1 and ξ_{it} is small; that is, the error is highly persistent over time. This assumption is plausible because of the way I calculate factor stocks. A household's level of capital, for example, is computed by backward depreciating the most recent purchase of each asset type and aggregating. Any error introduced by associating latest purchase price with productive value would be carried through to previous years.

The primary regression equation is then

$$\begin{aligned}\Delta \log y_{it} &= \alpha \Delta \log K_{it} + \Delta \varepsilon_{it} \\ &= \alpha \Delta \log K_{it}^o + \Delta \varepsilon_{it} - \alpha(1 - \rho)\varsigma_{i,t-1} + \xi_{it}.\end{aligned}$$

The first-stage regression is

$$\begin{aligned}\Delta \log K_{it} &= \pi_K \log K_{i,t-1} + \epsilon_{it} \\ \Rightarrow \Delta \log K_{it}^o &= \pi_K \log K_{i,t-1}^o + \epsilon_{it} + \pi_k \varsigma_{i,t-1}.\end{aligned}$$

The estimate of π_k will be attenuated because of the correlation between $K_{i,t-1}^o$ and the lagged measurement error terms. But this will not matter as long as the first-stage remains strong, and the results suggest it is.

More important is that the lagged level of capital is correlated with $-\alpha(1 - \rho)\varsigma_{i,t-1} + \xi_{it}$. But under the assumptions, this lagged error term is small (because $1 - \rho$ and ξ are small), so the bias of the IV estimator will be small as long as the first-stage is strong.

The intuition is simple: when the measurement error is persistent, it is captured by the fixed-effect so the parameter estimates are unbiased. What then happens to the estimates of efficiency when these fixed-effects make up some portion of anticipated TFP? Assuming the measurement error is independent of the true fixed component of anticipated TFP, it will almost certainly bias estimated misallocation upwards. To see why, assume the extreme case where markets are complete and contingent, so the observed outcome is perfectly efficient. My noisy estimates of TFP will produce a noisy computation of the efficient outcome; this will differ from the realized outcome, so the measure will erroneously mark the perfectly efficient allocation as inefficient. The true misallocation in rural Thailand may be even lower than I find.

Although measurement error in capital probably will be very persistent, that need not hold for labor. If a household member changes his actual time in the fields often the error in measured labor might fluctuate from year to year. The reader should keep that caveat in mind while evaluating the results.

A.5 Specification Tests (For Online Publication)

This appendix runs simple tests of whether the production function is approximately isoelastic (Cobb-Douglas) and whether my measure of anticipated productivity actually captures the farmer's ability to farm.

Table A.I
CES Production Function Estimates

	(1)
	NLS
	b/se
ϵ	1.013 (0.0068)
σ	0.767 (0.0408)
α	0.311 (0.0086)
β	0.369 (0.0119)
ϑ	0.000 (0.0000)
N	775.000
NT	6230.000
Pval: $\epsilon = 1$	0.061

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

5.1 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

$$y_{it} = A_{it}\phi_{it} \left[\alpha K_{it}^{\frac{\epsilon-1}{\epsilon}} + \beta T_{it}^{\frac{\epsilon-1}{\epsilon}} + (1 - \alpha - \beta)L_{it}^{\frac{\epsilon-1}{\epsilon}} \right]^{\sigma \frac{\epsilon}{\epsilon-1}} \quad (16)$$

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.²⁴

Column 1 of Table A.I reports the results of estimating the transformed equa-

²⁴The ideal way to estimate the equation is the nonlinear equivalent of Anderson-Hsiao: applying GMM to the first-differenced form of (16) using legged factors as instruments. Unfortunately, it does not converge.

tion with nonlinear least squares. The test of interest is whether ϵ differs substantially from 1. As (1) of Table A.I indicates, this null is actually rejected. However, the point estimate is almost identical to one ($\hat{\epsilon} = 1.013$) 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.

5.2 Is A_{it} Really a Valid Measure of Productivity?

Anticipated productivity A_{it} almost wholly determines a farmer's allocation with perfect markets. If \hat{A}_{it} is a bad measure of true productivity, the optimal allocations I compute could be completely wrong. If \hat{A}_{it} actually measures a farmer's productivity it should be correlated with individual and village characteristics that make a farmer more productive.

For example, if \hat{A}_{it} is capturing actual productivity it should be higher in areas with climate and soil better suited to growing rice. And if years spent in school raise productivity, it should also be higher for more educated farmers. Table A.II shows a regression of anticipated TFP \hat{A}_{it} on province dummies and two proxies: an index of agro-climactic suitability for rice and the farmer's schooling. For schooling I focus on primary education, as basic skills like literacy matter most in farming. The results indicate that TFP is indeed strongly positively correlated with climactic suitability, and also correlated with the farmer's education.

A.6 Data Appendix (For Online Publication)

6.1 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

Table A.II
Partial Correlations of Anticipated TFP

	(1) log \hat{A} b/se
Primary Schooling (years)	0.058*** (0.02)
Beyond Primary	0.026 (0.08)
Rice Suitability	0.082*** (0.02)
N	775.000
NT	6216.000

Note: The regression shows the correlation between predictors of productivity and my estimates of anticipated productivity. I cluster standard errors by household and include province dummies. Primary Schooling (years) represents the number of grades of primary school the household head has completed, where someone who has completed any grades in secondary school (Matthayom) is assumed to have completed six years of primary school (Prathom). Beyond Primary is a dummy for additional schooling, and Rice Suitability a district-level agro-climactic index.

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.

6.2 Household-Level Variables

6.2.1 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 depreciation 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, sub-district, 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 in-

terest 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.

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

6.2.3 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.