STEG WORKING PAPER

MISALLOCATION AND PRODUCT CHOICE

Stepan Gordeev and Sudhir Singh

APRIL 2024
STEG WP094

Misallocation and Product Choice

Stepan Gordeev* Sudhir Singh†

April 3, 2024

click for latest version

Abstract

We study the costs of misallocation of inputs between multi-product firms that endogenously choose among heterogeneous products. Misallocation of inputs between firms has been shown to be a significant drag on aggregate productivity: it is especially severe between farms in the agricultural sectors of low-income economies. Existing estimates of its costs have relied on models of single-product firms using a single aggregate production function. Using rich farm-crop-level data from India, we estimate product-level production functions and find that they are meaningfully different from one another and from the aggregate one. We build a general equilibrium model of firm-level misallocation in which multi-product firms (or farms) are able to choose the set and mix of heterogeneous products. Misinterpreting product heterogeneity as evidence of distortions and missing the endogenous product choice response to real distortions biases single-product models to overstate misallocation, while ignoring returns-to-scale heterogeneity and within-firm productivity dispersion biases them to understate it. On net, the single-product model understates the aggregate productivity cost of misallocation between Indian farms by 28%.

Keywords: misallocation, multi-product firms, agriculture, productivity, distortions

JEL classification: O4, O11, O13, Q1

We would like to thank Yan Bai, George Alessandria, Mark Bils, Sandro Steinbach, Armen Khederlarian, and Marcos Mac Mullen for their suggestions. We are also grateful to the participants of the Stockman Conference, Midwest Macro, and seminars at the University of Connecticut for their comments and suggestions.

*University of Connecticut, stepan.gordeev@uconn.edu.

†Michigan State University, singhsud@msu.edu.

1 Introduction

Misallocation of production inputs between firms is an important driver of cross-country income differences (Hsieh and Klenow 2009; Restuccia and Rogerson 2017). Various market frictions and distortionary government policies prevent inputs from being allocated to their most productive uses: some productive firms would prefer to hire more land, labor, or capital, while some less productive ones would be happy to supply these inputs. However, both parties find themselves unable to do so due to imperfect markets or institutional obstacles, limiting the aggregate productivity of the entire economy. Misallocation is especially severe between farms within the agricultural sectors of low-income countries: allowing farmers to efficiently trade land could triple the aggregate agricultural productivity in the most distorted markets (Chen, Restuccia, and Santaeulalia-Llopis 2022). Some of the estimated aggregate cost of misallocation can be explained my mismeasurement and model misspecification (Bils, Klenow, and Ruane 2021; Gollin and Udry 2021; Asker, Collard-Wexler, and De Loecker 2014). At the same time, a large portion of this cost reflects real malign frictions: poor land property rights enforcement, inefficient land distribution institutions, distortionary government policies, poor market access, and more (Gottlieb and Grobovšek 2019; Chen 2017; Krishnaswamy 2018; Le 2020; Morando 2023). Understanding both the costs and the causes of misallocation is crucial for understanding the obstacles to aggregate development, especially in agriculture, where cross-country productivity differences are particularly stark (Gollin, Lagakos, and Waugh 2014).

The misallocation literature has focused on studying single-product firms. The literature on agricultural misallocation has further modeled all farms as producing the same homogeneous agricultural good. Farms in the data, however, have wildly different selections of crops: we find that even the most commonly grown crop in India is produced by only one-quarter of all farmers. Moreover, two-thirds of Indian farmers grow two or more different products within a year.

In this paper, we study the aggregate costs of misallocation of inputs between multi-product farms that choose among heterogeneous products. When products have different production functions, they will be exposed differentially to misallocative market distortions. This causes conventional single-product models to overestimate the frictions present in the economy and to underestimate the ability of productive farms to expand output if frictions were removed.

We explore these mechanisms in rich farm-crop-level data from India. We document significantly heterogeneous product choice across farmers, across seasons, and frequent multi-product behavior by individual farmers. Next, we estimate crop-specific production functions using instruments derived from shocks to other plots within a farm to address the simultaneity bias. We find that many crops have significantly different input elasticities from one another—and from the traditional aggregate agricultural good production function—suggesting that products in the data are meaningfully heterogeneous for the purpose of quantifying misallocation.

We build a general equilibrium model of heterogeneous firms facing misallocative distortions. The main novelty of the model is that firms (or farms) choose among multiple products (or crops) with heterogenous production functions and can elect to produce any number of products at once. The model provides a mapping from the observable farm-crop-level input and output choices to the unobservable fundamental distortions that are the model's stand-in for all kinds of institutional and market frictions. Once these distortions are recovered from the data, the model provides a quantification of the aggregate cost of misallocation induced by the frictions.

The multi-product model highlights four mechanisms by which conventional single-product models misestimate the cost of misallocation when the data is generated by firms producing heterogeneous products. First, single-product models erroneously perceive heterogeneous product choice as evidence of underlying frictions even when the heterogeneity is driven purely by differential productivity draws: this leads single-product models to *overstate* the cost of misallocation in the data. Second, single-product models ignore firms' endogenous product choice response to the frictions they face. In the presence of misallocative distortions and a choice between products with heterogeneous production functions, model firms are able to partly reduce their exposure to distorted input markets by shifting their product choice away from goods intensive in the relatively distorted input toward goods intensive in the relatively undistorted one. This endogenous response reduces the aggregate cost of misallocation for given fundamental distortions: because single-product models are blind to this response, they *overstate* the cost of misallocation. Third, single-product models ignore productivity heterogeneity between products within a single firm, understating the underlying productivity dispersion: this leads them to *understate* the cost of misallocation. Finally, single-product models ignore the returns-to-scale heterogeneity between products, missing that, if frictions were removed, some firms would be able to scale production up more easily than others and consumers would be able

to substitute toward high returns-to-scale crops to take advantage of this difference: this leads them to *understate* the cost of misallocation.

We map the firm-product-level model to Indian farm-crop-level data. We conduct quantitative reallocation exercises in both the multi-product and the single-product models to evaluate the net effect of the four mechanisms on aggregate misallocation within India's agricultural sector. In our preferred calibration, the single-product model *understates* the cost of misallocation of inputs across farms by 28%: accounting for product heterogeneity and multi-product behavior, an efficient reallocation of land, labor, and intermediate inputs across farms promises an almost four-fold increase in agricultural productivity, compared to the three-fold increase estimated by the single-product model. This effect is mainly driven by the returns-to-scale mechanism, which overpowers the mechanisms that push the single-product model in the direction of overstatement. However, the sign and the magnitude of the single-product model's net error depend on the exercise being conducted. A partial reallocation that would bring the severity of distortions across India down to the levels of the least-distorted state would increase India's agricultural productivity by 35% according to the multi-product model: the single-product model *overstates* this figure by a quarter, with the bulk of the error explained by its erroneous view of heterogeneous product choice as evidence of frictions.

Accounting for the effect of product heterogeneity and multi-product behavior on the cost of misallocation is thus not a matter of applying a simple correction to the results of conventional single-product models: the severity and the direction of the single-product model's error will depend on the exercise of interest. When most estimated dispersion is seen to reflect real frictions that ought to be removed, the resulting expansive estimate of misallocation will be understated by the single-product model due to its missing heterogeneous returns to scale allowing some farms to expand output more easily. But when most estimated dispersion between farms is seen to be not feasible or even desirable to remove, the resulting conservative estimate of misallocation will be overstated by the single-product model due to its treatment of product heterogeneity as evidence of frictions.

Estimating the cost of agricultural misallocation in farm-level data is the subject of a growing literature. The foundational result on the severity of misallocation in the agriculture of low-income countries was established by Chen, Restuccia, and Santaeulalia-Llopis (2022). Some papers focus on identifying a particular institutional source of distortions in the data and then using a model to quantify the aggre-

gate cost of those distortions (Chen 2017; Gottlieb and Grobovšek 2019; Adamopoulos et al. 2022). Others seek to capture all sorts of market frictions and government distortions generally, and make methodological contributions to estimating their cost (Gollin and Udry 2021; Aragón, Restuccia, and Rud 2022; Ayerst, Brandt, and Restuccia 2023). Our paper contributes to the literature by departing from its assumption of a single homogeneous agricultural product. We estimate crop-specific production functions and build a model of multi-product farms endogenously choosing among those crops. We thus contribute both to the question of measuring distortions in the data and to the question of quantifying the aggregate cost of these distortions.

The effects of government policy or technological changes on country-level crop allocation have been explored before. Blanco and Raurich (2022) show that falling capital costs cause the aggregate allocation to shift toward capital-intensive crops. Le (2020) finds that land use policies in Vietnam generate aggregate misallocation by favoring rice over other crops. Krishnaswamy (2018) documents that Indian price support for staple crops distorts the market in favor of rice and wheat. We contribute to this literature by exploring the interplay between distortions and product choice not at the aggregate level, but at the level of individual farmers. The model we devise then allows us to map the micro-level distortions and farm responses to aggregate productivity: this mapping would apply to any individual friction the literature has studied before. A related contribution is by Ayerst, Brandt, and Restuccia (2023), who incorporate crop choice into their model of farm-level misallocation with farm dynamics. However, they model single-product farms producing the same homogeneous agricultural good: crops are heterogeneous only in their productivities and the productivity-maximizing choice is prevented only by direct land use restrictions. Our model instead incorporates multi-product farms choosing between crops with heterogeneous production functions, opening up several channels of interaction between product choice, estimated distortions, and resulting misallocation.

Multi-product firms have been discussed in the manufacturing misallocation literature by Jaef (2018) and Wang and Yang (2023). In both papers, firms face distortions and can choose the number of products to manufacture among varieties that have the same production function but different productivities. Both papers then calibrate their models to match cross-sectional distribution moments in the data. First, we contribute to this literature by exploring product choice between products not just with different productivities but also with different input elasticities: it's the latter that ends up driving the differential predictions between the single- and the multi-

product model in our setup. Second, we contribute by extending Hsieh and Klenow (2009)'s framework of mapping firm-level observables to model-implied firm-level fundamentals to the case of multi-product firms. This allows us to reproduce each observed firm (or farm) with its model analog exactly, without having to calibrate distributions of firms to aggregate moments.

Our contribution on the interaction between misallocation and endogenous product choice is not conceptually restricted to agriculture. However, there are several reasons why agriculture is the perfect setting to begin studying this interaction. First, firm-product-level (or farm-crop-level) inputs and outputs are far more feasible to measure in agriculture than in other sectors. Farmers are better able to estimate the amount of land, labor, and intermediate inputs applied to each crop. Such data is also more commonly collected: we use India's REDS, but multiple other farm-level surveys gather crop-level information. Second, products (crops) are far more homogeneous in agriculture than in other sectors. Outside of very specific sub-sectors, varieties produced by different manufacturers operating in the same sub-sector can be vastly different. In agriculture, however, specimens of the same crop grown on different farms can be reasonably approximated as homogeneous. This facilitates the task of estimating product-level production functions and exploring systematic differences in input intensities between products. For these reasons, we choose to focus our initial exploration of endogenous product choice on agriculture. But the mechanism applies more broadly, and insights gleaned from the agricultural setting of a developing country will apply also to the manufacturing sectors of more developed countries.

The paper proceeds as follows. Section 2 discusses the farm-level data we use and the institutional context of Indian agriculture. Section 3 documents the heterogeneous product choice among Indian farmers. Section 4 estimates crop-level production functions. Section 5 describes the multi-product firm model of misallocation. Section 6 quantifies misallocation and compares it across models. Section 7 concludes.

2 Background and Data

2.1 Background

Indian agriculture is very diverse and one of the largest in the world in terms of arable lands.[1] Around half of Indian agriculture is rainfed, although the share of the irrigated area has improved over the years. The agricultural sector employs 58 percent of households and is a major source of income and livelihood for the majority of rural households (Gulati and Juneja 2022). It is worth noting that women make up 30 percent of total cultivators and 41 percent of agriculture laborers. However, women only manage and operate 14 percent of the total operational holding (Chanana-Nag and Aggarwal 2020).

Land market and tenancy structure. The Indian agricultural system was a zamindari system during the colonial era, wherein lands were in the hands of the few who had the right to lease them out and collect revenue (Banerjee and Iyer 2005). However, after 1947, the government of India brought several land and tenancy-related reforms that changed the structure of ownership and leasing rights. These reforms varied in scale, timing, and implementation across Indian states (Deininger, Jin, and Nagarajan 2008).

The majority of agricultural land in India is privately owned by households. Smallholder farming is widespread. In 1999-2000, the average farm size was 5.3 acres per household, which declined to 4.4 acres in 2007-08 and further to 2.6 acres in 2015-16. This decrease in land holding size can be attributed to family splintering (Andrew D Foster and Mark R Rosenzweig 2002). Land is seldom sold and bought: more than 94 percent of land owners inherited their land. Furthermore, only 12 percent of agricultural cultivators participate in the land rental market by renting land either in or out.[2] Low participation in the land market hinders productivity as it limits the reallocation of land to productive farmers (Deininger et al. 2003; Bolhuis, Rachapalli, and Restuccia 2021).

Agricultural seasons. There are three major agricultural seasons in India, namely Kharif (monsoon), Rabi (spring/winter), and Zaid (summer/dry). The Kharif sea-

1. India possesses more than 150 million hectares of arable land and comes second after the USA.
2. Source: Rural Economic and Demographic Survey (REDS) 2007-08.

son typically spans from May/June to early October, during which crops like paddy (rice), maize, sugarcane, jute, cotton, groundnut, and turmeric are cultivated. The Rabi season takes place during winter, with crops being sown between October and December and harvested from March to April. Major crops of the Rabi season are wheat, barley, peas, gram, and mustard. Finally, the Zaid season, the shortest of the three, occurs from March/April to May/June; fruits and vegetables are grown during this season.

2.2 Data

Our primary data source is the Rural Economics and Demographic Survey (REDS), which was conducted by India's National Council of Applied Economic Research in 1971, 1981, 1999, and 2007-08.[3] The first round was collected in 1971. On each subsequent round, all households (or their descendants) living in the same village were surveyed again, along with a random selection of new replacement households (Vashishtha 1989).[4] The survey is nationally representative of rural India.

For this analysis, we use REDS 2007-08 (hence REDS07), covering a sample of 8,659 households across 242 villages in 17 major states in 2007/08. REDS07 collects detailed information on agricultural inputs and outputs at the level of plot-season-crops. We restrict our attention to this round because it is the only one in which agricultural data was collected at the plot level, which is necessary for our production function identification strategy. Household and village-level modules of the survey were administered during 2007-08. Of the surveyed households, 4,803 cultivated land: we treat each such household as a farm and restrict our sample to them. In total, these farmers cultivated 10,318 plots.

3 Heterogeneous Product Choice by Indian Farms

Crop choice is heterogeneous. Indian farmers vary greatly in which crops they choose to produce. Figure 1 shows the 10 most commonly produced crops in India. Rice is the most popular crop—yet it is grown only by 26% of farmers. Wheat is a close second, but all other crops are far behind: Indian farmers have little overlap in

3. Prior to 1981-82, the survey was known as the Additional Rural Income Survey (ARIS).

4. In 1971, the sampling frame included one district per state in the Intensive Agricultural District Program (IADP) and a random sample of other districts.

crops they choose to grow.[5]

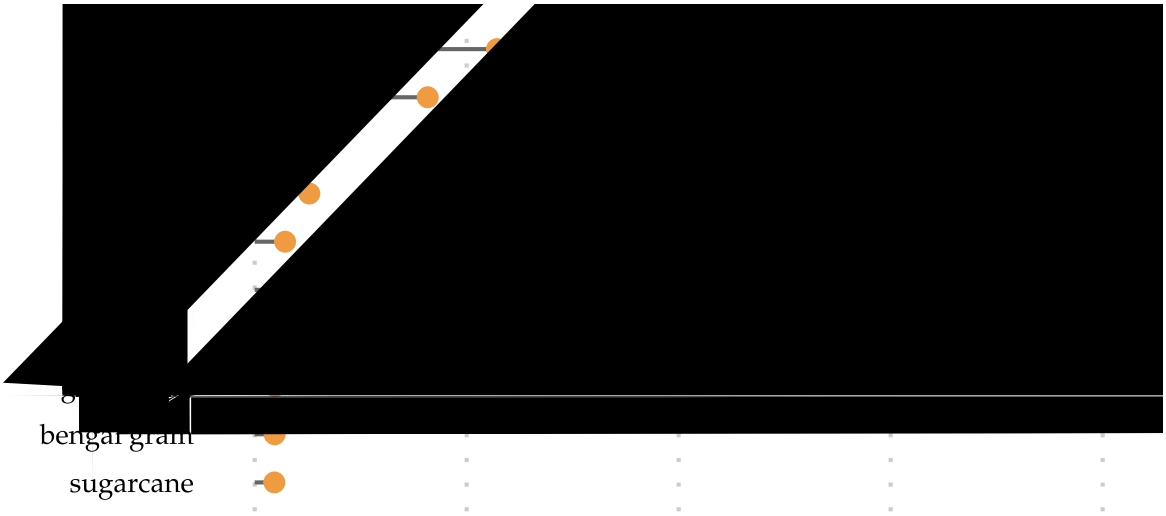

FIGURE 1: Top 10 crops by % of farm-seasons growing them

MANY FARMS GROW MULTIPLE PRODUCTS. Figure 2 displays the share of output produced by farms that grow different numbers of crops. The top bar counts crops that were grown by each farm within the survey year, even if the crops were grown in different seasons: over 80% of output is produced by farmers that grow multiple crops within a year. The bottom bar of Figure 2 treats farm-season combinations as separate entities, thus only counting crops that were all grown in the same season: two-thirds of output is produced by farmers that grow nothing else in the same season, and one-third is produced by multi-crop farmers. Farmers are more specialized within each season than across, although the extent of multi-crop behavior is significant in both cases.

These results have two potential implications for modeling and quantifying misallocation. First, they suggest that production should be split by season when quantifying misallocation, since crops grown in different seasons may be meaningfully different and would affect misallocation estimates if this heterogeneity is not taken into account. Second, the magnitude of multi-crop production is non-negligible both

5. Appendix Figure A.1 shows that the ranking of crops by their share in aggregate output has a similar shape, although the order of crops does change.

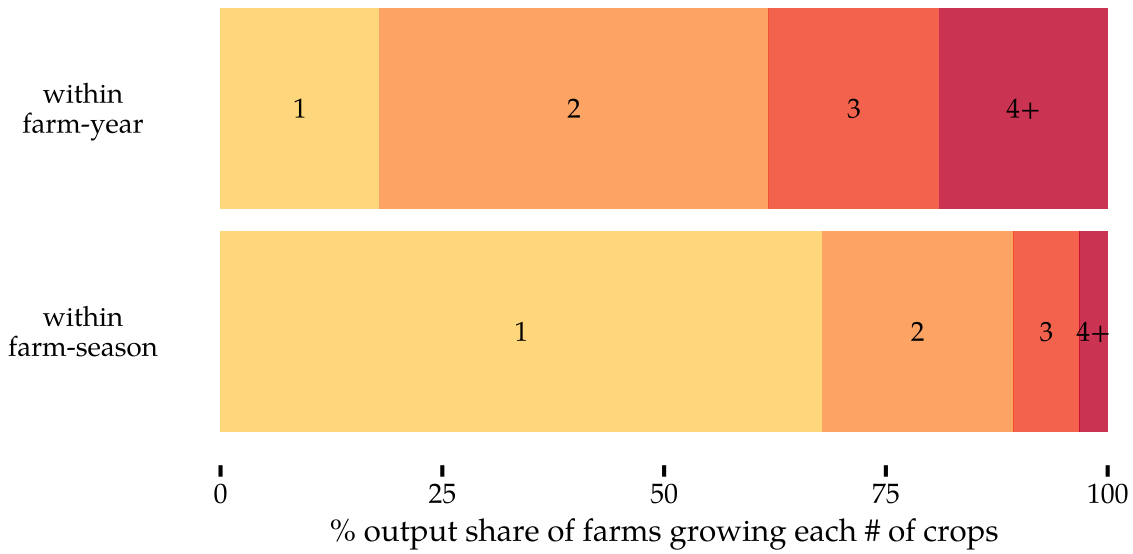

FIGURE 2: Share of output produced by farms growing each number of crops (within a year or within each season alone)

within and across seasons, suggesting that not only modeling heterogeneous crop choice but also the simultaneous cultivation of multiple crops may be important.

At the same time, the apparent heterogeneity of product choice across Indian farmers, across seasons, and the presence of multi-cropping may be of little consequence for estimating misallocation if different crops have similar production functions and would thus interact with distorted input markets the same way. In the next section, we estimate crop-specific production functions to test whether this is the case.

4 ESTIMATING PRODUCT-SPECIFIC PRODUCTION FUNCTIONS

The preceding section documented that Indian farmers engage in heterogeneous multi-product behavior. In this section, we investigate whether these products are meaningfully different in their production functions.

4.1 Specification

To estimate input elasticities, we estimate the following crop-specific Cobb-Douglas production function:

$$y_{f,i,t} = z_{f,i,t} l_{f,i,t}^{\gamma_i} x_{labor,f,i,t}^{\alpha_{labor,i}} x_{inter,f,i,t}^{\alpha_{inter,i}}$$

$y_{f,i,t}$ is the gross physical output of farm f for crop i in season t, weighted with market prices (this weighting is irrelevant at the estimation stage).[6] l is the land input with crop-specific elasticity γ_i. The land input is adjusted for quality, as described in Section 4.3 below. x_{labor} is the labor input with elasticity $\alpha_{labor,i}$, measured in the number of days worked, including both hired and family labor. x_{inter} is the intermediate input with elasticity $\alpha_{inter,i}$, measured as to

price of inputs on plot j, providing variation in plot-level input prices that can be used as instrumental variables for input allocations on plot j. The identification of the production function rests on two assumptions: first, that observed shocks to farm f's plot k affect the shadow price of inputs on f's plot j; second, that observed shocks affecting input demand on plot k are not correlated with unobserved shocks affecting input demand on plot j, conditional on observed shocks to j. The survey collects information on a wide array of agricultural, health, and social shocks at the household level. Gollin and Udry (2021) posit that these shocks will have different effects on individual plots depending on the physical features of each plot. Therefore, the household-level (or farm-level) shocks interacted with observed objective plot-level soil features provide the necessary plot-level variation, which serve as instruments for plot-level input choices. Of the three inputs, land deserves special treatment due to its centrality to agricultural production: we turn to it next.

4.3 Land Input Measure

There are two commonly used ways of measuring the quantity of the land input in agricultural survey data similar to India's REDS. The first is to measure it with the physical area of cultivated land. This quantity can be measured very precisely and objectively but lacks any information on the quality of land. The second is to measure it with the reported market price of cultivated land, which should capture quality differences between plots. However, land markets in India (and many other developing countries) are extremely under-developed, and so the market price reported by farmers is likely to be exceptionally noisy and mismeasured: in the absence of a market, farmers' ability (and need) to estimate the market price of their land is likely to be poor.

Instead of following either of these two approaches, we settle on a compromise solution that attempts to combine the best of both. We measure the land input on each plot with its physical area weighted by a quality index: price per acre predicted by a statistical model using only the observed objective characteristics of the plot as predictor variables. The objective characteristics measured in the data and included as predictors are various qualities of soil (its type, color, salinity, depth, percolation, and ease of drainage) and objective measures of irrigation access (presence and features of nearby irrigation wells and canals).

The task of constructing the land quality is effectively a regression. Instead of an

ordinary least squares regression (OLS), however, we employ random forests (RF): a supervised machine learning algorithm developed by Breiman (2001). The random forest model "grows" a large number of decision trees, with each tree using a random bootstrap sample of the data, and aggregates the predictions of individual trees.[8] Unlike OLS, RF does not require the econometrician to pre-specify the interactions and non-linearities expected in the data: it searches for them on the fly and adapts to what it finds in the data. Furthermore, random forests tend to outperform OLS and many other machine learning techniques in out-of-sample predictive performance, reducing the danger of overfitting the noise in reported market price (Varian 2014; Mullainathan and Spiess 2017).

To evaluate both methods, we estimate an OLS regression and an RF regression using a training sample that randomly selected two-thirds of our data, with the remaining one-third reserved for the test sample. Both methods regress log price per acre on the observed physical characteristics of the plot. The OLS regression includes a complete set of two-way interactions. The RF's parameters m_{try} and *nodesize* were tuned with k-fold cross-validation (with $k = 3$), although tuning provides only a small improvement in predictive performance relative to conventional defaults.[9]

Once the two methods have been estimated, we evaluate their performance on the reserved testing sample: results are presented in Table 1. The random forest has both a lower mean squared error and a higher R^2 than OLS. Due to RF's better out-of-sample performance, we elect to use it to construct the land quality index: predicted price per acre computed for the whole sample.

Because the random forest captures over half of the observed variation in reported price using nothing but observed objective plot features, it likely captures much of the underlying variation in land quality. Because we used predicted prices and took several steps to minimize overfitting (bootstrapping within RF, k-fold cross-validation, and train/test data split), this captured variation is driven by population-level relationships between plot characteristics and market price, rather than the noise in any

8. See Biau and Scornet (2016), Schonlau and Zou (2020), or Ziegler and König (2014) for excellent reviews.

9. Each tree within a random forest is constructed by repeatedly finding the predictor variable and its cutoff value at which to split the data into two branches that maximize the goodness of fit of the branching. m_{try} sets the number of candidate predictors randomly drawn to be considered for branching at each node when constructing each tree. *nodesize* sets the cutoff number of observations per node at which further splitting stops. The third free parameter in a random forest is the number of trees: we set it high enough that further increases do not improve predictive performance. See any of the RF reviews cited above for details on their operation and the role of these parameters. See Bischl et al. (2021) for an overview of k-fold cross-validation and other tuning techniques.

TABLE 1: Out-of-sample performance of OLS and RF in predicting land price

	Ordinary Least Squares	Random Forest
Mean Squared Error	0.61	0.49
R^2	0.39	0.51

Note. The table presents two metrics of out-of-sample performance of Ordinary Least Squares (OLS) and Random Forest (RF) regressions in predicting log price per acre of a plot with observed physical characteristics of the plot. Two-thirds of plot-level observations were used for estimation and one-third was reserved as a test sample to compute the presented out-of-sample performance metrics.

individual farmer's measurement.

4.4 PRODUCTION FUNCTIONS ARE HETEROGENEOUS ACROSS PRODUCTS

Complete results from production function estimation are shown in Appendix Table A.1. Figure 3 visually summarizes the estimated production function elasticities, including 95% confidence intervals. The first panel displays the input elasticities when a single production function is estimated on all crops pooled together: the commonly used approach. In this case, the results are in line with the literature's findings: land elasticity γ is the highest at 0.43, intermediates elasticity α_{inter} is slightly lower at 0.35, labor elasticity α_{labor} is 0.19. The sum of the three elasticities is 0.97, insignificantly different from 1, indicating that the production function exhibits slightly decreasing or roughly constant returns to scale.

As the following panels indicate, however, crop-specific production functions are heterogeneous across crops. Rice is significantly more land-intensive than the aggregate production function would indicate. Wheat is more intermediates-intensive than rice or the aggregate function, and its returns to scale are increasing. Cereals other than wheat are more in line with the aggregate elasticities but have significantly decreasing returns to scale. For pulses, all three inputs are similarly important. Vegetables, fruits, and oils are more labor-intensive than any other crop, and exhibit decreasing returns to scale.

Pairwise comparisons of a given input's elasticity across crops indicate that many pairs of crops have significantly different elasticities from one another: see Figure 4. For each input elasticity, there are $\binom{5}{2} = 10$ crop pairs that can be formed out of five crops. Out of $3 \times 10 = 30$ such combinations across three input elasticities, a third

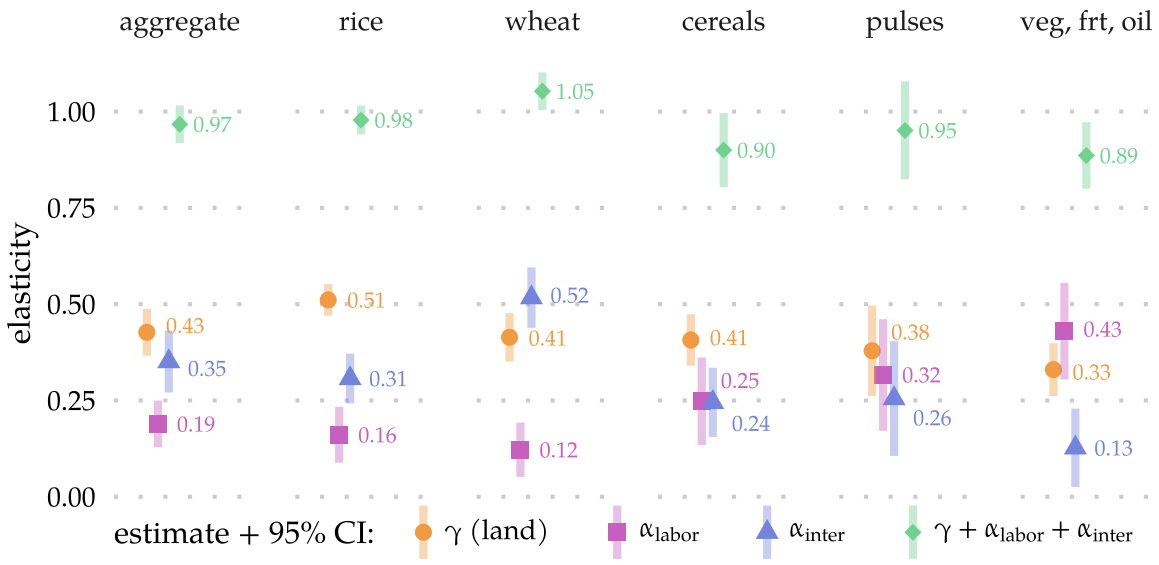

FIGURE 3: Estimated production function input elasticities

Note. The figure visualizes the estimated input elasticities reported in Table A.1. γ is the land, α_{labor} the labor, and α_{inter} the intermediates elasticity. The sum of elasticities $\gamma + \alpha_{labor} + \alpha_{inter}$ reflects the returns to scale.

are significantly different at the 1% level and half at the 10% level. Production functions are heterogeneous across products, and a single aggregate production function provides only a poor approximation of any one individual crop. To quantify the relevance of this heterogeneity for quantifying the cost of misallocation, we turn to the model.

5 MODEL

The model pursues three objectives. The first is to provide a framework for thinking about the production decisions of multi-product farmers in the presence of misallocative distortions. The second is to provide a mapping from observable farm-level input and output choices to unobserved distortions these farms must be facing. The third objective—and the ultimate goal of the paper—is to quantify the aggregate cost of misallocation induced by distortions and to tease out the role of product choice in explaining this cost.

To achieve these objectives, we build a model of farm-level misallocation in which

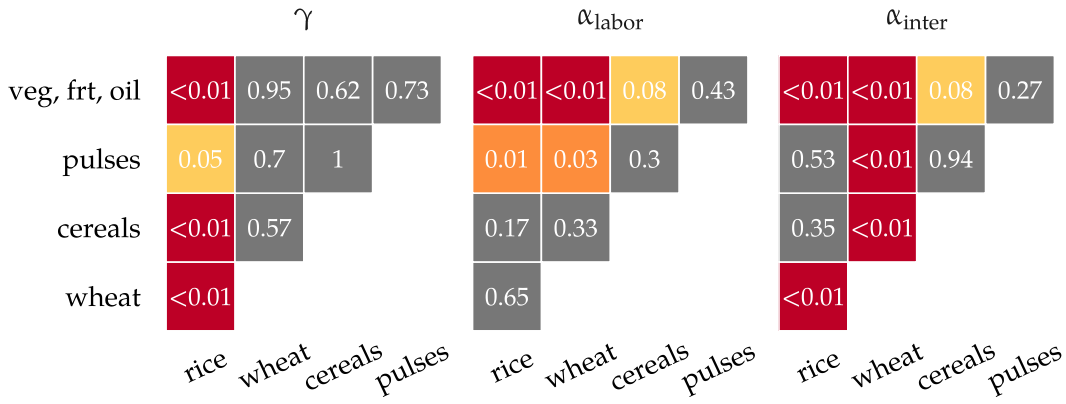

FIGURE 4: P-values of pairwise coefficient comparisons across crop-level production functions

Note. P-values are obtained from tests of equality of input elasticities between each pair of crops in Appendix Table A.1. Color-coded significance levels for rejecting coefficient equality within each crop-pair-input: $p < 0.1$, $p < 0.05$, $p < 0.01$.

multi-product farmers make production decisions over heterogeneous products. It builds upon the single-product firm model of Hsieh and Klenow (2009).[10] Because their model was applied to data on manufacturing firms, it assumed monopolistic competition and a constant-returns-to-scale production function. Because agricultural markets are populated with small producers growing fairly homogeneous crops, we assume perfect competition instead, in line with the simpler single-product agricultural misallocation framework of Chen, Restuccia, and Santaeulalia-Llopis (2022). Instead of assuming constant or decreasing returns to scale for all crops, we rely on our estimated returns to scale for each crop.

5.1 FARM PROBLEM

The economy is populated by F heterogeneous farmers making production decisions over N heterogeneous crops.

Consider the problem of a profit-maximizing farm f. The farm produces quantity $y_{f,i}$ of each crop i and sells it at market price p_i. The farm has idiosyncratic productivity $z_{f,i}$ in each crop. It chooses how much of each input to hire or allocate to each crop.

10. Which in turn is closely related to Restuccia and Rogerson (2008).

Inputs. Inputs $g = 1, \ldots, G$ are flexible: the farm can hire each at market rate r_g. The flexible input elasticity $\alpha_{g,i}$ is input-crop-specific. In our calibration, there are two flexible inputs ($G = 2$): labor and intermediates. The farm chooses quantity $x_{f,g,i}$ of each input g to allocate to each crop i.

Land input l is available to the farm in fixed supply L_f: the farm needs to decide how to split this endowment between available crops, assigning $l_{f,i}$ to each. The crop-specific land elasticities are γ_i. There are three reasons for modeling land as a fixed input. The first is that it is a reasonable approximation of the state of the land market in India. Land is rarely purchased, sold, or rented: it is usually simply inherited. The second is that having one input be in fixed supply generates interdependent crop production: a change in parameters of crop i will lead to a change in the shadow cost of the fixed input (land) and thereby impact the input and output decisions the farm makes on crops $j \neq i$.[11] The third is that an input in fixed supply justifies the production function identification assumptions we made in Section 4: land l and its shadow cost λ_f provides the link through which observable shocks to one plot can affect the shadow cost of inputs on other plots.

Inputs are combined with Cobb-Douglas production technology to produce each crop i: $y_{f,i} = z_{f,i} l_{f,i}^{\gamma_i} \Pi_g \left(x_{f,g,i}^{\alpha_{g,i}} \right)$.

Distortions. Misallocative frictions in this class of models are captured by distortion terms $\tau_{f,g}$. These are idiosyncratic to the farm and operate as taxes or subsidies. If farm f's $\tau_{f,g}$ is higher for input g than for input h, the farm is facing stronger frictions in input g's market, and will hire a relatively higher quantity of h instead. If farm f's frictions $\tau_{f,g}$ for all inputs are higher than those of farm e, then farm f will hire less of every input and will operate at an inefficiently small scale compared to farm e. This simple framework has been used in the misallocation literature to capture all kinds of market frictions and government distortions that firms may be suffering from. Any variation in input ratios or firm scales that is unexplained by the observed production functions and productivities can be replicated by these distortion terms.

Furthermore, once we add multi-product firms into the model, we have to add another layer of distortions. $\tau_{f,g,i}$ and $\tau_{f,l,i}$ are farm-input-crop-specific distortions that can fit unexplained variation in input ratios between crops *of a single farm*.

11. The importance of modeling interdependent crop production with a fixed input rather than any alternative mechanisms is discussed in Just, Zilberman, and Hochman (1983) and Shumway, Pope, and Nash (1984).

Finally, while land input l does not have an explicit distortion term like $\tau_{f,l}$ in the model, land is still de facto distorted since its quantity is fixed. Any variation in the shadow cost of land λ_f between farms reflects implicit τ-like frictions that prevent the farmers from operating an efficient amount of land. As long as land is not distributed in precisely the right way to equalize λ_f between farmers, it is distorted, with the dispersion of λ_f reflecting the severity of distortions.

Additional concavity. When crop-specific production functions have decreasing returns to scale, model farms choose to grow multiple crops at once. However, most of the crop-specific production functions estimated in Section 4.4 exhibit returns to scale that are close to constant, generating little multi-cropping in the model unless frictions are extreme: production concavity alone struggles to explain the extent of multi-product behavior in the data. To allow for a range of other mechanisms that may incentivize farmers to grow multiple crops, we raise the crop-level revenue within the farmer's objective function to some power $\eta \leq 1$. When $\eta < 1$, this generates additional concavity in crop-level revenue. Introducing this parameter allows the model to parsimoniously capture mechanisms such as risk (farmers multi-crop to hedge against uncertain revenue yields in each individual crop), diet diversification by subsistence farmers (subsistence farmers grow multiple crops to satisfy their love of variety), and market power (farmers with some market power in their local market face downward sloping demand for each crop, generating revenue that is concave in inputs). Appendix B shows that the latter two mechanisms, if modeled explicitly, generate a farm's problem that is exactly equivalent to the $\eta < 1$ case. Another reason to allow for $\eta < 1$ is that it takes the burden of generating the observed distribution of farm-crop revenues away from τ frictions, leading to more conservative misallocation estimates, as Section 5.4 discusses below.

Fixed cost. For each crop i that the farm chooses to produce ($y_{f,i} > 0$), it pays a fixed cost ω. Without this feature, farmers would produce all possible crops, with some of them grown in negligible quantities: introducing the fixed cost makes such marginal crops unprofitable and leads the farm to drop them entirely. We introduce the fixed cost to reproduce the observed distribution of the number of crops grown per farm, discussed in Section 3.

Problem. The complete problem of the farm is:

$$\max_{\{y_{f,i},l_{f,i}\}_i,\{x_{f,g,i}\}_{g,i}} \sum_{i=1}^{N} \left(p_i y_{f,i}\right)^{\eta} - \sum_{g=1}^{G} r_g \tau_{f,g} \sum_{i=1}^{N} \tau_{f,g,i} x_{f,g,i} - \sum_{i=1}^{N} \omega \cdot 1\left[y_{f,i} > 0\right] \tag{2}$$

s.t.

$$y_{f,i} = z_{f,i} l_{f,i}^{\gamma_i} \Pi_g \left(x_{f,g,i}^{\alpha_{g,i}}\right) \tag{3}$$

$$\sum_{i=1}^{N} l_{f,i} \tau_{f,l,i} = L_f \qquad (\lambda_f) \tag{4}$$

The solution and the computational algorithm are discussed in Appendix C.1.

5.2 General Equilibrium

Consumer. To close the model, we introduce a representative consumer who has constant elasticity of substitution preferences over N crops with elasticity of substitution σ and crop-specific taste weights φ_i. The consumer purchases quantity C_i of each crop from farms at price p_i. The consumer owns flexible inputs and rents out the endowment X_g^{agg} of each input g to the farms at rental rate r_g. The consumer also owns all farms and receives their profits as dividends Π.

The consumer's problem is:

$$\max_{\{C_i\}_{i=1}^N} \left(\sum_i \varphi_i C_i^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{5}$$

s.t.

$$\sum_i p_i C_i = \sum_g r_g X_g^{agg} + \Pi \tag{6}$$

For conciseness, we adopt the following formulation of the dividend Π:

$$\Pi = \sum_f \left[\sum_{i=1}^{N} p_i y_{f,i} - \sum_{g=1}^{G} r_g \sum_{i=1}^{N} x_{f,g,i} \right]$$

In this formulation, farmers act *as if* they faced distortions $\tau_{f,g}$ and $\tau_{f,g,i}$, fixed cost ω, and additional concavity η (since these show up in farmers' maximization problems), but these are ultimately non-monetary frictions and obstacles and so do not show up in the dividends Π that farmers send to the consumer. An exactly equiva-

lent formulation would be to interpret distortions, fixed costs, and additional concavity as monetary taxes or subsidies administered by the consumer: in this case, they would be included in Π and would likewise show up (with an opposite sign) in the consumer's budget constraint as government revenue or expense.[12]

Market clearing. Crop prices p_i and input prices r_g need to be such that all crops and all inputs markets clear:

$$C_i = \sum_f y_{f,i} \quad \forall i \tag{7}$$

$$\sum_f \sum_i x_{f,g,i} = X_g^{agg} \quad \forall g \tag{8}$$

Land needs to be distributed between farms in such a way that land used in cultivation equals the total amount of land available:

$$\sum_f \sum_i l_{f,i} = L^{agg} \tag{9}$$

The solution and the computational algorithm are discussed in Appendix C.2.

5.3 Mapping Observables to Fundamentals

For the model to be useful for quantifying the cost of misallocation in the data, unobserved distortions τ need to be estimable from observed characteristics. In the class of misallocation models following Hsieh and Klenow (2009), the fundamental distortions faced by each firm can be mapped to marginal revenue products of inputs chosen by the firm. With the assumed production structure, unobserved marginal revenue products can be extracted from observed inputs and outputs.

This result extends to our multi-product model as well, for both flexible inputs g and the land input l:

$$r_g \tau_{f,g} \tau_{f,g,i} = \alpha_{g,i} \eta \frac{\left(p_i y_{f,i}\right)^{\eta}}{x_{f,g,i}} \tag{10}$$

12. Extending this logic, distortion terms τ can be interpreted as a mixture of non-monetary frictions and monetary government taxes/subsidies. In this case, a fraction (representing the tax/subsidy portion) of each τ would show up in Π and the consumer's budget constraint. The choice between all these interpretations is completely arbitrary as they collapse to identical equilibrium conditions.

$$\lambda_f \tau_{f,l,i} = \gamma_i \eta \frac{\left(p_i y_{f,i}\right)^{\eta}}{l_{f,i}} \tag{11}$$

While only the total cost of the input can be extracted from the data in this way, the split of the total costs $r_g \tau_{f,g} \tau_{f,g,i}$ and $\lambda_f \tau_{f,l,i}$ into their components is arbitrary from the point of view of the farm. All that matters for quantifying misallocation is the *dispersion* in these costs, which is driven entirely by the dispersion in distortion τ terms. The average *level* of each distortion is then split from the level of market rental rates r_g and the average land shadow cost by the general equilibrium market clearing conditions. Furthermore, goods taste weights φ_i are picked to ensure that the current allocation is consistent with a general equilibrium: these details are discussed in Appendix C.2.

The productivity of the farm in each grown crop $z_{f,i}$ can be extracted from the data directly using the assumed production function, exploiting the availability of physical input and output quantities in the data:

$$z_{f,i} = \frac{y_{f,i}}{l_{f,i}^{\gamma_i} \Pi_g \left(x_{f,g,i}^{\alpha_{g,i}}\right)} \tag{12}$$

Seasons. As discussed in Section 3, viewing farms' production at the annual level may be deceptive as farmers tend to produce different products in different seasons. Bundling production across seasons may lead to nonsensical counterfactual predictions in the model like the model farm reallocating three times its total amount of land (summed across three seasons) to a single crop that's only ever grown in the most productive of the three seasons. Therefore, we map each model farm f not to an empirically observed farm, but to a farm-season combination. This ensures that different seasons have distinct productivity distributions. Section 6.1 will discuss the treatment of seasons at the aggregate level.

5.4 Calibration

Additional concavity parameter η. Farm-crop revenue can be written as:

$$p_i y_{f,i} = \Bigg(\underbrace{\left(\frac{1}{\lambda_f \tau_{f,l,i}}\right)^{\gamma_i} \Pi_g \left(\frac{1}{\tau_{f,g} \tau_{f,g,i}}\right)^{\alpha_{g,i}}}_{\text{composite distortion, } dist_{f,i}} \Bigg)^{\frac{1}{1-\eta\left(\sum_g \alpha_{g,i} + \gamma_i\right)}}$$

$$\cdot \underbrace{\left(p_i z_{f,i} \gamma_i^{\gamma_i} \eta^{\sum_g \alpha_{g,i} + \gamma_i} \Pi \left(\frac{\alpha_{g,i}}{r_g} \right)^{\alpha_{g,i}} \right)^{\frac{1}{1-\eta(\sum_g \alpha_{g,i} + \gamma_i)}}}_{\text{“objective” factors}} \tag{13}$$

where the composite distortion term summarizes the effects of distortions on revenue and the “objective factors” term captures effects driven by productivity, production parameters, and prices.

As Section 5.3 above discussed, the model can rationalize any observed heterogeneity across and within farms using τ distortions. However, the variance of distortions needed to achieve this (summarized in the variance of the $\log dist_{f,i}$ term) depends on the chosen additional concavity parameter η, as Figure 5 illustrates. When η is low, the degree of concavity in farm-crop revenue is high: in the frictionless optimum, farm-crop revenues are distributed comparatively uniformly, implying a uniform size distribution across farms and significant crop mixing within farms. The current observed allocation is highly dispersed in comparison, and the model rationalizes it with extreme frictions, producing a high variance of $dist_{f,i}$. When η is high, on the other hand, there is little concavity: in the frictionless optimum, farm-crop revenues are dispersed, implying high dispersion in farm sizes and significant crop specialization within farms. The current observed allocation is highly uniform in comparison, and the model rationalizes it with extreme frictions yet again. Somewhere in between is an intermediate value of η that allows the model to reproduce the observed level of dispersion in farm-crop revenues with the most conservative possible distribution of distortions.[13] Upon estimation, the value that minimizes the variance of $\log dist_{f,i}$ in Equation 13 is $\eta = 0.63$, and this is the value we elect for our headline calibration.

Misallocation estimates in general are sensitive to the degree of concavity in the model, whether that concavity comes from monopolistic competition (as in Hsieh and Klenow 2009) or decreasing returns to scale in production (as in Chen, Restuccia, and Santaeulalia-Llopis 2022). By introducing a flexible parameter η, we are able to explicitly explore the dependence of our estimates on this parameter (as sensitivity analysis will do in Section 6.2) and to intentionally pick a conservative estimate for the headline calibration.

13. Distortions are demeaned within the season.

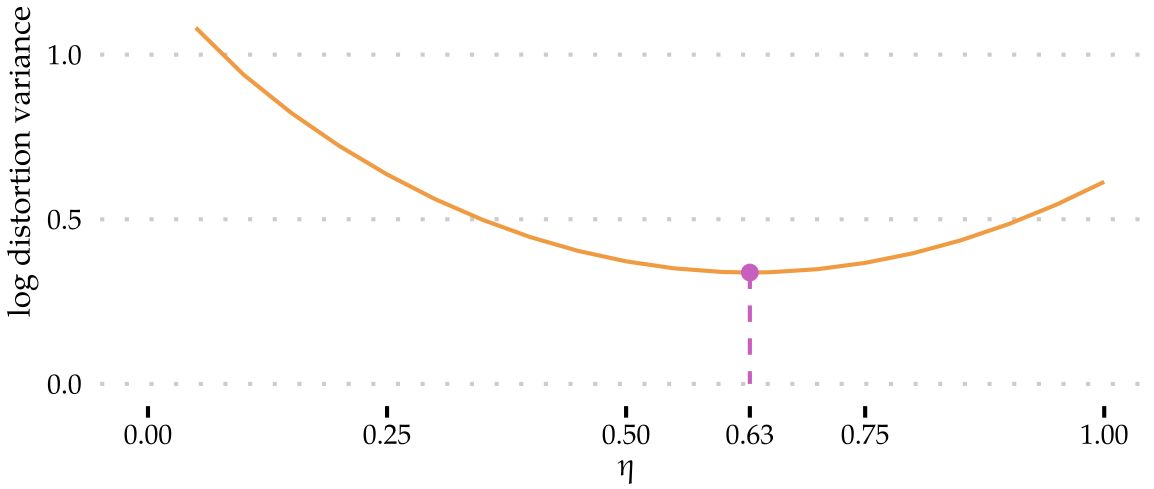

FIGURE 5: Variance of log extracted composite distortion term vs additional concavity parameter η

Note. The figure visualizes the variance of $\log dist_{f,i}$ from Equation 13 at different values of η. The highlighted value of η minimizes the variance.

ELASTICITY OF SUBSTITUTION σ. Representative consumer's demand for product i is

$$C_i = \left(\frac{\varphi_i}{p_i}\right)^{\sigma} \frac{\sum_g r_g X_g^{agg} + \Pi}{\sum_j \varphi_j^{\sigma} p_j^{1-\sigma}} \tag{14}$$

The log of product i's share in consumption expenditure can then be written as

$$\log\left(\frac{p_i C_i}{\sum_j p_j C_j}\right) = -\log\left(\sum_j \varphi_j^{\sigma} p_j^{1-\sigma}\right) + (1-\sigma)\log p_i + \sigma \log \varphi_i \tag{15}$$

Adapting the expression for household-level consumption data with households indexed by h and denoting crop i's share in h's expenditure with $s_{h,i}$, the expression can be written as a regression:

$$\log s_{h,i} = \beta_0 + \beta_1 \log p_{h,i} + \gamma_i + \varepsilon_{h,i} \tag{16}$$

Then the elasticity of substitution can be extracted from the coefficient on log

price: $\sigma = 1 - \beta_1$. Furthermore, product taste weights can be extracted (up to a normalization) from the estimated product fixed effects: $\varphi_i = \exp\left(\frac{\gamma_i}{\sigma}\right)$.

Estimating Equation 16 as-is will produce biased estimates of β_1 due to potential measurement errors in price creating a mechanical correlation between the left- and the right-hand-side of the equation, and because some unobservables at the household or village level could also be correlated with both crop prices and expenditure shares (Sotelo 2020; Kebede 2020). Therefore, we use an instrumental variable approach to address these concerns, instrumenting for crop prices using the interaction of village-level measures of elevation and ruggedness with rainfall, as well as the availability of pucca roads (all-weather roads). The rationale for using these instruments is that villages with high ruggedness may experience lower agricultural productivity when faced with negative rainfall shocks, consequently leading to higher prices. Similarly, the availability of roads will affect the shadow prices of inputs and hence affect agricultural production. Villages with bad road connectivity may experience higher labor prices and other input prices, raising the cost of crop production and, therefore, negatively affecting the prices of crops. Using these instruments, we estimate the equation 16.

Results of estimation are reported in Appendix Table A.2. We find that the F-statistic for the first stage is 230, which establishes the instrument's relevance. The estimation yields $\beta_1 = -0.7$, implying $\sigma = 1.7$. This is in between the estimates obtained in a similar setting by Kebede (2020) for Ethiopia, 1.3, and by Sotelo (2020) for Peru, 2.4.

5.5 Product Choice and Misallocation: Mechanisms

Allowing for product heterogeneity, endogenous product choice, and multi-product behavior in a model of firm-level misallocation impacts the estimates of misallocation obtained in the model. There are four major mechanisms through which a conventional single-product model misestimates the severity of misallocation in the data generated by farms growing heterogeneous crops. This section explores these mechanisms qualitatively.

Single-product models misinterpret product heterogeneity as frictions. Even in the absence of any fundamental frictions, the fact that different farms choose to produce different products will lead the single-product model to mistake heteroge-

neous crop choice for evidence of frictions.

Suppose that farm 1 draws a high productivity $z_{1,rice}$ in rice and farm 2 draws a high productivity $z_{2,veg}$ in vegetables, causing each farm to specialize in their respective high-productivity crop (so that $y_1 = y_{1,rice}$ and $y_2 = y_{2,veg}$). In the absence of frictions, marginal revenue products will be equalized across farms, accounting for the fact that rice is more land-intensive but vegetables are more labor-intensive:

$$\frac{\alpha_{g,rice}\eta(p_{rice}y_{1,rice})^{\eta}}{x_{1,g}} = \frac{\alpha_{g,veg}\eta(p_{veg}y_{2,veg})^{\eta}}{x_{2,g}} \quad \forall g$$

However, the marginal products evaluated by the single-product model that assumes a single aggregate production function for both crops will generally not happen to be equal:

$$\frac{\alpha_{g,agg}\eta(p_{agg}y_1)^{\eta}}{x_{1,g}} \neq \frac{\alpha_{g,agg}\eta(p_{agg}y_2)^{\eta}}{x_{2,g}} \quad \forall g$$

The heterogeneity in input choices between the two farms, and the resulting heterogeneity in single-product marginal revenue products will lead the single-product model to impute frictions where there are none and thus to *overstate* the true cost of misallocation.

Figure 6 illustrates this mechanism for a set of farms that suffer from no distortions ($\tau_{f,g} = 1 \ \forall f,g$; $\lambda_f = \bar{\lambda} \ \forall f$) but have different relative productivities in rice vs vegetables. Farmers who are relatively productive in rice specialize in this land-intensive crop, while farmers who are relatively productive in vegetables specialize in this labor-intensive crop. The single-product model explains the heterogeneous input choices with heterogeneity in friction terms, such as the pictured single-product model's estimated $\hat{\tau}_{f,labor}$.

Single-product models miss the endogenous response of farms to frictions. The existence of choice between heterogeneous products changes how farmers respond to distortions. Figure 7 shows how crop-specific outputs change as the labor market distortion is altered for a farm that has access to two crops: rice (relatively land-intensive) and the combined "vegetables, fruits, and oilseeds" crop (relatively labor-intensive). When the labor distortion $\tau_{f,labor}$ is below average (the farm's labor is effectively subsidized), the farm chooses to specialize heavily in vegetables, fruits, and oilseeds: this combined crop allows the farm to take advantage of the relatively cheaper labor input. As the distortion term grows and the subsidy diminishes, the

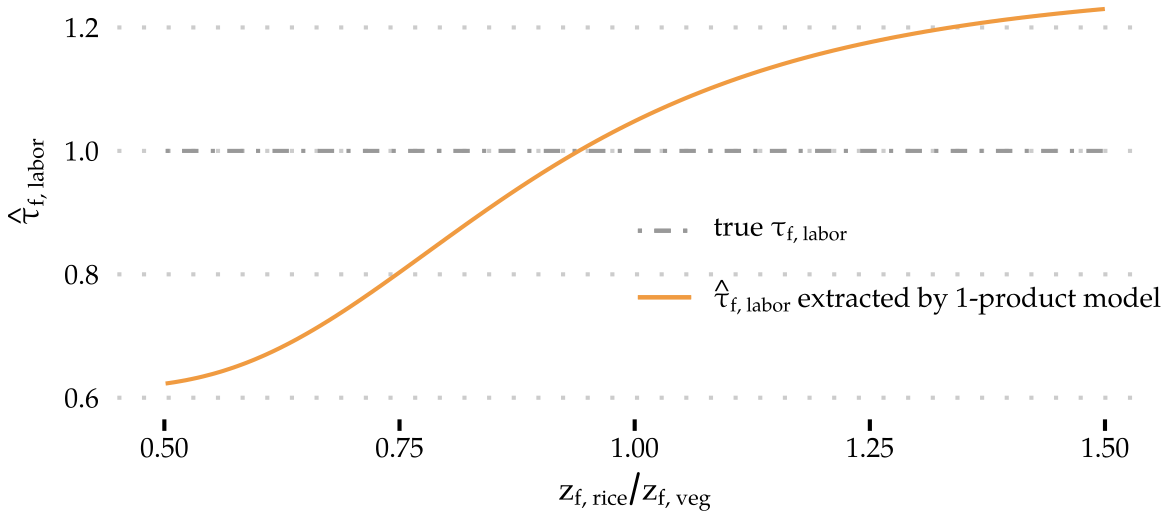

FIGURE 6: Distortion extracted by 1-crop model from multi-crop model data when productivities are heterogeneous: model illustration

Note. The figure displays the labor friction estimated by the single-product model when the data is generated by the multi-product model populated with farms that face no true frictions but have different productivities in rice and vegetables.

farm reduces its output as it is unable to hire as much labor. At a certain point, the subsidy is so minor that the farm starts growing both crops at the same time. As the subsidy turns into a tax and labor becomes unattractive relative to land, the farm changes its specialization to rice: the relatively less labor-intensive crop.

This endogenous adjustment of a farm's product choice to the distortions it faces allows the farm to lower its exposure to market frictions. Because traditional single-product models miss this margin of adjustment, they overstate the underlying distortions. A modest tax-like labor distortion would cause the farm to lower its $\frac{\text{labor}}{\text{output}}$ ratio both in the single- and the multi-product models of misallocation. But in the multi-product model, the distortion further leads the farm to shift part of its production to less labor-intensive crops, lowering the overall $\frac{\text{labor}}{\text{output}}$ ratio even more. This endogenous product choice response to a modest underlying friction would be mistaken by the single-product farm as evidence of severe underlying frictions. Figure 8 plots the labor distortion estimated by the single-crop model, $\hat{\tau}_{f,labor}$, when applied to observed inputs and outputs generated by the simulated multi-crop model. If

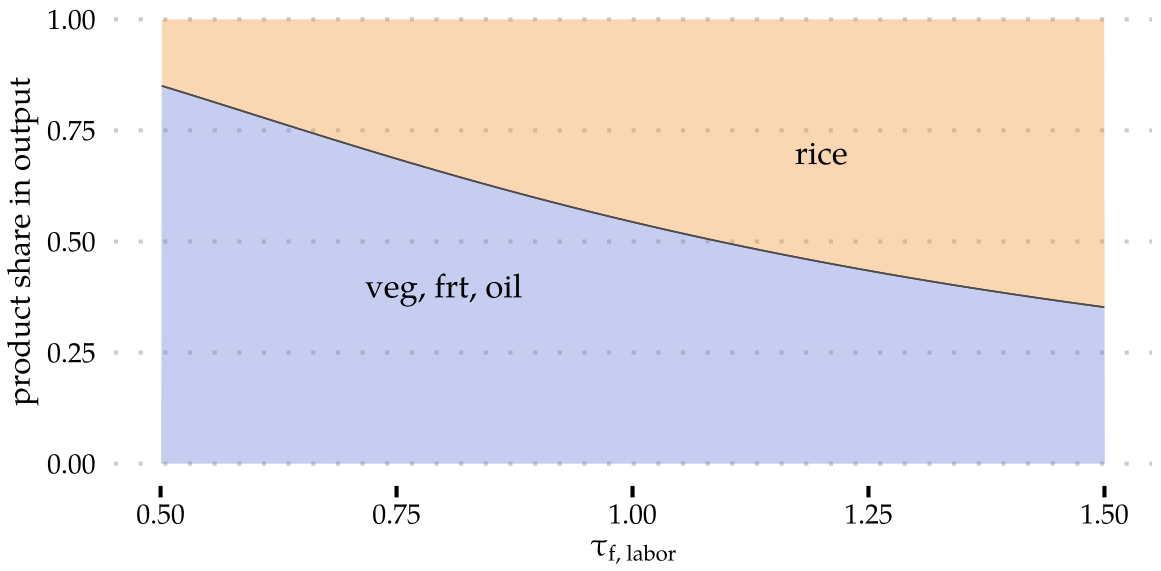

FIGURE 7: Product choice vs labor distortion: model illustration

Note. The figure illustrates the changing product choice of a farm as the labor friction it faces is varied.

the single-crop model could extract the fundamental friction correctly, its estimate would lie along the 45° line. Instead, the single-crop model significantly overstates the severity of both tax-like and subsidy-like distortions, and therefore *overstates* the cost of misallocation.

SINGLE-PRODUCT MODELS UNDERSTATE TFP DISPERSION. By ignoring heterogeneity across products within a single farm, single-crop models miss productivity differences between crops that farmers may exploit once frictions are removed.

Suppose that a given farm draws a low productivity in rice, $z_{1,rice} = z_L$, and a high productivity in vegetables, $z_{2,veg} = z_H$. If the farm faces severe frictions, they may distort the farm's product choice so that instead of specializing in vegetables, the farm also grows a significant quantity of the less productive rice. If the farm is treated as producing a homogeneous agricultural good, its farm-level TFP will be some average $z_L < z_M < z_H$. Applying the multi-product model to the farm would show that if frictions were removed, the farm would specialize in vegetables, raising its farm-level productivity from z_M to z_H. Applying a single-product model, however, would suggest that the farm-level productivity would remain at z_M even if frictions were removed, since the single-product model is blind to productivity differences between

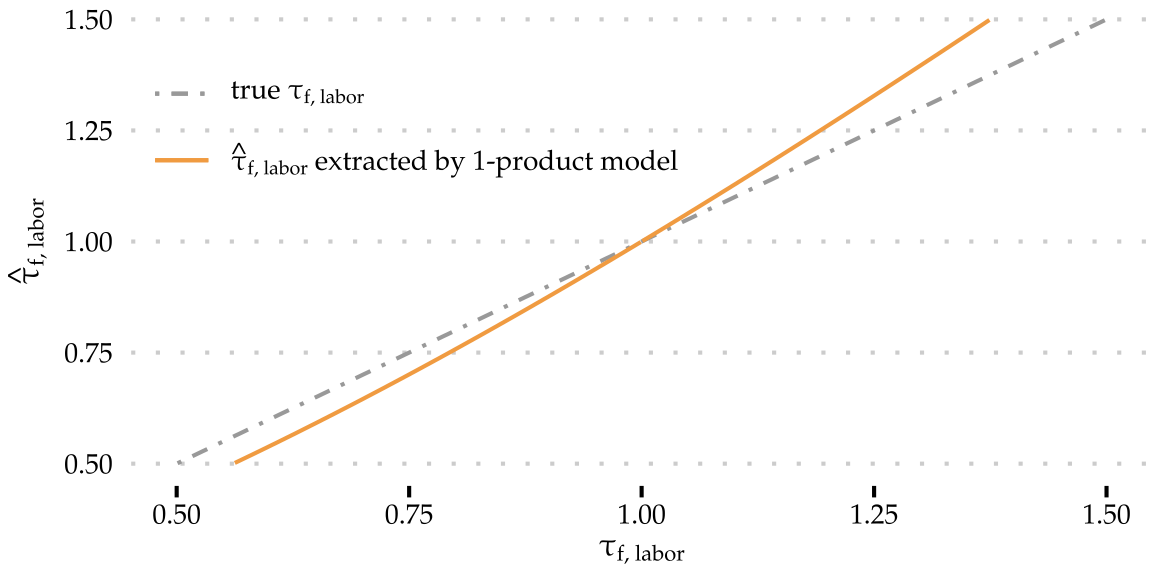

FIGURE 8: Distortion extracted by 1-crop model from multi-crop model data when distortions are heterogeneous: model illustration

Note. The figure displays the labor friction estimated by the single-product model from data generated by the multi-product model when farms are able to adjust product choice in response to frictions as in Figure 7.

crops within a single farm. Therefore, the single-product model would understate the underlying productivity differences within the economy and *understate* the true cost of misallocation.

SINGLE-PRODUCT MODELS MISS RETURNS-TO-SCALE HETEROGENEITY. Estimates of misallocation depend on the returns to scale in the production function. When returns to scale are small, the most productive firms are not able to grow their output by much even when the distortionary frictions are removed, reducing the gains from reallocation. Section 6.2 will illustrate the sensitivity of misallocation estimates to the degree of concavity in the model, of which decreasing returns to scale is one potential source. Single-product models impose the same level of returns to scale on the whole economy. But as Section 4.4 showed, we estimate returns to scale to be different across products: some, like "vegetables, fruits, and oilseeds", have significantly decreasing returns to scale. Others, like rice, have returns to scale that are slightly decreasing but close to constant. The estimates for wheat, on the other hand, suggest increasing returns to scale. A single-product model will miss the fact that, in

the event of a reallocation, farmers growing rice and wheat may be able to scale their production up more easily than vegetable farmers.

Figure 9 illustrates the mechanism. Suppose an economy is populated by two farms: a farm that has equally high productivity in all crops and faces equally high frictions in all inputs, and a farm that has equally low productivity in all crops and faces equally low frictions in all inputs. Consider a single-product model with one "aggregate" crop and a multi-product model with two crops: "rice" (slightly higher returns to scale than the "aggregate" crop) and "vegetables, fruits, and oilseeds" (lower returns to scale than the "aggregate" crop). Once the frictions are removed, output of all crops goes up in both models due to the reallocation of inputs from the previously subsidized unproductive farm to the previously taxes productive farm. However, the magnitudes are different: in the multi-product model, rice responds more strongly than vegetables due to the higher returns to scale in the production function of rice, while in the single-product model, the response of the aggregate crop is in between the two. Moreover, the consumer is able to adjust the consumption bundle to take advantage of this heterogeneity: the more substitutable different crops are for the consumer, the more the consumer reallocates consumption toward the higher returns-to-scale crop in the optimum, maximizing the gains from letting the most productive farm expand its output. Because the single-product model lacks this margin of adjustment, it will *understate* the true cost of misallocation, especially if crops are substitutable.

6 Quantitative Results

The preceding section showed that the way underlying distortions affect farm behavior and the way these distortions are extracted from observational data are both altered by the presence of heterogeneous products and multi-product farms. The present section quantifies the importance of these differences for estimating the aggregate cost of misallocation in the economy.

To estimate the aggregate cost of misallocation, we conduct counterfactual reallocation exercises in the model. Each exercise reduces or completely removes distortions between farms, computes their optimal counterfactual outputs, and aggregates the output to obtain a counterfactual reallocation output gain. The % difference in output between the counterfactual exercise and the currently observed one constitutes the headline measure of the aggregate cost of misallocation: it captures the

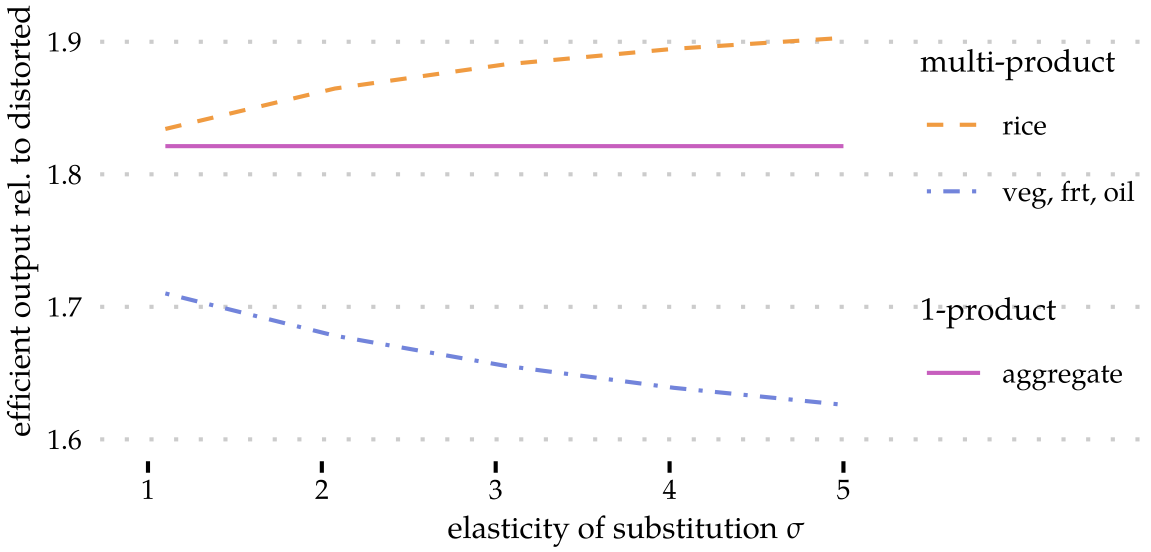

FIGURE 9: Efficient output relative to distorted output in multi-product vs single-product model: model illustration

Note. The figure displays how efficient output (attained when all friction terms are the same between farms) varies with elasticity of substitution in the single-product ("aggregate" crop) and the multi-product ("rice" and "veg, frt, oil" crops) model. Efficient output is rescaled relative to current distorted output in each crop.

magnitude of output being lost due to the market frictions currently present in the economy.

Note that the aggregate input quantities are not changed in the reallocation exercises: only the distortions are. Therefore, any aggregate output gain from reallocation is due purely to the increase in aggregate TFP. Furthermore, because the idiosyncratic farm-crop-level TFPs are not changed in the counterfactual either, the increase in aggregate TFP is driven not by any increase in farm-level TFPs but rather by an improved allocation of existing inputs to the more productive farms.

6.1 BENCHMARK REALLOCATION EXERCISE

In our benchmark reallocation exercise, we equalize all farm-level distortions while keeping the crop set of each farm restricted to what it is observed growing in the data. Furthermore, we leave farm-crop-level distortions intact to preserve idiosyncratic product choice motives.

Equalizing distortions. We impose equality of farm-input-level distortions: $\tau_{f,g} = \bar{\tau}_g \ \forall g, \lambda_f = \bar{\lambda}$. This equalization effectively makes input markets frictionless across farms. The levels of average distortions are fixed by the general equilibrium conditions. We leave the farm-input-crop-level distortions $\tau_{f,g,i}$, $\tau_{f,l,i}$ at their current levels to allow the model to capture idiosyncratic product choice motives, as will be discussed below.[14]

Keeping crop sets fixed. In this exercise, farmers will change the total quantity purchased of each input as well as the allocation of inputs between the crops they grow. However, farmers are not allowed to change the set of crops they are growing: if a particular farm is observed growing only rice and wheat in the data, it will not be allowed to start growing oilseeds in the counterfactual. This restriction permits us to only use productivities and distortions that are extracted directly from the data: we don't need to make any assumptions on what each farm's counterfactual productivity or farm-crop distortion in an unobserved crop would be.

Preserving idiosyncratic product choice motives. The model of a profit-maximizing farm is a simplification. Real-world farmers may choose to grow a particular crop for reasons other than pure profit maximization based on observed total factor productivity and input costs. For instance, semi-subsistence farmers may prefer a crop that is less profitable if sold on the market but serves their family's dietary needs better. The model is blind to these motives and would explain this observed preference for a non-profit-maximizing crop with a combination of an inefficient crop set and farm-input-crop-level distortions $\tau_{f,g,i}$ and $\tau_{f,l,i}$ making inputs relatively cheaper for the subsistence crop. Therefore, allowing farmers to change crop sets and equalize farm-input-crop-level distortions may result in the model overstating the true costs of misallocation as it would overstate the degree to which farmers are likely to switch to profit-maximizing crops. Avoiding this is the main reason we elect not to equalize farm-input-crop-level distortions and the second reason we keep crop sets fixed in the benchmark counterfactual exercises. Even though the model does not explicitly

14. Aragón, Restuccia, and Rud (2022) argue that estimates of agricultural misallocation are more reliable when the unit of analysis is a farm rather than a plot. We rely on plot-level data for production function identification, but in our misallocation exercise, we use variation across plots within a farm only when they produce different crops. Furthermore, leaving farm-input-crop-level distortions intact ensures that any remaining unexplained variation across plots is not interpreted as a source of misallocation, effectively making the exercise capture misallocation between farms and crops, but not between plots.

include subsistence and other idiosyncratic product choice motives, making these two restrictions makes the counterfactual reallocation exercise only capture the output gain from efficient input markets allowing farmers to more efficiently grow whatever they currently prefer, without making unrealistic predictions on how they could "improve" their existing product preferences.

Seasons. The treatment of seasons in the model is irrelevant for its ability to reproduce the observed allocation but becomes important for counterfactual exercises. If model farms f, mapped to farm-seasons in the data, are all bundled together into the same economy, a naive counterfactual reallocation may feature double-booking of inputs: land and labor inputs from less productive seasons may be reallocated to the more productive one, implying that a single physical piece of land may be allocated to three different uses in a single season and to no uses in the other two seasons. To preclude this nonsensical possibility, we split the model economy into three sub-economies by season, and conduct counterfactual reallocations only within each season: this ensures that no inputs are reallocated across seasons. The aggregate counterfactual output in each exercise is then a summation of three seasonal outputs.

Comparing multi-product and single-product models. To understand the effect of the product choice channel on aggregate productivity, we compare the estimated reallocation gain between the multi-product farm model and a traditional single-product farm model. The single-product farm model uses inputs and outputs measured at the farm level rather than the farm-crop level. Consequently, the single-product model uses a single production function with estimated "aggregate" elasticities from Figure 3.

6.2 Cost of Misallocation

Benchmark reallocation. The "benchmark" row of Table 2 shows the results of the benchmark counterfactual reallocation exercise. According to the multi-product model, allowing input markets to operate efficiently would raise the aggregate output of India's agricultural sector by 294%.[15] Repeating the same exercise in the single-product model promises a still formidable but smaller gain of 212%. For our preferred

15. Enormous costs of misallocation in the agricultural sectors of developing countries have been estimated previously, most notably in Chen, Restuccia, and Santaeulalia-Llopis (2022).

calibration and exercise, the single-product model *understates* the cost of misallocation by 28% (or 82 percentage points of current output).

TABLE 2: Reallocation exercises

	Gain	Gain (1 product)	Error (1 product)
benchmark	**294%**	**212 %**	**-28%**
best state	35%	45 %	26%
land only	30%	36 %	20%
complete	311%	212 %	-32%

Note. Column "Gain" displays the % agricultural output gain from each counterfactual exercise, estimated using the multi-product model. Column "Gain (1 product)" displays the same but using the single-crop model. Column "Error (1 product)" displays the relative difference between the two: positive values indicate an overestimation of reallocation gain by the single-product model. "Benchmark" exercise equalizes all frictions between farms. "Best state" exercise reduces population-level distortion variances to match those of the least-distorted state. "Land only" exercise equalizes land frictions between farms. "Complete" exercise additionally equalizes farm-crop frictions.

"BEST STATE" REALLOCATION. The benchmark exercise above predicts the effects of a blanket removal of distortions across farms. Such a removal is likely to be unattainable in practice as misallocation of inputs is significant even in developed countries. Furthermore, it may not even be desirable as some of the frictions picked up by the model likely reflect model misspecification rather than truly malign distortions.[16] Motivated by this, we conduct a conservative exercise that improves the input allocation not to the perfectly efficient level but to the level of the least-distorted Indian state. We rank Indian states by the variance of log composite distortion $dist_{f,i}$ discussed in Section 5.4 and identify the "least-distorted state", which is Tamil Nadu. We then rescale the distortions faced by individual farmers to equalize the population-level variance of each friction type with the state-level variance in Tamil Nadu.[17] This

16. Some of the innocuous causes of measured frictions explored in the literature are mismeasurement, late season productivity shocks, and dynamic input adjustment costs (Bils, Klenow, and Ruane 2021; Gollin and Udry 2021; Asker, Collard-Wexler, and De Loecker 2014).

17. Variances are computed and equalized separately for five types of distortions: $\hat{\lambda}_f$, $\tau_{f,g}$ (separately for each $g = 1, 2$), $\tau_{f,g,i}$, $\tau_{f,l,i}$. Letting $V_{d,s}$ be the variance of log distortion of type d in the least distorted state and $\hat{V}_d$ be the corresponding population variance, each distortion is rescaled to $\tau_{d,f}^{new} = \exp\left(\sqrt{\frac{V_{d,s}}{\hat{V}_d}} \cdot \log \tau_{d,f}\right)$.

"best state" exercise promises a 35% gain in agricultural TFP if the severity of distortions throughout India is brought down to the levels of its least-distorted state. In this case, the single-product model *overstates* the gain by 26% or 10 p.p.

Alternative reallocations. Our benchmark reallocation removes estimated frictions in all three modeled input markets: land, labor, and intermediates. The agricultural misallocation literature, however, has mainly focused on institutional frictions in the land market. In the alternative "land only" reallocation, we remove land frictions between farms (equalize λ_f) but keep their labor and intermediates frictions in place ($\tau_{f,g}$ stay at current levels). The expected reallocation gain in this case is far more modest but still sizeable: 30% of current agricultural output. The single-product model estimates a slightly higher reallocation gain of 36%: in contrast to the benchmark exercise, the single-product model *overstates* the potential gain in this incomplete reallocation by 20% (or 5 p.p.). Note also that Bolhuis, Rachapalli, and Restuccia (2021) arrived at a very similar gain of 38% from efficiently reallocating land in India, using a conventional single-product model.

The "land only" results are instructive also because of the gap between them and the benchmark exercise: efficiently reallocating all three inputs is much more effective than reallocating land alone, implying that labor and input market frictions impose a significant drag on Indian agricultural productivity.[18]

The benchmark reallocation removes frictions between farms but keeps farm-crop-level distortions intact in order to preserve any idiosyncratic product choice motives that the model does not explicitly capture. However, some of the farm-crop-level $\tau_{f,g,i}$ distortions are likely to reflect malign frictions rather than benign unmodeled motives. To place an upper bound on their importance, we conduct an alternative "complete" reallocation that equalizes these farm-crop-level frictions. Their contribution is relatively minor, increasing the reallocation gain from the 294% of the benchmark exercise to 311%. Because the single-product model is completely blind to the within-farm variation picked up by these frictions, its underestimation becomes even more severe, going from 28% to 32%.

Sensitivity of misallocation estimates. The multi-product model has two free parameters: additional concavity parameter η and elasticity of substitution between

18. See Andrew D. Foster and Mark R. Rosenzweig (2022) for an exploration of labor market frictions in India.

crops σ. Figure 10 displays the sensitivity of the benchmark reallocation exercise outcomes to alternative parameter values, centered on the benchmark calibration.

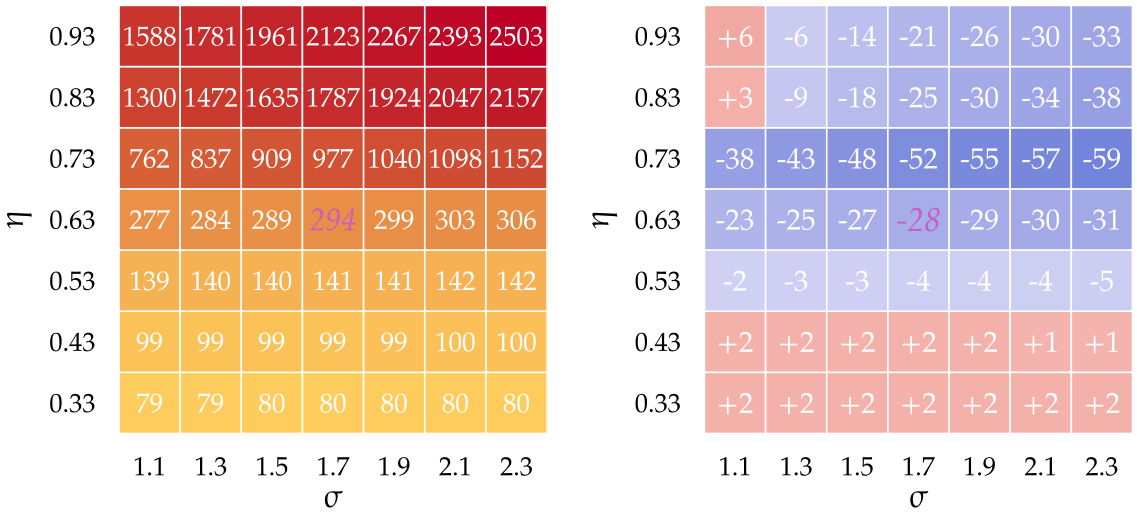

(A) Reallocation gain, %

(B) Single-product model error, %

Note. Panel (A) displays estimated reallocation gain using the multi-product model at each set of parameters σ, η, expressed as a ratio (1 = no gain). Panel (B) displays the error in the single-product model's reallocation gain ratio relative to the one displayed in Panel (A): positive values denote the overstatement of misallocation by the single-product model. The benchmark calibration is highlighted in the center.

FIGURE 10: Reallocation gain and single-product model error for alternative parameterizations

Panel (A) of Figure 10 shows how the reallocation gain predicted by the multi-product model changes with the two parameters. Our preferred calibration, discussed in Section 5.4, and its 294% reallocation gain is highlighted. The figure illustrates how sensitive misallocation estimates are to the degree of concavity in the model. The smaller the additional concavity (higher η), or the elasticity of substitution between crops (σ), the greater the reallocation gain predicted by the model. Productive farms grow in the efficient reallocation, but the lower the concavity, the greater their optimal size, driving up the benefits of reallocating inputs to the most productive farms. The figure tells a cautionary tale against basing any analysis of misallocation on a single parameterization: models of misallocation used in the literature are highly sensitive to the extent of concavity, whether it is driven by market power, decreasing returns to in production, or some other channel. This is also the

reason why in Section 5.4 we elected not to calibrate η to some moment in the data but instead to pick the value that implies the most conservative underlying distribution of frictions.

Panel (B) of Figure 10 shows the response in the single-product model's error to the changing parameter values. While at the benchmark calibration, the single-product model understates the reallocation gain by 28% (or 82 p.p.), neither the magnitude nor the sign of its error is set in stone. At some alternative calibrations, the single-product model *overstates* the reallocation gain instead, as the contribution of understating mechanisms diminishes. For instance, the less substitutable the crops are (the lower the σ), the smaller is the ability of the consumer to reallocate toward high returns-to-scale crops in the multi-product model (see Figure 9) and the smaller is the single-product model's understatement of misallocation.

These results imply that accounting for product heterogeneity and product choice in estimating misallocation is not a matter of applying a simple correction to single-product estimates: the sign and the magnitude of the needed correction depend on the exercise and calibration used by the researchers.

Role of seasons and states. Finally, we conduct two additional sets of exercises to clarify the role of seasons and states. Appendix Table A.4 repeats the counterfactual reallocation exercises above but treats different Indian states as separate economies, preventing inputs from moving across states in a reallocation. This "within state" exercise provides another conservative take on the extent of misallocation in Indian agriculture. Reallocation gains decrease substantially but remain significant: the benchmark reallocation promises a 107% productivity gain in the multi-product model, with the single-product model overstating this gain by 15%. Appendix Table A.5 judges the significance of our decision to prohibit input reallocation across seasons in the main exercises: this "no split by season" exercise allows inputs, even land, to be moved between seasons in the counterfactual reallocation. Estimated reallocation gains do rise in this case, but not drastically so.

6.3 Mechanism decomposition

Section 5.5 discussed several prominent mechanisms that make the single-product model's estimates of misallocation diverge from those of the multi-product model. In this section, we decompose the difference between the two models identified above

into contributions by each of these mechanisms.

Exercises. The mechanisms are conceptual (rather than driven by a single parameter) and highly interacting, and so do not allow a clean decomposition of the single-product model's error into additive components. However, we can get a sense of the contribution of each mechanism through a series of exercises.

To evaluate the extent to which the single-product model misinterprets product heterogeneity as frictions, we take the counterfactual allocation produced by the multi-product model once frictions are removed and use the single-product model to extract frictions from it as if it were real data. Then we conduct an additional counterfactual reallocation, removing these "fake" frictions from the single-product model. Because the fake frictions reflect only the fact that the single-product model misinterprets production function heterogeneity as evidence of frictions even when there are none, the additional reallocation gain obtained in this exercise provides an estimate of this mechanism's contribution to the single-product model's overall error.

The endogenous product choice mechanism can be shut down by precluding farms from changing their product sets and the relative allocation of inputs across crops. The former is already prohibited in the benchmark allocation (see Section 6.1). The latter can be shut down by solving an auxiliary model in which farms can change the total quantity of each input hired but not the allocation of this quantity between crops: product shares are fixed at exogenously given levels. This model is described in Appendix D. To get an estimate of the endogenous product choice response to frictions currently present in the economy, we re-introduce currently observed frictions into the counterfactual "efficient" economy, but don't let farmers change their product choice from the counterfactual ratios. Comparing the predicted output in this case to the observed actual output isolates the contribution of the endogenous mechanism.

The TFP dispersion mechanism can be shut down by applying the single-product model to current data in which all farm-season-crop combinations are treated as separate farm-seasons, effectively turning it from a farm-level model into a plot-level model. In this case, the single-product model still applies the same production function to all farms but it sees a more fine-grained distribution of productivities.[19]

19. We artificially preserve the farm-crop-level frictions $\tau_{f,g,i}$, $\tau_{m,i}$ when treating farm-crop-level data as farm-level data: otherwise the farm-crop-level frictions would be interpreted as farm-level frictions in this exercise and dropped, leading to an unfair comparison.

We isolate the contribution of the single-product model's blindness to returns-to-scale heterogeneity by shutting this heterogeneity down. We conduct an auxiliary counterfactual reallocation in the multi-product model after equalizing the returns to scale in every product's production function to those in the aggregate production function used by the single-product model. Elasticities of all three inputs are rescaled accordingly, to preserve their relative quantities. The difference in reallocation gain between this auxiliary exercise and the benchmark exercise indicates the contribution of returns-to-scale heterogeneity to the single-product model's error.

Decomposition of the benchmark reallocation. Figure 11 presents the contribution of each of the three mechanisms (excluding the TFP dispersion mechanism since it doesn't apply to the benchmark exercise) to the single-product model's misallocation error in the benchmark exercise, in percentage points. The contributions are visualized sequentially and additively, but note that the sum of the three mechanisms does not add up to the 82 percentage point error exactly, since the three identified mechanisms are neither exhaustive nor independent of one another. Still, the decomposition provides a useful overview of their relative importance.

The fact that the single-product model can treat efficient product heterogeneity as evidence of non-existing frictions leads the model to overstate misallocation by 9 p.p. The endogenous product choice mechanism makes a minor 0.5 p.p. contribution to the model's overstatement.[20] The understatement by the single-product model stemming from its blindness to within-farm TFP dispersion chips 3 p.p of this error. A far bigger understatement, however, is caused by the single-product model ignoring the heterogeneous returns to scale between products, which leads it to understate misallocation by 93 p.p.

Decomposition of the "best state" reallocation. Figure 12 presents the mechanism decomposition for the "best state" exercise, in which frictions across India are not removed but simply reduced to the levels of the least-distorted state. The effects of this modest reallocation depend less on the ability of most productive farmers to greatly expand their output, and thus the "returns to scale" channel is far less prominent, allowing the "product heterogeneity as frictions" channel to dominate.

20. Note that in this exercise, the endogenous product choice mechanism being identified works only through the relative quantities of crops grown by the farm rather than through the choice of *which* products it grows, since the latter is fixed at its current level.

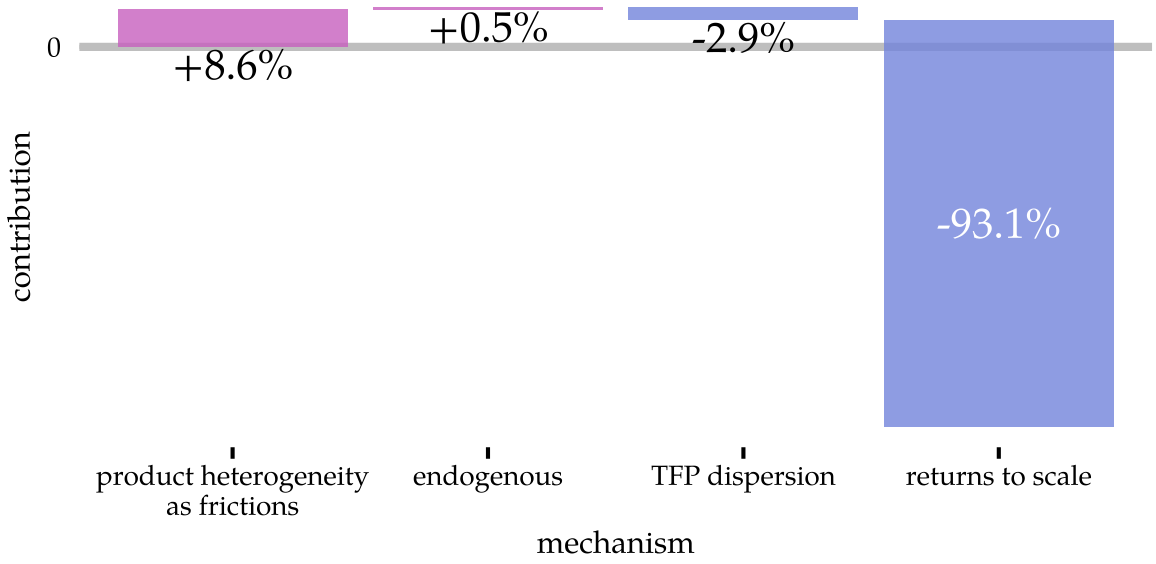

FIGURE 11: Mechanism contributions to the single-product model's misallocation error, benchmark reallocation

In this reallocation, the single-product model's biggest error is mistaking heterogeneous product choice for evidence of frictions, leading it to overstate the cost of misallocation on net, as Table 2 showed.

SINGLE-PRODUCT MODEL OVERSTATES CONSERVATIVE REALLOCATIONS, UNDERSTATES EXPANSIVE ONES. The difference between the "benchmark" reallocation and the "best state" reallocation in the single-product model's error and its decomposition into mechanism extends more broadly. We conduct a series of reallocation exercises in which the frictions present in the economy are progressively rescaled to smaller and smaller values. Figure 13 presents the resulting single-product model error and its decomposition at select points into the four mechanisms. The exercise that should be preferred depends on one's interpretation of the causes of observed dispersion in the data. If most of the dispersion extracted by the model is seen as driven by benign mechanisms and not reflecting true malign frictions, then exercises that treat a smaller fraction of this dispersion as misallocative frictions to be removed and achieve a correspondingly lower overall reallocation gain should be preferred as a measure of misallocation. In this case, the single-product model tends to overstate the cost of misallocation, driven by its treatment of benign product heterogeneity as evidence of frictions. If, on the other hand, most extracted dispersion is seen as evidence of

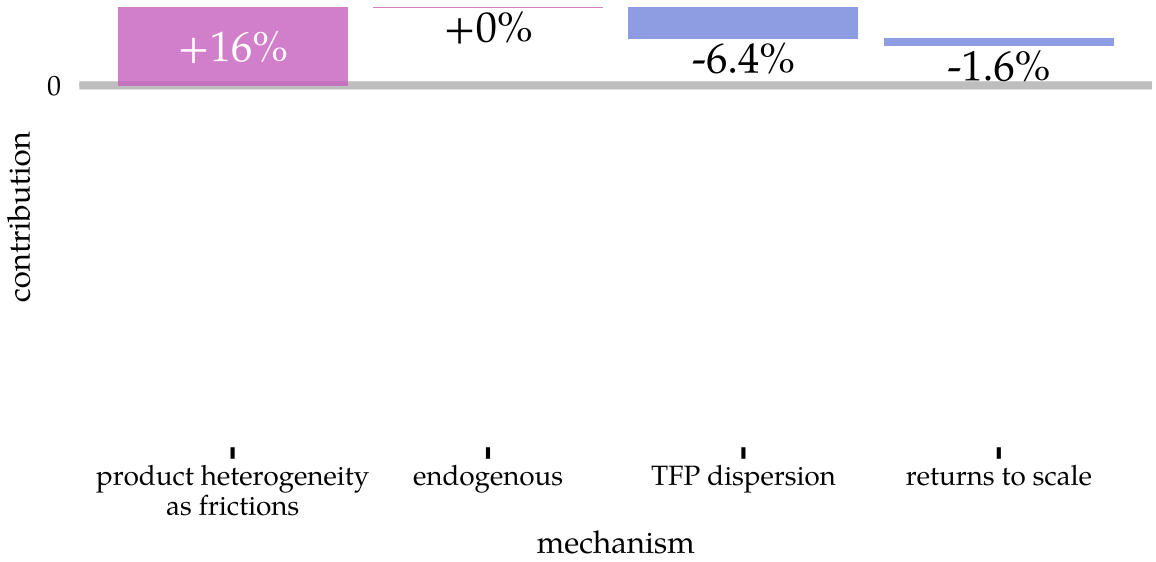

FIGURE 12: Mechanism contributions to the single-product model's misallocation error, "best state" reallocation

true malign frictions, then exercises that remove most of those frictions and achieve a correspondingly high reallocation gain should be preferred. In this case, the single-product model tends to understate the cost of misallocation: when misallocative frictions are seen as significant, correctly capturing the returns to scale heterogeneity is crucial for predicting the ability of firms to expand output once frictions are removed, and the single-product model understates this ability.

7 CONCLUSION

Misallocation of inputs between firms imposes a significant drag on the productivity of developing countries. The cost of misallocation in agriculture is especially large and contributes to explaining the enormous agricultural productivity gap between developed and low-income countries.

Estimates of misallocation in non-agricultural sectors have been limited by the treatment of all firms as producing only a single good or multiple homogeneous goods. Estimates of agricultural misallocation have been further limited by assuming that all farms produce the same agricultural product. We use Indian farm-level data to estimate production functions for individual crops and find that their input elasticities and returns to scale are meaningfully heterogeneous for the purposes of

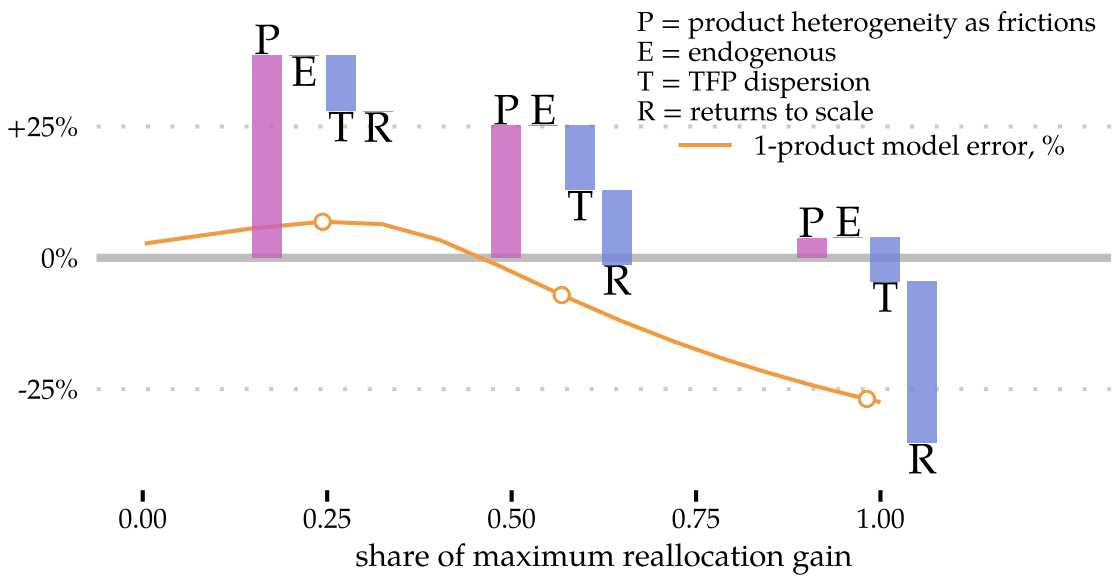

FIGURE 13: Single-product model error and decomposition for increasingly complete reallocations

Note. The figure presents the single-product model reallocation gain error (in relative % terms) in a series of reallocation exercises in which frictions are rescaled to progressively smaller values. Exercises are arranged on the x-axis by their reallocation gain as a share of maximum attainable gain (when all frictions are removed). Single-product model error at three select points is decomposed into four mechanisms, akin to Figures 11, 12, but rescaled to relative % terms. Mechanisms' additive contributions do not necessarily add up exactly to the overall error.

estimating misallocation.

We build a model of multi-product firms choosing among products with heterogeneous production functions. We find that conventional single-product models misestimate the cost of misallocation in several ways, some of which push the single-product models to overstate misallocation while others push them to understate it. The net error of a single-product model depends on the exercise of interest and chosen calibration. This implies that accounting for product heterogeneity and multi-product behavior is not a matter of applying a simple correction to a single-product model's output. When most observed dispersion between firms is interpreted to be driven by benign mechanisms, the remaining real misallocation tends to be overstated by the conventional single-product model due to its treatment of product heterogeneity as evidence of frictions. When, on the other hand, most observed dispersion is seen to reflect real frictions that ought to be removed, the single-

product model understates the significant real misallocation, driven by its blindness to heterogeneity in firms' ability to expand output.

We conducted our analysis on Indian farm-level data, since agricultural micro-level data provides several advantages in exploring the interplay of product heterogeneity and misallocation. However, the model developed in the paper—and the mechanisms it illustrates—would apply equally to non-agricultural firms. Future research can extend this analysis from agriculture to non-agriculture by applying the model to firm-product-level datasets.

Another future direction is to consider the role of dynamics in the interaction of product choice and misallocation.[21] Crop rotation can lead to input choices that appear inefficient in the short run but preserve the long-run productivity of land. Moreover, the heterogeneous suitability of different crops to different seasons can produce—when input rental markets are underdeveloped—input choices that appear inefficient in one season but are made to have the input in place for the next season. Capturing product choice complementarities across seasons and even years empirically and theoretically is thus a fruitful avenue of further research.

21. Dynamic choices have been shown to be important for misallocation using single-product models by, for instance, Asker, Collard-Wexler, and De Loecker (2014) and Kehrig and Vincent (2019).

References

Adamopoulos, Tasso, Loren Brandt, Jessica Leight, and Diego Restuccia. 2022. "Misallocation, Selection, and Productivity: A Quantitative Analysis With Panel Data From China." *Econometrica* 90 (3): 1261–1282.

Aragón, Fernando, Diego Restuccia, and Juan Pablo Rud. 2022. *Assessing Misallocation in Agriculture: Plots versus Farms.* Working paper w29749. Cambridge, MA: National Bureau of Economic Research, February.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Dynamic Inputs and Resource (Mis)Allocation." *Journal of Political Economy* 122 (5): 1013–1063.

Ayerst, Stephen, Loren Brandt, and Diego Restuccia. 2023. "Distortions, Producer Dynamics, and Aggregate Productivity: A General Equilibrium Analysis." *SSRN Electronic Journal.*

Banerjee, Abhijit, and Lakshmi Iyer. 2005. "History, institutions, and economic performance: The legacy of colonial land tenure systems in India." *American economic review* 95 (4): 1190–1213.

Biau, Gérard, and Erwan Scornet. 2016. "A Random Forest Guided Tour." *TEST* 25 (2): 197–227.

Bils, Mark, Peter J. Klenow, and Cian Ruane. 2021. "Misallocation or Mismeasurement?" *Journal of Monetary Economics* 124:S39–S56.

Bischl, Bernd, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, et al. 2021. "Hyperparameter Optimization: Foundations, Algorithms, Best Practices and Open Challenges." Preprint, November 24, 2021. Accessed March 25, 2023. arXiv: 2107.05847 [cs, stat]. http://arxiv.org/abs/2107.05847.

Blanco, Cesar, and Xavier Raurich. 2022. "Agricultural Composition and Labor Productivity." *Journal of Development Economics* 158:102934.

Bolhuis, Marijn, Swapnika Rachapalli, and Diego Restuccia. 2021. *Misallocation in Indian Agriculture.* Working Paper 709. University of Toronto Department of Economics, October.

Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45 (1): 5–32.

Chanana-Nag, Nitya, and Pramod K Aggarwal. 2020. "Woman in agriculture, and climate risks: hotspots for development." *Climatic Change* 158 (1): 13–27.

Chen, Chaoran. 2017. "Untitled Land, Occupational Choice, and Agricultural Productivity." *American Economic Journal: Macroeconomics* 9, no. 4 (1, 2017): 91–121.

Chen, Chaoran, Diego Restuccia, and Raul Santaeulalia-Llopis. 2022. *Land Misallocation and Productivity.* Working paper w23128. Cambridge, MA: National Bureau of Economic Research, March.

Deininger, Klaus, Songqing Jin, and Hari K Nagarajan. 2008. "Efficiency and equity impacts of rural land rental restrictions: Evidence from India." *European Economic Review* 52 (5): 892–918.

Deininger, Klaus W, et al. 2003. *Land policies for growth and poverty reduction.* World Bank Publications.

Foster, Andrew D, and Mark R Rosenzweig. 2002. "Household division and rural economic growth." *The Review of Economic Studies* 69 (4): 839–869.

———. 2022. "Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size." *Journal of Political Economy* 130, no. 3 (1, 2022): 636–680.

Gandhi, Amit, Salvador Navarro, and David Rivers. 2017. *How Heterogeneous Is Productivity? A Comparison of Gross Output and Value Added.* Working Paper.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "Agricultural Productivity Differences across Countries." *American Economic Review* 104, no. 5 (1, 2014): 165–170.

Gollin, Douglas, and Christopher Udry. 2021. "Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture." *Journal of Political Economy* 129 (1).

Gottlieb, Charles, and Jan Grobovšek. 2019. "Communal Land and Agricultural Productivity." *Journal of Development Economics* 138:135–152.

Gulati, Ashok, and Ritika Juneja. 2022. "Transforming Indian Agriculture." In *Indian Agriculture Towards 2030: Pathways for Enhancing Farmers' Income, Nutritional Security and Sustainable Food and Farm Systems*, 9–37. Springer Nature Singapore Singapore.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *The Quarterly Journal of Economics* 124 (4): 1403–1448.

Jaef, Roberto N. Fattal. 2018. "Entry and Exit, Multiproduct Firms, and Allocative Distortions." *American Economic Journal: Macroeconomics* 10, no. 2 (1, 2018): 86–112.

Just, Richard E., David Zilberman, and Eithan Hochman. 1983. "Estimation of Multicrop Production Functions." *American Journal of Agricultural Economics* 65 (4): 770–780.

Kebede, Hundanol A. 2020. *The gains from market integration: The welfare effects of new rural roads in Ethiopia.* Technical report. Mimeo, University of Virginia.

Kehrig, Matthias, and Nicolas Vincent. 2019. *Good Dispersion, Bad Dispersion.* Working paper w25923. Cambridge, MA: National Bureau of Economic Research, June.

Krishnaswamy, Nandita. 2018. *At What Price? Price Supports, Agricultural Productivity, and Misallocation.* Working paper.

Le, Kien. 2020. "Land Use Restrictions, Misallocation in Agriculture, and Aggregate Productivity in Vietnam." *Journal of Development Economics* 145:102465.

Morando, Bruno. 2023. "Subsistence Farming and Factor Misallocation: Evidence from Ugandan Agriculture." *The World Bank Economic Review* (6, 2023).

Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives* 31, no. 2 (1, 2017): 87–106.

Restuccia, Diego, and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics* 11 (4): 707–720.

———. 2017. "The Causes and Costs of Misallocation." *The Journal of Economic Perspectives* 31 (3): 151–174.

Schonlau, Matthias, and Rosie Yuyan Zou. 2020. "The Random Forest Algorithm for Statistical Learning." *The Stata Journal: Promoting communications on statistics and Stata* 20 (1): 3–29.

Shumway, C. Richard, Rulon D. Pope, and Elizabeth K. Nash. 1984. "Allocatable Fixed Inputs and Jointness in Agricultural Production: Implications for Economic Modeling." *American Journal of Agricultural Economics* 66 (1): 72–78.

Singh, Inderjit, Lyn Squire, and John Strauss. 1986. "A survey of agricultural household models: Recent findings and policy implications." *The World Bank Economic Review* 1 (1): 149–179.

Sotelo, Sebastian. 2020. "Domestic trade frictions and agriculture." *Journal of Political Economy* 128 (7): 2690–2738.

Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28, no. 2 (1, 2014): 3–28.

Vashishtha, Prem S. 1989. "Changes in Structure of Investment of Rural Households: 1970-71-1981-2." *Journal of Income and Wealth* 10 (2): 21–45.

Wang, Wenya, and Ei Yang. 2023. "Multi-Product Firms and Misallocation." *Journal of Development Economics* 163:103102.

Ziegler, Andreas, and Inke R. König. 2014. "Mining Data with Random Forests: Current Options for Real-World Applications: Mining Data with Random Forests." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 4 (1): 55–63.

Appendix

A Additional Figures and Tables

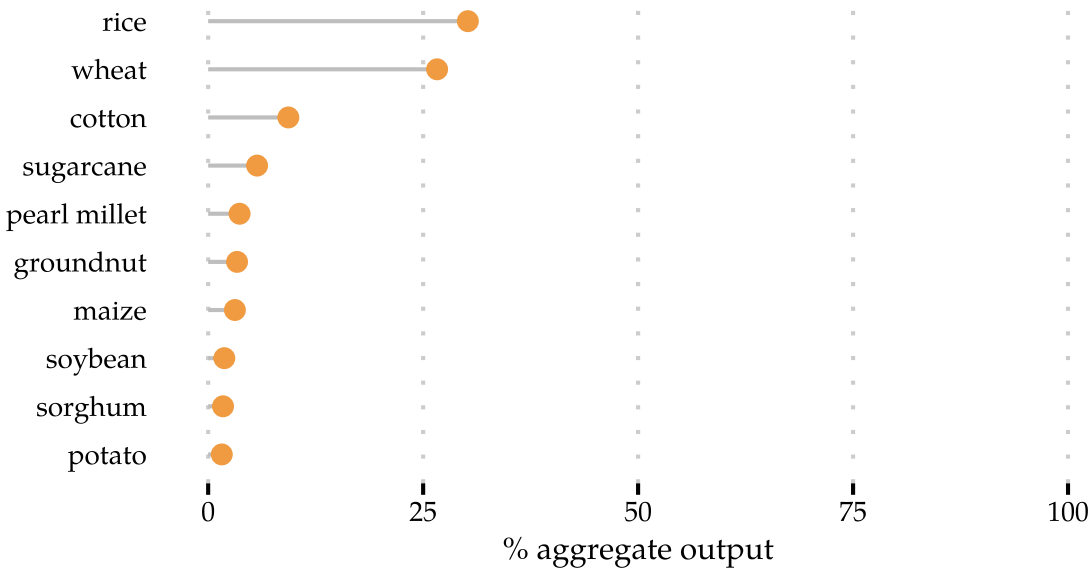

Figure A.1: Top 10 crops by % of aggregate output

TABLE A.1: Production function estimates

	(1) Aggregate	(2) Rice	(3) Wheat	(4) Other Cereals	(5) Pulses	(6) Veg, Frt, Oil
Land	0.427	0.511	0.414	0.407	0.379	0.330
	(0.031)	(0.021)	(0.032)	(0.034)	(0.060)	(0.035)
Labor	0.189	0.161	0.122	0.248	0.316	0.430
	(0.031)	(0.037)	(0.036)	(0.058)	(0.074)	(0.064)
Intermediates	0.351	0.307	0.517	0.245	0.255	0.127
	(0.041)	(0.033)	(0.040)	(0.046)	(0.076)	(0.052)
Observations	14,705	4,807	3,566	2,779	1,128	2,338
R^2	0.624	0.742	0.713	0.590	0.417	0.572
Village FEs	Y	Y	Y	Y	Y	Y
Season FEs	Y	Y	Y	Y	Y	Y
			First Stage: F statistics			
Land	77.0	62.0	40.3	37.8	15.7	19.3
Labor	49.3	34.7	17.7	25.2	12.9	14.8
Intermediates	35.8	31.7	21.5	19.9	8.9	11.8
K-Paap Wald F statistic	51.1	40.4	16.0	30.8	12.4	12.7

Note. Each column represents a separate regression. We use 2SLS to estimate the coefficients, using instruments described in Section 4.2. Robust standard errors are presented in parentheses. The "Other Cereals" group include barley, maize, jawar, bajra, ragi, and millets. The "Veg, frt, oil" group includes vegetables, fruits, oilseeds, and fiber crops. The "Aggregate" column pools all crops together. In panel B of the table, we report the F-stat from the first stage regression for each input and each regression separately. We also report the joint F-stat at the bottom line of the table.

TABLE A.2: Estimate of σ

	(1) $\log s_{h,i}$	(2) $\log p_{h,i}$
σ	1.699	
$\log p_{h,i}$	-0.699	
	(0.067)	
Elevation × rain		-0.004
		(0.000)
Ruggedness × rain		0.016
		(0.001)
Pucca roads availability		-0.023
		(0.001)
Observations	40,833	40,833
Kleibergen-Paap F stat	230.9	

Note. The table presents the estimation of Equation 16 using 2SLS. Column 1 presents second stage results and column 2 presents first stage results. Elevation, ruggedness, rain, and availability of pucca roads are village-level measures. "Pucca roads" are all-weather roads made of concrete and tar. Standard errors are clustered at the household level. The "σ" row shows the value of σ implied by β_1, the coefficient on $\log p_{h,i}$.

TABLE A.3: Crop list

Rice	**Wheat**	**Other Cereals**	**Pulses**	**Oilseeds, Fruits and Vegetables**		
		Barley	Black gram	**Oilseeds**	**Vegetables**	**Fruits /Condiments**
		Maize	Green peas	Sesame	Ash gourd	Mango
		Sorghum	Pigeon peas	Groundnut	Beet root	Papaya
		Pearl millet	Horse gram	Castor	Bitter gourd	Grapes
		Finger millet	Cowpea	Sunflower	Bottle gourd	Plum
		Others	Kidney bean	Niger	Eggplant	Cardamom
			Lentil	Soybean	Board bean	Chilli
			Chickpeas	Safflower	Cabbage	Cumin
			Others	Rapseed	Cauliflower	Dill seed
				Linseed	Carrot	Indian mustard
				Others	Potato	Other
					Cucumber	
					Peas	

TABLE A.4: Reallocation exercises: within state

	Gain	Gain (1 product)	Error (1 product)
benchmark	**107%**	**124 %**	**15%**
land only	20%	26 %	29%
complete	122%	124 %	1%

Note. Unlike the main exercises in Table 2, the present exercises only reallocate inputs within each state: effectively, separate states are treated as separate economies. The nomenclature follows that of Table 2.

TABLE A.5: Reallocation exercises: no split by season

	Gain	Gain (1 product)	Error (1 product)
benchmark	**314%**	**260 %**	**-17%**
best state	41%	46 %	11%
land only	30%	36 %	19%
complete	330%	260 %	-21%

Note. Unlike the main exercises in Table 2, the present exercises reallocate inputs across seasons: separate seasons are *not* treated as separate economies. The nomenclature follows that of Table 2.

B ADDITIONAL CONCAVITY MECHANISMS

Additional concavity governed by the parameter η, described in Section 5.1, is a reduced-form stand-in for multiple mechanisms that could incentivize farmers to produce multiple crops at once.

In this section, we explore two such mechanisms explicitly. Both generate first-order conditions that are functionally equivalent to those in the original problem. At the same time, different micro-foundations of η may carry with them different natural ways of aggregating crops and clearing markets at the aggregate level, generating slightly different general equilibrium implications.

B.1 SUBSISTENCE

Suppose that instead of operating a profit-maximizing farm, the household is operating a subsistence farm and consuming its output. The farmer is minimizing costs

of providing utility constant elasticity of substitution utility U to the family:

$$U = \left(\sum_i \varphi_i y_{f,i}^{\eta} \right)^{\frac{1}{\eta}}$$

In this interpretation, the farmer has an incentive to mix crops in order to provide a diverse diet for the family. The resulting first-order conditions are equivalent to those in Section 5.1.

B.2 Market Power

Suppose that instead of selling the harvest at a perfectly competitive market, the farmer has some market power in a monopolistically competitive market. This can be represented by introducing an intermediate crop aggregator operating within each crop i's market, which buys farms' varieties $y_{f,i}$ of this crop at $p_{f,i}$ and sells the combined i product to the consumer at price P_i. The aggregator uses the following production function:

$$Y_i = \left(\sum_f y_{f,i}^{\eta} \right)^{\frac{1}{\eta}}$$

This generates downward-sloping demand for farm f's variety of crop i, creating an incentive for the farmer to mix crops (to avoid excessive price drops in any one crop). The resulting first-order conditions are equivalent to those in Section 5.1

C Model Solution and Computational Algorithm

C.1 Solving the Problem of a Single Farm

For N products, there are $2^N - 1$ non-empty product sets that each farm could be producing. Let I_f be the set of products in a single such combination, s.t. $y_{f,i} > 0 \; \forall i \in I_f$ and $y_{f,i} = 0 \; \forall i \notin I_f$. If $\omega > 0$, the profit-maximizing product set might have $y_{f,i} = 0$ for some crops. Each combination has to be solved separately, and then their profits are compared to pick the best one.

For each product set I_f, the following non-linear equation needs to be solved for

the shadow cost of land λ_f:

$$\sum_{i \in I_f} \lambda_f^{\frac{\eta \sum_g \alpha_{g,i} - 1}{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}} \left((p_i z_{f,i})^\eta \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \left(\frac{\gamma_i}{\tau_{l,i}} \right)^{1 - \eta \sum_g \alpha_{g,i}} \right)^{\frac{1}{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}} \tau_{l,i} = L_f \tag{C.1}$$

Each term within the sum is decreasing in λ_f. Using this fact, the following bounds can be derived on λ_f that solves the condition, simplifying the numerical cost. For each crop, define:

$$A_i = \left((p_i z_{f,i})^\eta \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \left(\frac{\gamma_i}{\tau_{l,i}} \right)^{1 - \eta \sum_g \alpha_{g,i}} \right)^{\frac{1}{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}} \tau_{l,i} \tag{C.2}$$

Then the upper and lower bounds on λ_f are

$$\lambda_f^{UB} = \max_i \left(\frac{L_f}{|I_f| \cdot A_i} \right)^{\frac{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}{\eta \sum_g \alpha_{g,i} - 1}} \tag{C.3}$$

$$\lambda_f^{LB} = \min_i \left(\frac{L_f}{|I_f| \cdot A_i} \right)^{\frac{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}{\eta \sum_g \alpha_{g,i} - 1}} \tag{C.4}$$

Once the λ_f for this product set is known, input choices *for products within* I_f can be obtained from:

$$x_{f,h,i} = \frac{\alpha_{h,i}}{r_h \tau_{f,h} \tau_{f,h,i}} \left((p_i z_{f,i})^\eta \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \left(\frac{\gamma_i}{\lambda_f \tau_{l,i}} \right)^{\eta\gamma_i} \right)^{\frac{1}{1 - \sum_g \eta \alpha_{g,i} - \eta\gamma_i}} \tag{C.5}$$

$$l_{f,i} = \left((p_i z_{f,i})^\eta \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \left(\frac{\gamma_i}{\lambda_f \tau_{l,i}} \right)^{1 - \eta \sum_g \alpha_{g,i}} \right)^{\frac{1}{1 - \eta \sum_g \alpha_{g,i} - \eta\gamma_i}} \tag{C.6}$$

Set $x_{f,h,i} = l_{f,i} = 0 \ \forall i \notin I_f$.

With input choices computed, the profit for this product set can be evaluated using Equation 2. Iterate over $2^N - 1$ possible product sets and choose the product set yielding the highest profit.

C.2 Solving for the General Equilibrium

Merge consumer's product demands, farms' input demands, farms' optimal output, and market clearing conditions to obtain market clearing conditions in terms of fundamentals, shadow costs of land λ_f, goods prices p_i, inputs prices x_g, and farms' product sets I_f (letting the set of farms producing each product to be denoted by F_i):

$$\sum_f \sum_{i \in I_f} \left((p_i z_{f,i})^{\eta} \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \cdot \left(\frac{\gamma_i}{\lambda_f \tau_{f,l,i}} \right)^{1-\eta \sum_g \alpha_{g,i}} \right)^{\frac{1}{1-\eta \sum_g \alpha_{g,i} - \eta \gamma_i}} = L^{agg} \quad \text{(C.7)}$$

$$\sum_f \sum_{i \in I_f} \frac{\alpha_{h,i}}{r_h \tau_{f,h} \tau_{f,h,i}} \left((p_i z_{f,i})^{\eta} \cdot \eta \cdot \Pi_g \left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\eta \alpha_{g,i}} \cdot \left(\frac{\gamma_i}{\lambda_f \tau_{l,i}} \right)^{\eta \gamma_i} \right)^{\frac{1}{1-\eta \sum_g \alpha_{g,i} - \eta \gamma_i}} = X_h^{agg} \quad \text{(C.8)}$$

$$p_i = \left(\varphi_i^{-\sigma} \frac{\sum_j \varphi_j^{\sigma} p_j^{1-\sigma}}{\sum_j p_j \sum_{f \in F_j} y_{f,j}} \cdot \sum_{f \in F_i} \left(\eta^{\sum_g \alpha_{g,i} + \gamma_i} z_{f,i} \Pi_g \left[\left(\frac{\alpha_{g,i}}{r_g \tau_{f,g} \tau_{f,g,i}} \right)^{\alpha_{g,i}} \right] \left(\frac{\gamma_i}{\lambda_f \tau_{f,l,i}} \right)^{\gamma_i} \right)^{\frac{1}{1-\eta \sum_g \alpha_{g,i} - \eta \gamma_i}} \right)^{\frac{1-\eta \sum_g \alpha_{g,i} - \eta \gamma_i}{\left(1-\eta \sum_g \alpha_{g,i} - \eta \gamma_i\right)(1-\sigma)-1}} \quad \text{(C.9)}$$

The p_i condition can be simplified in two ways, to abstract from its dependence on all other goods' prices p_j. First, the price level can be normalized: $\sum_j \varphi_j^{\sigma} p_j^{1-\sigma} = 1$. Second, the nominal output can be substituted with a variable Y_{nom}, adding its identity to the list of market-clearing conditions:

$$Y_{nom} = \sum_j p_j \sum_{f \in F_j} y_{f,j} \quad \text{(C.10)}$$

Splitting distortions and ensuring the current allocation is consistent with general equilibrium. Once frictions have been extracted from the data (Section 5.3), additional steps need to be taken to ensure that the model representation of the current allocation is in general equilibrium.

Splitting $r_g \tau_{f,g} \tau_{f,g,i} = mrpg_{f,i}$ into r_g and $\tau_{f,g}$ is arbitrary. To ease the interpretation of frictions, we impose that $\tau_{f,g} \tau_{f,g,i}$ only distort the farm-level input demands and do not affect the aggregate input demand X_g^{agg}. This implies the following market price of g:

$$r_g = \frac{\sum_i \eta \alpha_{g,i} \sum_f (p_i y_{f,i})^\eta}{X_g^{agg}} \tag{C.11}$$

Extending the logic, we impose that $\tau_{f,g,i}$ only distort the within-farm allocation of inputs between crops and do not affect the farm-level input demand, implying the following:

$$\tau_{f,g,i} = \frac{\eta \alpha_{g,i} (p_i y_{f,i})^\eta}{\sum_j \eta \alpha_{g,j} (p_j y_{f,j})^\eta} \frac{\sum_j x_{f,g,j}}{x_{f,g,i}} \tag{C.12}$$

Finally, the unobserved taste weights φ_i need to be consistent with the observed allocation being in general equilibrium. To simplify notation, we map the price-weighted outputs in the data to real output $y_{f,i}$ in the model, implying that all prices in the current model allocation are: $\{p_i\} = 1$. Any such rescaling does not affect the real behavior of the model, neither in the current allocation nor in the counterfactual reallocations, since units of output in different products are arbitrary. This rescaling simplifies extracting the taste weights consistent with the current allocation being in general equilibrium to:

$$\varphi_i = \left(\frac{\sum_f y_{f,i}}{\sum_j \sum_f y_{f,j}} \right)^{\frac{1}{\sigma}} \tag{C.13}$$

Note that this procedure makes extracting taste weights from estimated fixed effects in Equation 16 unnecessary.

When Product Sets are Fixed. Although the current allocation extracted from the data can be easily made to be consistent with the general equilibrium in the model, market-clearing prices have to be solved numerically for any counterfactual reallocation. When product sets are fixed at the currently observed level, the problem of finding market-clearing prices using Equations C.7-C.9 is simplified by the fact that farmers cannot switch their product sets in response to changing prices.

If at least land is being efficiently reallocated between farms (λ_f is equalized between farms), the market-clearing prices can be solved without having to re-solve each individual farm's problem.[22] This is the case in the benchmark reallocation exercise discussed in Section 6.1.

In this case, farms' product sets I_f are known from the current allocation. Since land is being reallocated, $\lambda_f = \bar{\lambda}\ \forall i$. This allows the land market condition C.7, the flexible input markets conditions C.8 and the input market identity C.10 to be solved as a system with $G + 2$ dimensions in $G + 2$ unknowns: $\bar{\lambda}$, $\{r_g\}_g$, and Y_{nom}. On each guess of these variables, goods prices consistent with the guess can be computed from C.9.

When Product Sets are Allowed to Change. When farms are allowed to change their product set I_f from its current level, the need to solve all individual farms' problems on each price guess for new product sets can no longer be avoided.

By analogy with flexible inputs (Section 5.3), let each farm's land shadow cost deviation be fixed at $\hat{\lambda}_f = \lambda_f / \bar{\lambda}$.

The brute-force approach to solving for market-clearing prices is to solve conditions C.7, C.8, C.9 as a system with $G + N + 1$ dimensions in the same number of unknowns: $\bar{\lambda}, \{x_g\}_g, \{p_i\}_i$. Within each price guess:

1. Solve individual farms' problems. For each farm:
   (a) For each potential product set I_f, find the land endowment L_f consistent with $\lambda_f = \bar{\lambda} \cdot \hat{\lambda}$ using Equation C.1.
   (b) Solve the farm's problem for the optimal product set I_f, following Appendix C.1.
2. Evaluate conditions C.7, C.8, C.9.

The brute-force algorithm is guaranteed to find market-clearing prices but is exceptionally computationally costly due to its high-dimensionality and the need to re-solve all farms' problems on each price guess. However, there is a much faster sequential heuristic algorithm:

1. Pick an initial guess of $\bar{\lambda}, \{x_g\}_g, \{p_i\}_i$.

22. This approach works whether or not any of the flexible input frictions $\tau_{f,g}$, $\tau_{f,g,i}$ are being equalized.

2. Solve individual farms' problems as in the brute-force algorithm.

3. Solve the system of market conditions as in the fixed crop set algorithm for new price guesses: this assumes that product sets do not respond to changing prices.

4. Solve individual farms' problems as in the brute-force algorithm with price guesses obtained from the fixed crop set GE system.

5. Evaluate convergence in product sets. If all farms' optimal product sets obtained in step 4 are the same as in step 2, stop. If not, return to step 2.

This algorithm exploits the low-dimensionality and low computational cost of solving for general equilibrium when product sets are fixed. By sequentially updating the farms' product sets, this algorithm converges to the same answer as the brute-force approach in just a few iterations.

D Fixed Input Shares Model

An auxiliary model in which the relative allocation of inputs across crops within a farm is fixed is useful for shutting down the effect of the endogenous product choice mechanism (Section 6.3). In this model, the farmer can choose the overall quantity of each input hired, but the share of the input g allocated to each crop i is fixed at exogenously given $s_{f,g,i}$. The set of produced crops is likewise fixed at I_f. The setup of the model is otherwise similar to that in Section 5.1:

$$\max_{\{y_{f,i}\}_i, L_f^d, \{X_{f,g}\}_g} \sum_{i \in I_f} \left(p_i y_{f,i}\right)^{\eta} - \sum_{g=1}^{G} r_g \tau_{f,g} \sum_{i \in I_f} \tau_{f,g,i} s_{f,g,i} X_{f,g} \tag{D.1}$$

s.t.

$$y_{f,i} = z_{f,i} \left(s_{f,l,i} L_f^d\right)^{\gamma_i} \Pi_g \left[\left(s_{f,g,i} X_{f,g}\right)^{\alpha_{g,i}}\right] \tag{D.2}$$

$$\sum_{i \in I_f} s_{f,l,i} \tau_{f,l,i} L_f^d = L_f^s \qquad (\lambda_f) \tag{D.3}$$