Distributional Effects of Indian Agricultural Interventions

Shresth Garg[*] Sagar Saxena[†]

July 2023

Abstract

How do government programs that distort prices in agricultural markets affect producers and consumers along the income distribution? We study the distributional effects of three such programs in Indian agricultural markets: fertilizer subsidies, procurement of crops at minimum support prices (MSP), and sale of subsidized grains to households. These interventions directly impact hundreds of millions of people and cost about 1.2% of India's GDP. To examine their effects, we estimate a structural model of supply and demand with heterogeneous risk-averse producers, who choose a portfolio of crops and crop-specific inputs, and heterogeneous households who make consumption decisions. Using the estimated structural parameters, we solve for counterfactual equilibria in which these interventions are phased out. On the demand-side, we find these programs to be progressive. In their absence, consumption and expenditures of lower-income households would be affected more adversely. On the supply-side, we find these programs to be (weakly) regressive. Higher fertilizer prices, in the absence of subsidies, would be compensated by higher output prices so impact on farmer welfare would be minimal. Under no government-procurement at MSP, richer farmers would experience a greater welfare loss, while some of the poorest farmers would gain – a result driven partly by the inequitable implementation of the procurement program.

[*]Harvard University, Department of Economics. Email: garg@g.harvard.edu
[†]Yale University, Economic Growth Center. Email: sagar.saxena@yale.edu

1 Introduction

Government programs that distort prices in agricultural markets, such as input subsidies and price supports, are ubiquitous.[1] A key objective of such programs is redistribution,[2] which leads one to ask: *how do these price interventions affect market participants along the income distribution?* Yet, assessing these distributional effects is difficult. These programs are typically executed at a large scale and their equilibrium impact may amplify or dampen any direct effects on market participants. For instance, by lowering costs, input subsidies may positively affect farm profits but the equilibrium increase in aggregate output may decrease output prices enough to hurt profits.[3]

In this paper, we propose a structural model that enables us to examine large-scale agricultural interventions while accounting for spillover and equilibrium effects. We use our model to study the distributional effects of multiple price interventions along the agricultural supply chain in India. At the start of this supply chain, the government sells subsidized fertilizers to farmers. Next, upon harvest, the government buys a substantial share of the total output of key crops such as rice and wheat at prices known as *minimum support prices* (MSP); all other sales are made at market prices to private traders. Finally, the government sells subsidized foodgrains to households, subject to progressive income-based quotas, through the *public distribution system* (PDS). Jointly, these programs cost about 1.2% of India's GDP and impact nearly 800 million people.[4]

Our structural model closely follows the setup of the Indian agriculture sector. In our model, risk-averse farmers choose a portfolio of crops to plant and make crop-specific input allocations. Post-harvest, they sell output either to government agencies or to private traders. On the demand-side, households receive PDS entitlements from the government and make consumption decisions in the private market. In equilibrium, total PDS entitlements equal the sales made to government agencies and total household demand in the private market equals the sales made

[1]Input subsidies lower the costs of farm inputs such as fertilizers and seeds, while price supports typically serve as price floors at which farmers may sell output to government agencies. All 54 countries studied in OECD (2022), including the 38 OECD countries, have programs which provide support to the agriculture sector. The study excludes African nations; for an overview of similar programs in Africa, see Holden (2019).

[2]See Acemoglu and Robinson (2001) and OECD (2012).

[3]Large programs, in a variety of contexts, often generate equilibrium effects. Examples of papers studying such effects include Duggan and Morton (2006), Imbert and Papp (2015), Cunha et al. (2019), Egger et al. (2019), Rotemberg (2019), Breza and Kinnan (2021), Muralidharan et al. (2022), and Khanna (2022).

[4]See World Bank (2019).

to private traders. We estimate model parameters by matching simulated moments with empirical moments from publicly-available farmer- and household-level microdata. Finally, using estimated parameters, we simulate counterfactuals in which we phase out (1) fertilizer subsidies, and (2) government procurement at MSP and PDS entitlements.

We find these interventions to be progressive on the demand-side, and (weakly) regressive on the supply-side; in their absence, lower-income consumers and higher-income producers are affected more adversely. By raising output (and procurement) and lowering prices, these programs greatly benefit lower-income consumers who we find to be more price elastic and more reliant on PDS entitlements. For producers, we find that direct gains from fertilizer subsidies are nearly offset by equilibrium changes in market prices, so any impact on farmer welfare is minimal. In contrast, large-farmer bias in government-procurement at MSP accords greater benefits of this program to wealthier farmers, thus making it regressive. In an additional counterfactual where within-region bias for larger farmers is eliminated, we find substantial gains for smaller farmers.

Below, we summarize the main sections of this paper in more detail.

Motivating evidence. In Section 2, we present a mix of causal and descriptive evidence that motivates our model. First, using a natural experiment wherein subsidies for non-urea fertilizers were partially phased out, we show that subsidies affect production decisions. We also provide suggestive evidence that minimum support prices (MSP) influence production decisions. Correspondingly, in our model, we allow farmers' planting decisions to be determined by these programs. Next, we provide descriptive evidence to show that a large share of farmers sell at prices well below the MSP, and that there are stark income and spatial inequities in sales made to government buyers – larger farmers and farmers located in some regions are more likely to sell to government buyers.[5] Importantly, MSP appears to have an impact on production decisions only when sales to government buyers are likely. We add these findings to our model by allowing the likelihood of encountering a government buyer to depend on crop, location, and farmer size; this likelihood, in turn, determines how MSP affects farming decisions.

[5] Regional differences may be due to two reasons. First, procurement is still heavily reliant on infrastructure set up in the 1960s when only a few states produced surplus rice and wheat that could be procured. Second, in recent years, some states have introduced independent procurement schemes which only benefit farmers located in those states. There's little systematic evidence to explain the bias in favor of larger farmers. Conversations with local researchers suggest that corruption and bribery may explain part of this bias.

Model. In Section 3, we develop a model of multiproduct producers with endogenous product and input choice.[6] Farmers differ by productivity, location, and wealth (proxied by farm size). Given fertilizer subsidies and minimum support prices, they choose a set of crops to plant and make crop-specific area and input allocations subject to a farm size constraint. Farmers make these choices to maximize a mean-variance payoff function mediated by a farmer-specific risk aversion parameter. Finally, farmers also pay a fixed cost for the set of crops they plant.

At the time of planting, farmers face both output and price risk. Output risk shows up in crop-specific production functions in the form of idiosyncratic output shocks which scale output in a Hicks-neutral sense. Consequently, higher input usage yields greater output variance – a force which leads risk-averse farmers to moderate demand for inputs.[7] Importantly, input subsidies help offset this force and promote greater input usage.

Price risk arises from uncertainty over the price offered by private buyers and the uncertainty in accessing government buyers. Private buyer offers are distributed around an average private market price for each crop; the realized offer depends on post-harvest realization of an output quality shock (e.g. dust and moisture content).[8] Upon harvest, farmers encounter government buyers with a likelihood that depends on farmer size, location, and crop while a private buyer is always accessible. If a government buyer is present, the farmer sells to the government buyer if the minimum support price (MSP) is greater than the private buyer offer; else sales are made to the private buyer. Thus, for farmers who are more likely to find a government buyer, MSP provides greater insurance against downside price risk.

We allow both risk and risk aversion to differ by farmer. At the time of planting, these differences induce different choices on both the extensive and intensive margins: farmers may choose different bundles of crops and, even for the same bundle, may allocate different shares of their land to each crop. This risk channel, therefore, is an important determinant of how supply-side price interventions affect aggregate production and individual farmer welfare.

[6] Similar in spirit to Wollmann (2018) which considers a setting with oligopolistic producers.

[7] Presence of output risk combined with a lack of risk-mitigating technologies is a known source of underinvestment in farm inputs. See Rosenzweig and Binswanger (1993), Rosenzweig and Wolpin (1993), Mobarak and Rosenzweig (2013), Karlan et al. (2014), Cole et al. (2017), and Donovan (2021)

[8] In addition to output quality, other reasons such as transportation costs, storage costs, and intermediary market power/bargaining power may also explain the cross-sectional variance in prices in the private market but we do not model these. We assume that quality shocks are the only source of variance in private buyer offers; further, these only affect processing costs – lower quality crops have higher processing costs and therefore receive a lower private buyer offer. Finally, processed crops purchased by households in the private market do not differ in quality.

On the demand-side, households differ by income and entitlements from the public distribution system (PDS).[9] Quantity procured by the government is redistributed to households.[10] Residual demand, which depends on both income and PDS entitlements, is satisfied in the private market where households pay the average private market price, determined in equilibrium.

Estimation. We estimate the supply-side of our model in three steps. We rely primarily on publicly-available data from Cost of Cultivation Surveys (CCS) from 2008-2016, conducted each planting season by the Department of Agriculture in India. These include detailed information on prices, crop portfolio, and crop-specific input allocations for each farmer-season.

We begin by estimating parameters governing the distribution of price risk. Two challenges arise. First, private market prices below MSP are only observed if a government buyer is not found. Second, whether a government buyer is found is unobserved. While we can construct the likelihood of *selling* to government buyers – by crop, region, and farmer size – from administrative datasets,[11] this is not equal to the likelihood of *finding* a government buyer since farmers may sell to a private buyer if his offer is greater than MSP. We proceed with the help of a simulation-based estimator, described in detail in Section 4, which yields parameters that determine the likelihood of finding a government buyer and the distribution of private buyer offers. These allow us to assess the ex ante crop-, location- and farmer size-specific price risk faced by farmers.

Next, we estimate crop-specific production functions and the distribution of risk aversion which may depend on farmer size. Given a set of crops, these affect how farmers allocate land, labor, capital, and fertilizer to each crop in the set. Observed input choices and output are also influenced by unobserved farmer productivity (which we account for using farmer fixed effects) and the distribution of output shocks.[12] We estimate parameters using method of simulated moments: for each farmer, we solve the optimal portfolio choice problem for the observed set of crops planted and match simulated choices with observed moments that summarize crop-specific

[9] In addition to income, PDS entitlements may depend on household location. In our counterfactuals, we hold the targeting of the PDS system fixed.

[10] We treat these as in-kind transfers at zero cost to households.

[11] Data on the likelihood of selling to government buyers by farmer size, crop, and location are from the 77th round of the National Sample Survey (NSS) conducted in 2019.

[12] Note that standard production function estimation techniques fail given risk-averse farmers. These usually rely on a monotonicity assumption between productivity shocks and input demand. Productivity shocks, however, increase variance of output. For a risk-averse farmer, this yields a non-monotonic relationship between productivity and input demand; depending on the size of a positive shock, the farmer may choose to increase or decrease input demand.

output, land share patterns, and input usage.

Finally, we estimate the fixed cost of planting. The fixed cost for a crop is independent of the level of area allocated to that crop and depends on whether the crop is a *new* crop for a farmer. If it was part of his portfolio in the previous year, this cost is discounted by a parameter we estimate. The choice of which set of crops to plant is akin to a discrete choice problem where the choice set is composed of sets of crops. Once distribution of prices, production function parameters, output risk, and risk-aversion are known, given some guess of fixed cost parameters, we can simulate choice probabilities for each set of crops. We match these with the probability of observing a given set of crops in the data to estimate fixed cost parameters.

Once supply-side parameters are known, we can compute, for each farmer, the optimal set of crops and crop-specific inputs for any given input and (distribution of) output prices. This allows us to trace out aggregate supply curves for the private market and the government stockpile ("PDS supply") as a function of private market prices.[13] To pin down equilibrium private market prices, we require aggregate demand curves for the private market. For PDS crops, rice and wheat, we estimate demand using household-level data, accounting for PDS entitlements and income.[14] This yields an aggregate private market demand curve for each level of government stockpile (or PDS entitlements). For non-PDS crops, we use demand elasticities from Deaton (1997).

Main results. In Section 5, we evaluate the distributional effects of fertilizer subsidies, government procurement at minimum support prices (MSP), and in-kind transfers through the public distribution system (PDS) using two counterfactuals which shut down these interventions.[15]

In the first counterfactual, we phase-out fertilizer subsidies. In the data, we approximate an average subsidy rate of 50% for all fertilizer products. As such, in our counterfactual, we double fertilizer prices and solve for a vector of private market prices which clear all markets.

[13] We abstract away from modeling how MSP is set. Motivated by data, we assume that MSP tracks prices in the private market. Specifically, we assume a level of MSP such that conditional on finding a government buyer, only 35% of farmers would sell to private buyers i.e. MSP is set at the 65th percentile of the private buyer offer distribution.

[14] Data are from the 68th round of the National Sample Survey (2011). To deal with potential endogeneity of prices in our estimation, we construct Hausman-style price instruments (Hausman et al., 1994) by computing state-level average prices excluding the district in which each household resides. These are valid instruments under the assumption of idiosyncratic district-level demand shifters but correlated state-level supply shifters (e.g. processing costs).

[15] While useful for evaluating these programs, these counterfactuals also help us understand the impact of some of the proposed reforms to these programs in public and political debates. Both fertilizer subsidy program and government procurement at MSP are highly contentious topics with many in favor of shrinking their scale and scope.

We find equilibrium output of all crops falls and private market prices go up ($\approx 5\%$ for rice and wheat). For farmers, the impact of higher fertilizer prices is dampened by higher equilibrium private market prices. Therefore, we find a minimal impact on farmer welfare across the size distribution.

In the second counterfactual, we shut down government-procurement at MSP. Correspondingly, PDS entitlements of households also go to zero.[16] Now, farmers can only sell in the private market and households must satisfy all demand in the private market. Thus, both supply and demand in the private market go up. But farmers are also exposed to greater price risk now, especially those who were previously more likely to find government buyers. For these farmers, MSP was a meaningful price floor that protected against low private buyer offers. We find that in the absence of procurement, price of rice goes up ($\approx 5\%$) while there is a negligible change in the private market price of wheat. This differential response is driven by a greater fall in the total output of rice as some farmers switch to other crops due to greater estimated variance in private buyer offers for rice. These switchers are also large (they were previously more likely to find government buyers) and therefore have a noticeable impact on the aggregate output of rice. Despite switching, we find larger farmers have significantly larger welfare losses. Some of the smallest farmers experience modest gains as they now receive higher private market prices for rice.[17]

On the demand-side, in both counterfactuals, lower income households are more adversely affected as prices go up and PDS entitlements fall. These effects are driven by higher estimated price elasticities and higher observed PDS entitlements for lower-income households. Using a Laspeyres index, we show that without fertilizer subsidies, the lowest income households pay 3%-4% more to consume the old bundle of rice, wheat, and a numeraire good. In contrast, without government procurement at MSP, households pay 15%-20% more, which highlights the value of in-kind transfers. The impact on highest income households, in both scenarios, is close to zero.

Three implications for policy emerge. First, the incidence of large-scale government programs may depend crucially on the equilibrium channel. In our setting, fertilizer subsidies lower input

[16] This is an assumption. We can consider, for instance, a counterfactual where the government only shuts down procurement and offers consumption vouchers to households that can be redeemed in the private market.

[17] This result is partly driven by our assumption of perfect passthrough of output prices to farmers. The presence of intermediary market power would dampen this feedback effect.

costs for farmers but ultimately benefit only low-income consumers through lower output prices.

Second, in settings with multiple programs, a joint evaluation may be necessary to understand potentially important interactions. For example, we find that fertilizer subsidies not only lower output prices but also help increase government procurement by raising aggregate output. This highlights important complementarities in the two programs for the objective of improving food security of lower-income households.

Finally, implementation matters. Unequal access to government procurement at minimum support prices directly disadvantages small farmers. But it also indirectly hurts small farmers as greater procurement raises PDS entitlements, lowers private market demand and therefore, lowers prices in the private market. A more equitable program could, for example, match government buyers with farmers at random. We test the equilibrium impact of this alternative by eliminating within-region bias for larger farmers and find positive gains for smaller farmers.[18]

Related literature. Our work relates to a growing literature that uses structural models to study the agriculture sector in developing countries (Costinot et al., 2016; Sotelo, 2020; Allen and Atkin, 2022; Bergquist, Faber, et al., 2022; Chatterjee, 2022; Hsiao, 2022). We add to this literature by introducing a general simulation-based approach that integrates observational microdata on farmer- and household-level decisions in the estimation of structural models. Second, our paper relates to a large body of work that studies subsidies and transfers in agricultural markets, both on the supply-side (Duflo et al., 2008; Duflo et al., 2011; Karlan et al., 2014) as well as the demand-side (Banerjee et al., 2018; Banerjee et al., 2019; Gadenne, 2020; Gadenne et al., 2022). We contribute to this literature by jointly studying interventions that directly affect both producers and consumers. Finally, we provide new and timely empirical evidence on the impact of the largest agricultural interventions in India that are actively being discussed in public and political debates (Meenakshi and Banerji, 2005; Krishnaswamy, 2019; Gupta et al., 2021; Chatterjee et al., 2022).

[18] We do so by using the share of farmers in each region who found government buyers in the baseline as the uniform likelihood of finding government buyers in that region. This does not hold fixed the procurement in baseline since smaller farmers would be relatively more likely to find government buyers now. Finding a probability that holds procurement fixed is computationally non-trivial since it also affects farmers' production decisions.

2 Institutional Details: The Indian Agriculture Sector

The agriculture sector in India directly impacts the well-being and survival of over a billion people. On the supply side, nearly 300 million people rely on it for their livelihoods.[19] These agricultural households generally own small plots of farm land – average farm size in India is 2.8 acres compared to 445 acres in the United States USDA, 2021 – and have lower incomes (proxied by consumption expenditures in Figure 1a) relative to non-agricultural households. On the consumption-side, the agriculture sector supports a population of 1.4 billion, over 200 million of which are undernourished.[20]

Against this backdrop, several government-sponsored programs exist to support agricultural households and bolster food security in the country.[21] In this paper, we focus on three of the largest and longest-running such programs. These include fertilizer subsidies for agricultural use, government procurement of staple crops at pre-announced minimum support prices (MSP), and the redistribution of these crops at highly subsidized rates to low income households through the public distribution system (PDS).

These interventions date back to at least the mid-1960s. Newly-independent India faced severe food shortages, exacerbated by two successive drought years, and relied heavily on imports and foreign food aid to feed its rapidly growing population. To encourage greater production of foodgrains, the government started supplying farmers with high-yield variety seeds and heavily-subsidized fertilizers. In addition, the government promised attractive purchase prices for staples such as rice and wheat. These policies marked the beginning of the Green Revolution of the 1960s in India, during which yields increased many-fold and domestic production increased enough to allow India to become self-sufficient in food. Six decades after their introduction, these policies remain in place and make up a large share of the total budget of the central government – between 2010 and 2019, they amounted to 10% of annual government spending on average.[22]

[19]Estimated as the weighted sum of number of family members in households which report farming as their principal source of income in the 68th round of the National Samply Survey (2013); excludes agricultural labor.

[20]See FAO, IFAD, UNICEF, WFP and WHO (2021). Despite tremendous gains in agricultural production in the last few decades, malnutrition remains an issue. In the 2021 Global Hunger Index, India ranks 101st out of 116 countries. Rankings depend on the prevalence of undernourishment, childhood wasting, childhood stunting, and child mortality.

[21]Other programs, in addition to those studied in this paper, include subsidized crop insurance under Pradhan Mantri Fasal Bima Yojana (PMFBY), minimum income support for small and marginal farmers under Pradhan Mantri Kisan Samman Nidhi Yojana (PM-Kisan Yojana) launched in 2018, pension scheme for small and marginal farmers under Pradhan Mantri Kisan Maan-Dhan Yojana (PM-KMY) launched in 2019 etc.

[22]Central government spending on these programs went up in 2020-2021 due to COVID-19, and again in 2022 after the Russian invasion of Ukraine. In Figure A.1, we present these annual budget shares over time.

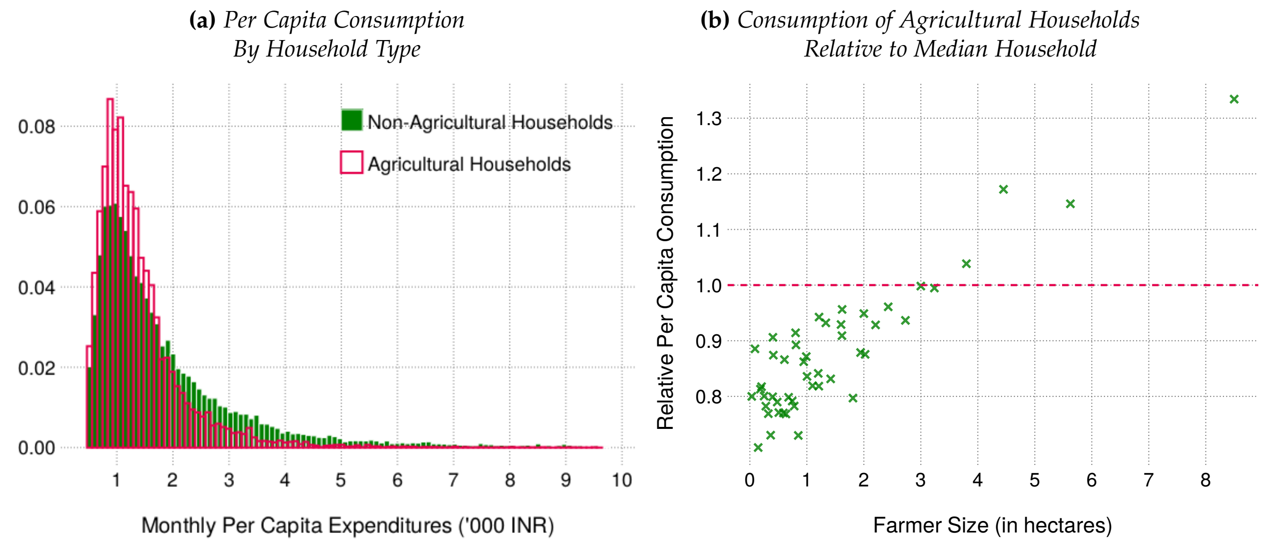

Figure 1: *Consumers and Producers in the Indian Agriculture Sector*

Notes. The left panel shows histograms of per-capita monthly expenditures of households who are identified as self-employed in agriculture against per-capita monthly expenditures of all other households in the 68th round of the NSS (2011-12). These values include home production valued at market prices. The plots only include households with reported per-capita expenditures greater than the 1st and lower than the 99th percentiles. In the right panel, we show median relative consumption of agricultural households binned by land size; relative consumption is defined as the ratio of per capita monthly expenditures and the median per capita expenditures in the data.

Before discussing these interventions in more detail, we briefly describe the datasets used in our study.

2.1 Data

We bring together several publicly-available administrative datasets for our analysis.

To estimate the supply side of our model, we require detailed farmer-level micro-data on planting decisions including crop and input choice. We obtain these data from three rounds (2008-09 to 2016-17) of Cost of Cultivation Surveys (CCS) conducted by the Department of Agriculture in India. In each round, a sample of farmers is followed for all planting seasons for three consecutive years, and plot-level data on output and input usage are recorded. In particular, input data are recorded very well. For a given farmer-year-season, we observe not only input expenditures but also physical quantities (e.g. hours of labor/machinery) of all inputs, logged separately for each crop grown in the season. We provide additional details in Appendix B.1.

While CCS surveys include information on realized output prices, the identity of the buyer is unknown. These data are critical for understanding which farmers have access to government buyers and are able to avail MSP. To get at the identity of buyers, we rely on the 77th round of the National Sample Survey (NSS) conducted in 2019 which surveyed a nationally representative sample of agricultural households in India.[23] These survey data allow us to map farmer size, region, and crop to the likelihood of making sales to a government agency.

On the household-side, we leverage a nationally-representative consumer expenditure survey conducted from July 2011 to June 2012 as part of the 68th round of the NSS. Relevant variables include household size and income, and quantities and values of rice and wheat purchased. Household purchases of these crops are broken down by source, so for each household, we observe the share of consumption that comes from PDS shops.

In addition to the above datasets, we rely on two sources of aggregate agricultural data. These include the ICRISAT District Level Database (DLD) from 1966-2016, which provide annual district-level statistics on cropping patterns, fertilizer consumption, and output prices. We use these data in our reduced-form analyses of the impact of fertilizer subsidies and minimum support prices on production decisions. Second, we use an agricultural census of all farm holdings conducted in 2016 to construct a nationally representative farm size distribution by crop.

2.2 Fertilizer Subsidies

Prior to economic liberalization in India in 1991, the government controlled the prices of all fertilizer products in India. It set the price at which it procured fertilizers from fertilizer producers and importers, and it set the price at which fertilizer products were sold to farmers; the difference between these prices was borne by the taxpayer.[24] All fertilizer products continue to be subsidized, but over the years, the government has taken steps to progressively decontrol non-urea fertilizers, in 1991 and then again in 2010.[25] In contrast, the price of urea, the most popular fertilizer product in India, continues to be tightly controlled and set directly by the government.[26]

[23] The NSS is a nationally-representative repeated cross-sectional survey.

[24] Fertilizers were procured from producers under the Retention Price Scheme; producer prices were specific to production plant and based on plant-specific costs of production.

[25] Though non-urea fertilizers are decontrolled, non-urea fertilizer producers still receive production subsidies; however, the producers now have more control over the sale price of their products.

[26] While prices paid to producers are not publicly available, we can estimate subsidy rates based on prices paid by farmers, total consumption of fertilizers in the country, and the total fiscal costs of fertilizer subsidies. For example, in 2019, the government spent USD 232 per tonne of urea, and set the controlled price at USD 76 per tonne, which amounts to a subsidy of 75% on the price of urea.

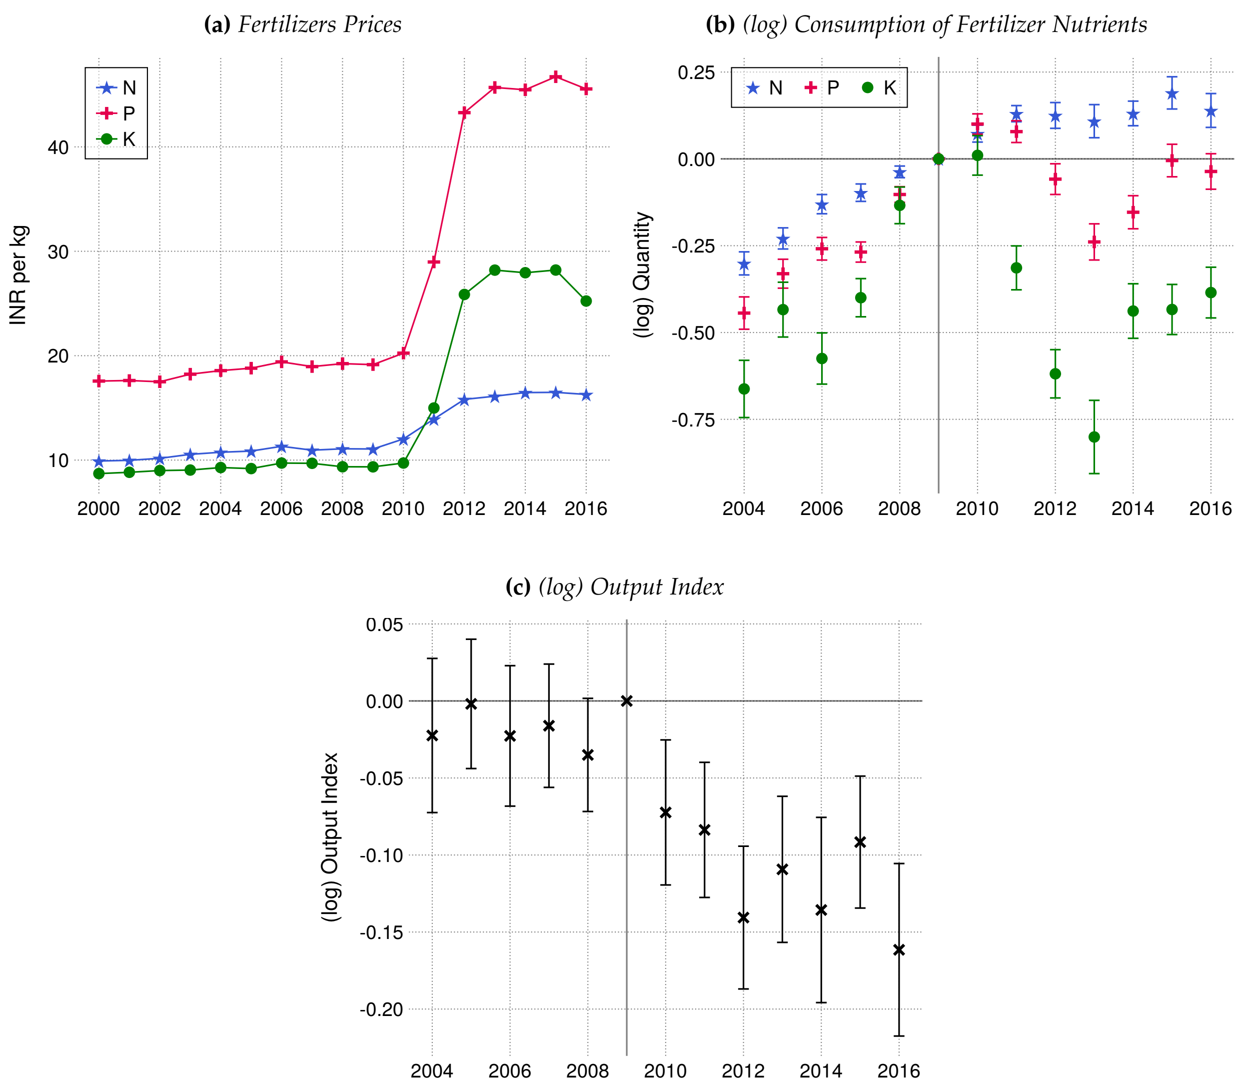

Figure 2: *Impact of Fertilizer Subsidies on Production Decisions and Output*

Notes. In the top-left panel, we plot (weighted) average reported prices of fertilizer nutrients N, P, and K in the Cost of Cultivation Surveys. In the top-right panel we show estimated coefficients from an event-study regression using district-level ICRISAT panel data. The dependent variable is (log) reported consumption of fertilizer nutrients (N, P, or K) at the district-level. The controls are year dummies (excluding 2009) and district fixed effects. In the bottom panel, we plot the estimated coefficients from a difference-in-differences specification with a continuous treatment variable using district-level ICRISAT panel data. Treatment intensity is defined as the per-unit area consumption of fertilizer nutrients P and K (aggregated using prices as weights) in the period 2004-2009, before prices of these nutrients increased sharply. The dependent variable is (log) output index at the district-level, constructed using output of all crops grown in that district aggregated using national median prices of those crops in the period 2004-2009. The controls are year and district fixed effects.

Do fertilizer subsidies affect production decisions? To study whether farmer behavior responds to these subsidies, we rely on a natural experiment. Starting 2010, subsidies for non-urea fertilizers, which are the only source of nutrients phosphorus (P) and potassium (K), were partially phased-out. In Figure 2a, we show that this resulted in a rapid increase in the price of fertilizer nutrients P and K, relative to nitrogen (N), as reported in the Cost of Cultivation Surveys.[27] Correspondingly, we find a decline in district-level consumption of nutrients P and K as shown in Figure 2b: this plot shows coefficients from a regression of (log) consumption on district fixed effects and year dummies (excluding 2009) estimated using ICRISAT District Level Database.

Next, to test how this partial phase-out of subsidies for nutrients P and K affected output, we construct a district-level measure of treatment intensity which captures the intensity with which these nutrients were used in each district prior to 2010. We use this measure of usage intensity to run the following (continuous) difference-in-differences specification

$$\log Y_{dt} = \beta_0 + \sum_{k \neq 2009} \beta_k \log \text{Avg. Usage Intensity}_d \cdot \mathbb{1}\{k = t\} + \phi_d + \gamma_t + \epsilon_{dt},$$

where ϕ_d and γ_t are district and year fixed effects.[28] Our main outcome of interest is a district-level (price-weighted) output index, which captures the value of agricultural output in each year.[29]

The estimated coefficients, shown in Figure 2c, suggest that districts where nutrients P and K were used more intensively experienced a greater decline in output post-2010 when prices of these fertilizer nutrients increased sharply. We take findings from this natural experiment as strong evidence that fertilizer subsidies not only affect farmers' fertilizer usage decisions but also have an impact on final output.

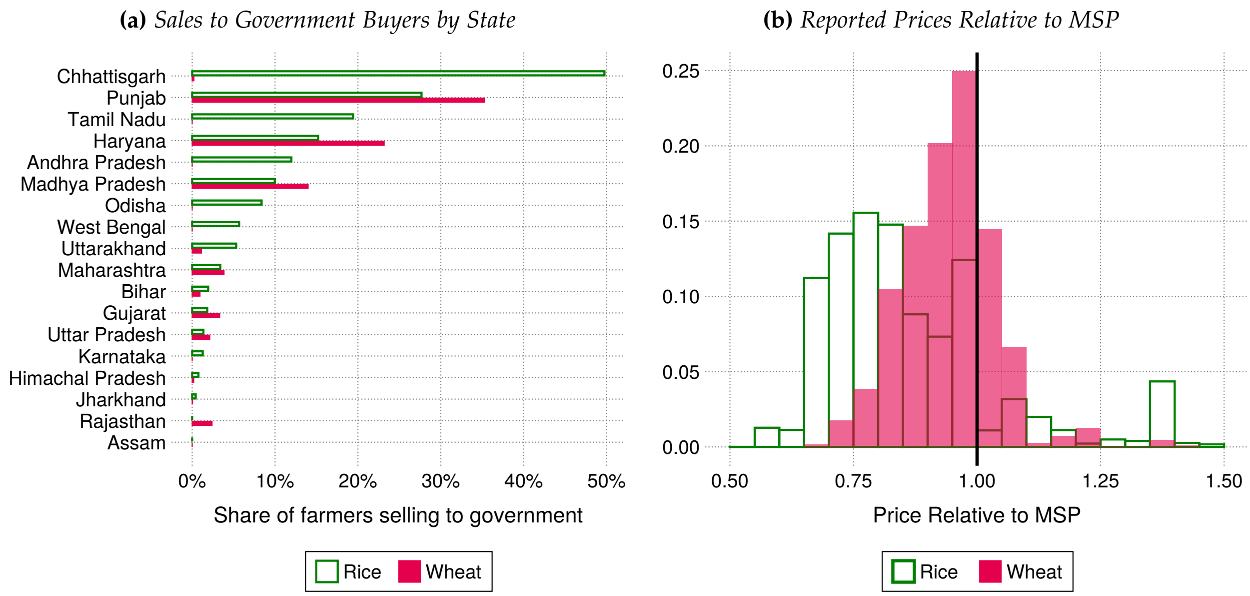

Figure 3: *Government-Procurement and Access to Minimum Support Prices (MSP)*

Notes. The left panel shows the share of rice (wheat) farmers in a state who made sales to government buyers. The right panel shows the distribution of prices received by farmers relative to the minimum support price (MSP) for that season. Source: 77th round of the National Sample Survey (NSS) conducted in 2019.

2.3 Procurement at Minimum Support Prices (MSP) & Redistribution Through the Public Distribution System (PDS)

India has two main planting seasons for crops – kharif (monsoon) and rabi (winter). At the start of each planting season, the government announces minimum support prices (MSP); these are prices farmers can expect to receive at the time of harvest if sales are made to government agencies. These prices are based on government-administered surveys known as Cost of Cultivation Surveys (CCS) designed to estimate the average costs of growing various crops in the country.[30]

[27]Urea only contains nitrogen (N), but non-urea fertilizers may also contain some nitrogen which might be why we see a small spike in the price of N as well.

[28]We construct the (price-weighted) average usage intensity of P and K in district d as

$$\text{Avg. Usage Intensity}_d = \frac{1}{6}\sum_{t=2004}^{2009} \frac{r_P^F F_{Pdt} + r_K^F F_{Kdt}}{\text{Total Area Planted}_{dt}}$$

where F_{Pdt} and F_{Kdt} are quantities consumed of nutrients P and K, respectively, while prices r_P^F and r_K^F are national median prices of the nutrients in the period 2004-2010.

[29]Note that the prices used to construct the output index are national-level median crop prices in the period 2004-2009, and only serve as weights to combine output of different crops

[30]While these surveys inform minimum support prices, the prices are also subject to political considerations.

While minimum support prices are announced for almost all major crops in India, only rice and wheat are subject to substantial procurement by the central government at minimum support prices.[31] In addition, there is substantial geographic variation in how intensively government agencies procure these crops in a region. In Figure 3a, we plot, by state, the share of farmers growing rice and wheat that report selling their output to government buyers.

Importantly, minimum support prices are not a legal price floor. Upon harvest, when a farmer brings the output to a regional market, they may only encounter private traders who are free to make price offers below the MSP. As shown in Figure 3b, a large share of farmers report receiving prices below the MSP. At the same time, the likelihood of selling to government buyers and therefore availing MSP is strongly correlated with the size of a farmer. We show this with the help of NSS data where farmers report whether sales were made to government agencies. As shown in Figure 4a, we find that larger farmers – proxied by total sales made – are more likely to sell their output to government buyers.[32] This relationship is robust to conditioning on farmer state.

The output procured by the government is fed into the public distribution system (PDS), which is a network of over half a million fair price or "ration" shops throughout the country where households can purchase staples rice and wheat at highly subsidized rates subject to income-based quotas.[33] Like fertilizer subsidies and government-procurement at MSP, the PDS has been in place since the 1960s, and is currently the largest food distribution program in the world (George and McKay 2019).[34]

The program assigns higher quotas to lower-income households. We confirm this in the data and also show, in Figure 4b, that lower-income households, proxied by total monthly expenditures, derive a larger share of their total consumption of rice and wheat from the PDS. This figure also highlights that PDS entitlements are inframarginal and that households across the income distribution rely on the private market for some share of their consumption.

[31] From 2011-2019, on average, the government procured over 30% of total annual output of rice and wheat in the country.

[32] See Footnote 5 for a discussion of why these patterns emerge.

[33] Depending on the region, these shops may sell other commodities but rice and wheat are sold almost everywhere.

[34] About 70% of Indian households interact with the PDS (Gadenne, 2020); 800 million people receive subsidized grains through the system (World Bank, 2019).

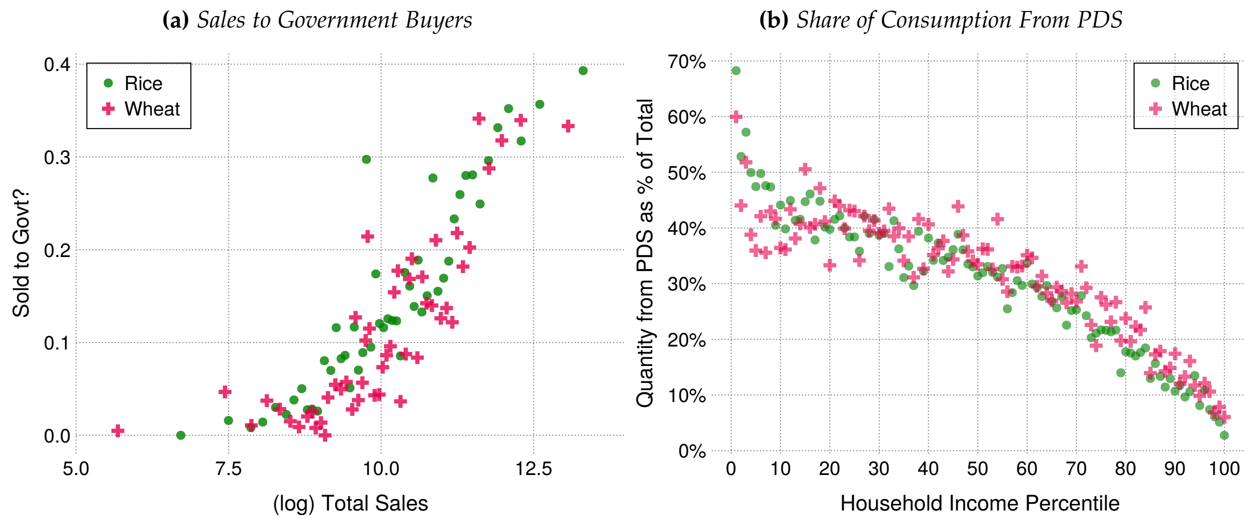

Figure 4: *Distributional Differences in the Impact of Government Interventions*

Notes. The left panel shows binned means of an indicator variable denoting whether sales were made to a government buyer against total sales made by the farmer, as reported in the 77th round of the NSS survey, conducted in 2019. The right panel shows binned means of the share of monthly consumption of rice and wheat obtained through PDS against total monthly expenses per capita as observed in the NSS Consumer Expenditure Survey, 2011.

Do minimum support prices affect production decisions? We provide evidence which suggests that farmers respond to higher MSP by increasing the share of area allocated to MSP crops, but only if the government actively procures in their state.

Specifically, let $X^1_{cs(d)t}$ be an indicator variable that equals one if in period $t-1$ the central government procured a nonzero quantity of crop c in state s of district d. Using the ICRISAT District Level Database, we estimate the following regression

$$\Delta\text{Share Area}_{cdt} = \underbrace{\{a^0_{cd} + a^0_{MSP} \cdot \Delta MSP_{ct}\}}_{\text{no procurement}} \times (1 - X^1_{cs(d)t}) + \underbrace{\{a^1_{cd} + a^1_{MSP} \cdot \Delta MSP_{ct}\}}_{\text{procurement}} \times X^1_{cs(d)t} + u_{cdt}$$

where $\Delta\text{Share Area}_{cdt}$ is the change in share of area allocated to crop c in district d relative to the previous year, while ΔMSP_{ct} is the change in minimum support price for the crop expressed in hundreds of rupees, deflated using a consumer price index. The intercepts a^0_{cd} and a^1_{cd} are crop $\times$ district fixed effects.

We plot coefficients a^0_{MSP} and a^1_{MSP} in Figure 5. Our estimates suggest that if the government procured a nonzero quantity of output of an MSP crop in a state, farmers in those states respond

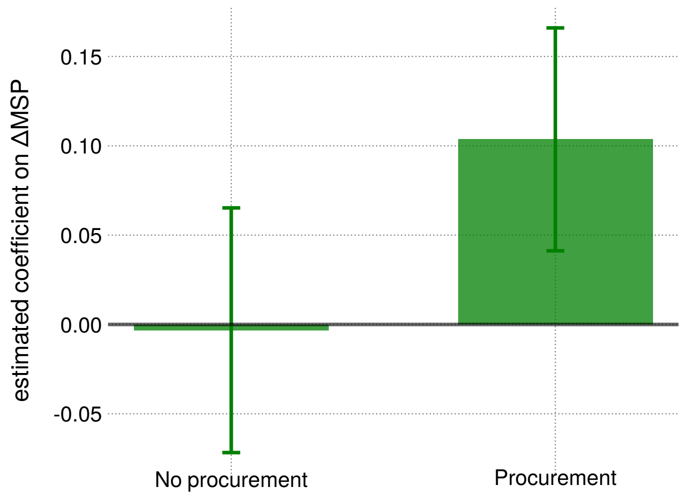

Figure 5: *Change in Share of Area Allocated to a Crop Responds to Minimum Support Prices If Government Buyers Active in the Region*

Notes. This figure plots estimated coefficients from a regression where the dependent variable is the change in share of area allocated to a given crop in a district in two consecutive years. The independent variable of interest is the change in (deflated) minimum support price for that crop expressed in hundreds of rupees. We interact this with an indicator variable for whether, in the previous year, the central government actively procured the given crop in the state in which the district is located. We also control for district $\times$ crop fixed effects. The data are from ICRISAT District Level Database (DLD) and reported at district $\times$ crop $\times$ year level. Only observations on rice and wheat are included in the regressions as these are the primary crops procured by central agencies. The plot shows the 95% confidence interval; standard errors are clustered at the district-level.

to a higher MSP for that crop by increasing the share of area allocated to it. Specifically, based on our estimate of a^1_{MSP}, increasing (deflated) MSP by 10 INR (mean = 15 INR) increases share of area allocated to that crop by 1 percentage point. In states with no government procurement in the previous year, change in MSP does not have a statistically significant effect on the change in cropping patterns.

We incorporate these findings in our model, which we present in the next section.

3 Model

The structural model consists of a supply-side with farmers who make production decisions, and a demand-side with households who consume agricultural output.

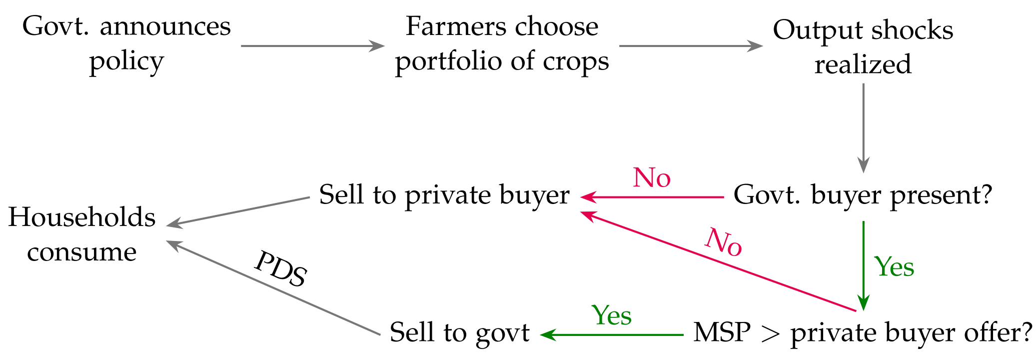

Notes. This figure provides an overview of the model. Before planting decisions are made, the government announces fertilizer subsidies and minimum support prices. Farmers take these into account and make planting decisions. Upon harvest, output shocks are realized. Farmers bring their output to the market where a government buyer may be present. If the government buyer is present, the farmer sells his crop to the government buyer if MSP is greater than the price offered by the private buyer. Otherwise, sales are made to the private buyer. Quantity procured by the government is distributed to households through the public distribution system (PDS). Household satisfy residual demand in the private market.

Figure 6: *Model Timeline & Overview*

3.1 Model Timeline

At the start of a planting season, the government announces fertilizer subsidies and crop-specific minimum support prices (MSP). Farmers observe these policy announcements and make planting decisions. After production decisions are made, idiosyncratic shocks are realized which affect output quantity as well as the price offer made by a private buyer. Farmers sell their output either to the government buyer or to the private buyer. Finally, households receive their PDS entitlements and make purchases from the private market. We summarize this in Figure 6.

3.2 Supply: Farmer's Problem

Planting decisions involve choices on both the extensive and intensive margins. Farmers choose which set of crops to plant and, for each crop in this set, they make crop-specific input allocations.

Farmer j in region r is endowed with a farm of total size A_j. In season t, he chooses a set s of crops to plant which maximize utility V_{jst}. This utility consists of two components: a mean-variance payoff U_{jst} and a fixed cost of planting κ_{jst}, and is expressed as

$$V_{jst} = U_{jst} - \kappa_{jst}, \quad \forall s \in \mathcal{S}_j$$

where $\mathcal{S}_j$ are all possible sets of crops farmer j can grow.

Farmers are risk averse with risk-aversion γ_j. Given a set of crops s, farmers choose how to allocate plots A_{jct} and inputs $\boldsymbol{X}_{jct}$ to each crop $c \in s$ in order to maximize the difference between expected total profits and risk-aversion weighted variance of total profits. The optimal allocation gives rise to the mean variance payoff U_{jst} defined as

$$U_{jst} = \max_{\{A_{jct}, \boldsymbol{X}_{jct}\}_{c \in s}} \mathbb{E}\left(\Pi_{jst}\right) - \gamma_j \operatorname{Var}\left(\Pi_{jst}\right)$$

$$\text{where } \Pi_{jst} = \sum_{c \in s} \pi_{jct}(A_{jct}, \boldsymbol{X}_{jct}) \text{ s.t. } \sum_{c \in s} A_{jct} \leq A_j \tag{1}$$

This optimization problem is only subject to an area constraint which requires that the sum of crop-specific area allocations is (weakly) less than the total land endowment A_j.[35]

Total profits Π_{jst} are the sum of crop-specific profits π_{jct} given by

$$\pi_{jct}(A_{jct}, \boldsymbol{X}_{jct}) = P_{jct} q_{jct}(A_{jct}, \boldsymbol{X}_{jct}) - \sum_{x \in \boldsymbol{X}_{jct}} w_t^x x$$

where P_{jct} is risky output price, q_{jct} is risky output, and w_t^x is the price of input x. One of the inputs is fertilizer F_{jct} and fertilizer subsidies directly affect the post-subsidy price of fertilizer w_t^F.

Next, we describe in detail (1) output, (2) output price, (3) risk aversion, and (4) fixed costs.

Risky Output

Output of crop c is given by $q_{jct}(A_{jct}, \boldsymbol{X}_{jct})$ which depends on plot area A_{jct} and inputs $\boldsymbol{X}_{jct}$. These inputs include labor L_{jct}, capital K_{jct}, and fertilizers F_{jct}. Further, output depends on a farmer-specific unobserved productivity term ω_j which is known at the time of planting. Output risk arises from an idiosyncratic output shock ε_{jct} that is realized at the time of harvest; only its

[35]We do not model credit constraints, which may also determine input choices, and assume that farmers can purchase any level of non-area inputs. In India, institutional credit for input purchases, through banks and cooperatives, might be readily available due to government support as agriculture is classified as a priority sector by the central bank, and banks and financial institutions are required to lend at least 18% of their credit to the agriculture sector. Further. all farmers are eligible for Kisan ("Farmer") Credit Cards that can be used to purchase inputs at low interest rates. In the 77th round of the NSS, we see that institutional lenders are responsible for 82% of agriculture loans, both by amount lent and number of loans extended. In Figure A.3 we show that the reported interest rates for farm loans are much lower than for consumption loans and these rates are similar across the farmer size distribution. Also see 2014 which finds agricultural risk to be a more important determinant of production decisions than input credit constraints in northern Ghana.

mean and variance are known ex ante.[36]

We express this risky output as follows

$$
\begin{aligned}
q_{jct}(A_{jct}, \boldsymbol{X}_{jct}) &= q_c(A_{jct}, \boldsymbol{X}_{jct}) \exp\{\omega_j + \varepsilon_{jct}\} \\
&= A_{jct}^{\beta_{ac}} L_{jct}^{\beta_{lc}} (1 + K_{jct})^{\beta_{kc}} (1 + F_{jct})^{\beta_{fc}} \exp\{\omega_j + \varepsilon_{jct}\}
\end{aligned}
\tag{2}
$$

where $\boldsymbol{X}_{jct} = \{L_{jct}, K_{jct}, F_{jct}\}$. In estimation of crop-specific production functions $q_c(A_{jct}, \boldsymbol{X}_{jct})$, we also account for the impact of location and season using relevant fixed effects.

The unobserved productivity term, ω_j, captures the average fertility of farmer j's land as well as any technological know-how and ability. Importantly, it does not differ by crop. It does, however, affect input choices and therefore gives rise to the standard input endogeneity concern (Marschak and Andrews, 1944; Hoch, 1962; Griliches and Mairesse, 1995). We explain how we deal with it when we discuss estimation in the following section.

Risky Prices

Upon harvest, farmers bring their output to the market where they may or may not encounter a government buyer. A private buyer is always present. Government buyer offers to buy PDS crops, rice and wheat, at the pre-announced minimum support price (MSP). If a government buyer is found and MSP for a crop is greater than the price offered by the private buyer, farmer sells all output of that crop to the government; otherwise, the farmer accepts the private buyer's offer.[37]

What does the private buyer offer? The price offered by the private buyer depends on an idiosyncratic output *quality* shock that is realized post-harvest.[38] This output quality shock, η_{jct}, captures factors such as dust and moisture content and only affects processing costs of the crop in the private market.[39] High quality crops have low processing costs and therefore receive higher private buyer offers. The expected price offer equals the equilibrium price in the private market,

[36] Output quantity shocks are uncorrelated across farmers. Therefore, the model features no aggregate shocks. Aggregate shocks can be added to the model at the expense of significantly larger computation requirements. An important consequence of aggregate shocks would be the negative covariance between output and private market prices which might lower the total revenue risk faced by farmers (Allen and Atkin, 2022). In the absence of aggregate shocks, we may overestimate risk and therefore, underestimate risk aversion.

[37] Non-PDS crops are always sold to private buyers at the offered price.

[38] There are several alternative justifications for cross-sectional variance in private buyer offers. We discuss some of these in Footnote 8.

[39] Processed crops that are purchased by households are homogeneous; there are no quality differences.

given by P_{ct}.[40] Now, we can express private buyer offers as

$$\widetilde{P}_{jct} = P_{ct} \cdot \exp\{\eta_{jct}\}$$

where $\eta_{jct} \sim N\left(-\frac{\sigma^2_{\eta c}}{2}, \sigma^2_{\eta c}\right)$. We assume that farmer j knows the distribution of private buyer offers. That is, he knows the distribution of quality shocks and the equilibrium private market prices for all crops at the time of planting.[41]

If a government buyer is present, the farmer only accepts private buyer offers if they are greater than MSP. Let $Z_{jct} = 1$ if farmer j encounters a govt buyer for crop c; $Z_{jct} = 0$ otherwise. Price received for crop c by farmer j is

$$P_{jct} = \mathbb{1}\{Z_{jct} = 1\} \max\left\{MSP_{ct}, \widetilde{P}_{jct}\right\} + \mathbb{1}\{Z_{jct} = 0\}\, \widetilde{P}_{jct} \tag{3}$$

Importantly, we assume that the farmer is uncertain about meeting a government buyer at the time of planting.[42] Motivated by data, the probability of finding a government buyer is a function of farmer size, crop, and location. Specifically,

$$\rho_{jct} = \Pr(Z_{jct} = 1) = \Phi(\alpha_{0rc} + \alpha_{1rc} \cdot \log A_j)$$

where $(\alpha_{0rc}, \alpha_{1rc})$ are crop- and region-specific coefficients, and A_j is the total area of farmer j. This probability, along with (3) and the distribution of private buyer offers, gives rise to a farmer size-, location-, and crop-specific mixture distribution which the farmer uses to compute mean-variance payoff given in (1). For non-zero ρ_{jct}, the mean of this distribution of prices is increasing in MSP.

[40] Note that we do not model intermediary market power. The average price paid by private buyers to farmers equals the price households pay for private market purchases. Several studies document and analyze trader market power in agriculture in India (Meenakshi and Banerji, 2005; Mitra, Mookherjee, Torero, and Visaria, 2018; Chatterjee, 2022) and elsewhere (Bergquist and Dinerstein, 2020). To check the robustness of our results, we plan to re-do our counterfactual analyses at different levels of calibrated passthrough.

[41] Knowing the equilibrium private market price requires solving a very complex problem. Alternatively, we can assume that farmers extrapolate equilibrium private market prices from the average prices in the previous year. This specification is easy to incorporate and we plan to add it as a robustness check.

[42] This can be relaxed and we can check the robustness of our results to this assumption; we plan to do this in the next iteration of estimation.

Risk Aversion

Risk aversion is parameterized to be a function of total land holdings A_j (a proxy for wealth).

$$\ln \gamma_j = \gamma_0 + \gamma_A \ln A_j + \psi_j, \quad \psi_j \sim N(0, \sigma_\gamma^2).$$

where ψ_j is an idiosyncratic component of risk aversion.

Fixed Costs

Fixed costs depend on the set of crops planted and do not scale by area. These help us rationalize low crop diversification observed in the data: most farmers grow at most 3 crops in a season.

Let $s_{j,t-1}$ be the set of crops planted by farmer j in the same season but in the previous year. Fixed cost κ_{jst} of planting a set s of crops jointly is given by

$$\kappa_{jst} = \sum_{c \in s} \kappa_{jct} \text{ where } \kappa_{jct} = \begin{cases} \kappa_c & c \notin s_{j,t-1} \\ \lambda \cdot \kappa_c & c \in s_{j,t-1} \end{cases}$$

where κ_c is a constant crop-specific parameter, and λ is a discount on fixed costs for repeating crops. In our estimation, we allow λ to differ by staple crops (rice and wheat), and all other crops.

3.3 Demand

PDS crops: rice and wheat

Households differ by income and PDS entitlements.[43]

Let q_{ch}^{PDS} denote the per-capita quantity of crop c received by household h through the PDS system.[44] Motivated by data, we assume this quantity is inframarginal to the total per-capita

[43] In addition to household income, these entitlements may also depend on where the household is located and how easily it can access PDS shops. To hold targeting fixed, we calibrate the total share of government procurement that a household receives using the 68th round of the NSS which is a nationally-representative survey of households. We hold these calibrated PDS shares as fixed in our counterfactuals.

[44] We assume that the households do not pay for these entitlements. In reality, households may pay a small amount depending on their income level.

quantity of crop c consumed by the household, which is given by

$$q_{ch} = q_{ch}^{PVT} + q_{ch}^{PDS}$$

where q_{ch}^{PVT} is the per-capita quantity of crop c purchased in the private market. The total per-capita demand depends on the equilibrium price in the private market, P_c (time subscript t is suppressed). In addition, it also depends on per-capita income y_h. For tractability, we assume a log demand function given by

$$\log(1 + q_{ch}) = a_{cp} \log P_c + a_{cy} \log y_h + a_{cpy} \log P_c \cdot \log y_h + u_{ch} \tag{4}$$

which can be approximated using a utility function discussed in Appendix C.1. Importantly, this function is compatible with Engel's law (and our data) that higher income households spend a lower share of their income on food.

Non-PDS crops

For non-PDS crops, we consider an aggregate demand function given by

$$q_{ct} = \mu P_{ct}^{e_c} \quad \forall c \notin \{\text{rice, wheat}\} \tag{5}$$

where e_c is the price elasticity of demand for crop c.

3.4 Equilibrium

The total quantity procured by the government can be expressed as

$$Q_{ct}^{govt}(P_{ct}) = \sum_j \mathbb{E}\left[q_{jct}(P_{ct}) \cdot \mathbb{1}\left\{Z_{jct} = 1\right\} \cdot \mathbb{1}\left\{MSP_{ct} \geq P_{ct} \cdot \exp\{\eta_{jct}\}\right\}\right]$$

We assume that MSP is set to track the equilibrium price P_{ct} in the private market. In particular, motivated by our estimates, MSP is set at the 65th percentile of the private buyer offer distribution. This implies that conditional on finding a government buyer, 65% of farmers sell to the government. All other sales are made to private buyers. Total equilibrium quantity in the private market, therefore, is

$$Q_{ct}^{pvt}(P_{ct}) = \sum_j q_{jct}(P_{ct}) - Q_{ct}^{govt}(P_{ct})$$

Our notion of equilibrium is a vector of average private market prices which farmers and consumers take as given, and which clears all markets. More precisely, a static competitive equilibrium is a vector of private market prices, $\{P_{ct}\}_c$ such that

1. Government procurement equals sum of PDS entitlements received by households.

$$Q_{ct}^{govt}(P_{ct}) = \sum_h q_{cht}^{PDS} \quad \forall c$$

2. Total purchases by private buyers equals total private market demand for all crops.

$$Q_{ct}^{pvt}(P_{ct}) = \sum_h q_{cht}^{PVT}(y_h, P_{ct}, q_{cht}^{PDS}) \quad \forall c$$

3. Sum of government procurement and private buyer purchases equals total output.

$$Q_{ct}^{govt}(P_{ct}) + Q_{ct}^{pvt}(P_{ct}) = \sum_j q_{jct}(P_{ct})$$

4 Estimation

4.1 Supply

We estimate the supply-side of the model in three stages. First, we estimate the parameters governing the distribution of prices at the time of planting. Next, we estimate the production function and risk aversion parameters. Finally, we estimate the fixed costs. All stages rely on simulation-based estimators (Pakes, 1986; McFadden, 1989; Pakes and Pollard, 1989).

The Distribution of Output Prices

Three sets of parameters determine the distribution of prices at the time of planting. These are

1. the equilibrium (average) private market price, P_{ct} for all crops and years,

2. the variance of output quality shocks, $\sigma^2_{\eta c}$ for all crops, and
3. crop- and region-specific parameters governing the likelihood of finding a government buyer, $\{\alpha_{0rc}, \alpha_{1rc}\}$ the latter of which is the coefficient on farmer size.

Our analysis is complicated by the fact that private buyer offers are observed only if a government buyer is absent or if the offers are higher than the minimum support price (MSP).[45] To get an unbiased estimate of the mean and variance of the private buyer distribution, we need to condition on the presence of a government buyer. However, whether a government buyer is present is not known. Our data only includes information on realized sales.[46] Since farmers may choose to sell to private buyers even when government buyers are present, this measure is an imperfect proxy for the likelihood of finding government buyers.

We estimate these parameters as follows. Let $\theta_c = \{\{P_{ct}\}_t, \sigma^2_{\eta c}, \{\alpha_{0rc}, \alpha_{1rc}\}_r\}$ and θ^g_c be a guess of these parameters. For each θ^g_c, we can simulate whether a farmer found a government buyer given his location, size, and crop, for all farmers in the CCS data. We can also draw a private buyer offer given a guess of average private market price and the variance of output quality shocks. Simulated *realized* price is the private buyer offer if a government buyer is not found or if the private buyer offer is greater than the MSP. Otherwise, the simulated price equals the MSP for that crop. This generates a distribution of simulated prices that farmers receive.

For each simulated distribution of prices, we compute the mean price by year, $\mathbb{E}[P_{jct}|\theta^g_c]$; recall that this does not necessarily equal the average private market price. We also compute the variance of this distribution. Finally, using the simulated data, we estimate the following probit model:

$$\Pr\Big(\text{Sold to government}_{jct} = 1|\theta^g_c\Big) = \Phi(\delta^g_{0rc} + \delta^g_{1rc} \cdot \log A_j) \tag{6}$$

which gives us region- and crop-specific coefficients δ^g_{0rc} and δ^g_{1rc}.

We construct empirical counterparts of these three sets of moments (mean, variance, and coefficients from probit model) using the Cost of Cultivation Surveys and the 77th round of the NSS. The former reports, in addition to farmer size and location, the realized output price for each crop planted. The latter (NSS) includes information on farmer size, location, and whether sales

[45] This is only an issue for PDS crops. For non-PDS crops, estimation is straightforward.

[46] In the 77th round of the NSS, agricultural households report whether sales were made to government agencies or to private buyers.

Table 1: *Standard Deviation of Output Quality Shocks Which Determine Private Buyer Offers*

	(1)	(2)
	σ_{η_c}	95% conf. interval
chickpea	0.124	[0.121, 0.126]
cotton	0.078	[0.076, 0.080]
finger millet	0.182	[0.177, 0.187]
groundnut	0.163	[0.159, 0.170]
maize	0.103	[0.101, 0.104]
mustard and rapeseed	0.082	[0.079, 0.084]
pearl millet	0.111	[0.109, 0.113]
pigeonpea	0.148	[0.143, 0.151]
rice	0.227	[0.225, 0.229]
sesamum	0.274	[0.263, 0.281]
sorghum	0.285	[0.282, 0.289]
sugarcane	0.133	[0.130, 0.136]
wheat	0.089	[0.088, 0.091]

Notes. This table shows the estimated standard deviation of output quality shocks, by crop, that determine private buyer offers. Column (2) is the 95% confidence interval estimated using bootstrap.

were made to government agencies or private buyers; this allows us to construct the auxiliary parameters δ_{0rc} and δ_{1rc} of (6) in the data.

For each crop, we estimate parameters θ_c by matching these empirical moments with the simulated moments.[47] We weigh the difference between empirical and simulated moments using inverse of the variance of empirical moments, which we estimate using bootstrap.

For non-PDS crops, identification is straightforward since the observed prices all come from the distribution of private buyer offers. For PDS crops, rice and wheat, identification is guaranteed if a positive mass of private buyer offers lies on both sides of MSP. Offers on the right of MSP guarantee that a non-trivial share of farmers sell to private buyers; thus, movements in the mean and variance of private buyer offers would shift the distribution of realized prices. Offers on the left ensure that if a government buyer is present, some sales would be made to government buyers. All else equal, a higher guess of the region-specific intercept α^g_{0rc} would uniformly

[47]Our approach is similar to an expectation–maximization (EM) algorithm (Dempster, Laird, and Rubin, 1977) commonly employed to estimate parameters of mixture distributions. However, instead of maximizing a likelihood, we match moments. Our approach also relies on the literature on indirect inference (see Gourieroux et al., 1993)

(across farmer size) increase the share of farmers selling to government buyers in that region, and therefore yield a higher δ^{g}_{0rc}. Similarly, the slope coefficient α^{g}_{0rc} would directly affect the auxiliary coefficient δ^{g}_{1rc}. We present estimated parameters in Tables 1 and 2.[48]

Production Function and Risk Aversion

Production function parameters include crop-specific input elasticities for area (β_{ac}), labor (β_{lc}), capital (β_{kc}), and fertilizers (β_{fc}).[49] Additionally, we need to recover unobserved farmer productivities ω_j for all j. Finally, we also require the mean and variance of output quantity shocks ε_{jct} which enter the production function in (2). In logs, output of crop c can be written as

$$\begin{aligned}\log q_{jct} &= \log q_c(A_{jct}, \boldsymbol{X}_{jct}) + \omega_j + \varepsilon_{jct} \\ &= \beta_{ac} \log A_{jct} + \beta_{lc} \log L_{jct} + \beta_{kc} \log\left(1 + K_{jct}\right) + \beta_{fc} \log\left(1 + F_{jct}\right) + \omega_j + \varepsilon_{jct}\end{aligned} \tag{7}$$

The unobserved productivity term, ω_j, is constant across time and across crops.[50] This is in contrast to standard production function specifications which usually allow productivity to vary over time.[51] However, their estimation requires a monotonicity assumption between productivity and input demand which fails in our setting with risk-averse farmers as positive productivity draws increase the variance of output which may lead farmers to reduce input demand.[52] We cannot calibrate elasticities using cost shares either since that too relies on profit-maximizing choices.

Our approach involves jointly estimating production function and risk-aversion parameters using farmer's optimization problem in (1); risk aversion parameters include intercept γ_0, the coefficient on farmer size γ_A, and the variance of the mean-zero risk aversion draw ψ_j, denoted by σ^2_γ. We proceed as follows. Let $\theta_\beta = \{\beta_{ac}, \beta_{lc}, \beta_{kc}, \beta_{fc}\}_c$ and $\theta_\gamma = \{\gamma_0, \gamma_A, \sigma^2_\gamma\}$. For each guess of parameters $(\theta^g_\beta, \theta^g_\gamma)$, where g denotes a candidate vector, we take the following sequence of steps.

1. Get $\xi^g_{jct} \equiv \omega^g_j + \varepsilon^g_{jct}$ by differencing out observed inputs from observed output using θ^g_β in

[48] Not included in the interest of space: crop $\times$ year mean private market prices.

[49] Note that we treat fertilizers as a single composite input. As a robustness check, we plan to also estimate this production function using data on farmer-level consumption of fertilizer nutrients N, P, and K.

[50] Since productivity does not differ by crop, selection into crops is not a concern for us.

[51] See Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020

[52] The monotonicity assumption allows researchers to construct control functions, using observed levels of intermediate inputs, which may account for unobserved productivity.

Table 2: *Parameters Governing the Likelihood of Finding a Government Buyer by State*

	(1) α_0 (Rice)	(2) α_1 (Rice)	(3) α_0 (Wheat)	(4) α_1 (Wheat)
Andhra Pradesh	-0.002	-0.019		
	[-0.004, -0.001]	[-0.036, -0.011]		
Bihar	-8.673	0.924		
	[-11.897, -8.105]	[0.893, 0.999]		
Chhattisgarh	-4.116	0.526		
	[-4.370, -3.351]	[0.480, 0.725]		
Gujarat	-6.323	0.580	-4.421	0.304
	[-7.812, -5.844]	[0.516, 0.738]	[-4.643, -4.113]	[0.268, 0.368]
Haryana	-3.921	0.373	-5.399	0.500
	[-5.469, -3.594]	[0.351, 0.468]	[-5.664, -5.311]	[0.482, 0.540]
Karnataka	-7.005	0.554		
	[-8.387, -6.578]	[0.496, 0.752]		
Madhya Pradesh	-7.311	0.665	-7.410	0.704
	[-8.731, -6.997]	[0.626, 0.786]	[-7.518, -7.259]	[0.676, 0.721]
Maharashtra	-5.719	0.462	-4.303	0.330
	[-5.946, -5.576]	[0.420, 0.665]	[-4.437, -4.116]	[0.263, 0.371]
Odisha	-8.063	0.745		
	[-11.297, -7.493]	[0.719, 0.832]		
Punjab	-3.110	0.307	-1.297	0.157
	[-4.298, -2.873]	[0.292, 0.344]	[-1.527, -1.181]	[0.148, 0.182]
Rajasthan			-5.223	0.373
			[-5.308, -5.121]	[0.357, 0.381]
Tamil Nadu	-4.038	0.417		
	[-4.571, -3.818]	[0.339, 0.441]		
Uttar Pradesh	-7.805	0.610	-7.586	0.649
	[-10.745, -7.321]	[0.596, 0.650]	[-7.676, -7.523]	[0.639, 0.659]
Uttarakhand	-12.485	1.171	-5.216	0.498
	[-15.250, -11.706]	[1.103, 1.376]	[-5.845, -4.411]	[0.439, 0.801]
West Bengal	-6.553	0.615		
	[-6.608, -6.500]	[0.595, 0.680]		

Notes. This table shows the estimated parameters governing the likelihood of finding a government buyer by state. Column (1) shows the intercept for rice. Column (2) shows the coefficient on (log) total farmer area for rice. Column (3) shows the intercept for wheat. Column (4) shows the coefficient on (log) total farmer area for wheat. Blank cells correspond to crop-states for which a negligible share ($<$ 1%) of farmers reported selling to government buyers. Confidence intervals are in square brackets below each point estimate and are estimated using bootstrap.

(7); then regress ξ^g_{jct} on farmer fixed effects to get ω^g_j and ε^g_{jct}.

2. Compute mean and variance of output shocks ε^g_{jct} by crop.

Table 3: *Risk Aversion Parameters*

	(1) estimate	(2) 95% conf. interval
Intercept, γ_0	-9.911	[-9.939, -9.895]
Coefficient on farmer size, γ_A	-0.118	[-0.125, -0.116]
Std. dev. of distribution, σ_γ	0.946	[0.929, 0.954]

Notes. This table shows the estimated parameters governing farmer risk aversion. Column (1) is the estimated parameters. Column (2) shows 95% confidence intervals estimated using bootstrap.

3. Draw risk aversion γ_j for each farmer using θ^g_γ.

4. For the observed set of crops for each farmer-season, solve the portfolio choice problem in (1) using θ^g_β, ω^g_j, γ_j, mean and variance of output shocks for each crop, and the previously estimated parameters which govern the distribution of output prices.

The last step is computationally intensive; it gives us crop-specific input allocations of area, labor, capital, and fertilizers for the observed set of crops planted by each farmer in each season in the data. Since these input choices maximize the mean-variance utility for a given set of crops, they do not depend on fixed costs.

Using these simulated choices, we construct, by crop, first and second moments of simulated output, area, share of area conditional on planting, labor, capital, and fertilizer. These moments are sensitive to the guess of input elasticities and risk aversion parameters as both govern how farmers allocate inputs to different crops in a given set of crops.[53] We jointly identify these parameters by matching these simulated moments with their empirical analogs in the CCS data.[54] Estimated parameters are presented in Tables 3 and 4.

Fixed costs

The fixed costs parameters to be estimated are crop-specific constants κ_c, and the discount parameter on repeated crops λ.[55] We denote these by $\theta_\kappa = \{\{\kappa_c\}_c, \lambda\}$.

[53] For example, in Figures A.6 and A.7 we show how changing risk-aversion changes fertilizer usage and land share allocations.

[54] The differences are weighted by the inverse-variance weighted before summing.

[55] We allow crop-specific constants to differ by season (monsoon or winter). The discount parameter is also allowed to be different for staples (rice and wheat) and all other crops.

Table 4: *Production Function Parameters*

	(1) land	(2) labor	(3) capital	(4) fertilizer
chickpea	0.455	0.437	0.259	0.059
	[0.444, 0.487]	[0.421, 0.471]	[0.240, 0.329]	[0.056, 0.069]
cotton	0.321	0.899	0.099	0.203
	[0.314, 0.334]	[0.898, 0.899]	[0.090, 0.136]	[0.192, 0.209]
finger millet	0.754	0.628	0.261	0.111
	[0.699, 0.857]	[0.610, 0.667]	[0.243, 0.333]	[0.100, 0.172]
groundnut	0.517	0.506	0.258	0.135
	[0.505, 0.549]	[0.492, 0.517]	[0.229, 0.302]	[0.132, 0.139]
maize	0.586	0.435	0.194	0.102
	[0.581, 0.592]	[0.428, 0.447]	[0.189, 0.208]	[0.098, 0.106]
mustard and rapeseed	0.581	0.290	0.183	0.078
	[0.571, 0.590]	[0.282, 0.309]	[0.178, 0.186]	[0.076, 0.080]
pearl millet	0.386	0.368	0.337	0.075
	[0.378, 0.396]	[0.359, 0.384]	[0.332, 0.358]	[0.071, 0.077]
pigeonpea	0.595	0.486	0.234	0.109
	[0.584, 0.611]	[0.478, 0.503]	[0.215, 0.250]	[0.104, 0.112]
rice	0.708	0.386	0.082	0.079
	[0.701, 0.741]	[0.380, 0.392]	[0.076, 0.086]	[0.074, 0.082]
sesamum	0.243	0.299	0.135	0.054
	[0.237, 0.265]	[0.289, 0.317]	[0.126, 0.140]	[0.046, 0.056]
sorghum	0.133	0.729	0.400	0.030
	[0.123, 0.146]	[0.722, 0.739]	[0.400, 0.400]	[0.023, 0.048]
sugarcane	0.551	0.665	0.126	0.091
	[0.527, 0.601]	[0.654, 0.675]	[0.118, 0.143]	[0.084, 0.104]
wheat	0.712	0.212	0.188	0.092
	[0.704, 0.747]	[0.207, 0.215]	[0.183, 0.192]	[0.088, 0.095]

Notes. This table shows the estimated production function parameters. Column (1) is the output elasticity of land. Column (2) is the output elasticity of labor. Column (3) is the output elasticity of capital. Column (4) is the output elasticity of fertilizer. Confidence intervals are in square brackets below each point estimate and are estimated using bootstrap.

Given a set $s \in \mathcal{S}_j$ of crops, the estimated parameters so far allow us to compute the mean-variance payoff U_{jst} for farmer j by solving the optimal portfolio choice problem. Farmer j computes this U_{jst} for all possible sets of crops he can plant, and then uses the fixed costs κ_{jst} to

Table 5: *Crop-Specific Fixed Costs by Season*

	(1) season	(2) $\log(\kappa_c)$	(3) 95% conf. interval
chickpea	kharif (monsoon)	11.622	[10.235, 16.646]
chickpea	rabi (winter)	10.501	[9.985, 10.948]
cotton	kharif (monsoon)	11.356	[11.036, 11.641]
cotton	rabi (winter)	15.425	[12.664, 18.997]
finger millet	kharif (monsoon)	12.062	[10.670, 14.092]
finger millet	rabi (winter)	12.238	[10.075, 14.160]
groundnut	kharif (monsoon)	10.232	[9.906, 10.460]
groundnut	rabi (winter)	15.624	[13.281, 18.272]
maize	kharif (monsoon)	10.565	[10.353, 10.970]
maize	rabi (winter)	11.456	[10.455, 14.836]
mustard and rapeseed	rabi (winter)	9.843	[9.561, 11.401]
pearl millet	kharif (monsoon)	12.090	[8.683, 13.659]
pearl millet	rabi (winter)	10.373	[8.507, 12.854]
pigeonpea	kharif (monsoon)	10.315	[10.019, 10.543]
pigeonpea	rabi (winter)	15.869	[12.432, 20.763]
rice	kharif (monsoon)	10.605	[10.473, 10.981]
rice	rabi (winter)	11.328	[10.664, 12.296]
sesamum	kharif (monsoon)	8.808	[8.664, 8.987]
sesamum	rabi (winter)	10.660	[9.493, 11.542]
sorghum	kharif (monsoon)	10.196	[9.608, 12.015]
sorghum	rabi (winter)	11.353	[10.181, 12.932]
sugarcane	kharif (monsoon)	13.268	[12.064, 15.825]
sugarcane	rabi (winter)	13.029	[12.247, 13.968]
wheat	rabi (winter)	10.172	[9.700, 10.696]

Notes. This table shows the estimated crop-specific fixed cost constants κ_c by season. Farmers grow some crops in both seasons, while others are only grown in one season in our data. The last column reports the 95% confidence interval estimated via bootstrap.

determine the set of crops which yield the highest utility $V_{jst} = U_{jst} - \kappa_{jst}$. This exercise is similar to a discrete choice problem where the choice set is a set of sets $\mathcal{S}_j$.

To estimate θ_κ, we proceed as follows. First, we build the set $\mathcal{S}_j$ for each farmer j. To do so, for each farmer, we take the set of crops planted by farmers in his state, and then take all combinations of up to length 3.[56] Then, we compute U_{jst} for all $s \in \mathcal{S}_j$ for each farmer in our data. This is a computationally intensive exercise but only needs to be done once.

[56] Almost all farmers in our data grow 3 or fewer crops in a season.

Table 6: *Discount Parameters for Repeating Crops*

	(1) estimated λ	(2) 95% conf. interval
rice and wheat	-0.143	[-0.183, -0.091]
all other crops	-0.006	[-0.024, -0.001]

Notes. This table shows the estimated discount parameter on repeated crops. We estimate these separately for staple crops (rice and wheat), and all other crops. The last column reports the 95% confidence interval estimated via bootstrap.

Next, for each guess of parameters θ_κ^g, we find the set $s_j^*(\theta_\kappa^g)$ that maximizes utility V_{jst} for farmer j. Note that farmers are not forward-looking but their fixed costs depend on which crops they planted in the previous period; in our simulations, we take the set of crops in the previous period from the data, and then predict choices for this period given θ_κ^g. We take these simulated choices and compute simulated "market" shares for all sets $s \in \bigcup_j \mathcal{S}_j$. We also calculate, for each crop, the unconditional probability of being added to and dropped from a set between two consecutive periods.

We estimate θ_κ by matching these simulated market shares and switching probabilities with their empirical counterparts.[57] Moment conditions used in estimation are weighted by the inverse of the variance of empirical probabilities, estimated using bootstrap. We present the estimated parameters in Tables 5 and 6.

Our functional form assumption helps in the identification of the level of fixed costs. For example, consider the possible sets of crops with two crops c_1 and c_2: $\{c_1\}, \{c_2\}, \text{and} \{c_1, c_2\}$. If we increase the fixed cost associated with c_1 and c_2 by Δ, the relative attractiveness of $\{c_1\} \text{and} \{c_2\}$ will remain the same. However, $\{c_1, c_2\}$ will become relatively less attractive as costs go up by 2Δ and farmers will switch out of it. Thus, the share of farmers growing a given set of crops is informative about crop-specific constants κ_c. The discount on repeated crops, λ, is informed by the observed persistence of crop choice. We capture this persistence by its counterpart, the probability of dropping a crop grown in the last period. Lower λ lowers the fixed cost of repeated crops, and therefore lowers the probability of dropping a crop.

[57] In practice, we only match sets of crops with greater than 1% share in our data.

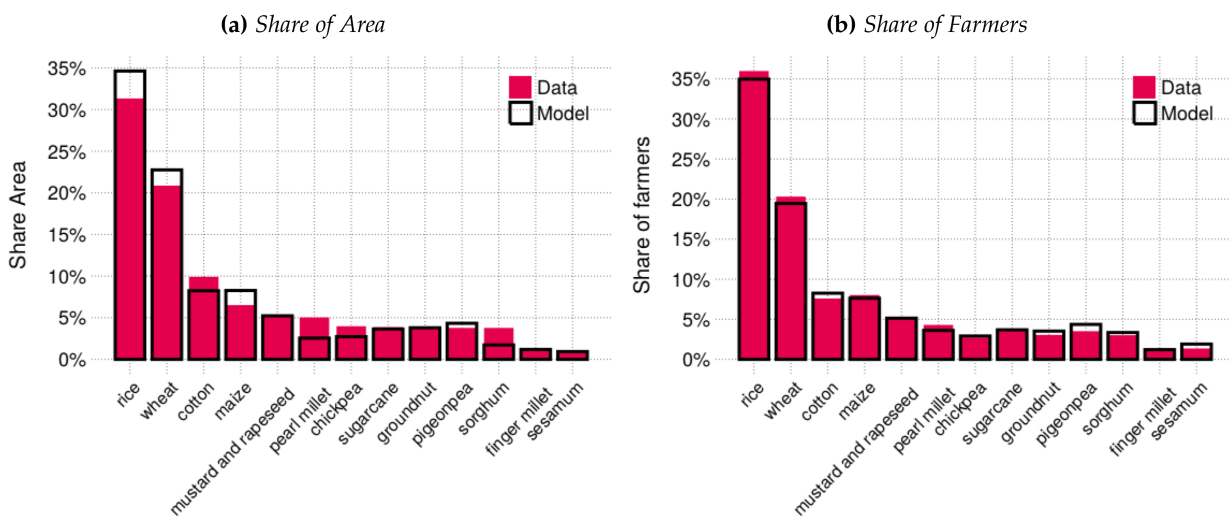

Figure 7: *Supply-Side Estimates: Comparing Model-Predictions with Data*

Notes. The left panel compares the share of area allocated to each crop as observed in the data and as predicted by the model. The right panel reports the same for the share of farmers growing a crop.

Supply Model Fit

To assess model fit, we sample a set of farmers and solve their crop and input choice problem, keeping prices fixed. We compare these simulated choices with the data in Figures 7a and 7b.

4.2 Demand

PDS Crops: Rice and Wheat

For rice and wheat, we estimate the specification in (4) using household-level consumption data from the 68th round of the NSS, conducted in 2011-12. We proxy for household income using total monthly expenditures. All variables are measured at per capita level.

To address potential endogeneity of prices, we rely on Hausman et al. (1994) and instrument prices using average price in the state excluding own district. These are valid instruments under the assumption of idiosyncratic district-level demand shocks, which may enter the error term u_{ch} in (4), but correlated state-level supply shocks such as processing costs and/or transportation costs. Note that these are outside of our model and only used for the estimation of demand parameters; in our counterfactuals, households would face a single average private market price for each crop. We present the estimated parameters in Table 7. We also show the implied demand

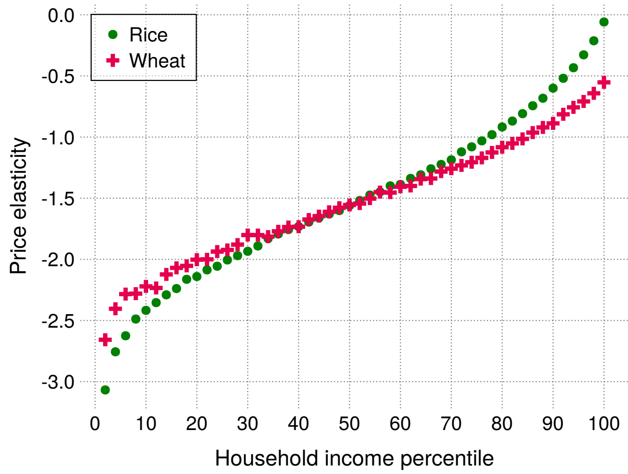

Figure 8: *Price elasticities of demand for rice and wheat by household income*

Notes. This figure shows simulated price elasticities using estimated demand parameters for households with different income levels – proxied using monthly household expenditures.

Table 7: *Estimated Demand Parameters for PDS Crops*

	$log\,(1+q)$
	(1)
(Rice) log price	-7.985***
	(0.268)
(Rice) log income per capita	-2.945***
	(0.105)
(Rice) log price × log income per capita	0.942***
	(0.034)
(Wheat) log price	-3.917***
	(0.150)
(Wheat) log income per capita	-1.055***
	(0.046)
(Wheat) log price × log income per capita	0.422***
	(0.016)
N	186,866
R^2	0.425

Notes. This table shows estimated parameters using the specification in (4) for PDS crops, rice and wheat. Standard errors are reported in parentheses.

elasticities by household income in Figure 8.

Table 8: *Calibrated Demand Elasticities for Non-PDS Crops*

	(1) elasticity
chickpea	-0.57
cotton	-0.85
finger millet	-3.29
groundnut	-0.28
maize	-3.29
mustard and rapeseed	-0.28
pearl millet	-0.45
pigeonpea	-0.57
sesamum	-0.28
sorghum	-0.45
sugarcane	-0.33

Notes. This table reports the calibrated demand elasticities for non-PDS crops.

Non-PDS Crops

We calibrate demand for non-PDS crops using estimates of price elasticities in Deaton (1997). These are given in Table 8.

5 Counterfactuals

In this section, we evaluate the distributional effects of fertilizer subsidies, government-procurement at minimum support prices (MSP), and redistribution of foodgrains through the public distribution system (PDS). We do so with the help of two counterfactuals in which we phase out these programs. These include: (1) no fertilizer subsidies, and (2) no government-procurement at minimum support prices; the latter also results in zero PDS entitlements for households.[58]

While these counterfactuals help us understand the effects of existing programs, they also help us study equilibrium effects of proposed reforms that aim to minimize government's role in the agriculture sector. These include proposals to end fertilizer subsidies (Gulati, 2014) as well as legislation to promote a greater role of private players and potentially smaller role of government

[58] In ongoing work, we also consider how to end government procurement without impacting household consumption through alternative programs such as consumption vouchers.

buyers in output markets (Mashal, Schmall, and Goldman, 2021).[59]

5.1 Solving for the baseline equilibrium

We begin by describing how we solve for the equilibrium in the prevailing regime of fertilizer subsidies, MSP procurement, and PDS entitlements using our estimated parameters. We compute this equilibrium for a sample of 20,000 farmers and all households in the 68th round of the NSS (weighted by sampling weights), that we hold fixed across counterfactuals.

Equilibrium consists of a vector of average private market prices, for the 13 crops we include in our sample, which clears all markets (see Section 3.4). We start with a guess of price vector, solve for optimal production and consumption decisions, and test if all equilibrium conditions hold. If not, we update our guess.[60]

On the supply-side, given this vector of prices, farmers choose which set of crops to plant and make crop-specific input allocations.[61]. We simulate whether sales are made to government buyers or to private traders based on the estimated likelihood of finding government buyers and the distribution of private buyer offers. This gives us the aggregate private market supply and the level of government stockpile of rice and wheat. We then redistribute the government stockpile to households in proportion to their observed entitlements.[62] Given the estimated demand function, and the guess of the price vector, we also know the total demand for each household. We can subtract their PDS entitlements from the total demand to get their private market purchases in the counterfactual. Summing across households gives the total private market demand for PDS crops. For non-PDS crops, private market demand is the predicted aggregate demand from (5).

[59] In 2020, the Indian government attempted to pass bills which would have paved the way for greater private sector involvement in output markets (where farmers sell their harvest). But this attempt was met with a large-scale farmers' protest which lasted for a year, and ended with the repeal of these bills and a demand for a legal guarantee for MSP.

[60] To update, we decrease prices for crops with excess private market supply and increase prices for crops with excess private market demand.

[61] MSP is set by the government taking into account cost of cultivation, and expected market prices. The announced MSP closely tracks the private market prices (see Figure A.2). We do not model this endogenous MSP setting process, but instead assume that the government announces an MSP based on the expected distribution of prices in the private market. On average, the announced MSP is at 59th percentile of the private market price distribution for wheat and 72nd percentile for paddy. We take the mid-point of the two, and assume that the announced MSP is at 65th percentile for wheat and paddy in the counterfactual.

[62] For each household, we compute the share of total entitlements received in the 68th NSS round. We hold these shares constant in each counterfactual simulation, and redistribute total quantities of rice and wheat procured by the government using these shares.

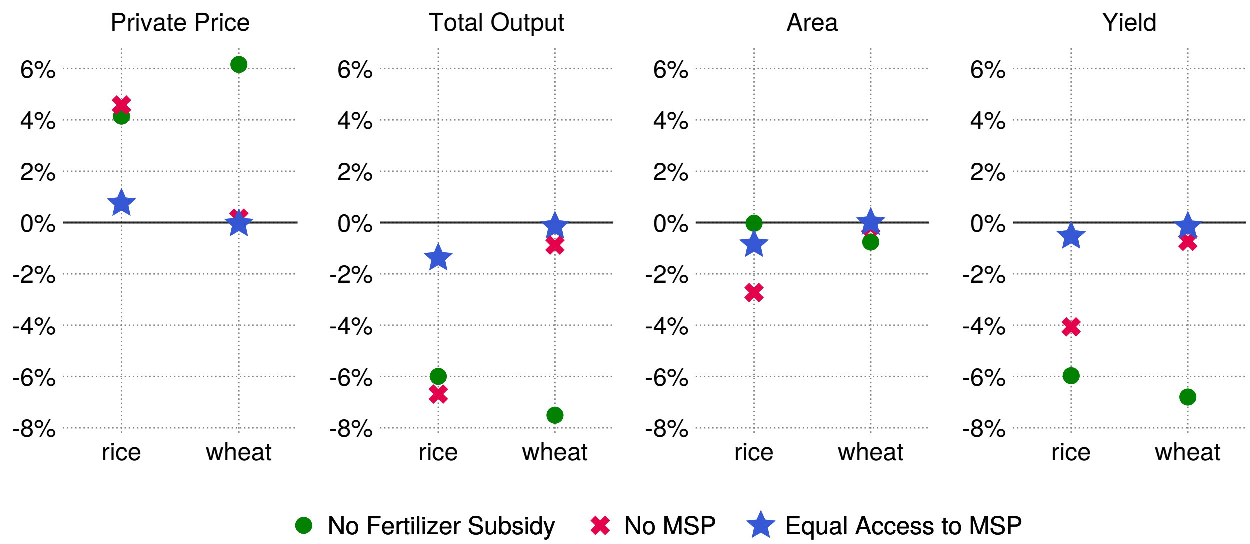

Figure 9: *Percent Changes in Key Variables Relative to Baseline*

Notes. The first panel shows percent change in equilibrium private market prices under different counterfactual policies relative to the baseline prices. The other three panels repeat this exercise for total output, total area, and average yield in the economy. For similar plots for all other crops in our data, see Figure A.4.

5.2 How do we measure distributional effects?

Before presenting results from our counterfactuals, we describe how we characterize the impact on farmer and consumer welfare along the income distribution.

For farmers, the net impact is captured by utility V_{jst} for each farmer j. We compute this utility for each farmer in each counterfactual and calculate changes relative to the baseline equilibrium described above. For example, when we consider the impact of ending government-procurement at MSP, negative changes in V_{jst} would imply that farmer j was relatively better off in the baseline. Then, we compute summary statistics of these changes grouped by income (or farmer size) bins and present them below.

Consumers, or households, differ along two dimensions – income and PDS entitlements; the latter is strongly correlated with the former. In our estimation, we find that lower-income households have relatively higher price elasticities. These households are also more reliant on the PDS for their consumption. Therefore, their consumption and their food expenditures are very sensitive to private market prices as well as the total size of the government stockpile.

While we present effects on consumption and expenditures along the income distribution on the demand-side, neither of these measures fully captures the impact of prices and in-kind transfers on consumers. To summarize this impact, we construct a Laspeyres index for each counterfactual, as described below.

Let the minimum food expenditures to consume a vector of quantities $\boldsymbol{q}^{TOT}$, given a vector of PDS entitlements $\boldsymbol{q}^{PDS}$, be given by

$$e(\boldsymbol{p}^{PVT}, \boldsymbol{q}^{TOT}; \boldsymbol{q}^{PDS}) = \left(\boldsymbol{q}^{TOT} - \boldsymbol{q}^{PDS}\right)' \boldsymbol{p}^{PVT}$$

where $\boldsymbol{p}^{PVT}$ is a vector of private market prices. In the baseline regime, household h consumes quantities $\boldsymbol{q}_{0,h}^{TOT}$ of crops rice and wheat, given by

$$\boldsymbol{q}_{0,h}^{TOT} = \boldsymbol{q}_{0,h}^{PVT} + \boldsymbol{q}_{0,h}^{PDS}$$

In addition, they consume quantity $c_{0,h}$ of the numeraire good, given by

$$c_{0,h} = y_h - e(\boldsymbol{p}_0^{PVT}, \boldsymbol{q}_{0,h}^{TOT}; \boldsymbol{q}_{0,h}^{PDS})$$

where y_h is the total monthly expenditures of household h. The modified Laspeyres index (MLI) under the two counterfactuals is given by

$$MLI_{\text{no fert subsidy},h} = \frac{c_{0,h} + e(\boldsymbol{p}_{\text{no fert subsidy}}^{PVT}, \boldsymbol{q}_{0,h}^{TOT}; \boldsymbol{q}_{\text{no fert subsidy},h}^{PDS})}{c_{0,h} + e(\boldsymbol{p}_0^{PVT}, \boldsymbol{q}_{0,h}^{TOT}; \boldsymbol{q}_{0,h}^{PDS})}$$

$$MLI_{\text{no msp},h} = \frac{c_{0,h} + e(\boldsymbol{p}_{\text{no msp}}^{PVT}, \boldsymbol{q}_{0,h}^{TOT}; \boldsymbol{0})}{c_{0,h} + e(\boldsymbol{p}_0^{PVT}, \boldsymbol{q}_{0,h}^{TOT}; \boldsymbol{q}_{0,h}^{PDS})}$$

In other words, this index captures the relative change in expenditures if the household were to continue to consume the baseline bundle of rice and wheat, and the numeraire good, in counterfactual regimes.[63]

Finally, in the main text below, we do not discuss the change in aggregate consumer surplus associated with non-PDS crops, but we present results in the appendix in Figure A.5.

[63]Since this excludes the impact of change in prices in other crops i.e. non-PDS crops, this is only an approximation to the actual relative change in expenditures.

5.3 The equilibrium without fertilizer subsidies

In a single-crop economy, the impact of removing fertilizer subsidies on downstream consumers is unambiguous. When fertilizer prices go up, fertilizer demand falls, individual and aggregate production falls, market price rises, and the consumption of downstream households falls. The impact on welfare of risk-neutral producers depends on the price elasticity of demand: if demand is inelastic, the decline in demand is low relative to the increase in price, so profits go up.

The conclusions in our setting with crop choice, on the extensive and intensive margins, and risk-averse farmers are more ambiguous. To understand equilibrium distributional effects, we solve for a new equilibrium without fertilizer subsidies. In our data, we approximate an average subsidy rate of 50% across all fertilizer products; as such, we double the price of fertilizer and solve for production and consumption decisions, accounting for government procurement, at different guesses of average private market prices. We stop when all equilibrium conditions are met.

Without fertilizer subsidies, aggregate output of all crops falls, driven by lower consumption of fertilizers, and private market prices go up (see Figure A.4 for all crops). For rice and wheat, prices go up by about 5%, output falls by about 7% (and so does yield) as shown in Figure 9. We estimate government savings to be approximately \$4.35 billion or \$70 per farmer.[64]

If we do not allow farmers to adjust input and crop choice, when fertilizer prices go up, farm profits would unambiguously fall. But since farmers are free to use less and produce less, in equilibrium, they receive higher prices in the private market. Our results suggest that these higher output prices are nearly enough to compensate for the higher per unit cost of fertilizers. As shown in Figure 11b, we find a minimal impact on farmer welfare in the absence of fertilizer subsidies.

On the demand-side, consumption falls. The fall in consumption of rice and wheat is greater for lower-income households who we estimate to be more price sensitive (see Figures 10a and 10b). Since total output falls and, therefore, government procurement is low, PDS entitlements go down. But this impact is small. Expenditures on rice and wheat, as a share of total expenditures, fall by 0.5 percentage points for the lowest-income households; this is just due to lower consumption of rice and wheat by these households. Finally, we summarize the net impact

[64] We use 2014 as our reference year and use the average exchange rate of 0.0164 USD = 1 INR.

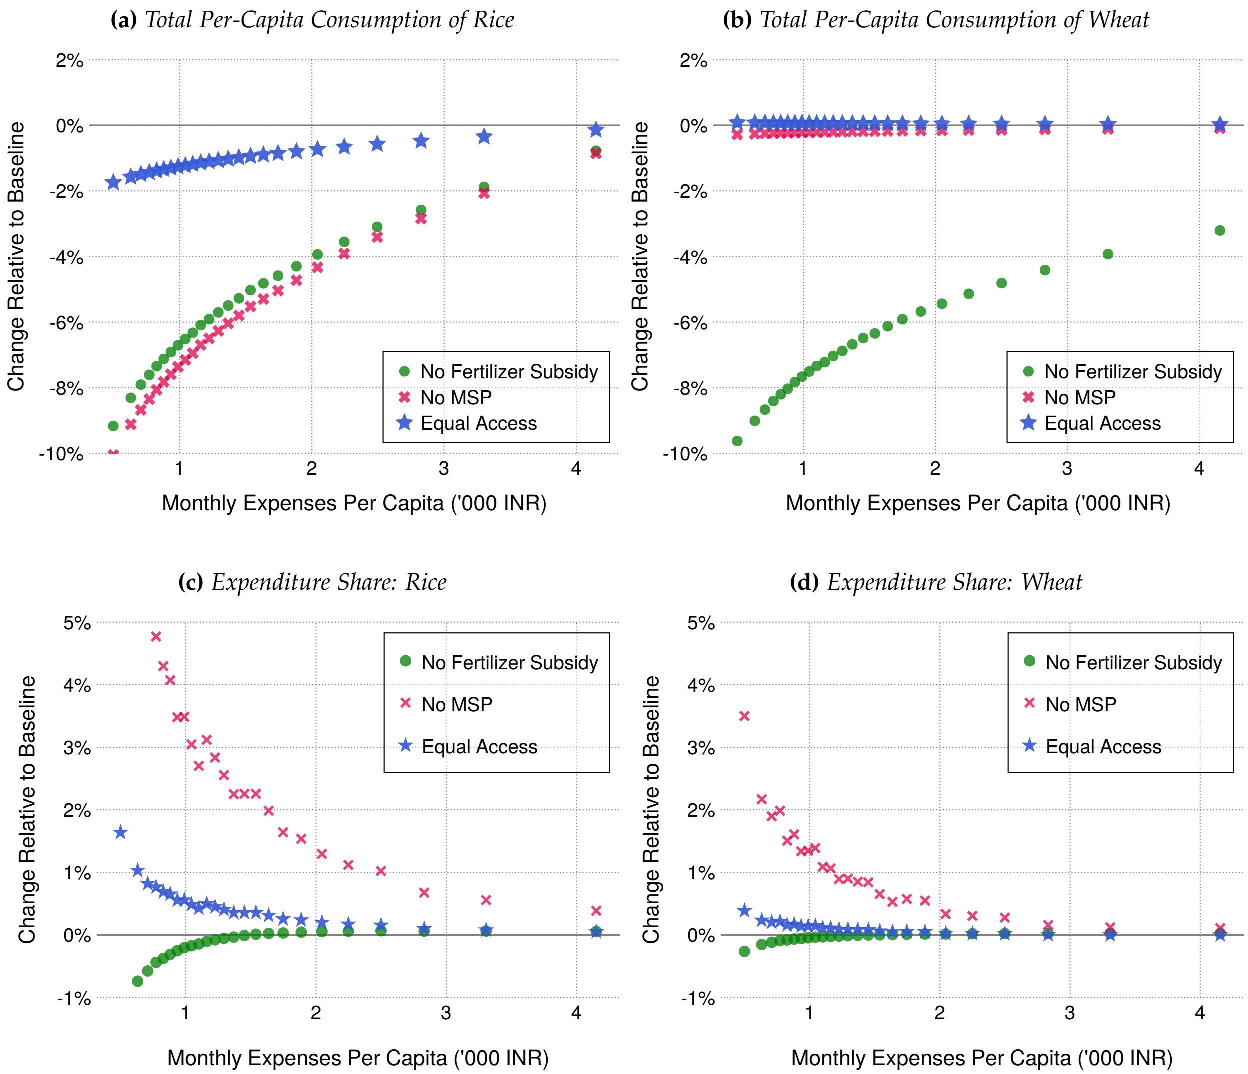

Figure 10: *Change in Household Consumption & Expenditures Relative to Baseline*

Notes. The top-left panel shows mean percent change in household consumption of rice relative to baseline, binned by total household expenses (proxy for household income) under different counterfactual policies. The top-right panel repeats the same for consumption of wheat. The bottom-left panel shows the change in expenditures on rice as a share of total household expenditures (in percentage points) relative to baseline. The bottom-right panel repeats the same for share of expenditures on wheat.

on consumers in Figure 11a using the Laspeyres index described above. To consume the same bundle of rice, wheat, and numeraire good as in the baseline, lowest-income households must now spend 3%-4% more relative to baseline.

5.4 The equilibrium without government-procurement at MSP

Let us consider again the scenario with a single crop. For a risk-averse farmer, minimum support price (MSP), if available, increases the mean and reduces the variance of output price. In the absence of MSP, the farmer would face greater price risk; to reduce exposure to this risk, he would lower input usage and produce less. On the household side, removing government procurement would take PDS entitlements to zero.[65] Consequently, household demand in the private market would go up. Low supply and high demand would give rise to an equilibrium with higher prices. With zero PDS entitlements and higher private market prices, lower-income households would suffer more given their higher reliance on PDS entitlements and their greater price sensitivity.

With multiple crops, farmer response would depend on the relative impact on mean and variance of prices across crops. For example, while rice and wheat may both become less attractive, wheat may become more attractive relative to rice. This could result in more output for wheat when MSP for rice and wheat is taken away as farmers switch from producing rice to wheat. To understand equilibrium effects in our setting, we set the probability of finding a government buyer to zero, simulate farmer and household decisions, and solve for a new vector of equilibrium private market prices. We find that the private market price of rice goes up by about 5% and aggregate output falls by over 6%. We find a minimal impact on the private market price and output of wheat. These differences are due to the differential price risk of rice and wheat in the absence of government procurement at MSP – the estimated variance of private buyer offers for rice is much greater than that of wheat. In the absence of government-procurement at MSP, we estimate government savings to be approximately $8.5 billion or $137 per farmer.

On the supply-side, we find that larger farmers experience a larger loss in welfare. This is because they were more likely to find a government buyer and avail MSP in the baseline. Some of the smallest farmers experience modest gains since they were less likely to sell to government buyers in the baseline and they now receive higher private market prices for rice, which is the preferred crop of small farmers.

On the demand-side, households must now satisfy all demand in the private market. As such, expenditures in the private market go up, as shown in Figures 10c and 10d. This increase

[65] This is an assumption of this counterfactual. We can also consider a scenario where the government stops procuring at MSP but continues to subsidize consumption of lower income households in the private market (e.g. through consumption vouchers).

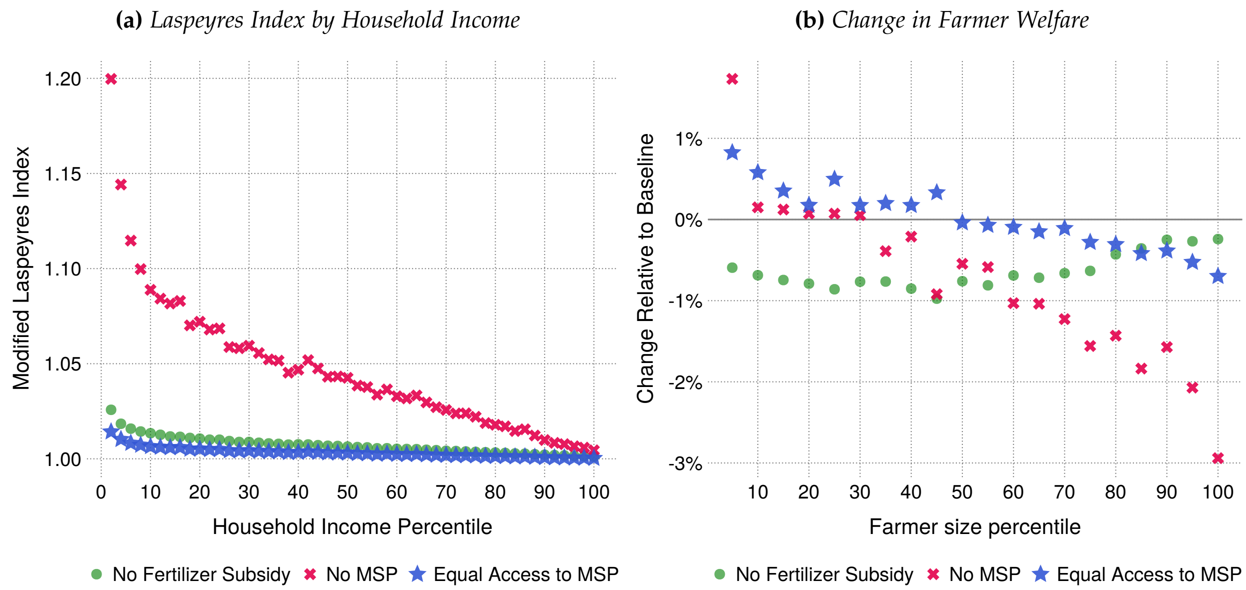

Figure 11: *Distributional Effects on Consumers and Producers*

Notes. The left panel shows a binscatter plot of Laspeyres index by household income under different policy regimes. The index accounts for changes in both prices and in-kind transfers. The right panel shows median percent change in farmer utilities, V_{jst}, relative to baseline, binned by farmer size under different counterfactual policies.

in expenditures is greater for lower-income households since they derived a greater share of their total consumption of these crops from the public distribution system (PDS). For rice, not only do PDS entitlements go to zero, private market prices also go up which exacerbates the adverse effects of this counterfactual on lower-income households. Again, we summarize the impact using our Laspeyres index described above. As shown in Figure 11a, the lowest-income households must now spend 15% to 20% more to consume the baseline bundle of rice, wheat, and a numeraire good.

5.5 What if there was no large-farmer bias in government procurement?

As an additional counterfactual, we consider the impact of a policy where the large-farmer bias in government-procurement at MSP is eliminated. To do so, we hold fixed the number of farmers in each state that the government procures from and randomly assign all farmers to government buyers. Note that this does not hold fixed the quantity of output procured by the government. Total procurement is expected to go down since government buyers would now match with smaller farmers with higher frequency than in the baseline. In fact, procurement of rice falls by about 17% and procurement of wheat falls by about 11%. As such, the government saves about \$1.3 billion or \$20 per farmer.

This alternative policy has minimum impact on private market prices and total output of rice and wheat, as shown in Figure 9, as well as all other crops shown in Figure A.4. Smaller farmers gain (not just the smallest) and the average gains are greater than the scenario where MSP procurement is phased out. Larger farmers are worse off but these losses are small. Importantly, on the demand-side, the impact on lower-income households is minimal. As shown in Figure 11a, under this counterfactual, lowest-income households only pay 1%-2% more to consume their baseline bundle of goods.

6 Conclusion

In this paper, we develop and estimate a structural model of the agriculture sector in India, accounting for the impact of various government-sponsored price interventions on production and consumption decisions. We estimate this model using observational microdata at the farmer and household level and run counterfactuals to characterize the distributional effects of these programs. On the demand-side, we find these interventions to be progressive – these accord greater benefits to lower-income households. In contrast, on the supply-side, we find these interventions to be (weakly) regressive due to inequities in implementation which favor wealthier farmers.

A Additional Figures

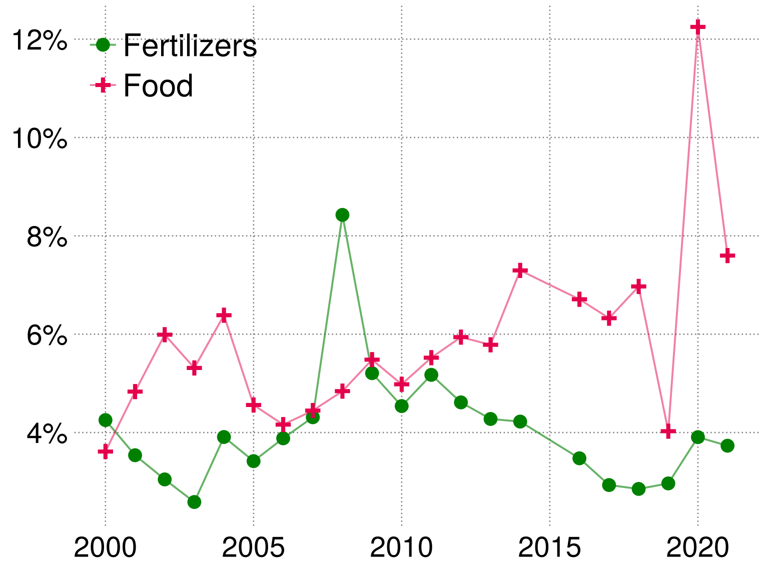

Figure A.1: *Program Costs As a Share of Total Government Spending*

Source. (Revised) budget estimates of the Government of India

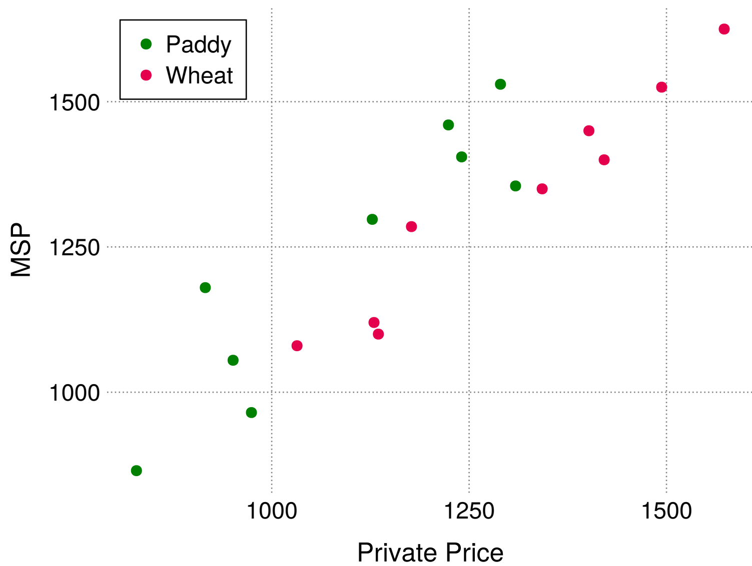

Figure A.2: *MSP Relative to Mean Private Market Price*

Notes. This figure plots the government announced MSP for rice and wheat on the private market mean prices recovered from our estimation.

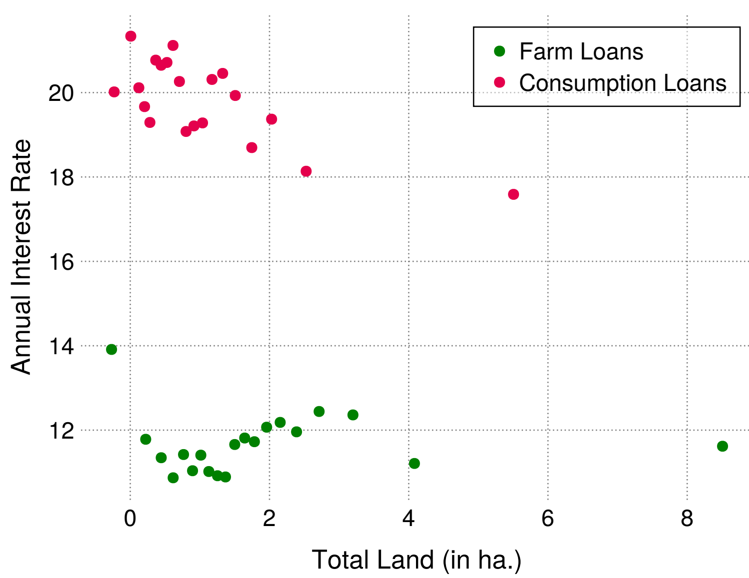

Figure A.3: *Annual Interest Paid for Farm and Consumption Loans*

Notes. The figure plots the average interest rate paid by farmers for farm and consumption loans on total land holdings (in ha.) of the farmer. The data are from the 77th round of the NSS (2019).

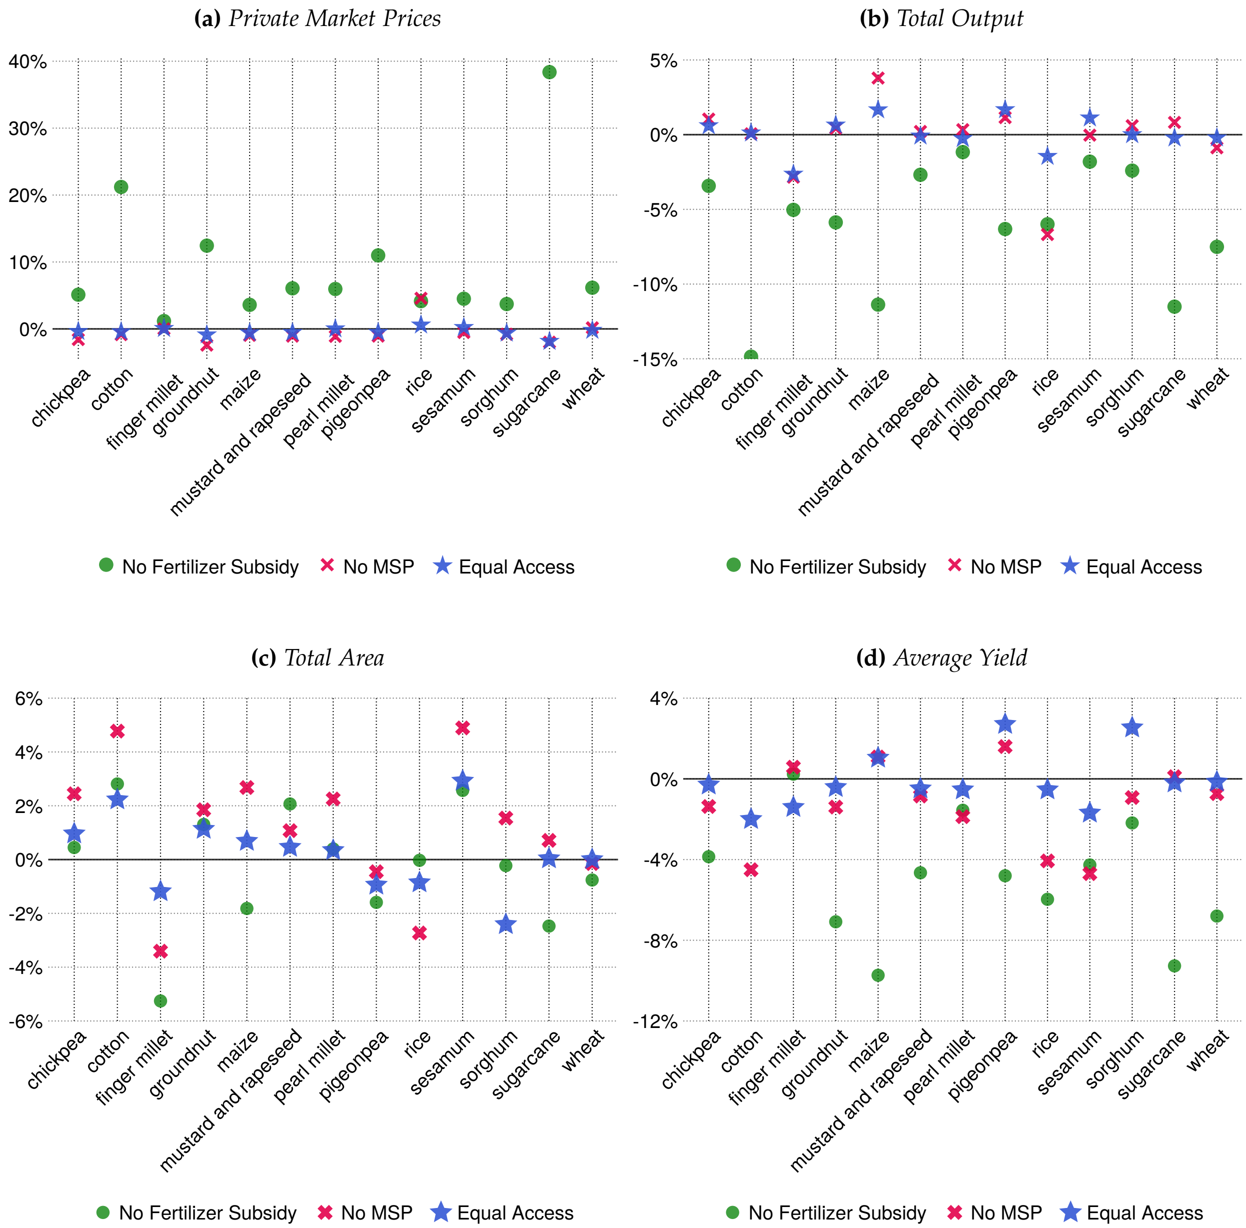

Figure A.4: *Counterfactuals: Percent Change Relative to Baseline, By Crop*

Notes. These plots show the relative change in key aggregate statistics in the various counterfactuals for each crop in our sample.

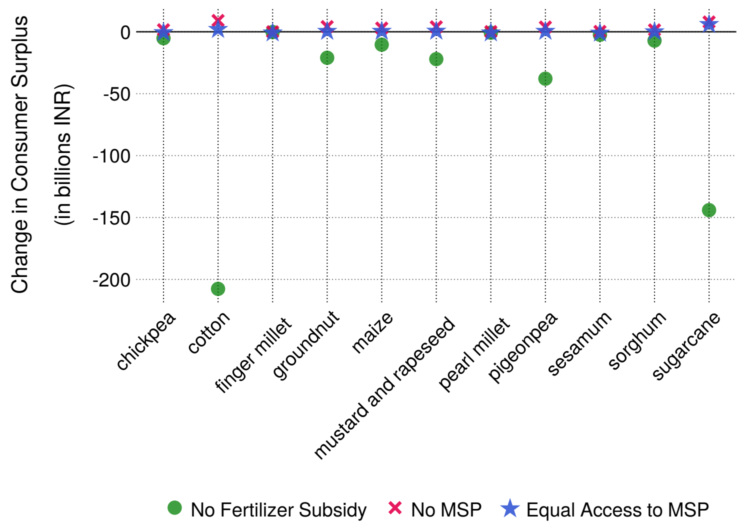

Figure A.5: *Non-PDS Crops: Change in Consumer Surplus*

Notes. This figure shows the change in consumer surplus in different counterfactuals relative to baseline for non-PDS crops. Change in consumer surplus is defined as the area under the demand curve in (5) between the baseline price and the counterfactual price.

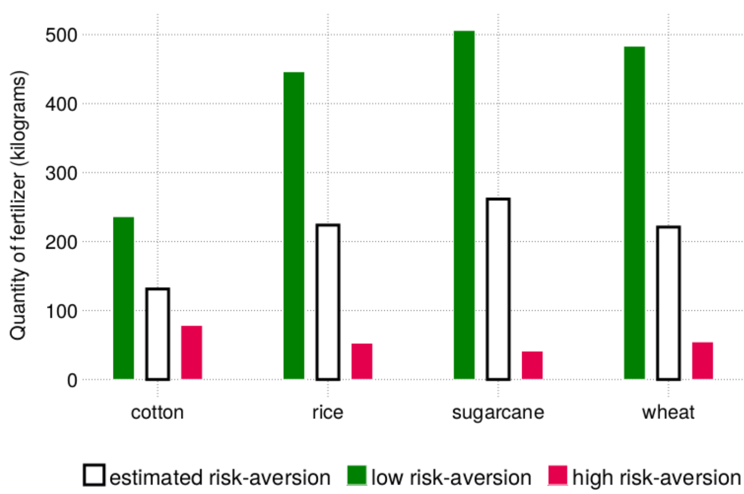

Figure A.6: *Impact of Risk Aversion on Fertilizer Usage*

Notes. This figure shows how predicted average fertilizer usage by crop would differ if risk-aversion were set to a very low or very high level, relative to the model-predicted level of risk-aversion. For all crops, average fertilizer usage falls as risk aversion goes up.

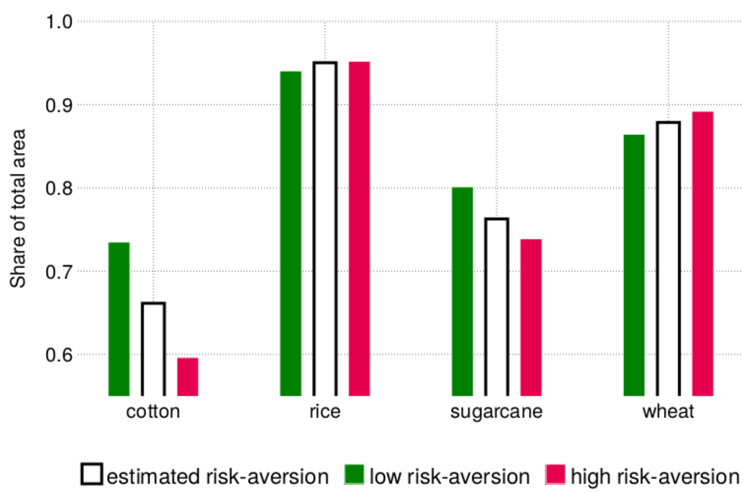

Figure A.7: *Impact of Risk Aversion on Crop Area Allocation*

Notes. This figure shows how predicted average area allocated to each crop would differ if risk-aversion were set to a very low or very high level, relative to the model-predicted level of risk-aversion. For staple crops such as rice and wheat, conditional on growing these crops, average area allocated goes up as risk aversion goes up. The converse is true for cash crops such as cotton and sugarcane.

B Data Appendix

B.1 Cost of Cultivation Surveys (CCS)

Table B.1: *Descriptive Statistics for Cost of Cultivation Surveys*

	(1) area share	(2) observation share	(3) fertilizer use / ha.	(4) labor use / ha.	(5) capital use / ha.
rice	29.70	34.84	150.59	829.96	11.34
wheat	19.76	20.24	156.74	380.14	12.63
cotton	10.08	7.54	184.87	915.09	16.09
maize	6.83	8.16	123.78	514.46	9.12
pearl millet	5.57	4.47	47.41	354.74	9.12
mustard and rapeseed	5.01	4.98	113.13	428.85	10.99
chickpea	4.37	2.98	46.03	299.57	13.91
pigeonpea	4.34	3.69	64.95	474.17	17.51
groundnut	4.20	3.29	93.57	634.47	12.72
sorghum	4.08	3.15	61.65	380.01	8.79
sugarcane	3.82	3.84	377.01	1688.76	10.21
sesamum	1.31	1.56	47.67	371.63	6.36
finger millet	0.93	1.27	88.10	767.92	7.93

Notes. This table shows some descriptive statistics from the Cost of Cultivation Survey, after resampling to match the agricultural census. Column (1) is the share of land allocated to different crops. Column (2) is the share of observations for different crops. Column (3), (4), and (5) are average fertilizer, labor, and capital per hectare for different crops. (3) is recorded in kgs per hectare, while (4) and (5) are hours of use per hectare.

Resampling CCS The government runs the Cost of Cultivation Surveys to get an unbiased estimate of the average cost of growing different crops in the country for farmers of different sizes. The sampling strategy makes the survey unrepresentative due to two reasons. First, within each primary survey unit (PSU - typically a village) the government will sample 2 farmers from each quintile of farm size distribution. Second, PSUs are sampled in proportion to area under cultivation instead of number of farmers in the PSU.

To get a representative sample at the national level, we reweight CCS using the 2016 agricultural census. Agricultural census gives us the proportion of farmers in each size-group $\times$ crop bin. For example, the proportion of farmers that have marginal land holdings ($<$ 0.5 ha.) and grow paddy. We reweight our sample to match this distribution as follows.

Let G denote a group defined by size-category and crop. Let $P_{\text{ag census}}(G)$ be the probability of the group in agricultural census and $P_{\text{ccs}}(G)$ be the probability of the group in Cost of Cultivation

Figure B.1: *Share of farmers growing a given number of crops*

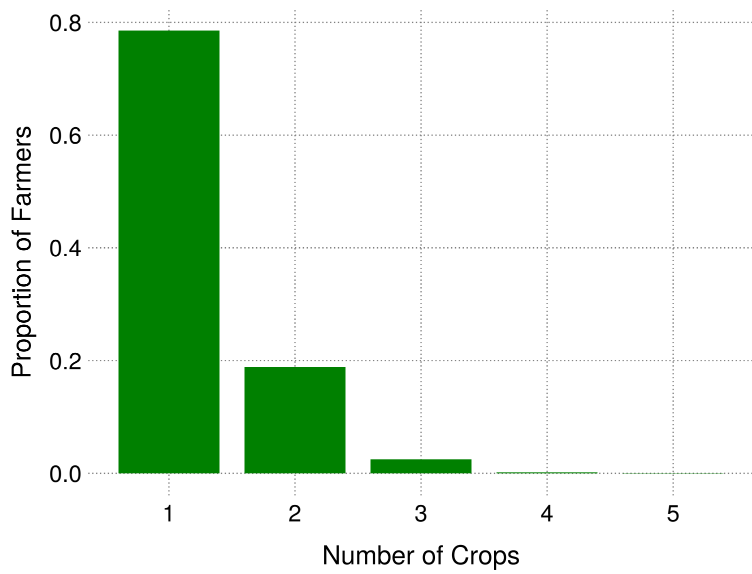

Notes. This figure plots the proportion of farmers that grow different number of crops within the same season. The data is from (resampled) Cost of Cultivation Surveys.

Surveys. The probability in CCS is computed as the proportion of G at the farmer-season-crop level, i.e., the share of farmer-season-crops that belongs to G. We assign a new weight for each farmer-season-crop observation in CCS as,

$$\text{weight} = \frac{P_{\text{ag census}}(G)}{P_{\text{ccs}}(G)}.$$

To compute the farmer weights, we take the mean over all season-crops for the farmer. We resample farmers according to these weights.

C Model Appendix

C.1 Household Demand for PDS Crops

Given income y_h, household h chooses private market quantities of rice and wheat, denoted by $q^{PVT}_{\text{rice},\,ht}$ and $q^{PVT}_{\text{wheat},\,ht}$ to maximize consumption utility given by

$$\begin{aligned}\max_{q^{PVT}_{\text{rice},\,ht},\,q^{PVT}_{\text{wheat},\,ht}} \mathcal{U}_{ht} = &\left(1 + q^{PVT}_{\text{rice},\,ht} + q^{PDS}_{\text{rice},\,ht}\right)^{\delta_{\text{rice},\,h}} + \left(1 + q^{PVT}_{\text{wheat},\,ht} + q^{PDS}_{\text{wheat},\,ht}\right)^{\delta_{\text{wheat},\,h}} \\ &+ \delta_{yh}\left(y_h - P_{\text{rice},\,t} \cdot q^{PVT}_{\text{rice},\,ht} - P_{\text{wheat},\,t} \cdot q^{PVT}_{\text{wheat},\,ht}\right) \end{aligned} \tag{8}$$

where $P_{\text{rice},\,t}$ and $P_{\text{wheat},\,t}$ are the equilibrium private market prices of rice and wheat.

Differentiating (8) with respect to private market quantity for crop c gives

$$\delta_{ch}\left(1 + q^{PVT}_{cht} + q^{PDS}_{cht}\right)^{\delta_{ch}-1} = \delta_{yh} P_{ct}$$

Taking logs and re-arranging gives

$$\begin{aligned}\log\left(1 + q^{PVT}_{cht} + q^{PDS}_{cht}\right) &= \log\left(1 + q_{cht}\right) \\ &= \frac{\log\delta_{ch} - \log\delta_{yh}}{1-\delta_{ch}} - \frac{\log P_{ct}}{1-\delta_{ch}}\end{aligned}$$

where q_{cht} is the total consumption of crop c.

Consider the approximation

$$\begin{aligned}\frac{\log\delta_{ch} - \log\delta_{yh}}{1-\delta_{ch}} &\approx \alpha_{cy}\log y_h \\ \frac{1}{1-\delta_{ch}} &\approx -(\alpha_{cp} + \alpha_{cpy}\log y_h)\end{aligned}$$

Plugging it back in gives

$$\log\left(1 + q_{cht}\right) = a_{cp}\log P_{ct} + a_{cy}\log y_h + a_{cpy}\log P_{ct}\cdot\log y_h$$

which is the specification proposed in (4).