Distributional Effects of Agricultural Interventions in India*

Shresth Garg[†] Sagar Saxena[‡]

First Draft: Nov 13, 2022

Current Draft: Dec 15, 2025

Abstract

India's primary agricultural and food support programs — fertilizer subsidies, procurement at Minimum Support Prices (MSP), and the Public Distribution System (PDS) — cost over 1.5% of GDP and directly impact hundreds of millions of producers and consumers. Yet there is limited understanding of how these large programs interact in equilibrium to shape welfare across the farm-size and income distributions. We develop an equilibrium model featuring heterogeneous risk-averse farmers who choose crop portfolios and input allocations, and heterogeneous households who make consumption decisions. We simulate counterfactuals that either scale back existing interventions or alter their design. Reducing access to procurement generates welfare losses for larger farmers and the poorest consumers. Conversely, the negative impact of cutting fertilizer subsidies on farmers is largely mitigated by an equilibrium increase in market prices. In contrast, redesigns such as equalizing procurement access or targeting fertilizer subsidies to smaller farmers shift benefits toward the poorest producers while generating net welfare gains. More broadly, our analysis reveals that these policies are deeply interconnected: reforms to one propagate through procurement, prices, and food transfers to alter the incidence of others, highlighting the need for a unified framework to assess the impact of any individual program.

*We are grateful to Pol Antras, Emily Breza, Edward Glaeser, Myrto Kalouptsidi, Robin Lee, Marc Melitz, Ariel Pakes, and Gautam Rao for their guidance and support. We also thank Lauren Bergquist, Kevin Donovan, Amit Khandelwal, Robbie Minton, Rohini Pande, Ljubica Ristovska, Mark Rosenzweig, Nick Ryan, Giorgio Saponaro, Roman Sigalov, and seminar participants at Harvard University, Yale University and several other institutions for helpful comments. We gratefully acknowledge financial support from STEG. All errors are our own.

[†]University of Pennsylvania, Department of Economics. Contact: shresth@sas.upenn.edu

[‡]University of Pennsylvania, Wharton School. Contact: sagarsxn@wharton.upenn.edu

1 Introduction

Governments intervene heavily in agricultural input and output markets.[1] In developing countries, where these interventions command a substantial share of the public budget, policymakers often champion them as essential instruments for poverty alleviation and food security.[2] Yet despite the scale and stated redistributive goals of these programs, there is limited evidence on their equilibrium *distributional* effects. Understanding these effects is especially challenging because beneficiaries often face multiple concurrent programs and the incidence of any single intervention may depend on the others in place.

In this paper, we study the distributional effects of agricultural interventions when multiple programs operate simultaneously. We do so in India, focusing on three of its oldest and largest agricultural programs: government procurement of staple grains at Minimum Support Prices (MSP), subsidized food distribution through the Public Distribution System (PDS), and fertilizer subsidies. These programs have been in place for over five decades, reach a large share of the population, and account for roughly 10 percent of government budget. To study their equilibrium effects, we develop a structural model with heterogeneous producers and consumers, estimate it using farmer and household-level microdata, and use it to quantify the incidence of existing policies and evaluate potential reforms.

We begin by documenting three facts that motivate our modeling choices. First, government procurement and food distribution have opposing distributional profiles: the probability of selling to the government rises sharply with farm size, while the share of consumption sourced from the PDS falls with household income. Second, the MSP does not function as a price floor. Roughly 85 percent of rice and wheat sales occur at prices below the announced MSP, and the share of farmers who sell to the government varies widely across states—from

[1] Total support to agriculture across 54 major economies monitored by the OECD reached USD 817 billion per year in 2019–21 (OECD, 2022).

[2] In sub-Saharan Africa, input subsidy programs take up roughly 30 percent of agricultural budgets on average, reaching as high as 70 percent of public agricultural expenditure in countries like Malawi. (Jayne & Rashid, 2013). These programs persist in part because politicians frame them as essential support for small and marginal farmers – a narrative that makes their actual distributional effects a first-order policy question (Anderson, Rausser, & Swinnen, 2013; Birner, Gupta, & Sharma, 2011; Holden, 2019).

over 35 percent of wheat farmers in Punjab to less than 1 percent in Bihar. Third, fertilizer subsidies matter for agricultural output. Using a 2010 reform, which raised non-urea fertilizer prices while holding urea prices fixed, we show that districts that had previously relied more on the newly expensive fertilizers saw larger declines in both fertilizer use and crop output. Together, these facts suggest that policy incidence depends on who can access each program and how farmers respond to the incentives they create.

To quantify how these programs jointly impact producers and consumers, we develop a structural model that links farmers' crop and input decisions to household demand in the presence of multiple government interventions. On the supply side, risk-averse farmers choose a portfolio of crops and then allocate land, labor, machinery, and fertilizer to each crop under yield and price risk. Crop selection is driven by two key sources of unobserved heterogeneity. The first is farmer-crop-specific productivity, which captures differences in skill, soil suitability, and local growing conditions. The second is access to MSP procurement. In our data, we observe whether a farmer sold to the government, but not whether he could have: a farmer with access may transact with private buyers when market prices are attractive. Motivated by the descriptive evidence, we model the probability of access as depending on crop, farm size, and region; for farmers with access, the MSP truncates downside price risk by guaranteeing a floor if private buyers offer less. On the demand side, households vary in income and PDS entitlements, and supplement PDS grains with market purchases. In equilibrium, prices adjust to clear both the private market and the government's procurement and distribution operations.

We estimate the model using detailed microdata from multiple nationally representative surveys: the Cost of Cultivation Surveys, which follow the same farmers over three consecutive years (6 planting seasons) and record inputs and outputs separately for each crop; a cross-sectional survey of agricultural households that records sales channels and prices; and a consumer expenditure survey that captures household purchases by source, including how much comes from the PDS. The panel structure of the production data allows us to estimate production function parameters from within-farmer-crop variation over time. Price data let us estimate the distribution of private market offers after accounting for MSP-induced censoring. Moreover, once we know how often offers fall below the MSP, the share of farmers selling to the government becomes informative about the probability of access to the MSP. The remaining supply-side parameters—risk aversion, access probabilities, productivity distributions, and fixed costs—jointly influence crop choice and are estimated together. We do so using simulated method of moments, matching observed moments to their model-implied counterparts: crop portfolio shares pin down fixed costs, yields identify productivity distributions, the share selling to government buyers identifies access probabilities, and input expenditure shares identify risk aversion. Finally, we calibrate demand parameters directly from observed expenditure shares across income groups after accounting for PDS transfers.

We use the estimated model to evaluate four policy reforms: reducing access to government procurement, cutting fertilizer subsidies uniformly, targeting subsidy reductions to larger farmers, and equalizing MSP access across farm sizes while holding total procurement fixed. A key finding from these exercises is that the three programs are interdependent. Reducing MSP access, for example, lowers fertilizer use because the guaranteed price floor encourages input use. Conversely, cutting fertilizer subsidies contracts agricultural supply, which reduces procurement and, in turn, PDS transfers. These linkages mean that reforms to any single policy propagate through the system to alter the incidence of the others.

These reforms produce different effects across the farm-size and income distributions. Scaling back MSP access hurts larger farmers, who have high baseline access to procurement, and the poorest consumers, who depend heavily on food transfers through the PDS. In contrast, uniformly raising fertilizer prices has a limited impact on farmers because equilibrium output prices rise to offset higher input costs, while the PDS partially insulates low-income households from these price increases. Redesigning programs, rather than simply cutting them, produces more progressive outcomes. When fertilizer subsidies are targeted only to smaller farmers, the smallest producers gain nearly three percent in welfare relative to baseline: they continue to receive subsidized inputs while the supply contraction by larger farmers pushes up market prices. Similarly, equalizing MSP access redistributes benefits from the largest farmers to smaller ones without affecting consumers. Of the scenarios we consider, equalizing access to procurement generates the largest net gain at approximately

1.3 percent of baseline government expenditure, while scaling back MSP access is the only reform that reduces aggregate surplus. Taken together, these results suggest that debates over agricultural policy in India need not be framed as a choice between free markets and state intervention. Reforms that improve targeting—whether by directing input subsidies to smaller producers or by expanding procurement infrastructure to underserved farmers—can shift benefits toward the poorest farmers and consumers while generating net gains.

Related literature. We add to a growing literature that uses structural methods to examine the equilibrium impact of government policies on agricultural markets in developing countries. Existing work typically considers one policy instrument at a time, whether input subsidies (Bergquist et al., 2025; Diop, 2025), trade restrictions (Chatterjee, 2023), or environmental regulation (Hsiao, 2025; Scott et al., 2025; Souza-Rodrigues, 2019). Our setting features multiple large-scale programs that jointly move both supply and demand, and our main contribution is integrating these policies into a single equilibrium model, estimated using detailed farmer-by-crop panel microdata, where reforms to one policy alter the incidence of the others. Chakraborty, Chopra, and Contractor (2025) also study the joint impact of MSP and input subsidies in India but their two-sector model takes a macro perspective on sectoral reallocation and the agricultural productivity gap rather than the precise mechanics of how these policies operate within the agricultural sector. Thus, a key payoff from our more granular and institutionally grounded approach is the ability to evaluate alternative policy designs rather than simply assessing whether these programs should exist.

We also contribute to the broader literature on structural models of agricultural production. Much of this work has examined the drivers of the large agricultural productivity gap in developing countries (Foster & Rosenzweig, 1995; Gollin, Lagakos, & Waugh, 2014; Gollin & Udry, 2021; Lagakos & Waugh, 2013; Suri, 2011). Within this literature, our work is most closely related to studies examining how risk distorts input (Donovan, 2021) and crop choice (Allen & Atkin, 2022; Mobarak & Rosenzweig, 2013; Rosenzweig & Binswanger, 1993). We build on this work by modeling the farmer's decision as a discrete-continuous portfolio choice problem: farmers select a set of crops (extensive margin) and then allocate inputs to them (intensive margin) under yield and price risk, and in response to the specific incentives created by government policies. Methodologically, we treat farmers as multiproduct firms with endogenous product portfolios in the spirit of frameworks from industrial organization (Eizenberg, 2014; Wollmann, 2018). This distinguishes our work from existing models of farm production that typically have farmers choose between broad sectors (e.g., agriculture vs. non-agriculture) or crop categories (e.g., staples vs. cash crops) rather than specific bundles of crops. This distinction is consequential because both policy exposure and household demand vary at the crop level, and ignoring the crop-choice margin may obscure the precise channels through which policy reshapes production and welfare.

Finally, our paper speaks to a literature on the distributional effects of large public programs. Work on food subsidies and transfer design shows that implementation details such as targeting, leakage, and delivery mechanisms shape who ultimately benefits (Banerjee et al., 2018, 2019; Gadenne, 2020). A complementary body of work documents that large-scale interventions generate equilibrium effects that can redistribute surplus between buyers and sellers (Cunha, De Giorgi, & Jayachandran, 2019; Imbert & Papp, 2015; Muralidharan, Niehaus, & Sukhtankar, 2023). These studies typically examine policies targeting one side of the market. We contribute to this literature by studying the joint distributional impact of consumer-side (PDS) and producer-side (MSP procurement and fertilizer subsidies) interventions and tracing their equilibrium incidence across farm-size and income distributions.

2 The Indian Agricultural Policy Landscape

Agriculture remains central to India's economy, even as its manufacturing and service sectors have grown. The sector is the primary source of livelihood for nearly half of the nation's workforce, approximately 600 million people, yet contributes only about 18% to national GDP.[3] The farming landscape is predominantly one of small and marginal holdings, with an average farm size of just over one hectare, which is a fraction of the average farm size

[3] Around 46% of India's population (roughly 600 million people) relies on agriculture as its primary source of livelihood (Government of India, 2025). This agricultural workforce comprises both land-owning cultivators and landless agricultural laborers. According to the 2011 Census, landless laborers account for about 55% of all agricultural workers in India, while cultivators make up the remaining 45% (Singh, 2025).

in other major agricultural producers like the United States (187 hectares) or Brazil (69 hectares).[4] For the many rural households that depend on cultivation, this limited scale often translates to low and volatile incomes, with average farm household incomes falling well below the national median. As a broader national challenge, despite being a major food producer, India is home to hundreds of millions of undernourished people, making food security a policy priority (FAO et al., 2024).

Against this backdrop, the Indian government has established a patchwork of policies to support farmers and ensure food security. The three longest-running and most significant of these are (1) government procurement at Minimum Support Prices (MSP) for select crops, (2) distribution of subsidized foodgrains through the Public Distribution System (PDS), and (3) input subsidies that reduce fertilizer costs for farmers.

The Minimum Support Price (MSP) is a form of market intervention intended to insulate farmers from price volatility and provide a measure of income stability. Each season, the government announces an MSP for several crops and procures select staples, primarily rice and wheat, at these pre-determined prices. The scale of this operation is vast; in recent years, the government has procured about 38% of the total rice and 35% of the total wheat produced in the country.[5] In theory, the MSP acts as a price floor that farmers can expect to receive for their crops. However, as we document in Section 2.2, access to this program is severely limited, meaning that in practice, most farmers sell their output in private markets at prices well below the MSP.

Complementing its price support for farmers, the government operates the Public Distribution System (PDS) to enhance food security for consumers. The PDS is the largest food subsidy program in the world, providing heavily subsidized rice and wheat to over 800 million eligible people through a vast network of over 500,000 "fair-price shops".[6] The program has its roots in the post-independence era of food shortages and has since evolved from a univer-

[4]The average farm size in India is 1.08 hectares based on the 2015-16 Agricultural Census (Press Information Bureau, Government of India, 2020).

[5]From Government of India (2020), in 2018-19, total output of rice was 116.48 million tonnes and total output of wheat was 103.60 million tonnes. In that same year, the government procured 44.40 million tonnes of rice and 35.80 million tonnes of wheat, which corresponds to 38.1% and 34.6% of total output, respectively.

[6]See Press Information Bureau, Government of India (2021) and World Bank (2019) for more details.

sal entitlement to a system targeted primarily at low-income households. Operationally, the grain procured from farmers under the MSP forms the central stock for distribution through the PDS, creating a direct link between producer support and consumer subsidies.

The third major intervention is on the input side: fertilizer subsidies. Introduced in the 1970s during the Green Revolution, the policy compensates manufacturers for the difference between the cost of production and a fixed, low maximum retail price paid by farmers. The subsidy is particularly pronounced for urea, the most common nitrogenous fertilizer, which receives a 75% subsidy per kilogram compared to approximately 35% for DAP and MOP fertilizers (Government of India, 2016). While fertilizer use has become widespread and is credited with significant gains in crop yields, the subsidy regime has also raised concerns about skewed nutrient consumption, fiscal unsustainability, and environmental degradation from overuse (Gulati & Banerjee, 2015).

These three policies collectively represent a massive fiscal commitment: between 2011-12 and 2021-22, combined expenditures on food and fertilizer subsidies averaged over 10.6% of annual government expenditures (see Appendix Figure A.1), or more than 1.5% of India's GDP.[7] This large fiscal footprint, combined with questions about efficiency and equity, has spurred persistent calls for reform. The tension is particularly acute for procurement policy: while policymakers have debated how to contain costs, farmers' groups have mobilized in the opposite direction, organizing large-scale protests to demand a legal guarantee for the MSP, aiming to transform it from a government procurement price into a mandatory, universal price floor (Reuters, 2021; The New York Times, 2024). Simultaneously, the ballooning fiscal burden of fertilizer subsidies and their environmental externalities have fueled a long-standing, separate debate on how to rationalize input support without harming farmer welfare and food security (Gulati & Banerjee, 2015). Quantifying the net impact of this complex policy environment on farmers, consumers, and government budgets, therefore, remains a first-order empirical question for policy reform.

[7] These figures are based on revised estimates from India's union budget documents.

2.1 Data Sources

Our empirical analysis of these policies draws on several publicly-available datasets.

Cost of Cultivation Surveys (CCS)

To model farm-level production, we use detailed micro-data from the Cost of Cultivation Surveys (CCS), conducted by the Department of Agriculture in India. The CCS employs a multi-stage stratified sampling design that is representative of the production conditions for major crops grown in India. States are first divided into agro-climatic zones, from which tehsils (sub-districts) and then village clusters are selected with probability proportional to the area under major crops. Within each cluster, operational holdings are stratified into five size classes, and a fixed number of farmers is randomly sampled from each class.[8]

A key feature of the survey is its panel structure. We use three waves of the survey, covering the years 2011-12 to 2019-20, during which each sampled farmer is followed for all planting seasons over three consecutive years. Importantly, these surveys record detailed data on inputs — both in physical quantities (e.g., hours of labor, kilograms of fertilizer, hours of machinery) and expenditures — separately for each crop cultivated by the farmer. This disaggregation allows us to estimate crop-specific production functions which feed into our supply-side model of crop and input choices. See Appendix C for more details.

NSS 77th Round: Survey of Agricultural Households

The 77th round of the National Sample Survey (NSS) of Agricultural Households, a nationally representative cross-sectional survey conducted in 2019, provides several key inputs for

[8] By stratifying holdings by size before the final stage of sampling, the survey ensures that it captures the full range of input intensities and technologies used across different scales of operation. A simple random sample of farmers would fail to capture enough of the medium and large farms that, while fewer in number, account for a significant share of total output and land use.

our model.[9] First, we use these data to observe the shares of farmers cultivating different bundles of crops. These empirical portfolio shares serve as important moments for identifying the crop-specific fixed costs in our structural estimation. Second, the survey provides detailed information on prices and sales channels. While the CCS provides rich production data, it does not identify the buyers of farm output. The NSS data fill this gap, allowing us to observe the probability that a farmer with certain characteristics (e.g., farm size, region) sells to a government agency. As we argue later, these characteristics-specific probabilities are essential for modeling access to the Minimum Support Price (MSP). Furthermore, the data on realized output prices received by farmers, when combined with consumer prices of key crops, allow us to estimate the parameters governing region-specific distributions of markdowns in the private market. We provide more details on these components of the model in Section 3 and present summary statistics for the NSS data in Appendix Table C.2.

NSS 68th Round: Consumer Expenditure Survey

To model demand, we leverage the 68th round of the NSS Consumer Expenditure Survey, which is a nationally representative survey of over 100,000 households conducted from July 2011 to June 2012. Unlike the 77th round, which focused on agricultural households in rural areas, the 68th round was a survey of both rural and urban households. These data provide information on household size, income, and detailed consumption of various commodities, including rice and wheat. Crucially, household purchases of these staple crops are broken down by source, allowing us to observe the share of a household's consumption that is met through the Public Distribution System (PDS). This allows us to construct household-specific PDS entitlements which we incorporate directly into our demand-side model.

[9] The survey is conducted by the National Sample Survey Office (NSSO) to generate a sample representative of the nation's agricultural households. The survey was conducted in 2019 in two visits to the same set of households, covering the two halves of the agricultural year 2018-19: the kharif (monsoon) season, covering July-December 2018, and the rabi (winter) season, covering January-June 2019. The sample we rely on includes approximately 59,500 rural households.

ICRISAT District Level Database (DLD)

Finally, we also rely on the ICRISAT District Level Database (DLD) for the period 1966-2016. This database provides annual, district-level statistics on key agricultural variables, including cropping patterns, fertilizer consumption, and output prices. We use these panel data for our reduced-form analyses of the impact of fertilizer subsidy policy changes on production decisions and outcomes.

2.2 Facts

We document three empirical patterns in India's agricultural policy landscape that inform our subsequent modeling choices.

Fact 1: Government procurement skews toward larger farmers while foodgrain distribution targets lower-income consumers.

We begin by documenting the opposing distributional outcomes of India's two interlinked foodgrain programs. As shown in Figure 1, these complementary policies create contrasting distributional profiles of their respective beneficiaries: procurement at Minimum Support Price (MSP) is more accessible to larger farmers, while the Public Distribution System (PDS) functions as a progressive transfer to low-income consumers.

On the production side, sales to government buyers are heavily skewed towards larger farmers.[10] Figure 1a plots the probability of selling to a government buyer against farm size. Conditioning on farmers who sell any output, we see a clear positive and roughly (log) linear relationship: as farm size increases, so does the likelihood of selling to a government agency.[11]

[10]While this pattern has been documented by several studies (e.g., Gupta, Khera, and Narayanan (2021)), we emphasize it here as it provides the direct empirical motivation for a core feature of our model: heterogeneity in access to government procurement by farm size.

[11]The positive relationship between the probability of selling to a government buyer and farm size is not just due to regional differences in land size distribution. The pattern persists even after controlling for state fixed-effects (see Figure A.3). Nor is the pattern generated by an intensive margin choice, where large farmers might be more likely to sell to both government and private buyers. The data show farmers who sell to the government tend to sell nearly

Although the pattern is clear, its causes are not: there is limited evidence that pinpoints why small farmers are less likely to sell to the government but several plausible mechanisms have been suggested. Small and marginal farmers typically have much lower marketable surpluses and face higher transaction costs (such as the time and expense of transporting grain to distant procurement centers), which often leads them to sell at the farm-gate to local traders instead of to government buyers (Basu, 2011; Das, 2020; G et al., 2024). Larger farmers, by contrast, are more likely to be aware of MSP opportunities and have the resources (including better transport and storage facilities) to reach procurement centers and wait for procurement agents, allowing them to take advantage of the MSP program. Frictions like corruption and political patronage may also tilt the system in favor of bigger producers and contribute to the pro-large-farmer bias in procurement (Jitendra, 2015).

On the consumption side, the Public Distribution System (PDS) appears to function as a progressive transfer program.[12] Figure 1b shows the share of total rice and wheat consumption that households receive from the PDS, plotted against household income percentiles. The program is well-targeted, with the lowest-income households sourcing 60-70% of their staple grain consumption from the PDS. As intended by the program, in absolute terms, these households also receive a greater quantity (in kilograms) of PDS grains, as shown in Appendix Figure A.2. These patterns highlight the heavy reliance of poor households on the PDS for food security.

Fact 2: MSP is not a price floor and access is highly geographically concentrated.

A potential explanation for the low likelihood of selling to government buyers is that private markets offer higher prices. The data, however, suggest the contrary: the MSP is an attractive price relative to the distribution of prices received in the private market. As shown in Figure 2a, the distribution of prices received by farmers is heavily skewed below the MSP,

their entire output to a single buyer (see Figure A.4). As such, our structural model also does not feature an intensive margin choice of *how* much to sell to the government.

[12] While PDS grains are sold at nominal prices (₹2-3/kg) rather than distributed free, these prices represent less than 10% of market prices. For modeling purposes, we treat PDS transfers as free, though incorporating the nominal prices would be straightforward and would not materially affect our results.

with a large mass of transactions ($\sim$85%) occurring at prices below the supposed floor for rice and wheat.

These national-level statistics mask stark geographic disparities in access to the MSP. In Figure 2b, we show the proportion of rice and wheat farmers in each state who report selling to a government buyer. Nationally, only 8.05% of rice farmers and 4.6% of wheat farmers report selling to the government. As the figure illustrates, likelihood of selling to government buyers varies sharply across Indian states. For example, in Punjab and Haryana, over 35% and 22% of wheat farmers sell their output to government buyers, respectively. In stark contrast, less than 1% of farmers in Bihar report selling wheat to the government, despite Bihar (2.16 million hectares) having a similar area under wheat cultivation as Haryana (2.55 million hectares) in 2018-19 (Government of India, 2020).

This geographic skew is a legacy of historical policy and infrastructure bias. The public procurement machinery for food grains was built up during the Green Revolution era in the late 1960s-70s, focusing on Punjab and Haryana due to their early adoption of high-yield crop varieties and irrigation infrastructure. These states received heavy investments in *mandi* (wholesale markets) networks, storage, and transport to support wheat and rice procurement for the Public Distribution System (Gupta, Khera, & Narayanan, 2021). Decades later, MSP procurement remains concentrated in those same states. By contrast, many other regions (especially eastern India) never saw comparable market infrastructure development. For example, eastern states like Bihar see minimal procurement, leaving most farmers to sell to local traders at farm-gate prices (Chaudhary et al., 2024). Fragmented agricultural markets further entrench these regional disparities. Regulations like the Agricultural Produce Market Committees (APMC) acts prevent farmers from easily crossing state borders to sell to government buyers in regions with better procurement infrastructure (Chatterjee, 2023).[13]

While structural barriers are significant, state-level policy can overcome them. Chhattisgarh, for instance, operates under a decentralized procurement system, offering substantial

[13] This is not to say that agricultural commodity markets themselves are fragmented. Evidence suggests that these markets are relatively well-integrated for traders and intermediaries, with significant trade occurring across states (Allen & Atkin, 2022). The frictions discussed here, stemming from APMC regulations, historically created barriers specifically for farmers seeking to directly sell produce outside their designated local markets.

bonuses above the central MSP and running a network of village-level procurement centers that enable the vast majority of the state's paddy farmers to sell at guaranteed prices (Puri, 2012). Such interventions, however, remain exceptions. The combination of limited procurement infrastructure and interstate trade barriers ensures that for most of India, the MSP functions more as a geographically-concentrated benefit than a universal price guarantee.

Fact 3: Fertilizer prices are an important driver of fertilizer use and agricultural productivity.

India's fertilizer subsidy regime, established through the retention price scheme in the late 1970s, has fundamentally shaped the country's agricultural landscape (Tanpure, 2011). Unlike Sub-Saharan Africa where fertilizer adoption remains persistently low, India's generous subsidies have driven widespread fertilizer use—so much so that policy concerns have shifted from encouraging adoption to preventing overuse (Chand & Pavithra, 2015; Gulati & Banerjee, 2015; Sheahan & Barrett, 2017).

To assess the role these subsidies play in driving fertilizer use and agricultural productivity, we exploit a natural experiment arising from a major policy reform. In April 2010, the Government of India launched the Nutrient-Based Subsidy (NBS) scheme, partially deregulating the fertilizer market. This reform replaced the prior system of product-specific concessions on non-urea fertilizers with subsidies fixed per-kilogram of phosphatic (P) and potassic (K) nutrients.[14] Critically, while urea prices remained controlled, the reform lifted retail price ceilings on non-urea fertilizers (primarily nutrients P and K) allowing manufacturers to set market prices. (Business Standard, 2012; Press Information Bureau, 2025).

This deregulation triggered immediate and substantial price increases for nutrients P and K relative to nitrogen (N) that is primarily supplied by urea (Figure 3a).[15] Correspondingly,

[14]For urea, the most heavily used fertilizer, the program operates by setting a low, uniform maximum retail price (MRP) for farmers and compensating manufacturers and importers for the difference between this controlled price and their costs of production and distribution.

[15]We obtain these inputs prices for these nutrients from the Cost of Cultivation Surveys (CCS). While urea contains only nitrogen (N), some non-urea fertilizers also contain nitrogen, which may explain the small increase observed in nitrogen prices.

district-level consumption of these nutrients declined sharply. This pattern is shown in Figure 3b, which plots coefficients from an event-study regression of (log) fertilizer consumption by nutrient, controlling for district fixed effects.

To investigate the impact of these price changes on agricultural output, we exploit district-level variation in pre-reform reliance on P and K fertilizers. We posit that districts that relied more heavily on P and K fertilizers before 2010 were more exposed to the ensuing price shock assuming inertia in production technologies and agronomic practices prevented easy substitution from these nutrients in the short run. We therefore construct a district-level measure of "treatment intensity" based on the average, price-weighted usage of P and K nutrients per unit of planted area in the pre-reform period (2004-2009). Using this continuous measure of exposure in a difference-in-differences framework, we estimate the causal effect of the price shock on agricultural production. The results show that districts with a higher historical intensity of P and K use experienced significantly larger declines in agricultural output following the 2010 reform. We show this for the district-level output of rice in Figure 3c, and for a district-level output index in Figure 3d.[16] Additional details on the empirical strategy are provided in Appendix B.

These findings support the idea that fertilizer prices play an important role in driving fertilizer use and that India's agricultural productivity is strongly tied to the fertilizer subsidy regime. Thus, any large-scale reforms to fertilizer policy would require careful design to avoid adverse impacts on food production and farm livelihoods.

***

The empirical facts presented above illustrate the key features of India's primary agricultural interventions. While we discussed one policy instrument at a time, these policies do not operate in isolation. They form an interdependent policy ecosystem. For instance, the volume of grain procured by the government at minimum support prices (MSP) directly determines the supply available for the public distribution system (PDS). Similarly, fertilizer

[16] Output index is constructed as a price-weighted sum of output of all crops grown in that district aggregated using national median prices of those crops in the period 2004-2009. Note that the prices used to construct the output index are held constant and therefore only serve as weights to combine output of different crops.

subsidies and price supports are mutually reinforcing: subsidies help ensure sufficient production of crops targeted for MSP procurement, while a guaranteed high price in the form of the MSP, in turn, may incentivize farmers to invest in higher fertilizer use.

Given these interdependencies, evaluating proposed reforms to any one of these policies requires a framework that can account for all major policy instruments. In the following section, we develop a structural model of the agricultural sector to trace how these policies jointly shape production choices, market prices, and farmer and consumer welfare. The framework is designed to capture distributional consequences of these policies, quantifying impacts both across groups (producers and consumers) and within them by income level.

3 Model

In this section, we outline an empirical model of supply and demand for multiple agricultural commodities which incorporates the impact of the suite of major agricultural programs that are currently in place in India.

3.1 Timeline

This is a static, single-period model of the agricultural sector. At the start of the planting season, the government announces policy decisions – 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 a government buyer or to the private buyer. Finally, households receive their PDS entitlements and make purchases from the private market. We summarize the setup of this model in Figure 4.

3.2 Supply: Farmers' Production Problem

In each season, farmers decide which bundle of crops to plant and how to allocate various inputs, including land, to each crop. Both extensive and intensive margin decisions are sensitive to government interventions, which we denote by $\boldsymbol{G} = \left\{\tau_f, \{MSP_c\}_c, \{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}\right\}$, where τ_f is the fertilizer subsidy rate and MSP_c is the minimum support price for crop c. The policy parameters $\{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}$ govern access to government buyers in region r for crop c; we provide more details about these parameters later in this section.

Farmers are indexed by j and are endowed with a farm of total size A_j. In each season, farmers choose a set s of crops to plant which maximize their expected utility over net profits. Their indirect utility, V_{js}, is given by

$$\begin{aligned} V_{js} &= \mathbb{E}\, U_j\left[\Pi_{js} - \kappa_{js}\right] \\ s^* &= \arg\max_{s \in \mathcal{S}_j} V_{js} \end{aligned} \tag{1}$$

where $\mathcal{S}_j$ is the set of all feasible bundles of crops, Π_{js} are the gross profits from planting set s, κ_{js} is the fixed cost of planting set s, and U_j is the utility function of farmer j.

Expected utility $\mathbb{E}\, U_j(\cdot)$ is derived from a nested optimization problem, where, given a set s of crops, farmers choose an allocation of land and inputs for each crop in this set so as to maximize a function of the mean and variance of total net profits. More precisely, let *realized* gross profits at the end of the season be given by

$$\Pi_{js} = R_{js} - C_{js} = \sum_{c \in s} \left(P_{jc} \cdot A_{jc} \cdot Y_{jc}\left(\boldsymbol{X}_{jc}\right) - C(\boldsymbol{X}_{jc}; \boldsymbol{w}_r, \tau_f)\right) \tag{2}$$

where R_{js} and C_{js} are the realized revenue and cost of planting set s, respectively; P_{jc} is the realized price of crop c, A_{jc} is the amount of land allocated to crop c, $Y_{jc}\left(\boldsymbol{X}_{jc}\right)$ is the realized yield of crop c given inputs $\boldsymbol{X}_{jc}$, and $C(\boldsymbol{X}_{jc}; \boldsymbol{w}_r, \tau_f)$ is the cost of inputs $\boldsymbol{X}_{jc}$ given region-specific vector of input prices $\boldsymbol{w}_r$ and fertilizer subsidy rate τ_f. The set of inputs $\boldsymbol{X}_{jc}$ includes amount of fertilizer, labor hours, and machine hours applied per unit of land for crop c. As we describe in the next section, both prices and yields are realized at the time of harvest and therefore treated as random variables at the time of planting. However, input costs are known to the farmers at the start of the season.

Now, the expected utility of farmer j from planting set s is given by

$$\begin{aligned} \mathbb{E}\, U_j \left[\Pi_{js} - \kappa_{js}\right] = \max_{\{A_{jc},\, \boldsymbol{X}_{jc}\}_{c \in s}} & \ \mathbb{E}\left(R_{js}\right) - \gamma_j \cdot \sqrt{\mathrm{Var}\left(R_{js}\right)} - C_{js} - \kappa_{js} \\ \text{s.t.} & \ \sum_{c \in s} A_{jc} \leq A_j \end{aligned} \tag{3}$$

where γ_j is a parameter that captures farmer j's level of risk aversion. This objective function penalizes expected revenue by its standard deviation, measuring risk in the same monetary units as revenue; γ_j scales this penalty. Since both fixed and input costs are known to farmers at the start of the season, they are treated as constants in the expected utility function. 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 area of the farm, A_j.[17]

Next, we provide more details on the components of the expected utility function.

Risky Output

Output of crop c depends on the area allocated to it, A_{jc}, and yield $Y_{jc}\left(\boldsymbol{X}_{jc}\right)$. Yield is a function of a vector of inputs per unit area, $\boldsymbol{X}_{jc} = \{L_{jc}, K_{jc}, F_{jc}\}$, where L_{jc} is the number of labor hours per unit of land A_{jc}, K_{jc} is the number of machine hours per unit of land, and F_{jc} is the amount of fertilizer per unit of land. Further, yield depends on a farmer-crop-specific productivity term, ω_{jc}, which is known to the farmer at the time of planting.

Finally, yield is also subject to an idiosyncratic shock, ε_{jc}, which is realized at the time of harvest. This could be a local weather shock (e.g. extreme heat, low rainfall) or a pest or disease shock. Only the mean and variance of this shock are known to the farmer at the time of planting. The uncertainty about its exact realization generates output risk.[18]

[17] We abstract from credit constraints in our analysis. This choice is motivated by government priority-sector lending policies (e.g., Kisan Credit Cards) and NSS data, which indicate that observed farm loans are predominantly sourced from low-cost institutional lenders and that credit for agricultural purposes is available at a significantly lower interest rate than for consumption loans. We discuss this further in Appendix E.

[18] We assume that yield shocks, ε_{jc}, are uncorrelated across farmers, so the model does not feature aggregate

Yield has a Cobb-Douglas form, given by

$$Y_{jc}(\boldsymbol{X}_{jc}) = L_{jc}^{\beta_l} \cdot K_{jc}^{\beta_k} \cdot F_{jc}^{\beta_f} \cdot \exp\{\omega_{jc} + \varepsilon_{jc}\} \tag{4}$$

where β_l, β_k, and β_f are the elasticities of labor, capital, and fertilizer, respectively.[19]

The unobserved productivity term, ω_{jc}, captures the suitability of farmer j's land for crop c as well as any technological know-how and ability, and directly affects both crop choice and input choices. Therefore, in addition to the standard input endogeneity problem (Griliches & Mairesse, 1995; Hoch, 1962; Marschak & Andrews, 1944), we also need to account for selection into crops since the distribution of productivity for the farmers who choose to grow a particular crop will be different from the non-selected distribution. We come back to this issue 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 the 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. Non-PDS crops are always sold to private buyers at the offered price.

What price does the private buyer offer? The price offered by the private buyer, $\widetilde{P}_{jrc}$, to farmer j in region r depends on the price of that crop c in the (national) consumer market, P_c, and a farmer-specific markdown, μ_{jr}. This markdown captures intermediary market

risk. While incorporating such shocks is not infeasible, it would require solving the model for each realization of the aggregate shock and having farmers form expectations across these realizations, which would significantly increase the computational burden. An important consequence of aggregate shocks would be a negative covariance between aggregate output and private market prices, which would act as a natural hedge and lower the total revenue risk faced by farmers. By assuming this covariance is zero, our model likely overestimates total revenue risk and, consequently, may underestimate the degree of farmer risk aversion. However, for this mechanism to be quantitatively significant, the shocks would need to be large and national in scope. Given India's well-integrated agricultural markets, locally correlated shocks are less likely to move national prices, a point emphasized by (Allen & Atkin, 2022).

[19] These elasticities are assumed to be constant across crops for parsimony although relaxing this assumption is straightforward.

power as well as the cost of transporting the crop to the consumer market (Chatterjee, 2023; Meenakshi & Banerji, 2005; Mitra et al., 2018), and is farmer-specific because it depends on which private buyer matches with the farmer at the time of sale. More concretely, the price offered by the private buyer is given by

$$\widetilde{P}_{jrc} = \mu_{jr} \cdot P_c$$

where μ_{jr} is drawn from a region-specific distribution, $F(\theta_r^\mu)$, which is exogenous to the model.[20] Note that, while these distributions differ by region, they do not differ by crop. Thus, a given farmer in region r faces the same markdown for all his crops.

We assume farmers have rational expectations and know the equilibrium price in the national market, P_c, when making planting decisions.[21] However, at the time of planting, farmers do not know the realization of the markdown for the private buyer they will encounter. This uncertainty about the realized price, and its interaction with whether a government buyer is accessible or not, generates price risk.

More precisely, the realized price for crop c for farmer j in region r is given by

$$P_{jrc} = \mathbb{1}\left\{Z_{jrc} = 1\right\} \cdot \max\left\{MSP_c, \widetilde{P}_{jrc}\right\} + \mathbb{1}\left\{Z_{jrc} = 0\right\} \cdot \widetilde{P}_{jrc} \tag{5}$$

where $Z_{jrc} = 1$ if farmer j can access a government buyer for crop c in region r; $Z_{jrc} = 0$ otherwise. Thus, with access to a government buyer, the farmer accepts the private offer only if it exceeds the minimum support price for that crop, MSP_c; otherwise, the realized price equals the MSP for that crop. Without access to a government buyer, the realized price equals the private buyer's offer. Note that for non-PDS crops $Z_{jrc} = 0$ for all j and r.

Whether a government buyer is accessible or not is known to the farmer at the time of planting. This access to government procurement at a pre-announced price provides the

[20]For tractability, we do not endogenize the markdown distributions. While markdowns charged by private intermediaries may respond to changes in government procurement levels (for instance, reduced government presence could increase intermediaries' market power), we hold these distributions constant across our counterfactuals. This simplification is unlikely to severely bias our welfare calculations because, as shown in Figure 8a, equilibrium price changes under our MSP counterfactuals are negligible (less than 0.2%). With such minimal price movements, there is limited scope for markdown adjustments to materially affect our welfare estimates.

[21]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 extension is easy to incorporate since we observe the same farmer multiple times.

farmer with a valuable price floor, analogous to holding a put option: the farmer has the right, but not the obligation, to sell their output at a "strike price" equal to the MSP for that crop, which effectively truncates the lower tail of the private price distribution. This unambiguously increases the farmer's expected price and reduces price variance.

However, not all farmers have access to government buyers. As documented in the previous section, the probability of selling to a government buyer varies significantly by farmer size, region, and crop. We do not endogenize the matching process between farmers and government buyers, but instead treat it as exogenously determined by farmer characteristics. Specifically, we assume that the probability of *finding* a government buyer is given by

$$\rho_{jrc} = \Pr(Z_{jrc} = 1) = \Phi(\alpha_{0rc} + \alpha_{1rc} \cdot \log(A_j)) \tag{6}$$

where $(\alpha_{0rc}, \alpha_{1rc})$ are crop- and region-specific coefficients, and A_j is the total area of farmer j. This specification, therefore, allows for heterogeneity in access to government buyers by farmer size and region, enabling us to match the patterns observed in the data.

Note that while the $(\alpha_{0rc}, \alpha_{1rc})$ parameters are part of the policy vector $\boldsymbol{G}$, they are not directly observable. We think of them as reduced-form proxies for the institutional environment governing procurement; they capture the joint impact of physical infrastructure, administrative capacity, and frictions like transaction costs and information asymmetries. This allows us to then simulate policy counterfactuals that alter this procurement landscape.

Risk Aversion

A central feature of our model is that farmers are not merely profit-maximizers; they are also sensitive to the risk inherent in agricultural production. We formalize this by assuming farmers make production decisions to trade off the mean and variance of total profits. This trade-off is governed by a parameter, γ_j, which summarizes farmer j's preference for risk.

We allow this preference to be heterogeneous across the population, reflecting the idea that farmers differ in their underlying attitudes toward risk, perhaps due to differences in wealth, access to informal insurance, or other unobserved characteristics. To capture this heterogeneity, we assume that each farmer's risk aversion parameter is an independent draw from a common distribution. Specifically, we assume

$$\gamma_j \sim Exponential(\theta^\gamma) \tag{7}$$

where θ^γ is the mean of the exponential distribution from which γ_j is drawn.

Accounting for risk aversion is particularly important in the context of Indian agriculture. The sector is dominated by small and marginal farmers who have limited capacity to absorb adverse income shocks (D'Exelle & Verschoor, 2015; Donovan, 2021; Emerick et al., 2016). This vulnerability is compounded by imperfect consumption credit markets, which make it difficult to smooth consumption in the face of a poor harvest. Furthermore, with agriculture being largely rain-fed, yields are subject to the vagaries of the monsoon, and formal risk-mitigation instruments that could insulate farmers from such shocks are not widely used; the penetration of crop insurance remains low, and while futures markets for agricultural commodities exist, they see negligible direct participation from farmers and are not integrated into policy instruments like crop insurance in the way they are in other contexts, such as the United States. In such an environment, production choices become a primary mechanism for managing risk, making risk aversion an important behavioral feature to model.

Fixed Costs

Despite having many crops in their choice sets, farmers in our data typically cultivate one or two crops per season. To rationalize this limited diversification, we introduce a fixed cost associated with cultivating a given set of crops. These fixed costs can be interpreted as the cost of acquiring crop-specific knowledge, the indivisible cost of specialized equipment, or the time and effort needed to establish relationships with new input suppliers or output buyers.

We model the fixed cost of planting a set of crops s for farmer j, denoted κ_{js}, as being additive and dependent on the farmer's overall scale of operation, A_j. Specifically, we divide farmers into two groups, small and large, based on the median farm size in the data. The fixed cost for a farmer in group $g \in \{\text{small, large}\}$ is the sum of the crop-specific costs for each crop in their chosen set

$$\kappa_{js} = \sum_{c \in s} \kappa_{g(j),c} \tag{8}$$

where $\kappa_{g(j),c}$ is the fixed cost for crop c for a farmer in size group $g(j)$. By making diversification costly, these fixed costs create a trade-off between the risk-reduction benefits of planting more crops and the additional costs incurred in doing so, thereby helping the model match the observed levels of crop diversification.

3.3 Demand: Households' Consumption Choices

The demand side of our model captures how Indian households make consumption decisions across agricultural commodities, accounting for the important role of government food transfers through the Public Distribution System (PDS). In our framework, we treat the supply (production) and demand (consumption) sides of the economy separately. Farmers maximize expected utility over profits (as specified in equation (1)), while households maximize utility from consumption subject to an exogenous budget constraint.[22]

Our assumption about the separability of production and consumption decisions stems from data limitations; specifically, our datasets do not link farm production decisions to household consumption budgets. Our production data (CCS and NSS 77th Round) provide detailed information on farm profits but reveal nothing about the household's total non-farm income. Conversely, our consumption data (NSS 68th Round) capture total household expenditure but do not disaggregate the income sources (farm vs. non-farm) that fund it. Without this income link, we cannot model how realized farm profits endogenously determine the household's consumption budget. Attempting to impute these missing linkages would require strong assumptions that may contaminate identification of supply-side parameters. By maintaining separation, we achieve clean identification while noting that our estimated

[22]This approach is equivalent to assuming that farming households sell all of their output at market prices and then make consumption decisions as pure consumers at some baseline income (i.e. the budget spent on consumption does not change as agricultural profits change).

parameters may absorb some effects of non-separation. For instance, our farmer-size-specific fixed costs of growing particular crops may partly reflect consumption benefits from home production, and this may show up as a lower effective fixed cost of growing staples for smaller farmers who value food security. Similarly, our risk aversion parameters and markdown distributions may capture not just preferences and market frictions, but also differences in consumption-smoothing opportunities across farmers.

Household Preferences

Households are indexed by h and consume a bundle of agricultural goods $\boldsymbol{q}_h$ and a numeraire good η_h (representing all other consumption) to maximize their utility. This consumption vector is the sum of purchases from the private market, $\boldsymbol{q}_h^{pvt}$, and transfers received from the PDS, $\boldsymbol{q}_h^{pds}$. Thus, $\boldsymbol{q}_h = \boldsymbol{q}_h^{pvt} + \boldsymbol{q}_h^{pds}$.[23]

Unlike farmers, who make production decisions under price and yield uncertainty, households make their consumption choices after market prices have been realized. Consequently, their problem is one of deterministic utility maximization, and therefore, we do not model consumers as risk averse.

We specify a nested Cobb-Douglas utility function to reflect plausible substitution patterns between different types of commodities. At the top level, we assume households distinguish between staple and non-staple goods, where staples are rice and wheat, which are procured by the government and distributed through the PDS. Non-staple goods are all other crops in our sample (e.g., maize, cotton, soybeans, etc.).

The overall utility for household h is given by

$$U_h\left(\boldsymbol{q}_h, \eta_h\right) = U_{h,\text{staple}}\left(\boldsymbol{q}_h^{staple}\right)^{\alpha_{h,\text{staple}}} \cdot U_{h,\text{non-staple}}\left(\boldsymbol{q}_h^{non-staple}, \eta_h\right)^{1-\alpha_{h,\text{staple}}} \tag{9}$$

where $\boldsymbol{q}_h^{staple}$ and $\boldsymbol{q}_h^{non-staple}$ are the consumption vectors of staple and non-staple crops, respectively. The parameter $\alpha_{h,\text{staple}}$ represents the household-specific preference for the

[23] As noted in footnote 12, PDS payments are negligible relative to market prices; we therefore treat these transfers as free.

staple composite good.

The staple group consists of rice and wheat, the two primary PDS commodities. The sub-utility for staples is given by:

$$U_{h,\text{staple}}(q_h^{staple}) = q_{h,rice}^{\alpha_{h,rice}} \cdot q_{h,wheat}^{1-\alpha_{h,rice}} \tag{10}$$

The non-staple group includes all other agricultural commodities in our model (e.g., maize, cotton, soybeans) and the numeraire good, η_h. The sub-utility for this group is

$$U_{h,\text{non-staple}}(q_h^{non-staple}, \eta_h) = \eta_h^{\alpha_\eta} \prod_{c \in \text{non-staple}} q_{h,c}^{\alpha_c} \tag{11}$$

where the expenditure shares are constrained to sum to one: $\alpha_\eta + \sum_{c \in \text{non-staple}} \alpha_c = 1$.

To account for differing consumption patterns across the income distribution, we allow for preference heterogeneity. Specifically, the preference parameters for staples, $\alpha_{h,\text{staple}}$ and $\alpha_{h,rice}$, are modeled as functions of household income. For the non-staple nest, we assume preferences are homogeneous across households. This is a simplifying assumption, motivated in part by data limitations, as we do not observe detailed household-level consumption for many non-staple industrial crops.

Budget Constraint

Household h is endowed with a cash income of y_h. The household's problem is to choose a consumption bundle to maximize utility subject to its budget but the problem is complicated by the fact that households also receive PDS transfers.

We assume that PDS transfers are fungible, meaning households treat them as a cash transfer equal to the market value of the grains. This simplification is exact when transfers are inframarginal—that is, when households supplement their PDS entitlements with purchases from the private market. This condition holds for the majority of households in our data; we apply the simplification to all households to maintain tractability.[24]

[24] Alternatively, we could brute force a solution to the demand system by setting $q_{hc}^{pvt} = \max\left\{0, q_{hc} - q_{hc}^{pds}\right\}$ and set

This fungibility allows us to simplify the household's problem significantly. We can model the household as maximizing its utility subject to a single linear budget constraint based on its *effective income*, $\tilde{y}_h$, which is the sum of its cash income and the market value of its PDS entitlements

$$\tilde{y}_h = y_h + \sum_c P_c \cdot q_{hc}^{pds} \tag{12}$$

Treating this as the household's budget constraint, the optimization problem can now be written concisely as

$$\max_{\boldsymbol{q}_h, \eta_h} U_h\left(\boldsymbol{q}_h, \eta_h\right) \quad \text{s.t.} \quad \sum_c P_c \cdot q_{hc} + \eta_h \leq \widetilde{y}_h \tag{13}$$

where the price of the numeraire good η_h is normalized to one, and U_h is given by (9).

PDS Entitlements

Finally, we specify how the quantity of PDS transfers, q_{hc}^{pds}, each household receives is determined. While official PDS entitlements are based on criteria like income and household size, the actual quantities received can vary due to imperfect implementation and local administrative capacity.

Rather than modeling this complex allocation process, we take an empirical approach that leverages our nationally representative consumption data. We calculate each household's share of the total observed PDS distribution for each staple crop. This share, ϕ_{hc}, is assumed to be fixed for the household. The PDS quantity for crop c received by household h is thus modeled as:

$$q_{hc}^{pds} = \phi_{hc} \cdot Q_c^{PDS} \tag{14}$$

where Q_c^{PDS} is the total quantity of crop c distributed through the PDS system (which, in equilibrium, is equal to the total quantity of crop c procured by the government). These

the household budget constraint to $\sum_c p_c \cdot q_{hc}^{pvt} \leq y_h$. However, this increases computational burden without changing the results meaningfully, as in our counterfactuals, at the equilibrium price vector, most households consume all of the transfers received from the food distribution program.

household-specific shares, ϕ_{hc}, are calculated from our data and held constant across counterfactual policy simulations.

3.4 Equilibrium

An equilibrium in this economy is characterized by a vector of national market prices, $\boldsymbol{P}^* = \{P_c^*\}_c$, that all agents take as given and which simultaneously clears the market for each agricultural commodity. Given this price vector and government policies, $\boldsymbol{G}$, farmers make their production decisions to maximize expected utility, and households make their consumption choices to maximize utility. For these individually optimal decisions to constitute an equilibrium, the resulting aggregate quantities must satisfy two key conditions:

1. **Overall Market Clearing:** For each crop c, the total amount produced must equal the total amount consumed:

$$Q_c^S(\boldsymbol{P}^*, \boldsymbol{G}) = Q_c^D(\boldsymbol{P}^*, \boldsymbol{G}) \quad \forall c \tag{15}$$

   where $Q_c^S(\boldsymbol{P}^*, \boldsymbol{G})$ is the aggregate supply, i.e., total expected output of crop c produced by all farmers resulting from their optimal choices of crop portfolio, land allocation, and input usage, and $Q_c^D(\boldsymbol{P}^*, \boldsymbol{G})$ is the aggregate demand for crop c from all households.

2. **Government Stockpile Clearing:** The quantity of each crop procured by the government must equal the quantity it distributes to households via the PDS:

$$Q_c^{govt}(\boldsymbol{P}^*, \boldsymbol{G}) = Q_c^{PDS} \quad \forall c \tag{16}$$

   where $Q_c^{govt}(\boldsymbol{P}^*, \boldsymbol{G})$ is the portion of aggregate supply that is sold to government buyers and Q_c^{PDS} is the total quantity of crop c distributed to households through the Public Distribution System (PDS).

If these two conditions hold, the private market for each crop also clears. The supply sold on the private market, $Q_c^S - Q_c^{govt}$, will exactly equal the demand purchased on the private market, $Q_c^D - Q_c^{PDS}$.

We conclude by highlighting two simplifications in our equilibrium analysis. First, our model represents a closed economy and abstracts from international trade. This is a reasonable approximation for our study period, as, with the exception of rice, only a small share of domestic production for the commodities we consider was traded internationally.[25] Second, we do not model the government's objective function. The goal of our model is to evaluate the impact of given policies, not to explain their endogenous determinants. We therefore treat policy levers as exogenous and focus entirely on modeling the production and consumption decisions of farmers and households in response.

4 Estimation

In this section, we describe our estimation strategy for recovering the parameters of the empirical model outlined above. We begin by describing the estimation of the supply-side of the model, and then move on to the demand-side.

4.1 Supply

We estimate the supply-side of the model in three steps. In the first two steps, we estimate the parameters governing the distribution of private prices (or markdowns) and the parameters governing the yield function. In the final step, we estimate the remaining parameters.

The Distribution of Markdowns

The price offered to farmer j in region r by a private buyer is a fraction of the national consumer price, given by $\widetilde{P}_{jr} = \mu_{jr} \cdot P_c$. The term μ_{jr} is the farmer-specific markdown, which we model as a random draw from a region-specific Beta distribution, $\mu_{jrt} \sim \text{Beta}(\alpha_r^\mu, \beta_r^\mu)$.

[25] In 2018-19, rice exports were ∼10% of domestic production. Other trade was negligible: rice imports (0%), wheat exports (0.2%), and wheat imports (0%) (Government of India (2020)).

The Beta distribution is a natural choice as its support is bounded between 0 and 1. Our goal is to estimate the parameters $(\alpha_r^\mu, \beta_r^\mu)$ that characterize this distribution for each region.

Note that markdowns are assumed to be constant across crops. This is due to a lack of national consumer price data for all crops in our sample; while such prices are readily available for rice and wheat in household survey data, consumer prices for crops like cotton and soybeans are not. We therefore use price data for only rice and wheat to estimate region-specific markdown parameters. Moreover, note that while the minimum support price (MSP) is set nationally, some states offer a bonus on top of the MSP, which creates a higher, state-specific effective MSP. In our estimation, we use these state-specific MSPs as the relevant price floor that censors the private price distribution.

The primary econometric challenge is that government procurement at the MSP creates a selection problem. For farmers with access to government buyers, any private offer below the MSP is rejected in favor of selling to the government. Consequently, these low offers are unobserved, and the empirical distribution of private prices is censored from below. A naive estimation using only the observed prices would yield biased parameters, systematically understating the true magnitude of private market markdowns.

Our estimation strategy addresses this censoring and differs based on the level of procurement in a region. For regions with negligible MSP procurement, the procedure is straightforward. We directly observe the complete distribution of markdowns and can estimate the parameters $(\alpha_r^\mu, \beta_r^\mu)$ using the empirical mean and variance of markdowns, which we calculate from farmer-level prices (NSS 77th round) and national consumer prices (NSS 68th round).

For regions with significant procurement, we employ a simulated method of moments (SMM) procedure. This procedure finds parameters of the underlying markdown distribution that, after accounting for reshuffling of mass induced by government procurement, best replicate the moments of the observed price data. In particular, government procurement moves some mass of farmers from below the MSP to the MSP as shown in Appendix Figure D.1. Our simulation replicates this reshuffling to recover the true underlying parameters. The procedure is as follows.

For a given set of candidate parameters $(\alpha_r^\mu, \beta_r^\mu)$, we first simulate a distribution of private price offers. We then censor this simulated distribution from below, randomly dropping price realizations below the MSP until the share of observations below the MSP in the simulated data matches the share of reported prices below the MSP in the data (NSS 77th round). Next, to ensure we are comparing moments from unambiguously private sales, we drop from both the real and simulated data all prices that fall at the MSP (or within a small bandwidth around it). Finally, we compute the mean and variance of the remaining prices in both the remaining real data and the remaining simulated data, and search for the parameters $(\alpha_r^\mu, \beta_r^\mu)$ that minimize the distance between these moments.

Identification of the markdown parameters comes from the shape restriction imposed by the Beta distribution. The estimation procedure matches the moments of the observed private price distribution after appropriately accounting for the reshuffling of mass induced by the MSP and dropping the observations clustered at the MSP. By fitting the Beta distribution to these remaining data, we identify the parameters that govern the entire underlying distribution. The estimated markdown distribution parameters are reported in Table 1.

Yield Function

In the second step of our supply-side estimation, we estimate the parameters of the yield function specified in equation (4). Taking logarithms, the yield function becomes:

$$\log(Y_{jc}) = \beta_l \log(L_{jc}) + \beta_k \log(K_{jc}) + \beta_f \log(F_{jc}) + \omega_{jc} + \varepsilon_{jc}$$

where L_{jc}, K_{jc}, and F_{jc} are crop-specific input intensities (input per unit area) for labor, capital, and fertilizer, respectively. Our goal is to estimate the production elasticities $(\beta_l, \beta_k, \beta_f)$ and the parameters governing the distributions of the two stochastic components: the farmer-crop specific productivity, ω_{jc}, and the idiosyncratic yield shock, ε_{jc}.[26]

[26] We apply a series of filters to the CCS data that we use to estimate the yield function. These filters are outlined in Appendix C.1. Further, given the Cobb-Douglas structure of the yield function, we estimate the yield function parameters using observations with positive input intensities. In the filtered data, $\approx 90\%$ of observations have positive input intensities.

We assume the yield shock is crop-specific and normally distributed as follows:

$$\varepsilon_{jc} \sim N(-\sigma_{\varepsilon c}^2/2, \sigma_{\varepsilon c}^2) \tag{17}$$

Thus, output risk varies systematically across crops. The farmer-crop productivity term is also assumed to be normally distributed as follows:

$$\omega_{jc} \sim N(\mu_{\omega rc}, \sigma_{\omega c}^2) \tag{18}$$

where the mean varies by region and crop. To maintain tractability and reduce the number of parameters to estimate, we decompose the mean as $\mu_{\omega rc} = \mu_{\omega r} + \mu_{\omega c} - \sigma_{\omega c}^2/2$, where $\mu_{\omega r}$ and $\mu_{\omega c}$ are region- and crop-specific means, respectively.

Our yield function estimation departs from the control function methods commonly used in the productivity literature (Ackerberg, Caves, & Frazer, 2015; Gandhi, Navarro, & Rivers, 2020; Levinsohn & Petrin, 2003; Olley & Pakes, 1996). These methods rely on inverting the input demand function to recover unobserved productivity, which requires a monotonic relationship between productivity and input use. However, this monotonicity assumption fails when farmers have mean-variance utility and productivity enters as a Hicks-neutral shifter. Higher productivity increases both expected output and output variance, creating opposing effects on input demand through the mean-variance utility framework. Risk-averse farmers may reduce input use in response to higher productivity if the variance effect dominates which violates the monotonicity required for control function approaches.

Given these limitations, we employ a fixed effects approach using our panel data. We estimate the production function by regressing log yield on log input intensities while including farmer-crop fixed effects. This specification identifies the elasticity parameters $\beta_l, \beta_k, \beta_f$ from within farmer-crop variation over time, using how changes in a farmer's input use for a specific crop correlate with changes in yield across years. The farmer-crop fixed effects absorb the time-invariant productivity term ω_{jc}, thereby addressing the endogeneity arising from farmers' knowledge of their own productivity when making input decisions. The estimated yield function parameters are reported in Table 2.

The residuals from this fixed-effects regression are the realized yield shocks, ε_{jc}. Because these shocks are realized after production decisions are made, they are uncorrelated with input choices and free from selection bias. We therefore use these residuals to directly estimate the crop-specific variances of the yield shocks, $\sigma^2_{\varepsilon c}$, via maximum likelihood. However, the estimated fixed effects, $\hat{\omega}_{jc}$, cannot be used to directly estimate the parameters of the productivity distribution because they are only observed for farmers who cultivate crop c, meaning its distribution reflects a selected sample. As such, we defer estimation of the productivity distribution parameters to the final step of our estimation procedure, which explicitly accounts for this selection.

Remaining Supply-Side Parameters

We next estimate the remaining supply-side parameters. These include: (i) the parameter θ^{γ} governing the distribution of risk aversion, (ii) the parameters $\{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}$ which determine the likelihood of finding a government buyer, (iii) the parameters $\{\{\mu_{\omega r}, \}_r, \{\mu_{\omega c}, \sigma_{\omega c}\}_c\}$ characterizing the distribution of farmer-crop productivity, and (iv) the fixed costs $\{\kappa_{g,c}\}_{g,c}$ of planting crops.[27]

All of these parameters jointly affect crop choice decisions and cannot easily be isolated. For example, a farmer's decision to plant rice depends on their risk aversion, the probability they can sell their rice to the government at the MSP, their inherent productivity in growing rice, and the fixed cost of entering rice cultivation. Thus, to identify the fixed costs, for example, one needs to know the other parameters. Relatedly, because we only observe outcomes for the crops farmers choose to plant, we also need to account for selection when estimating these parameters.

We address these challenges by employing a full solution approach where, given a guess of parameters, we solve the production problem for each farmer in our sample, and minimize the difference between observed and simulated moments (McFadden, 1989; Pakes, 1986; Pakes & Pollard, 1989). This is computationally challenging because the space of parameters we need

[27]We estimate the parameters governing MSP procurement, $\{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}$, for state-crop combinations with significant procurement and set procurement to zero for the rest. A state-crop has significant procurement if at least 2% of the farmers in the state grow that crop and sell it to the government.

to search over is high-dimensional. To make this high-dimensional search computationally feasible, we use a nested algorithm that proceeds as follows. In the outer loop, we search over the space of the main behavioral and distributional parameters: risk aversion, government procurement access, and productivity distribution. Then, for each guess of these outer-loop parameters, we solve for the optimal fixed costs in an inner loop. With these conditionally estimated fixed costs, we compare the simulated moments with their empirical counterparts and stop once the difference is below a threshold; otherwise, we update the outer-loop parameters and repeat the process. We provide further details in Appendix D.

The inner loop estimates the crop- and farmer size-group-specific fixed costs, $\kappa_{g,c}$, by matching the model's predicted shares of crop portfolios to those observed in the data. For a given set of outer-loop parameters, we compute the gross expected utility (exclusive of fixed costs) for every farmer and every feasible crop portfolio.[28] The inner loop then finds the vector of fixed costs that, when subtracted from these utilities, best rationalizes the observed portfolio choices. The level of crop-specific fixed costs is identified by the additive structure imposed on the fixed cost of a set of crops. To see this, consider two crops that can be in three possible portfolios: $\{c_1\}$, $\{c_2\}$, and $\{c_1, c_2\}$. If we increase the fixed cost associated with c_1 and c_2 by Δ, the relative attractiveness of $\{c_1\}$ 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, simultaneously matching single-crop and multi-crop shares allows us to identify the level of fixed cost for each crop.

In the outer loop, we estimate parameters related to risk aversion, government procurement access, and productivity distribution.

The parameter governing the distribution of risk aversion, θ^γ, is identified from farmers' input use decisions, specifically, their use of fertilizers. Risk aversion directly affects this intensive margin of production through the mean-variance trade-off in equation (3). Since fertilizer is a variable input that increases both expected yield and yield variance, risk-averse farmers will use less fertilizer than risk-neutral farmers, all else equal. We therefore include

[28] In practice, we limit the possible sets of crops to the 30 most common portfolios observed in the data. Allowing for all possible portfolios is straightforward but computationally expensive.

the average share of revenue spent on fertilizer as a moment in our estimation.[29]

The parameters $\{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}$ that determine the probability of finding a government buyer are identified from observed patterns of government sales. We match two sets of moments: (i) the unconditional share of farmers selling to the government by region and crop, and (ii) how these shares vary by farm size groups (small vs. large). The unconditional nature of these moments is crucial: we compute the probability that a farmer sells crop c to the government, not the probability conditional on growing that crop. This distinction matters because MSP policies create selection: government procurement incentivizes farmers to grow MSP crops, so the conditional probability of selling rice to the government (among those who grew rice) overstates true accessibility.

Finally, in the full solution approach, we also estimate the parameters characterizing the distribution of farmer-crop productivity, $\{\{\mu_{\omega r},\}_r, \{\mu_{\omega c}, \sigma_{\omega c}\}_c\}$. These are identified from the variation in yield observed in the data. We match three sets of moments: (i) average yield by crop, (ii) average yield by region, and (iii) the standard deviation of yield by crop. A key challenge is that we only observe yields for farmers who choose to grow each crop — a selected sample with likely higher average productivity than the population average. Our approach addresses this selection: by simulating the full model, we ensure that farmers in the simulated data also select into crops based on productivity (and other parameters), just as they do in the real data. Matching these selected moments allows us to recover the parameters of interest. We report the estimated supply-side parameters in Tables 2 through 6. The risk aversion parameter, along with yield function elasticities, is presented in Table 2. Fixed costs of cultivation and crop-specific yield shock variances are given in Table 3. Parameters governing access to government procurement are in Table 4. Finally, the parameters characterizing the farmer-crop productivity distribution are reported in Table 5 for crop-specific components and in Table 6 for region-specific components.

[29] This identification strategy follows Donovan (2021) that links risk attitudes to intermediate input use. In their framework, farmers facing uninsurable shocks reduce their use of inputs to limit their exposure to low consumption outcomes following a bad productivity shock. While our setup is different, the intuition is the same: input intensity is a relevant empirical moment that is sensitive to farmers' risk attitudes. In Appendix D.4, we provide further evidence to support this identification strategy by demonstrating the sensitivity of fertilizer expenditure shares to perturbations in the risk aversion parameter.

4.2 Demand

We calibrate the parameters of the household demand system using nationally representative consumption data from the 68th round of the National Sample Survey (NSS). The calibration leverages the properties of the nested Cobb-Douglas utility function specified in equations (9) to (11), where preference parameters correspond to expenditure shares.

First, we calibrate the preference parameters for the staple goods nest, $\alpha_{h,\text{staple}}$ and $\alpha_{h,\text{rice}}$. To capture heterogeneity in consumption patterns, we allow these parameters to vary across household income deciles. As such, we split households into deciles based on their per-capita consumption expenditure, which serves as a proxy for income.[30] For each household h in decile d, we calculate their effective income, $\tilde{y}_h$, as the sum of their cash expenditure and the market value of their PDS entitlements. The preference for the staple composite good for households in that decile, $\alpha_{d(h),\text{staple}}$, is then calculated as the total expenditure on staples divided by the total effective income for all households in the decile:

$$\alpha_{d,\text{staple}} = \frac{\sum_{h\in d}(p_{\text{rice}} \cdot q_{h,\text{rice}} + p_{\text{wheat}} \cdot q_{h,\text{wheat}})}{\sum_{h\in d} \tilde{y}_h}$$

where $q_{h,c}$ represents the total consumption of crop c by household h (from both private markets and PDS) and p_c is the private market price. Similarly, the preference for rice within the staple sub-utility, $\alpha_{d,\text{rice}}$, is calibrated as the total expenditure on rice divided by the total expenditure on staples for households in that decile:

$$\alpha_{d,\text{rice}} = \frac{\sum_{h\in d}(p_{\text{rice}} \cdot q_{h,\text{rice}})}{\sum_{h\in d}(p_{\text{rice}} \cdot q_{h,\text{rice}} + p_{\text{wheat}} \cdot q_{h,\text{wheat}})}$$

Next, we calibrate the preference parameters for the non-staple goods, $\{\alpha_c\}_{c\in\text{non-staple}}$, and the numeraire good, α_η, which are assumed to be homogeneous across all households. Directly calculating expenditure shares from the household survey is not possible for all non-staple crops, particularly industrial crops like cotton or soybeans, as their consumer prices are not observed. We therefore adopt an indirect approach. First, we infer the national

[30] Since PDS transfers involve small copayments rather than being entirely free, we adjust household income by removing these nominal PDS expenditures to obtain baseline income y_h and use this to define income deciles and in subsequent analysis.

consumer price for each non-staple crop by taking the average price received by farmers (from the NSS 77th round) and adjusting it upward using our previously estimated region-specific markdown distributions (see Appendix D.5). Second, under the closed economy assumption of our model, we set total expenditure on each crop equal to the total value of its production, which we calculate by multiplying the inferred national consumer price by the total quantity produced nationally. Finally, the preference parameter α_c for each non-staple crop is its share of the total value of the non-staple consumption bundle. The share of the numeraire good, α_η, is the residual which ensures that the shares within the non-staple nest sum to one. The calibrated demand parameters are reported in Figure 5.

4.3 Model Fit

We evaluate how well the model matches the data by comparing model predictions to observed outcomes. On the supply-side, to generate these predictions, we solve the optimization problem for every farmer in our sample using the estimated parameters and simulated productivity draws, holding consumer prices fixed. Figure 6b compares model-predicted aggregate output shares for each crop to observed shares. The model assigns a large share of output to rice and wheat, similar to the data, and it also captures the relative shares of other major crops reasonably well. Figure 6a shows that the model also fits farmers' crop choices: the predicted fraction of farmers growing each crop is close to what we see in the data.

On the demand side, Figure 7a shows that the model tracks the observed share of spending on rice and wheat across the income distribution. As income rises, both the data and the model show a smaller share of total spending going to these staple grains. Within the staples bundle, Figure 7b shows that the model captures the relative preference for rice versus wheat across income deciles. Overall, the model provides a close fit to both supply and demand patterns in the data.

5 Counterfactual Analysis

Having estimated the parameters of our structural model, we now use it to evaluate the effects of changing the scale and design of existing agricultural policies. For each counterfactual, we keep the same sample of farmers from NSS 77 and households from NSS 68 that we used for estimation and solve for the new equilibrium as follows. We begin with a guess of national consumer prices and solve for optimal crop and input choices for each farmer and the optimal consumption bundles for each household. We then aggregate supply and demand for each crop and update our price vector until market-clearing conditions are met. Additional details are in Appendix F.1.

To assess the impact of counterfactual policies, we simulate a baseline equilibrium under the existing policy environment and use it as a benchmark. For farmers, we measure welfare impacts by calculating the change in net utility (V_{js}), expressed in rupees, relative to this baseline. For households, welfare changes are assessed using compensating variation (CV), defined as the additional income needed to maintain baseline utility under the new prices and PDS transfers.[31]

We present distributional effects across farmer size categories and household income levels, as well as aggregate impacts on welfare and government expenditure. Since we do not explicitly model the agricultural labor market and treat wages as exogenous, our analysis excludes welfare impacts for landless laborers, a significant portion of the agricultural workforce.[32] However, we report changes in total labor demand, which may be used to interpret partial equilibrium outcomes for wage earners.

Policy experiments. We consider four counterfactuals. The first two counterfactuals change the scale of existing interventions: (1) lower access to government procurement at minimum support prices (MSP) and (2) lower fertilizer subsidies. The final two counter-

[31] In the figures and tables below, we report $-$CV so a value less than zero denotes a welfare loss and a value greater than zero denotes a welfare gain.

[32] We thank Mark Rosenzweig for emphasizing this caveat.

factuals change the design of existing interventions: (3) targeted fertilizer subsidies and (4) equal access to MSP for all farmers.

Under the first counterfactual, we uniformly reduce the probability of access to government buyers such that the average farmer in each state has a 50% lower chance of finding a government buyer, holding MSP levels fixed. In the second counterfactual, we increase the composite fertilizer price by 20% for all farmers.[33] In the third counterfactual, under targeted subsidies, we apply the 20% fertilizer price increase only to farmers with landholdings above 0.65 hectares; smaller farmers continue to face the subsidized price.[34] In the fourth counterfactual, we equalize access to MSP for all farmers in states with non-negligible procurement by assigning a common access probability (separately for rice and wheat) chosen to match baseline total procurement.[35] Thus, the total quantity of rice and wheat procured by the government is held fixed.[36]

5.1 Aggregate Impacts: Prices and Quantities

We begin by describing the impact on market-level prices and quantities. Figure 8 presents the results for rice and wheat: panel 8a reports percentage changes in prices and panel 8b reports percentage changes in total output, each relative to the simulated baseline for the four counterfactuals introduced above. For non-staple crops, Appendix Figure F.3 reports analogous outcomes.

Under lower access to MSP, equilibrium rice and wheat output falls slightly (by about 0.5%), yet the composition of sales changes dramatically. As shown in Figure 9, government procurement collapses (-46% for rice, -30% for wheat), while private market sales surge

[33] We approximate the average subsidy on the composite fertilizer input at about 50%, so the unsubsidized price is twice the observed price used in the model; the observed price is a weighted average of nutrient prices (N, P, K) in the cost of cultivation surveys.

[34] This is the unweighted median of landholdings in our sample. The weighted median is even lower at 0.47 hectares.

[35] For each staple (rice and wheat), restricting to states where at least 2% of farmers report selling to the government, we search over a single common access probability. For a given guess, we compute the equilibrium, compute total procurement, and update the guess until model-implied procurement matches the baseline. Farmers in states without meaningful procurement in the baseline remain without access; we do not open new procurement states. The resulting equal-access probabilities are 26.8% for rice and 34.9% for wheat.

[36] Note that farmers in states without much procurement in the baseline continue to not have access to MSP.

(+3.4% and +4.3%, respectively). On the supply side, reduced MSP access triggers two opposing effects: higher price risk leads farmers to contract total output, but because fewer farmers can sell to the government, a larger share of this output flows to private buyers. In our simulations, private supply expands on net. On the demand side, lower procurement means lower PDS transfers, which also triggers two opposing effects: the loss of transfer income exerts downward pressure on demand, while households must also replace lost PDS grains with private market purchases. In our simulations, the replacement effect dominates, so private demand also expands. With both private supply and private demand rising, equilibrium prices change by +0.05% for rice and -0.1% for wheat.

The direction of the net price change varies by crop. As shown in Appendix Figure F.5, reducing MSP access generates a larger increase in price risk for rice than for wheat. This is because rice is more likely to be grown in states with more dispersed markdown distributions (greater variance in private buyer offers). Thus lower MSP access induces a stronger supply contraction for rice which tempers the surge in private market supply relative to demand and pushes prices up slightly. For wheat, the weaker supply response means private supply outpaces demand, pushing prices down slightly.

In the second counterfactual, we lower fertilizer subsidies by increasing the observed price by 20%. Since a key input is now more expensive, the agricultural supply curve shifts inward, raising market prices and reducing overall output. When we lower fertilizer subsidies for all farmers, the price of rice and wheat goes up by approximately 1.8%, while total output falls by about 2%. Prices for all other crops also increase, accompanied by declines in their outputs. However, when targeting the subsidy reduction only to larger farmers, as in the third counterfactual, these supply-side impacts are dampened: rice and wheat prices still rise, but by less (around 1.2%) and total production decreases by about 1.3% for both staples. Non-staple crops similarly experience smaller price and quantity movements.

For the two fertilizer subsidy counterfactuals (uniform and targeted reductions), the contraction in agricultural supply generates a secondary feedback loop through the demand side. As output falls, government procurement declines by about 17% (Figure 9a).[37] This reduces the quantity of grain available for the public distribution system (PDS). As discussed above, lower PDS transfers trigger two opposing forces: the loss of transfer income exerts downward pressure on private demand, while the need to replace PDS grains pushes private demand outward. While we do not decompose these specific channels here, the aggregate market outcomes are primarily driven by the initial inward shift in the supply curve. Under the targeted subsidy policy, these demand-side feedbacks are more muted; because more than half of all farmers continue to receive subsidies, the aggregate supply contraction—and the associated fall in procurement—is smaller than under the uniform subsidy cut.

In the final counterfactual, all farmers in states with non-negligible procurement in the baseline are equally likely to find a government buyer. These common access probabilities (26.8% for rice, 34.9% for wheat) ensure that total government procurement matches the baseline. This shuts down the demand-side feedback channel through procurement and PDS transfers while redistributing MSP access from larger to smaller farmers. Smaller farmers find it more attractive to grow rice and wheat and larger farmers find it less attractive. As above, since price risk for rice is higher, the net impact is a reduction in rice output (about −0.1%) and an increase in its price (about +0.1%). On the other hand, for wheat, output goes up (about +0.05%) and price goes down (about −0.08%).

Before we wrap up our discussion on aggregate impacts, note that reducing access to MSP procurement not only decreases procurement volumes but also lowers fertilizer use as shown in Appendix Figures F.8b and F.9. In other words, government procurement at MSP incentivizes farmers to use more fertilizer, similar to direct fertilizer subsidies. Conversely, when fertilizer subsidies are scaled back, agricultural supply contracts, leading to reduced procurement and consequently lower PDS transfers (Figure 9a). These interactions highlight that agricultural policy reforms in India should account for complementarities across interventions; any evaluation of changes to a particular policy is incomplete without understanding how it interacts with other existing policies.

[37] Farmers still match with government buyers with the same probability as before but now they show up with lower quantities of rice and wheat. The underlying assumption is that the government does not change its procurement strategy to compensate for the lower supply.

Finally, we turn to the impact on agricultural labor demand to offer suggestive evidence on welfare effects for landless laborers. As shown in Appendix Figure F.8a, the decline in labor demand is most pronounced under the first counterfactual (lower MSP access), where labor demand falls by 0.7% as farmers contract production in response to greater price risk. For the remaining counterfactuals, the effect is much smaller, with labor demand falling by less than 0.2%. The smallest decline (0.15%) occurs under the targeted fertilizer subsidy. These results suggest that policy changes affecting production incentives have non-trivial implications for agricultural labor, although a complete welfare analysis would require modeling wage adjustments.

5.2 Distributional Impacts

The aggregate outcomes presented above result from choices made by heterogeneous agents, and we now zoom in on these agents, producers and consumers, and quantify the distributional impacts of the four counterfactual interventions above. We begin by discussing impacts on producers across the farm size distribution, then turn to impacts on consumers across the income distribution.

Impact on producers

Figure 10 shows the welfare impacts on producers across the farm size distribution. Panel 10a reports the change in net utility in rupees and panel 10b shows this change as a percentage of average baseline utility for each farmer size decile. In the first counterfactual, which reduces access to MSP, all farmers are worse off, with larger farmers experiencing the largest welfare losses. This is because larger farmers have greater baseline access to government buyers (Appendix Figure F.4) and thus experience the greatest increase in price risk when access is lowered. For the largest decile of farmers, welfare falls by about 1% relative to baseline utility. Though seemingly modest, note that this loss represents a national average across all farmers, the majority of whom have no access to MSP procurement even at baseline.

Moreover, these effects account for farmers' ability to adapt by adjusting their crop portfolios and input choices, which mitigates the welfare impact relative to a setting where cropping decisions are held fixed. In Appendix Figures F.6 and F.7, we show that farmers do indeed make these intensive and extensive margin adjustments.

The first counterfactual changed two things simultaneously: farmers' access to MSP and households' PDS transfers. To disentangle the direct impact of reduced MSP access from the indirect effect of lower PDS transfers, we run an additional policy experiment. Here, we isolate the demand channel by holding procurement fixed while allocating PDS grains randomly across households rather than targeting the poor (see Appendix F.2). Because lower-income households have a high marginal propensity to consume staples, this redistribution generates a negative aggregate demand shock that depresses private market prices. Small farmers, who rely on private sales, suffer welfare losses, while large farmers remain insulated by their MSP access. The contrast with the first counterfactual—where large farmers lost the most—confirms that those losses were driven by the removal of the MSP price floor, not by the contraction in consumer demand.

In the second counterfactual, uniformly raising fertilizer prices by 20% has a minimal impact on farmer welfare. While larger farmers, who use more fertilizer, are the most affected, even their welfare declines by at most 0.4%, as higher input costs are largely offset by the resulting increase in equilibrium output prices. However, the welfare effects are more pronounced under the third counterfactual which increases fertilizer prices for larger farmers. Now, the largest farmers lose about 1.2%, which is their largest loss in any scenario we consider. This is because they face higher input costs without the large output price increase observed under the second counterfactual. In contrast, small and medium-sized farmers (with landholdings below 0.65 hectares) benefit, with the smallest decile gaining nearly 3% in welfare. These farmers continue to receive subsidized fertilizer while also benefiting from higher market prices.

Under the fourth counterfactual, equalizing MSP access across farmers redistributes welfare from the largest farmers to smaller ones. The top decile loses about 0.7% in welfare as their access probability falls, while most other farmers gain, with the smallest decile seeing a welfare increase of about 0.7%.

Impact on consumers

We now turn to consumers. Figure 11 presents welfare impacts on consumers across the income distribution. Consistent with our sign convention, we plot $-$CV so negative values denote losses. Panel 11a reports per-capita CV in rupees by income decile, while Panel 11b scales CV by average income in each decile.[38]

Under the first counterfactual, which reduces MSP access, all households are worse off, but lower-income households bear substantially larger welfare losses. For the poorest decile, welfare falls by about 1.2% of average income, compared to negligible losses for higher-income households. This pattern reflects the differential impact of reduced PDS transfers: lower-income households, who rely more on subsidized food grains, experience a larger decline in effective income when procurement, and consequently PDS distribution, contracts.

In the second counterfactual (uniform 20% fertilizer price increase), all households are worse off as equilibrium prices rise for all crops. Interestingly, when measured in levels (panel 11a), higher-income households appear more adversely affected. This occurs because the PDS remains operational when fertilizer prices rise: lower-income households continue to receive PDS transfers that insulate a large portion of their consumption from market price increases (although these transfers are lower than baseline). In contrast, higher-income households receive little or no PDS support and are fully exposed to price increases. Thus, the PDS program dampens the adverse impact of fertilizer price increase on lower-income households. In proportional terms (panel 11b), however, the pattern reverses as expected: higher-income households lose less as a share of their income since they allocate a smaller fraction of their budgets to food. Under the third counterfactual (targeted fertilizer subsidies), welfare effects

[38] Households are ranked by per-capita monthly consumption expenditure (MPCE) from the NSS 68th round, which we use as a proxy for income. Consumption expenditure includes the imputed value of goods produced and consumed at home (e.g., own-produced food) and other in-kind items. We also use MPCE to normalize Panel 11b: values there are CV divided by the decile-specific average MPCE for each income decile.

are qualitatively similar but more muted, since equilibrium price increases are smaller.

Finally, equalizing MSP access in the fourth counterfactual has minimal impact on all households (approximately 0%). By construction, total procurement and therefore PDS transfers remain unchanged, and equilibrium prices of rice and wheat barely move, leaving consumer welfare largely unaffected across the income distribution.

5.3 Net impacts

In this section, we present the net change in welfare under each counterfactual. These calculations combine the impacts on producers, consumers, and government expenditure.[39] Figure 12 summarizes our findings and presents all values normalized by total baseline government spending on fertilizer subsidies and MSP procurement.

First, we consider the impact on the government budget. Savings are largest under the second counterfactual (uniform fertilizer price increase), amounting to about 24% of baseline expenditure. The targeted subsidy reform (counterfactual 3) also yields significant savings of about 18%. Reducing MSP access (counterfactual 1) saves about 15%, while equalizing access (counterfactual 4) only saves about 2%.

These budgetary savings, however, come largely at the expense of consumer welfare. The welfare loss for consumers is equivalent to 21% of baseline government spending when MSP access is reduced, and 23% and 17% for uniform and targeted fertilizer price increases, respectively. For farmers, the aggregate impacts are much smaller, as prices adjust to partially offset policy changes and most farmers lack access to MSP even at baseline. The largest aggregate loss for farmers occurs when MSP access is reduced (-1.2%), while the smallest loss occurs when it is equalized (-0.3%).

Summing these components gives the net change in welfare. Since we exclude landless laborers, these results should be interpreted with caution. We find that reducing MSP

[39] Our welfare calculations do not include the impact on landless laborers. Since, under all counterfactuals, demand for agricultural labor falls, our estimates likely represent an upper bound on net welfare gains.

access is welfare-reducing: the combined loss for farmers and consumers exceeds government savings, resulting in a net welfare loss of about 7% of baseline government expenditure. The other three scenarios generate net welfare gains. The reforms to fertilizer subsidies, both uniform and targeted, yield a modest net welfare gain of about 0.3%. The largest net welfare gain, approximately 1.3% of baseline spending, comes from equalizing MSP access.

5.4 Discussion

Our counterfactual analysis offers several insights into agricultural policy reform in India. First, the interventions are directly and deeply interconnected. For example, output price supports (MSP) and input subsidies are complementary in driving fertilizer demand; the existence of MSP dampens the impact that a fertilizer subsidy cut would have on fertilizer use. Our finding that reducing MSP access directly lowers fertilizer demand confirms this link. In the other direction, scaling back fertilizer subsidies contracts agricultural supply, which then reduces procurement and PDS transfers. These interactions show that changes to a single policy can affect incidence and incentives directly targeted by other policies.

Second, the counterfactual policies generate significant distributional consequences across the farm size distribution. The largest farmers are worse off under all counterfactuals, particularly when MSP access is reduced and fertilizer subsidies are targeted. In contrast, smaller farmers face minimal negative impacts and gain significantly from policies that are better targeted or more equitable. Note, however, that even farmers in our largest decile are small by global standards: their average landholding is about 6 acres (Appendix Figure F.10). Thus, while our results highlight meaningful distributional differences across the farm size distribution, the trade-offs occur among relatively small-scale producers.

Third, our analysis reveals how the PDS mediates the welfare impacts of upstream agricultural policies on consumers. When fertilizer prices rise, the PDS partially insulates lower-income households from food price increases, even as the program itself contracts due to lower procurement. This protective function means that reforms affecting procurement have large impacts on the poor, while impacts of reforms affecting input prices are partially mitigated for PDS beneficiaries. This finding also underscores the importance of jointly modeling policies that affect food prices and food transfers; studying these policies in isolation may bias our estimates of welfare changes for the poor and mischaracterize how welfare changes distribute across income groups.

Fourth, our analysis highlights a novel linkage between the design of consumer safety nets and producer welfare: targeting food transfers to lower-income households acts as an indirect price support for small farmers. As shown in Appendix F.2, redistributing PDS transfers away from lower-income households generates a negative demand shock in the aggregate that depresses private market prices. While large farmers can retreat to the safety of government procurement, small farmers, who are dependent on private markets, bear the full incidence of this price decline. This suggests that the political economy of PDS targeting extends beyond consumer welfare; directing transfers to the poor generates equilibrium spillovers that support the incomes of the smallest and most vulnerable producers.

Finally, while our results are specific to India, the policies we study are not; governments throughout the developing world rely on some combination of input subsidies, output price supports, and food aid. Moreover, the production environment we model—small-scale, risk-averse farmers making crop and input decisions under yield and price uncertainty—characterizes agriculture across much of the developing world. Our central finding, that these policies interact in equilibrium, therefore applies broadly, while our framework provides a template for analyzing bundled agricultural interventions in other contexts.

Several caveats merit emphasis. Our analysis excludes landless laborers, who comprise a significant portion of the agricultural workforce. Given that labor demand falls under all counterfactuals (by 0.15-0.7%), this omission likely leads us to overstate net welfare gains. Additionally, while equalizing MSP access yields the largest net welfare gain (1.3% of baseline spending), we do not account for the operational costs involved. If the operational costs of reaching smaller farmers substantially exceed those for current procurement operations, these gains could diminish or reverse. Finally, we hold private market markdowns constant, although one expects intermediaries to adjust their behavior in response to changes in procurement policies. However, given the minimal equilibrium price changes under MSP-related counterfactuals, this assumption likely does not significantly affect our conclusions.

6 Conclusion

This paper studies the distributional effects of India's major agricultural interventions: procurement at minimum support prices (MSP), food aid through the public distribution system (PDS), and fertilizer subsidies. We develop a structural model linking household demand for agricultural commodities to farmers' crop and input choices under risk and in response to government interventions, allowing for rich heterogeneity in observed and unobserved farmer characteristics. Our equilibrium model also captures how these interventions interact with private markets to clear residual supply and demand and allows us to quantify how these policies shape outcomes across the farm-size and consumer income distributions.

Our analyses point to meaningful differences in who benefits from each policy. While MSP procurement insures farmers against downside price risk, these benefits mostly reach larger farmers and those in certain regions with historically high procurement. Input subsidies lower production costs for all farmers, but because larger farmers use more fertilizer, they capture more of this support. For smaller farmers, these lower costs are largely offset by lower market prices for food resulting from greater aggregate supply. On the demand side, both fertilizer subsidies and PDS transfers improve welfare for lower-income households. PDS transfers, in particular, support aggregate demand for staples by providing food directly to the poor. The resulting income effect pushes private market prices up and helps small farmers (who are more likely to sell in private markets) earn higher prices for staples.

Taken together, these findings underscore how deeply the policies are intertwined. A change in one policy, such as procurement or subsidies, affects incentives and impacts of other policies. Therefore, evaluating a single program in isolation can lead to misleading conclusions about who benefits and by how much. Our analyses also show that improving the design of these programs yields better outcomes than simply scaling them back, implying meaningful reform is not just a choice between free markets and state intervention.

We conclude with caveats that point toward future research. First, our framework abstracts from labor markets, precluding an analysis of how interventions affect agricultural wages and the welfare of landless laborers, who are among the most vulnerable in the rural economy. Incorporating labor markets would give a more complete picture of welfare impacts for the rural poor. Second, we treat private trader margins as given. Future analysis could endogenize these markdowns to capture how intermediary market power adjusts in response to procurement. Finally, our analysis is static. By shoring up income through price supports and subsidies, the current regime may keep unproductive farmers in agriculture, potentially slowing structural transformation. Adding dynamics would allow us to weigh these long-term consequences against the immediate distributional gains documented here.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Allen, T., & Atkin, D. (2022). Volatility and the Gains From Trade. *Econometrica*, *90*(5), 2053–2092.

Anderson, K., Rausser, G., & Swinnen, J. (2013). Political economy of public policies: Insights from distortions to agricultural and food markets. *Journal of Economic Literature*, *51*(2), 423–477.

Banerjee, A., Hanna, R., Kyle, J., Olken, B. A., & Sumarto, S. (2018). Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia. *Journal of Political Economy*, *126*(2), 451–491.

Banerjee, A., Hanna, R., Kyle, J., Olken, B. A., & Sumarto, S. (2019). Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia. *Journal of Political Economy*, *127*(1), 101–137.

Basu, K. (2011). India's Foodgrains Policy: An Economic Theory Perspective. *Economic and Political Weekly*, *46*(5), 37–45.

Bergquist, L. F., Faber, B., Fally, T., Hoelzlein, M., Miguel, E., & Rodriguez-Clare, A. (2025, April). *Scaling agricultural policy interventions.*

Birner, R., Gupta, S., & Sharma, N. (2011). *The political economy of agricultural policy reform in india: Fertilizers and electricity for irrigation.* International Food Policy Research Institute (IFPRI).

Business Standard. (2012). Time to fix fertiliser policy. *Business Standard.*

Chakraborty, P., Chopra, A., & Contractor, L. (2025, July). *The equilibrium impact of agricultural support prices and input subsidies.*

Chand, R., & Pavithra, S. (2015). Fertiliser Use and Imbalance in India: Analysis of States. *Economic and Political Weekly*, *50*(44), 98–104.

Chatterjee, S. (2023). Market Power and Spatial Competition in Rural India. *The Quarterly Journal of Economics*, *138*(3), 1649–1711.

Chaudhary, S., Mehta, A., James, A., Arya, H. K., & Rana, K. (2024). Evaluation of agricultural price support systems: A comparative analysis of msp implementation in bihar and punjab states of india. *International Journal of Agriculture Extension and Social Development*, *7*(SP-Issue 11), 185–188.

Cunha, J. M., De Giorgi, G., & Jayachandran, S. (2019). The Price Effects of Cash Versus In-Kind Transfers. *The Review of Economic Studies*, *86*(1), 240–281.

Das, R. (2020). Minimum support price in India: What determines farmers' access? *Agricultural Economics Research Review*, *33*(1), 61.

D'Exelle, B., & Verschoor, A. (2015). Investment behaviour, risk sharing and social distance. *The Economic Journal*, *125*(584), 777–802.

Diop, B. Z. (2025). *Upgrade or migrate: The effects of fertilizer subsidies on rural productivity and migration*.

Donovan, K. (2021). The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity (M. Tertilt, Ed.). *The Review of Economic Studies*, *88*(5), 2275–2307.

Eizenberg, A. (2014). Upstream innovation and product variety in the u.s. home pc market. *The Review of Economic Studies*, *81*(3), 1003–1045.

Emerick, K., de Janvry, A., Sadoulet, E., & Dar, M. H. (2016). Technological innovations, downside risk, and the modernization of agriculture. *American Economic Review*, *106*(6), 1537–1561.

FAO, IFAD, UNICEF, WFP, & WHO. (2024). *The state of food security and nutrition in the world 2024: Financing to end hunger, food insecurity and malnutrition in all its forms* (Annual Report) (Joint publication by five UN agencies monitoring global progress towards SDG 2). Food, Agriculture Organization of the United Nations, International Fund for Agricultural Development, United Nations Children's Fund, World Food Programme, and World Health Organization. Rome.

Foster, A. D., & Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, *103*(6), 1176–1209.

G, B., Sarojini, S. G., Mathew, R., & M, M. A. (2024). Ensuring livelihoods: A critical review of the minimum support price's impact on small and marginal farmers. *International Journal of Agriculture Extension and Social Development*, *7*(11), 108–113.

Gadenne, L. (2020). Can Rationing Increase Welfare? Theory and an Application to India's Ration Shop System. *American Economic Journal: Economic Policy*, *12*(4), 144–177.

Gandhi, A., Navarro, S., & Rivers, D. A. (2020). On the Identification of Gross Output Production Functions. *Journal of Political Economy*, *128*(8), 2973–3016.

Gollin, D., Lagakos, D., & Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, *129*(2), 939–993.

Gollin, D., & Udry, C. (2021). Heterogeneity, measurement error, and misallocation: Evidence from african agriculture. *Journal of Political Economy*, *129*(1), 1–50.

Government of India. (2016). Reforming the fertiliser sector. In *Economic survey 2015-16* (pp. 130–139).

Government of India. (2020). Agricultural statistics at a glance 2020. *Ministry of Agriculture & Farmers Welfare, Department of Agriculture, Cooperation & Farmers Welfare, Directorate of Economics & Statistics.*

Government of India. (2025). Chapter 9. In *Economic survey.*

Griliches, Z., & Mairesse, J. (1995, March). *Production Functions: The Search for Identification* (NBER Working Paper No. w5067). National Bureau of Economic Research. Cambridge, MA.

Gulati, A., & Banerjee, P. (2015, August). *Rationalising Fertiliser Subsidy in India: Key Issues and Policy Options* (Working Paper No. 307). Indian Council for Research on International Economic Relations. New Delhi, India.

Gupta, P., Khera, R., & Narayanan, S. (2021). Minimum Support Prices in India: Distilling the Facts. *Review of Agrarian Studies*, *11*(1), 48–71.

Hoch, I. (1962). Estimation of Production Function Parameters Combining Time-Series and Cross-Section Data. *Econometrica*, *30*(1), 34–53.

Holden, S. T. (2019). Economics of Farm Input Subsidies in Africa. *Annual Review of Resource Economics*, *11*(1), 501–522.

Hsiao, A. (2025). *Coordination and Commitment in International Climate Action: Evidence from Palm Oil*.

Imbert, C., & Papp, J. (2015). Labor Market Effects of Social Programs: Evidence from India's Employment Guarantee. *American Economic Journal: Applied Economics*, *7*(2), 233–263.

Jayne, T. S., & Rashid, S. (2013). Input subsidy programs in sub-saharan africa: A synthesis of recent evidence. *Agricultural Economics*, *44*(6), 547–562.

Jitendra. (2015). At market's mercy. *Down To Earth*.

Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural Decisions After Relaxing Credit And Risk Constraints. *The Quarterly Journal of Economics*, *129*(2), 597–652.

Lagakos, D., & Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, *103*(2), 948–980.

Levinsohn, J., & Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, *70*(2), 317–341.

Marschak, J., & Andrews, W. H. (1944). Random Simultaneous Equations and the Theory of Production. *Econometrica*, *12*(3/4), 143–205.

McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica*, *57*(5), 995–1026.

Meenakshi, J., & Banerji, A. (2005). The unsupportable support price: An analysis of collusion and government intervention in paddy auction markets in North India. *Journal of Development Economics*, *76*(2), 377–403.

Mitra, S., Mookherjee, D., Torero, M., & Visaria, S. (2018). Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers. *The Review of Economics and Statistics*, *100*(1), 1–13.

Mobarak, A. M., & Rosenzweig, M. R. (2013). Informal Risk Sharing, Index Insurance, and Risk Taking in Developing Countries. *The American Economic Review*, *103*(3), 375–380.

Muralidharan, K., Niehaus, P., & Sukhtankar, S. (2023). General equilibrium effects of (improving) public employment programs: Experimental evidence from india. *Econometrica*, *91*(4), 1261–1295.

OECD. (2022, November). *Agricultural policy monitoring and evaluation 2022: Reforming agricultural policies for climate change mitigation* (Policy Brief). Organisation for Economic Co-operation and Development. Paris.

Olley, G. S., & Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, *64*(6), 1263–1297.

Pakes, A. (1986). Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. *Econometrica*, *54*(4), 755–784.

Pakes, A., & Pollard, D. (1989). Simulation and the Asymptotics of Optimization Estimators. *Econometrica*, *57*(5), 1027–1057.

Press Information Bureau. (2025, March). Under the nutrient based subsidy (nbs) scheme, a fixed amount of subsidy is provided on subsidized p&k fertilizers depending on their nutrient content.

Press Information Bureau, Government of India. (2020, March). Decrease in Agricultural Holdings.

Press Information Bureau, Government of India. (2021, July). 5.38 lakh Fair Price Shops (FPS) operational in the country.

Puri, R. (2012). Reforming the Public Distribution System: Lessons from Chhattisgarh. *Economic and Political Weekly*, *47*(5), 21–23.

Reuters. (2021). Indian farmers hold mass rally to keep pressure on modi despite climbdown. *Reuters*.

Rosenzweig, M. R., & Binswanger, H. P. (1993). Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments. *The Economic Journal*, *103*(416), 56–78.

Scott, P. T., Souza-Rodrigues, E., Rosenbaum, T., & Goel, S. (2025, November). *Cows and trees*.

Sheahan, M., & Barrett, C. B. (2017). Ten striking facts about agricultural input use in Sub-Saharan Africa. *Food Policy*, *67*, 12–25.

Singh, V. (2025). Systemic agrarian crisis & changing contours of farm workers. *NewsClick*.

Souza-Rodrigues, E. (2019). Deforestation in the amazon: A unified framework for estimation and policy analysis. *The Review of Economic Studies*, *86*(6), 2713–2744.

Suri, T. (2011). Selection and Comparative Advantage in Technology Adoption. *Econometrica*, *79*(1), 159–209.

Tanpure, S. S. (2011). A Study of Fertilizer Policy in India. *International Journal of Agriculture Sciences*, *3*(3), 145–149.

The New York Times. (2024). Why farmers are marching toward delhi again. *The New York Times*.

Wollmann, T. G. (2018). Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. *American Economic Review*, *108*(6), 1364–1406.

World Bank. (2019, February). Schemes to Systems: The Public Distribution System.

Figure 1: *Government Programs and their Beneficiaries Along the Income Distribution*

(a) *Sales to Government Buyers*

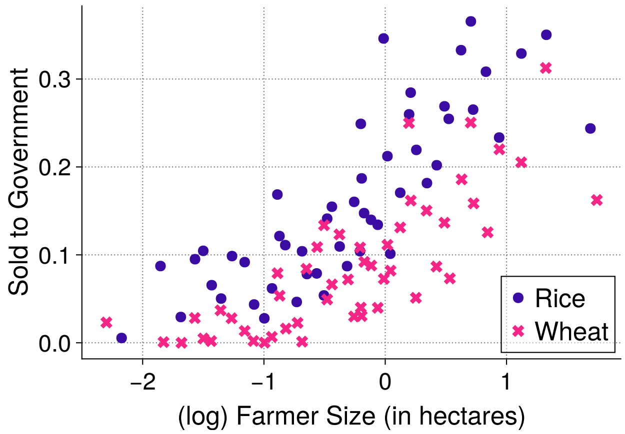

(b) *Share of Consumption From PDS*

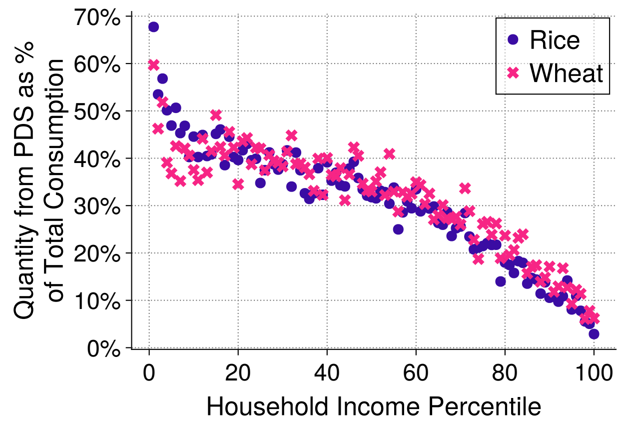

Notes. The left panel shows binned means of an indicator variable denoting whether sales were made to a government buyer against (log) total farm size of the reporting household (in hectares). These data are conditional on selling non-zero amount of output in the market to anyone. The right panel shows binned means of the share of monthly consumption of rice and wheat obtained through PDS against household income percentile, computed using per capita monthly expenditure as observed in the 68th round of the NSS Consumer Expenditure Survey, 2011-2012.

Figure 2: *Government-Procurement and Access to Minimum Support Prices (MSP)*

(a) *Reported Prices Relative to MSP*

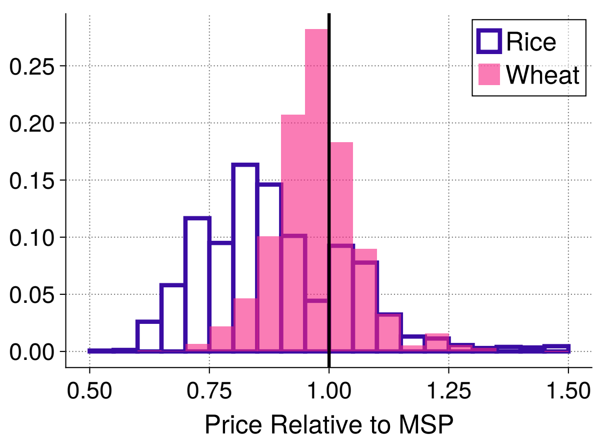

(b) *Sales to Government Buyers by State*

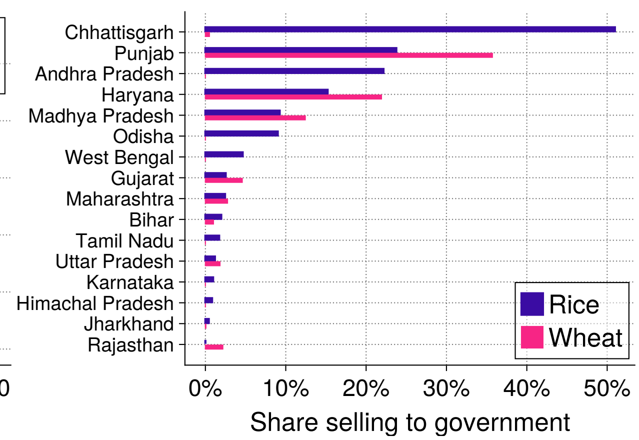

Notes. The left panel shows the distribution of prices of rice and wheat received by farmers relative to the minimum support price (MSP) for that season. The right panel shows the share of rice and wheat cultivators in a state who sold output to government buyers. Source: 77th round of the National Sample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

Figure 3: *Impact of Fertilizer Prices on Production Decisions and Output*

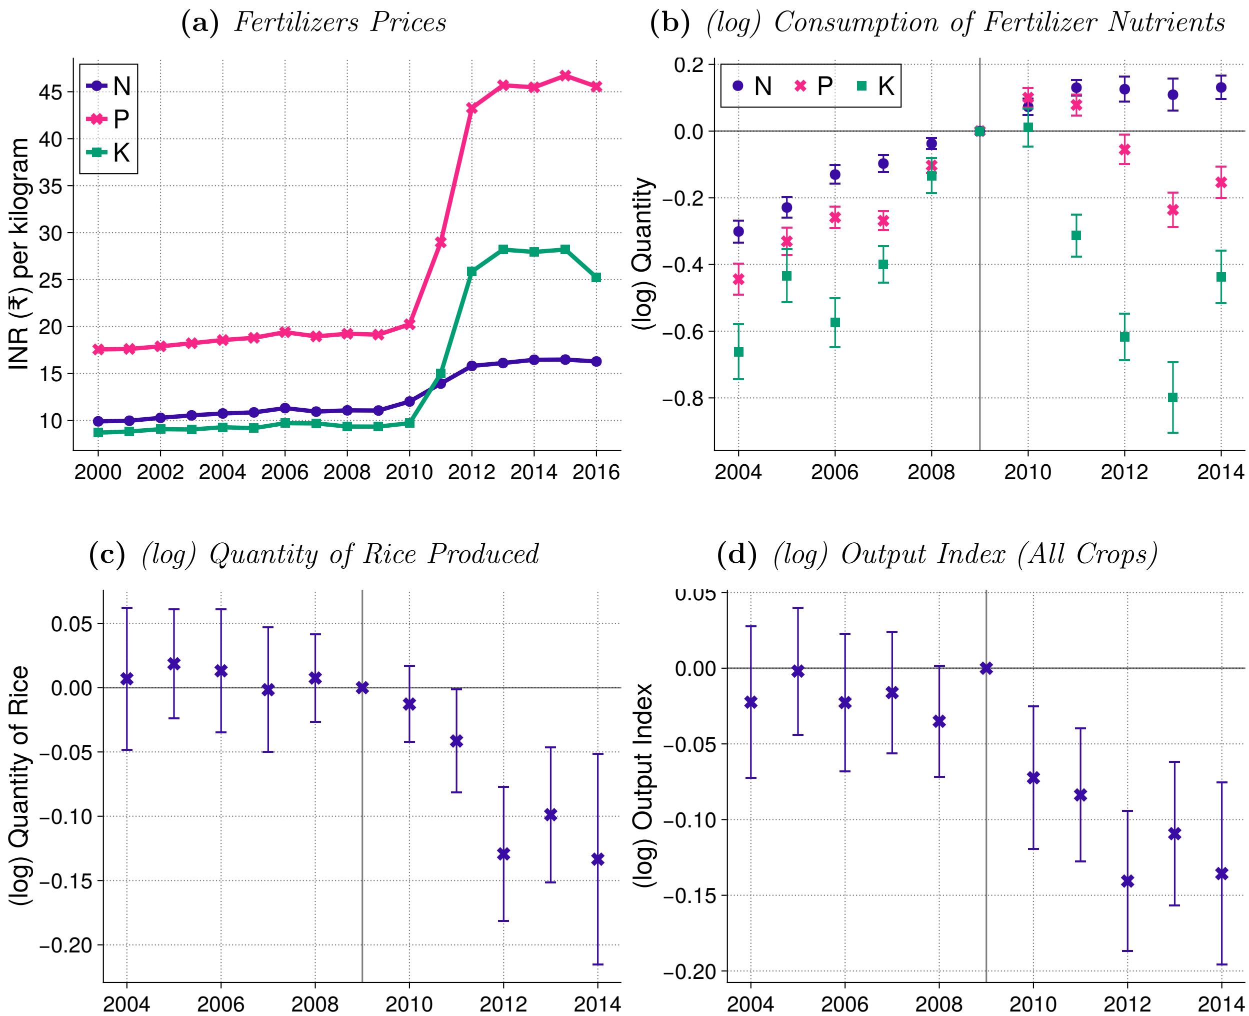

(a) *Fertilizers Prices* **(b)** *(log) Consumption of Fertilizer Nutrients* **(c)** *(log) Quantity of Rice Produced* **(d)** *(log) Output Index (All Crops)*

Notes. In the top-left panel, we plot (weighted) average reported prices of fertilizer nutrients nitrogen (N), phosphorus (P), and potassium (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-left 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 nutrient prices as weights) in the period 2004-2009, before prices of these nutrients increased sharply. The dependent variable is (log) output of rice at the district-level. The controls are year and district fixed effects. The bottom-right panel repeats this exercise using a (log) output index at the district-level as the dependent variable. The output index is constructed using output of all crops grown in that district aggregated using national median prices of those crops in the period 2004-2009.

Figure 4: *Model Timeline & Overview*

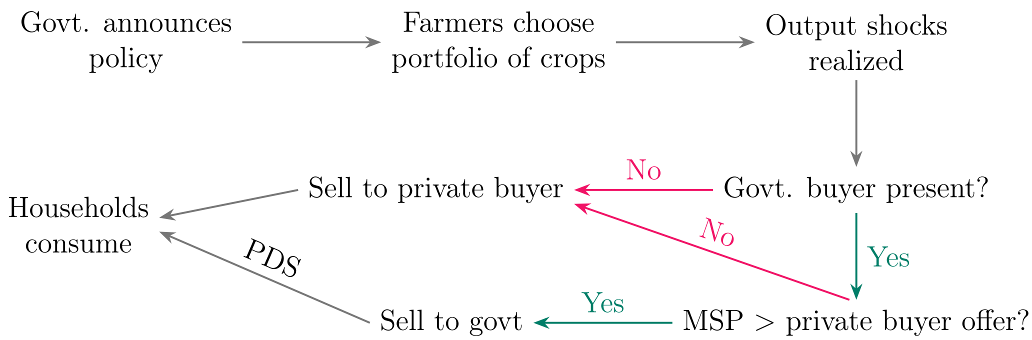

Notes. This figure provides an overview of the model. Before planting decisions are made, the government announces fertilizer subsidies and minimum support prices. Farmers take these into account and make planting decisions. Upon harvest, output shocks are realized. Farmers bring their output to the 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 5: *Calibrated Demand Parameters*

(a) *Parameters by Income Decile*

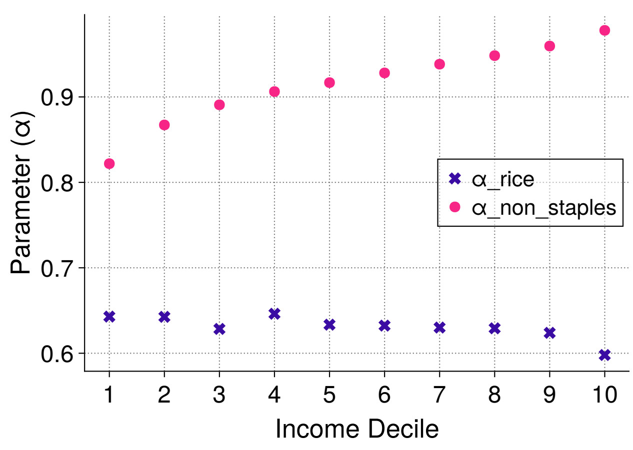

(b) *Other Crop Parameters*

	Value
Panel A: Other Crops	
Chickpea	0.005
Cotton	0.015
Groundnut	0.003
Maize	0.009
Mustard and Rapeseed	0.006
Pearl Millet	0.003
Pigeon Pea	0.003
Sorghum	0.003
Soybean	0.006
Panel B: Numeraire	
Numeraire	0.947

Notes. The figure and table display the calibrated expenditure share parameters (α) for the demand model. Panel (a) plots the expenditure shares for rice (α_{rice}) and for the non-staples nest ($\alpha_{non\text{-}staples}$) across household per-capita income deciles. These are calculated using household consumption data from India's National Sample Survey (NSS) 68th round. α_{rice} is the share of the staples budget (rice and wheat) allocated to rice, while $\alpha_{non\text{-}staples}$ is the share of total income allocated to goods other than rice and wheat. Panel (b) lists the expenditure shares for other non-staple crops (α_c) and a residual numeraire good, which are calibrated as shares of the total non-staple budget based on national production and price data.

Figure 6: *Supply-Side Estimates: Comparing Model-Predictions with Data*

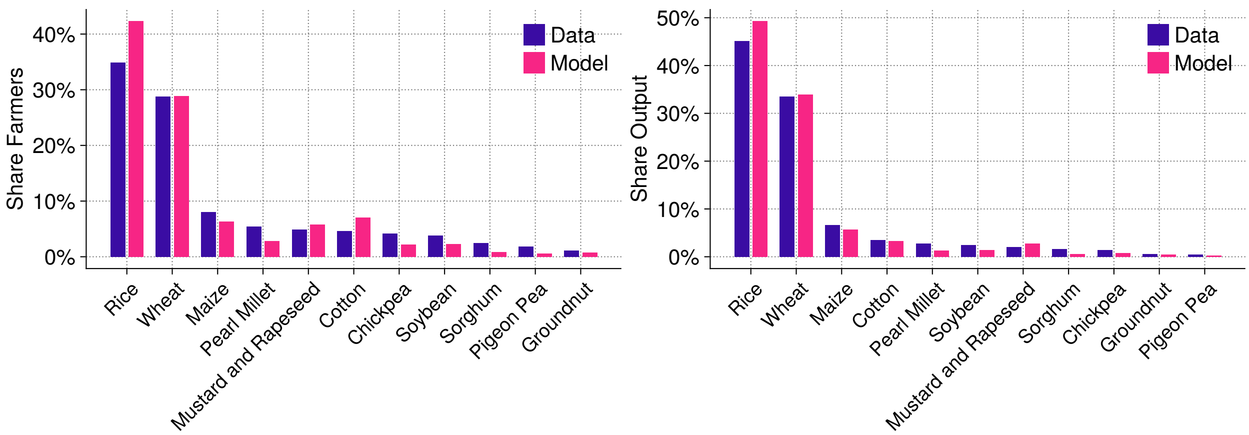

(a) *Number of Farmers Growing a Crop* **(b)** *Output*

Notes. This figure compares model-predicted aggregate supply outcomes with the data. The left panel compares the share of farmers growing each crop, while the right panel compares each crop's share of total output of all crops.

Figure 7: *Demand-Side Model Fit*

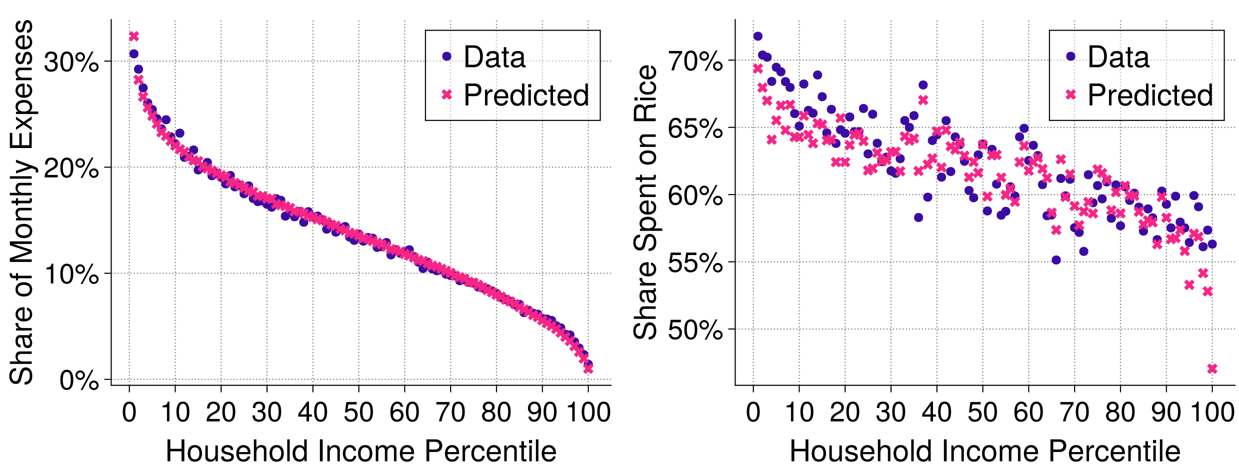

(a) *Expenditures on Rice and Wheat* **(b)** *Expenditures on Rice Relative to Wheat*

Notes. The left panel shows the share of total monthly expenditures allocated to rice and wheat as observed in the data and as predicted by the model. The right panel reports the same for the share of total expenditures on rice and wheat that is spent on rice.

Figure 8: *Impact on Prices and Output of Rice and Wheat*

(a) *Price* **(b)** *Quantity*

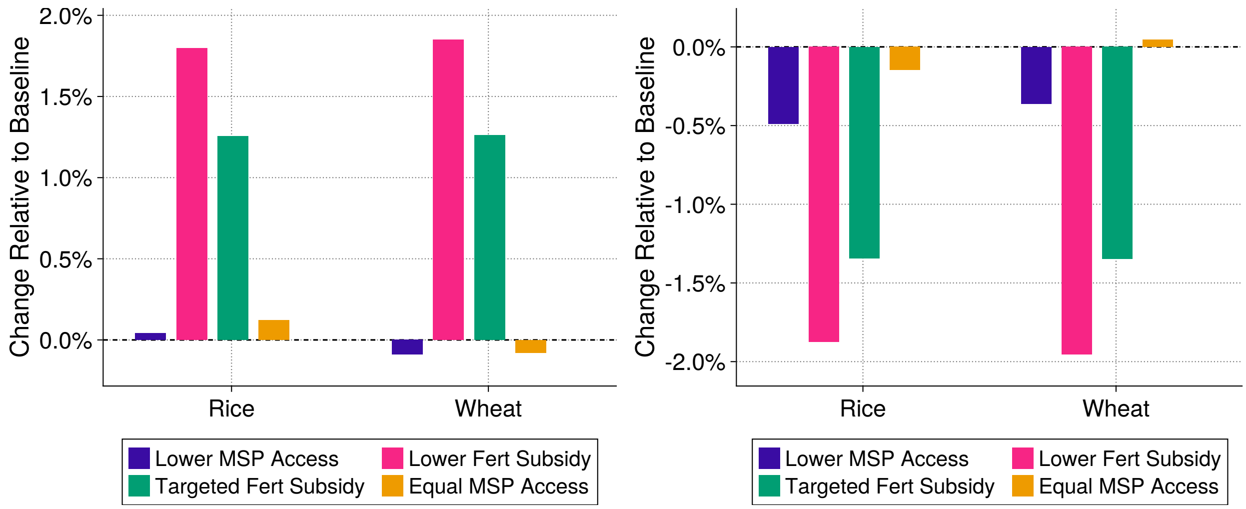

Notes. This figure shows the percentage change in equilibrium market-level outcomes for rice and wheat under the four counterfactual policies, relative to the simulated baseline. Panel A reports the percentage change in national-level market prices paid by consumers. Panel B reports the percentage change in total output, which is the sum of sales in the private market and procurement by the government.

Figure 9: *Impact on Quantity by Market*

(a) *Change in Government Procurement* **(b)** *Change in Private Market Sales*

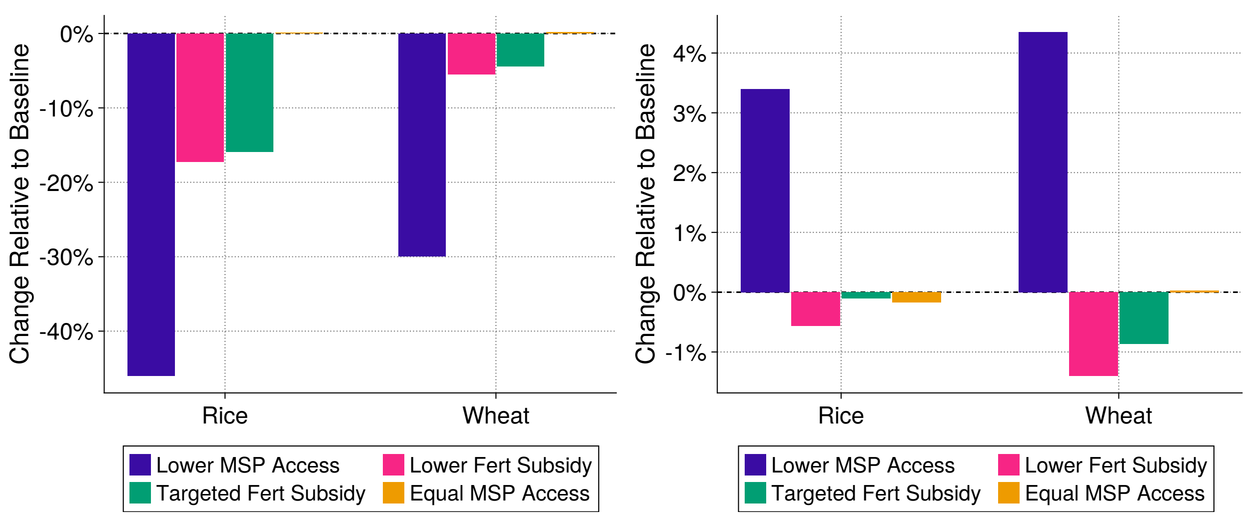

Notes. This figure shows the percentage change in the quantity of rice and wheat sold in different markets under the four counterfactual policies, relative to the simulated baseline. Panel A reports the percentage change in the quantity procured by the government. Panel B reports the percentage change in the quantity sold in the private market.

Figure 10: *Impact on Farmers*

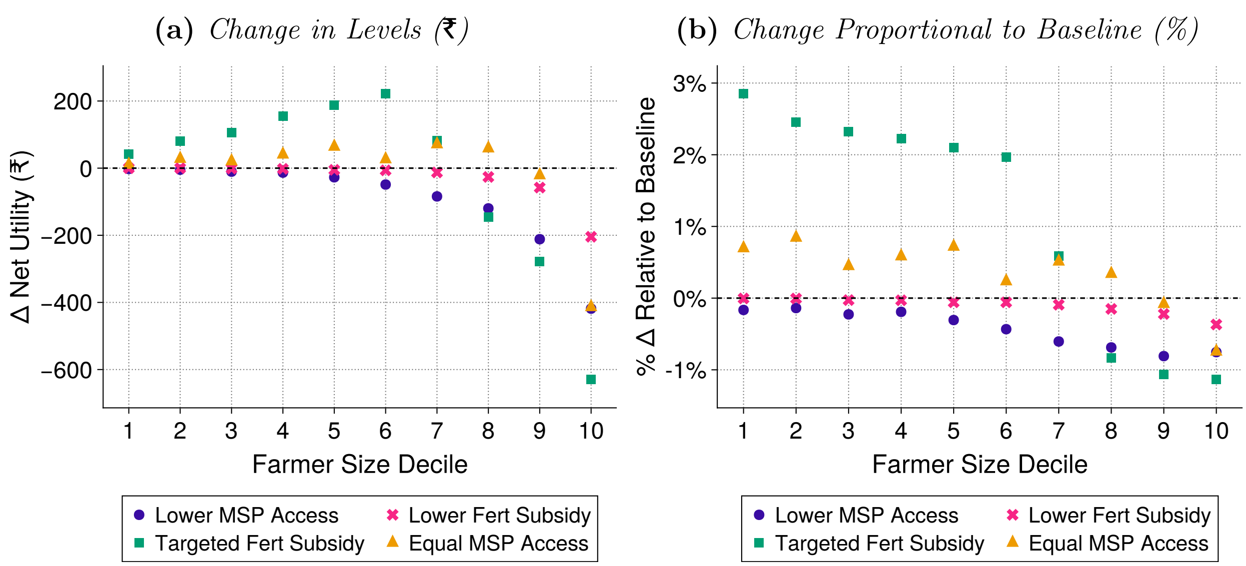

Notes. This figure shows the welfare impacts on farmers across the farm size distribution. The "net utility" corresponds to farmers' expected utility from production given in (3) in section 3. Panel A plots the average change in net utility, measured in rupees, for each farmer size decile. Panel B plots the same change as a percentage of the average baseline utility for each decile. Farmers are sorted into deciles based on their total farm area, using sampling weights from the NSS 77th round.

Figure 11: *Impact on Consumers*

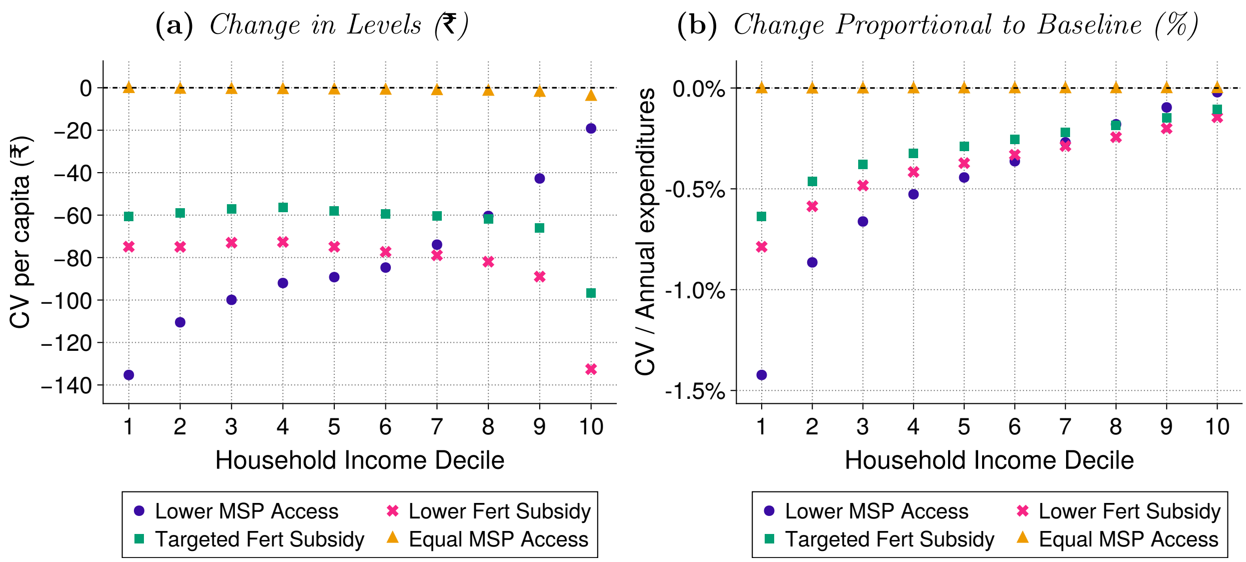

Notes. This figure shows the welfare impacts on consumers across the income distribution. We plot negative compensating variation ($-CV$), so values greater than zero represent welfare gains. Panel A plots the average per-capita CV in rupees for each income decile. Panel B plots the per-capita CV as a percentage of average baseline per-capita expenditure for each decile. Households are sorted into deciles based on their baseline monthly per-capita consumption expenditure (MPCE) from the NSS 68th round, using sampling weights.

Figure 12: *Comparing Net Effects Across Counterfactuals*

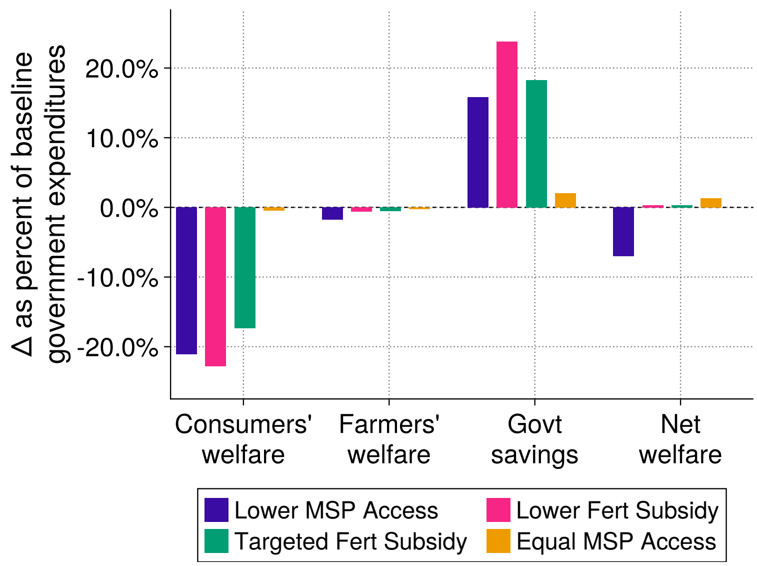

Notes. This figure presents the aggregate welfare impacts for each counterfactual policy, broken down by component. The components are the total change in consumer welfare (sum of CV), the total change in farmer welfare (sum of change in net utility), and government savings (the negative of the change in government expenditure). Net welfare is the sum of these three components. All values are normalized by total baseline government expenditure on fertilizer subsidies and MSP procurement.

Table 1: *Private Price Markdown Distribution Parameters*

State	Beta Distribution (α)	Beta Distribution (β)	Average Price Wedge
	(1)	(2)	(3)
Andhra Pradesh	46.42	26.02	0.36
	(2.34)	(1.18)	
Bihar	52.57	36.61	0.41
	(1.23)	(0.89)	
Chhattisgarh	19.60	7.98	0.29
	(1.80)	(0.91)	
Gujarat	16.64	8.99	0.35
	(1.36)	(0.89)	
Haryana	641.96	293.29	0.31
	(200.33)	(91.69)	
Himachal Pradesh	6.46	2.49	0.28
	(1.81)	(0.81)	
Karnataka	33.97	22.53	0.40
	(2.47)	(1.82)	
Madhya Pradesh	101.95	55.50	0.35
	(8.05)	(4.40)	
Maharashtra	8.01	3.25	0.29
	(1.06)	(0.46)	
Odisha	65.15	53.75	0.45
	(4.86)	(4.16)	
Punjab	793.94	386.91	0.33
	(73.67)	(35.87)	
Rajasthan	97.19	46.46	0.32
	(5.04)	(2.68)	
Tamil Nadu	11.47	6.69	0.37
	(1.78)	(1.24)	
Uttar Pradesh	37.37	23.08	0.38
	(1.16)	(0.78)	
West Bengal	72.81	58.85	0.45
	(5.57)	(4.10)	

Notes. This table presents the estimated parameters of the region-specific Beta distributions for private price markdowns, $(\alpha_r^\mu, \beta_r^\mu)$. These parameters are estimated using a simulated method of moments (SMM) procedure on price data for rice and wheat from the NSS 77th round, accounting for censoring due to government procurement. The final column reports the implied mean price wedge (1 - markdown), calculated as $1 - \alpha_r^\mu/(\alpha_r^\mu + \beta_r^\mu)$. Standard errors are reported in parentheses and computed using Bayesian bootstrap with 30 draws.

Table 2: *Production Function and Risk Aversion Parameters*

Variable	Estimate	Standard Error
Labor (β_l)	0.194	0.009
Capital (β_k)	0.088	0.005
Fertilizer (β_f)	0.098	0.005
Risk aversion (θ^γ)	1.483	0.273

Notes. This table reports the estimated parameters for the Cobb-Douglas yield function and the mean of the exponential distribution governing farmer risk aversion. The standard error is computed via Bayesian bootstrap with 30 draws.

Table 3: *Variance of Yield Shocks and Fixed Costs by Farmer Size*

Crop	Output shock	Small	Large
	(1)	(2)	(3)
Chickpea	0.229	460	181
	(0.009)	(156)	(224)
Cotton	0.314	3,523	2,138
	(0.007)	(794)	(515)
Groundnut	0.329	874	752
	(0.015)	(252)	(794)
Maize	0.249	722	1,697
	(0.007)	(184)	(775)
Mustard and Rapeseed	0.186	643	738
	(0.005)	(154)	(437)
Pearl Millet	0.235	703	1,474
	(0.006)	(201)	(926)
Pigeon Pea	0.332	716	255
	(0.010)	(303)	(378)
Rice	0.161	1,135	4,751
	(0.002)	(128)	(946)
Sorghum	0.318	705	582
	(0.023)	(188)	(557)
Soybean	0.358	802	673
	(0.008)	(359)	(539)
Wheat	0.147	2	0
	(0.003)	(35)	(197)

Notes. This table reports the estimated crop-specific parameters for the variance of idiosyncratic yield shocks ($\sigma^2_{\varepsilon c}$) in column (1) and the fixed costs of cultivation ($\kappa_{g,c}$) in columns (2) and (3). The variance is estimated from the residuals of the yield function regression. Fixed costs are estimated as part of the full structural model and are allowed to differ for small and large farmers, defined as those below or above the (unweighted) median farm size. Fixed costs are reported in thousands of rupees. Standard errors are reported in parentheses and computed using Bayesian bootstrap with 30 draws.

Table 4: *Parameters Governing Access to Government Procurement in States with Non-Negligible MSP Procurement*

State	Crop	MSP Parameter (α_{0rc})	MSP Parameter (α_{1rc})
		(1)	(2)
Andhra Pradesh	Rice	-0.62	0.14
		(0.14)	(0.17)
Chhattisgarh	Rice	0.20	0.96
		(0.22)	(0.38)
Haryana	Rice	-1.45	0.56
		(0.25)	(0.32)
Haryana	Wheat	-0.13	1.40
		(0.26)	(0.46)
Madhya Pradesh	Wheat	-1.27	0.77
		(0.15)	(0.25)
Odisha	Rice	-1.11	0.72
		(0.17)	(0.41)
Punjab	Rice	-0.85	1.14
		(0.98)	(1.32)
Punjab	Wheat	-0.27	0.03
		(0.22)	(0.31)
West Bengal	Rice	-1.00	0.69
		(0.26)	(0.24)

Notes. This table reports the estimated parameters governing the probability of a farmer accessing a government buyer, as specified in Equation (6). The parameters $\{\alpha_{0rc}, \alpha_{1rc}\}_{r,c}$ are only estimated for state-crop combinations with non-negligible MSP procurement levels. Standard errors are reported in parentheses and computed using Bayesian bootstrap with 30 draws.

Table 5: *Productivity Distribution Parameters: Crop-Specific Components*

Crop	Productivity ($\mu_{\omega c}$)	Std ($\sigma_{\omega c}$)
	(1)	(2)
Chickpea	0.905	0.310
	(0.058)	(0.032)
Cotton	0.946	0.389
	(0.082)	(0.032)
Groundnut	0.653	0.510
	(0.094)	(0.065)
Maize	2.224	0.359
	(0.120)	(0.039)
Mustard and Rapeseed	1.081	0.400
	(0.059)	(0.036)
Pearl Millet	2.510	0.099
	(0.183)	(0.054)
Pigeon Pea	0.581	0.452
	(0.053)	(0.053)
Rice	3.302	0.234
	(0.211)	(0.015)
Sorghum	1.524	0.315
	(0.106)	(0.042)
Soybean	1.629	0.167
	(0.132)	(0.056)
Wheat	3.355	0.190
	(0.200)	(0.013)

Notes. This table reports the estimated crop-specific components of the farmer-crop productivity distribution, where $\omega_{jc} \sim N(\mu_{\omega r} + \mu_{\omega c} - \sigma^2_{\omega c}/2, \sigma^2_{\omega c})$. Standard errors are reported in parentheses and computed using Bayesian bootstrap with 30 draws.

Table 6: *Productivity Distribution Parameters: Region-Specific Components*

State	Productivity ($\mu_{\omega r}$)	State	Productivity ($\mu_{\omega r}$)
	(1)	(2)	(3)
Andhra Pradesh	2.199 (0.148)	Maharashtra	1.189 (0.081)
Bihar	1.964 (0.131)	Odisha	1.777 (0.113)
Chhattisgarh	1.372 (0.084)	Punjab	2.461 (0.153)
Gujarat	1.382 (0.093)	Rajasthan	1.579 (0.101)
Haryana	2.121 (0.131)	Tamil Nadu	2.215 (0.170)
Himachal Pradesh	1.125 (0.079)	Uttar Pradesh	1.853 (0.109)
Karnataka	1.446 (0.110)	West Bengal	2.436 (0.154)
Madhya Pradesh	1.364 (0.085)		

Notes. This table reports the estimated region-specific component of the mean of the farmer-crop productivity distribution, $\mu_{\omega r}$. Standard errors are reported in parentheses and computed using Bayesian bootstrap with 30 draws.

Appendix

A Additional Figures

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

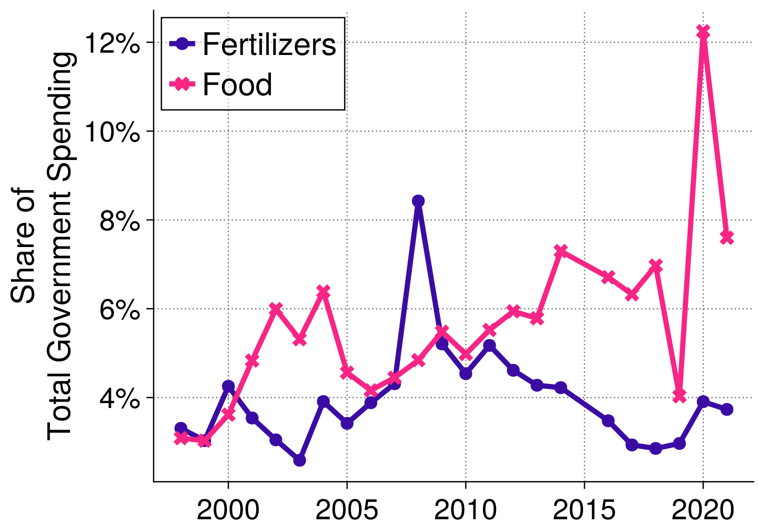

Notes. The figure plots combined central government expenditure on food and fertilizer subsidies as a share of total annual expenditures of the central government, from 1998-99 to 2021-22. Source: (Revised) budget estimates of the Government of India.

Figure A.2: *PDS Consumption by Household Income Percentile*

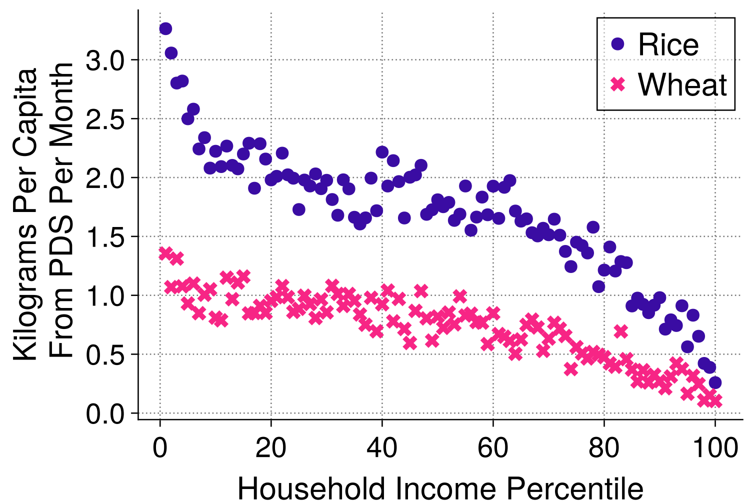

Notes. The figure shows binned means of the total monthly quantity (in kilograms) of rice and wheat received from the PDS against household income percentile. Income percentiles are computed using per capita monthly expenditure. Source: 68th round of the NSS Consumer Expenditure Survey, 2011-2012.

Figure A.3: *Sales to Government Buyers by Farmer Size (with state fixed effects)*

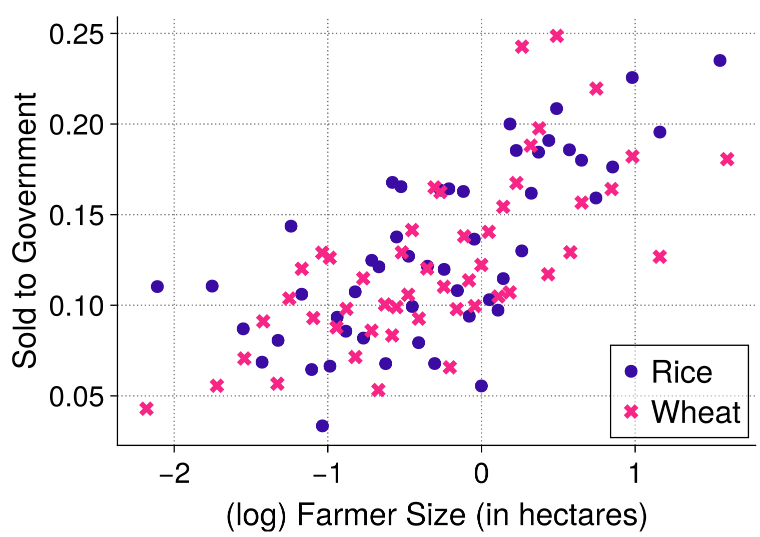

Notes. The figure shows binned means of an indicator variable for whether sales were made to a government buyer against (log) total farm size (in hectares), residualized by state fixed effects. The sample is conditional on the household selling a non-zero amount of output. Source: 77th round of the NSS, 2019.

Figure A.4: *Share of Output Sold to Government*

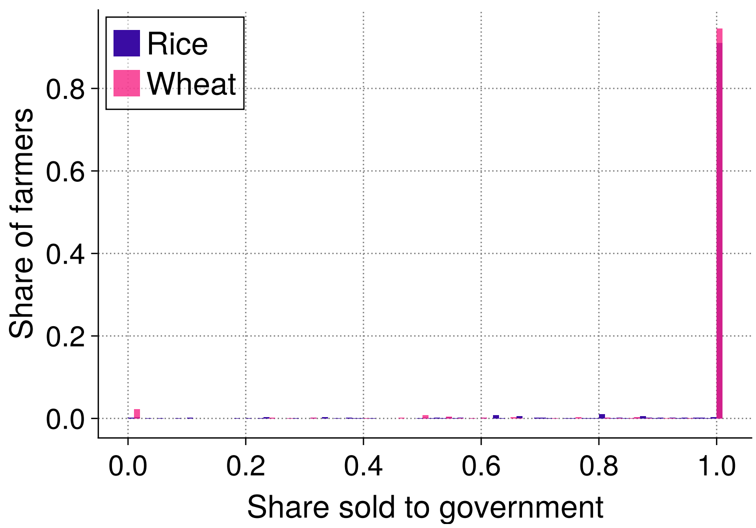

Notes. The figure plots the distribution (histogram) of the share of total output sold to government buyers. The sample is conditional on the household reporting any sale to a government buyer. Source: 77th round of the NSS, 2019.

B Reduced-form evidence on the impact of fertilizer prices

This appendix provides additional details on the difference-in-differences strategy used to study the impact of 2010 Nutrient-Based Subsidy (NBS) reform, as discussed in Section 2.2.

To measure each district's exposure to the policy change, we use a continuous treatment intensity variable. This measure is constructed for each district d based on its pre-reform fertilizer usage patterns. Specifically, we define treatment intensity as the average, price-weighted usage of P and K nutrients per unit of planted area in the pre-reform period (2004-2009).

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

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 2004-2009.

First, we estimate a simple event study to examine how fertilizer usage patterns changed following the reform:

$$\log F_{ndt} = \alpha_0 + \sum_{k \neq 2009} \alpha_k \cdot \mathbb{1}\{k = t\} + \phi_d + \epsilon_{dt},$$

where F_{ndt} is the quantity of nutrient $n \in \{N, P, K\}$ used in district d in year t, and ϕ_d are district fixed effects. The coefficients α_k capture the average change in fertilizer use in year k relative to 2009. The estimates are shown in Figure 3b.

Second, to examine whether districts more reliant on P and K fertilizers experienced differential changes in agricultural output, we use the treatment intensity measure defined above 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 Y_{dt} is the outcome of interest, and ϕ_d and γ_t are district and year fixed effects. We consider two outcomes of interest: (1) district-level output of rice, and (2) district-level output index. To construct this index, we aggregate the production of all major crops within a district using national median prices from the pre-reform period (2004-2009). The formula for the output index (OI_{dt}) in district d in year t is,

$$OI_{dt} = \sum_{c} p_c \cdot y_{cdt}$$

where p_c is the national median price of crop c in the period 2004-2009, and y_{cdt} is the quantity of crop c grown in district d in year t. The estimates are shown in Figures 3c and 3d.

Table B.1: *Summary Statistics of Treatment Intensity*

Statistic	Value
Number of Districts	310
Standard Deviation (Log)	1.103
Interquartile Range (Log)	1.033
25th Percentile (Rs/ha)	414
Median (Rs/ha)	761
75th Percentile (Rs/ha)	1,164

Notes. Treatment intensity is defined as the average price-weighted consumption of Phosphate (P) and Potash (K) fertilizers per unit of gross cropped area (in hectares) during the pre-reform period (2004-2009). Nutrient prices are fixed at their national median levels during the same period.

Table B.1 presents summary statistics for the treatment intensity variable. The median treatment intensity is 761 Rs/ha, with an interquartile range of roughly 750 Rs/ha (moving from 414 to 1,164 Rs/hectare). The distribution of log intensity has a standard deviation of 1.10 and an interquartile range of 1.03.

To interpret the magnitude of our estimates, we compute the implied percentage change in the outcome variable associated with a one standard deviation and an interquartile range increase in treatment intensity. Table B.2 displays these calculated effects for the output index.

Table B.2: *Interpretation of Event Study Coefficients (Output Index)*

Year	Coefficient	SE	1 SD Effect (%)	IQR Effect (%)
2010	-0.072	0.024	-7.672	-7.203
2011	-0.084	0.022	-8.818	-8.283
2012	-0.141	0.024	-14.367	-13.520
2013	-0.109	0.024	-11.358	-10.677
2014	-0.136	0.031	-13.900	-13.079

Notes. Columns (2) and (3) report the point estimate and standard error (β_k) from the event study regression where the outcome is the log of the output index. Column (4) computes the percentage change in the outcome associated with a one standard deviation increase in log treatment intensity: $(\exp(\beta_k \times SD) - 1) \times 100$. *Column (5) computes the effect of an interquartile range increase:* $(\exp(\beta_k \times IQR) - 1) \times 100$.

The results suggest economically significant effects. By 2012—two years after the policy change—districts with one standard deviation higher (log) treatment intensity experienced a 14.4% larger decline in aggregate agricultural output relative to 2009. Equivalently, districts at the 75th percentile of pre-reform fertilizer intensity experienced a 13.5% larger decline in output compared to those at the 25th percentile.

C Data

C.1 Supply-side data

We apply a series of filters to the CCS and the NSS 77th round data.

Cost of Cultivation Surveys (CCS)

The raw data are available at the plot-crop level, where individual farmers may operate multiple plots and cultivate the same crop across different plots. To align with our farmer-level framework, we aggregate this data by summing up inputs and outputs across all plots operated by each farmer for each crop.

Our empirical analysis focuses on the largest eleven crops. We restrict the sample to farmers who cultivate only these crops, and drop all observations for farmers who grow crops outside of these eleven. The estimation relies on within-farmer variation across time. Therefore, we impose a minimum threshold on the number of observations a farmer has, by dropping farmers with fewer than three observations.

To deal with outliers, we winsorize input intensities, input prices, and output prices at the 99% level. Finally, we retain only crop-state combinations with at least 20 observations in every year and states with at least 100 observations in every year. Summary statistics for the filtered CCS data are shown in Table C.1.

National Sample Survey (NSS) 77th Round

We drop observations with missing values for key variables, including crop name and state name. We also drop farmers with land sizes above the 99th percentile. We drop northeastern states (Assam, Meghalaya, Sikkim, and Nagaland), and union territories (Andaman and Nicobar Islands, Dadra and Nagar Haveli, and Daman and Diu, and Pondicherry).

Table C.1: *Summary Statistics: Cost of Cultivation Surveys (CCS)*

Variable	5th	25th	50th	75th	95th	N Farmers	N Observations
Total Area (hectares)	0.27	0.74	1.33	2.80	6.10	14,930	76,602
Number of Crops (s)	1.00	1.00	1.00	1.00	2.00	14,930	76,602
Labor / Land	125.50	301.30	453.41	706.33	1143.80	14,930	76,602
Machine / Land	4.40	8.12	11.91	16.67	29.17	14,930	76,602
Fertilizer / Land	40.25	91.67	142.94	197.82	289.32	14,930	76,602

Notes. This table presents summary statistics from the Cost of Cultivation Surveys (CCS) for the period 2011-2019. The statistics are calculated at the farmer-season level. The table reports the 5th, 25th, 50th, 75th, and 95th percentiles for total cultivated area, number of unique crops grown, and average input intensities (per hectare). The final two columns report the total number of unique farmers and the total number of farmer-crop-season observations in the sample.

We drop observations for crops that are rarely grown in a season or in a state. A crop is rare for a season if less than 20% of observations for that crop are in that season. A crop is rare for a state if less than 2% of farmers grow that crop in that state. Summary statistics for the filtered NSS 77th round data are shown in Table C.2.

We harmonize the states between the CCS and NSS 77th round data by keeping only those states that are present in both datasets after all the above filters.

C.2 Demand-side data

We use the 68th round of the NSS Consumer Expenditure Survey (2011-2012) to calibrate household demand for rice, wheat, and other crops. The data include household size, monthly per-capita expenditure, and quantities of rice and wheat purchased through the Public Distribution System (PDS) and private markets. Summary statistics for these variables are shown in Table C.3. Total quantities are the sum of PDS and private sources and are reported in

Table C.2: *Summary Statistics: NSS 77th Round Agricultural Households*

Variable	5th	25th	50th	75th	95th	N Farmers	N Observations
Total Area (hectares)	0.12	0.30	0.61	1.21	2.43	28,243	49,879
Number of Crops (s)	1.00	1.00	1.00	1.00	2.00	28,243	49,879
Sold to Govt	0.00	0.00	0.00	0.00	1.00	28,243	49,879

Notes. This table presents summary statistics from the 77th round of the National Sample Survey (NSS) of Agricultural Households (2019). The statistics are calculated at the household-visit level. The table reports the 5th, 25th, 50th, 75th, and 95th percentiles for total cultivated area and the number of unique crops grown. "Sold to Govt" is an indicator variable equal to one if the household sold any of its output to a government agency during the visit. The final two columns report the total number of unique farming households and the total number of household-crop observations in the sample.

monthly kilograms.

Table C.3: *Summary Statistics: NSS 68th Round Household Data*

Variable	5th	25th	50th	75th	95th	N Households
Household Size	1	3	4	6	8	99,637
PDS Rice (kg)	0.00	0.00	0.00	12.00	30.00	99,637
PDS Wheat (kg)	0.00	0.00	0.00	2.00	18.00	99,637
Total Rice (kg)	0.00	4.00	15.00	30.00	55.00	99,637
Total Wheat (kg)	0.00	1.00	6.00	20.00	50.00	99,637

Notes. This table presents weighted summary statistics from the NSS 68th round household consumption data (2011-2012). The statistics are calculated at the household level and report the 5th, 25th, 50th, 75th, and 95th percentiles for household size and monthly household quantities of rice and wheat. "PDS" quantities are purchases through the Public Distribution System; "Total" quantities sum PDS and private sources. "N Households" reports the unweighted number of households in the sample.

D Additional Details on Estimation of the Structural Model

D.1 Estimation: Markdown Distribution

The estimation of region-specific markdown distributions, characterized by the parameters $(\alpha_r^\mu, \beta_r^\mu)$, uses data from NSS 68th round for national consumer prices and the NSS 77th round for farmer-level prices.

National consumer prices for rice and wheat are constructed by taking the median of all household-reported consumer prices in the NSS 68th round. For farmer-level prices, we use the NSS 77th round. To mitigate the influence of outliers and potential measurement error, we trim the price data by dropping observations above the 95th percentile and below the 5th percentile within each state-crop combination. This results in a sample with observations ranging from a minimum of 237 to a maximum of $\approx 7,000$ in a given state-crop combination.

For regions without MSP procurement, the full distribution of private prices is observed. The parameters of the markdown distribution $(\alpha_r^\mu, \beta_r^\mu)$ can be expressed as a function of the moments of the observed markdown distribution.

$$\alpha_r^\mu = \bar{\mu}_r \left(\frac{\bar{\mu}_r (1 - \bar{\mu}_r)}{Var(\mu_r)} - 1 \right),$$
$$\beta_r^\mu = \alpha_r^\mu \frac{(1 - \bar{\mu}_r)}{\bar{\mu}_r},$$

where $\bar{\mu}_r$ is the mean of the markdown distribution and $Var(\mu_r)$ is the variance of the markdown distribution.

For regions with significant procurement, we implement the simulated method of moments (SMM) procedure described in the main text. In the simulation step, we use 10,000 draws to generate the private price distribution. When dropping observations that fall at the MSP to construct moments from unambiguously private sales, we define a narrow bandwidth of 1% of the MSP.

Figure D.1: *Simulated Markdowns with MSP Censoring*

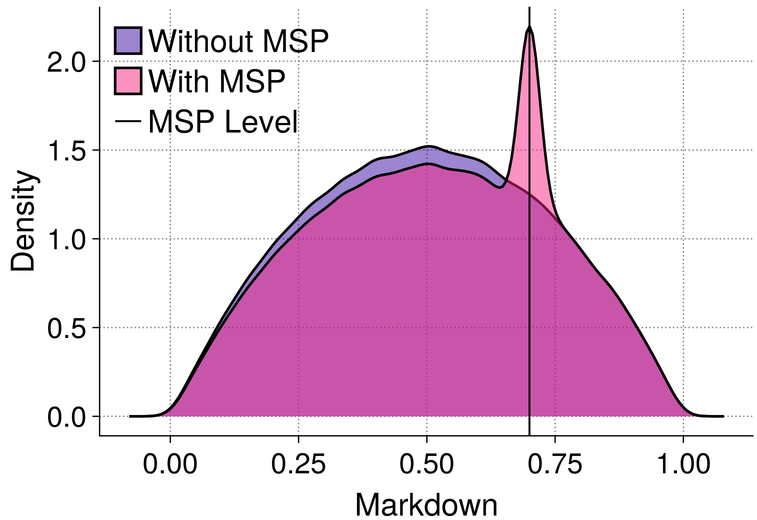

Notes. This figure illustrates how government procurement at the Minimum Support Price (MSP) affects the observed distribution of private price markdowns. The "Without MSP" density curve represents the true, latent distribution of markdowns, simulated here from a Beta distribution. The "With MSP" density curve represents the observed distribution that results after accounting for procurement. It is generated by replacing a fraction of the offers that fall below the MSP level (indicated by the vertical line) with the MSP itself. This censoring of the lower tail creates a point mass at the MSP.

D.2 Estimation: Yield Function

The yield function parameters are estimated using data from the CCS survey. After applying the filters detailed above, we are left with a panel dataset of $\approx 76,600$ farmer-year-crop observations.

As described in Section 4, we estimate the yield elasticities $(\beta_l, \beta_k, \beta_f)$ via a fixed-effects regression. The specification identifies the parameters from within-farmer-crop variation over time. Consequently, farmer-crop units with only a single observation do not contribute to identification and are dropped, reducing the effective estimation sample to $\approx 71,000$ observations.

The residuals from this regression, $\hat{\varepsilon}_{jc}$, are the estimates of realized idiosyncratic yield shocks. We use these residuals to estimate the parameters of the crop-specific shock distributions via maximum likelihood. The number of observations available for this estimation varies by crop, from the maximum of $\approx 27,000$ for rice to the minimum of 154 for sorghum.

D.3 Estimation: Remaining Supply-Side Parameters

The remaining supply-side parameters are estimated by matching moments from the NSS 77th round and the Cost of Cultivation Surveys (CCS). While most moments are derived from the NSS, the average revenue share spent on fertilizer is calculated using the 2017-2019 CCS wave. This wave was selected because its timing aligns most closely with the NSS 77th round. This CCS subsample has $\approx 25,000$ observations.

MSP parameter moments: For each state, farmer size group (small and large), and MSP crop (rice and wheat), we compute the unconditional probability of selling the crop to the government. This is calculated as the number of farmers selling the crop to the government divided by the total number of farmers in that state-group. The number of observations underlying these moments varies from a minimum of 78 for large farmers in Himachal Pradesh to a maximum of $\approx 6,500$ for small farmers in Uttar Pradesh.

Three sets of moments are used to identify the farmer-crop productivity distribution.

1. **Average yield by crop:** For each of the eleven crops, we calculate the average yield across all farmers who cultivate it. The number of observations per crop ranges from 987 for groundnut to $\approx 17,700$ for rice.
2. **Average yield by region:** For each region, we calculate the average yield by averaging across all crops and farmers within that region. The number of observations per region ranges from 681 for Himachal Pradesh to $\approx 11,300$ for Uttar Pradesh.
3. **Standard deviation of yield by crop:** This moment is constructed in two steps. First, we compute the standard deviation of yield for each crop within each state. Second, we take the mean of these standard deviations across states for each crop. The underlying sample sizes for the state-crop combinations vary significantly, from a minimum of 2 for Pigeon Pea in Chhattisgarh to a maximum of $\approx 4,700$ for Wheat in

Uttar Pradesh.

D.4 Identification of Risk Aversion

To validate our identification strategy, we perform a sensitivity analysis where we perturb the estimated risk aversion parameter (θ^{γ}) by a scaling factor ranging from 0.5 to 1.5, holding all other parameters constant. For each scaled value, we solve the full structural model and compute the resulting average fertilizer cost share.

Figure D.2 plots the results of this exercise. The figure demonstrates a clear monotonic relationship: as risk aversion increases, farmers reduce their exposure to risk by lowering the intensity of variable inputs like fertilizer. This strong sensitivity confirms that the fertilizer cost share is a highly informative moment for identifying the degree of farmer risk aversion.

Figure D.2: *Impact of Risk Aversion on Input Usage*

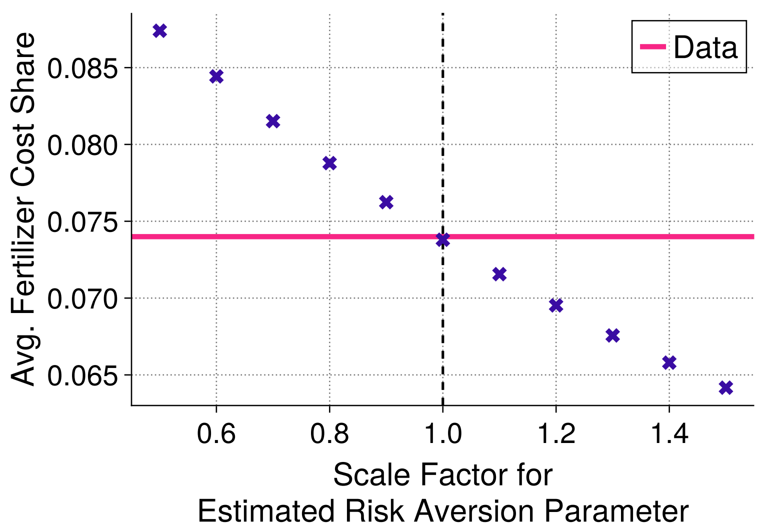

Notes. This figure plots the model-implied average fertilizer cost share (y-axis) against a scaling factor applied to the estimated risk aversion parameter (x-axis). The horizontal line represents the observed moment in the data. The steep slope indicates that input choices are highly sensitive to the level of risk aversion and provides support for the identification strategy.

D.5 Constructing National Consumer Prices

National consumer prices are not readily available for non-staple crops like cotton and soybeans, and we recover them using our estimated markdown distributions and the observed farmer prices from the NSS 77th round. This is performed in two steps.

First, for each crop c, we construct a national average markdown, $\bar{\mu}_c$. This is calculated as a weighted average of the estimated region-specific mean markdowns, $\bar{\mu}_r$, where the weight for each region is its share of the total number of farmers in the country cultivating that crop. This weighting scheme ensures that the national average markdown reflects the geographic distribution of production. Second, the national average markdown is combined with the observed national farmer price, $\bar{P}_{fc}$, to recover the implied national consumer price, $P_c = \frac{\bar{P}_{fc}}{\bar{\mu_c}}$.

E Credit Constraints

In our model, we abstract from credit constraints as a determinant of input choices, assuming farmers can finance their desired level of non-area inputs. This assumption is motivated by the fact that government policy actively promotes credit access for farmers. Agriculture is designated as a priority sector by the central bank, mandating that banks and financial institutions allocate at least 18% of their lending to the sector. Further, all farmers are eligible for Kisan ("Farmer") Credit Cards that can be used to purchase inputs at subsidized interest rates.

Data from the NSS 77th round corroborates this view. The data includes information on loan sources, purpose, amount, and interest rates. According to the data, institutional lenders, including banks and cooperatives account for 82% of the value and the number of loans for farm related revenue expenditures. The remaining share is sourced from non-institutional lenders, such as moneylenders. Credit for agricultural purposes is also available at a lower interest rate than for consumption loans. In Figure E.1 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. Our focus on risk over credit frictions is also consistent with findings in other contexts; for instance, Karlan et al. (2014) find agricultural risk to be a more important determinant of production decisions than input credit constraints among farmers in northern Ghana.

Figure E.1: *Annual Interest Paid for Farm and Consumption Loans*

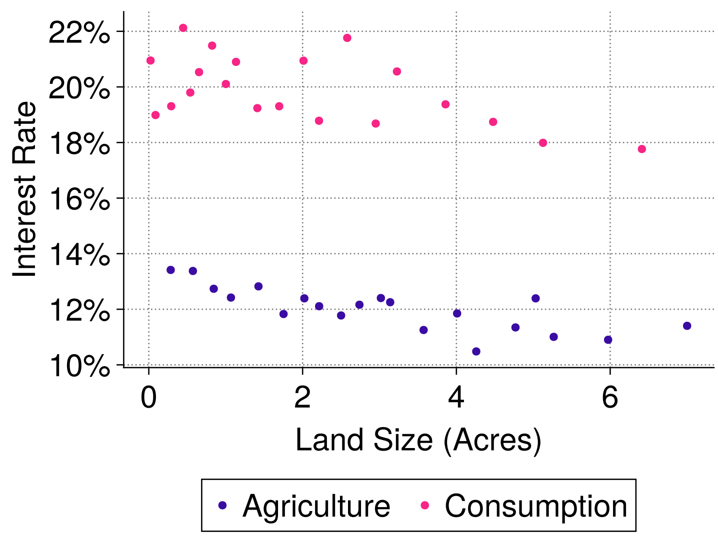

Notes. The figure plots the average interest rate paid by farmers for farm and consumption loans on total land holdings (in acres) of the farmer. The data are from the 77th round of the NSS (2019).

F Appendix for Counterfactuals

F.1 Computational Details

This appendix provides additional details on how we compute the equilibria in the baseline and counterfactual policy environments. An equilibrium in our model is defined by a vector of national market prices, $\boldsymbol{P}^*$, at which the markets for all agricultural commodities clear. We solve for this equilibrium iteratively using an algorithm inspired by tatonnement processes.

Before starting, we first compute a scalar multiplier, λ, to ensure that the aggregated supply lines up with aggregate demand at the observed national prices in the data. We take these observed national consumer prices as given and solve a one-dimensional optimization problem to find the value of λ that minimizes the sum of squared differences between aggregate private supply and aggregate private demand for the two main staple crops, rice and wheat. This calibrated multiplier is then held constant across all subsequent counterfactual simulations. This multiplier helps adjust for differences in sampling weights and population sizes between the farmer and household surveys, which were conducted in different years (2011-12 for households and 2019 for farmers). It also helps adjust for differences in post-harvest losses that occur during storage, transport, and processing, as well as net trade. These are all factors which are not explicitly included in our model.

Once a policy parameter is changed (e.g., the fertilizer price or the parameters governing MSP access), we solve for the new equilibrium price vector, $\boldsymbol{P}^*$. This is accomplished using an iterative price-adjustment algorithm which proceeds as follows:

1. **Initialization**: The algorithm begins with an initial guess for the national consumer price vector, $\boldsymbol{P}^{(0)}$, which we set to the observed baseline prices. A step-size parameter, δ, is initialized to a starting value (e.g., 0.1).
2. **Iterative Search**: The algorithm iterates for a fixed maximum number of iterations (T_{max}). In each iteration t:
   (a) **Solve Agent Problems**: Given the current price vector $\boldsymbol{P}^{(t-1)}$, we solve the

utility maximization problems for every farmer and household in our sample. This yields each farmer's optimal crop portfolio and input choices, and each household's optimal consumption bundle.

(b) **Aggregate Supply and Demand**: We aggregate these individual decisions using sampling weights (and the multiplier λ for supply) to compute the total quantity supplied and demanded for each crop. This gives us government procurement, Q_c^{govt}, and private market sales, $Q_c^{S,pvt}$, on the supply side. On the demand side, since government procurement determines PDS transfers ($Q_c^{PDS} = Q_c^{govt}$), we compute the residual private market demand, $Q_c^{D,pvt}$.

(c) **Calculate Excess Demand**: For each crop c, we calculate the excess demand in the private market as $ED_c^{(t-1)} = Q_c^{D,pvt} - Q_c^{S,pvt}$.

(d) **Update Prices**: The price of each crop is updated based on the sign of its excess demand. If demand exceeds supply, the price is adjusted upwards; if supply exceeds demand, it is adjusted downwards. The update rule is:

$$P_c^{(t)} = P_c^{(t-1)} \cdot (1 + \delta \cdot \text{sign}(ED_c^{(t-1)}))$$

This process is repeated for all crops, yielding a new price vector $\boldsymbol{P}^{(t)}$.

3. **Adaptive Step Size**: The algorithm features an adaptive step-size. We track whether the sign of excess demand for a crop has "flipped" relative to its sign in the first iteration. If the signs for all crops flip, it indicates that the prices are oscillating around an equilibrium. When this occurs, the step size δ is reduced by a factor of 10 (e.g., from 0.1 to 0.01) to facilitate finer adjustments and help with convergence.

4. **Termination**: The algorithm runs for a pre-specified number of iterations (e.g., 100). Throughout this process, it tracks the iteration that yields the minimum error, where error is defined as the maximum percentage deviation between private supply and private demand across all crops. The final counterfactual equilibrium price vector, $\boldsymbol{P}^*$, is the price vector from the iteration that achieved this minimum error.

F.2 Additional Counterfactual: Random Allocation of PDS Grains

To isolate the equilibrium effects of PDS targeting, we simulate a counterfactual where the total volume of PDS grain is distributed randomly across the population rather than being targeted based on income. In our baseline model, PDS entitlements are determined by household-specific shares, ϕ_{hc}, which are calibrated from the NSS 68th round of the survey. In this experiment, we replace these heterogeneous shares with a uniform allocation rule, effectively redistributing entitlements from lower-income households to higher-income households while keeping the total quantity of PDS grain fixed. We do this while holding the supply-side policy environment constant: farmers face the same MSP levels and access probabilities as in the baseline.

Figure F.1: *Impact of Random Allocation of PDS Grains on Welfare*

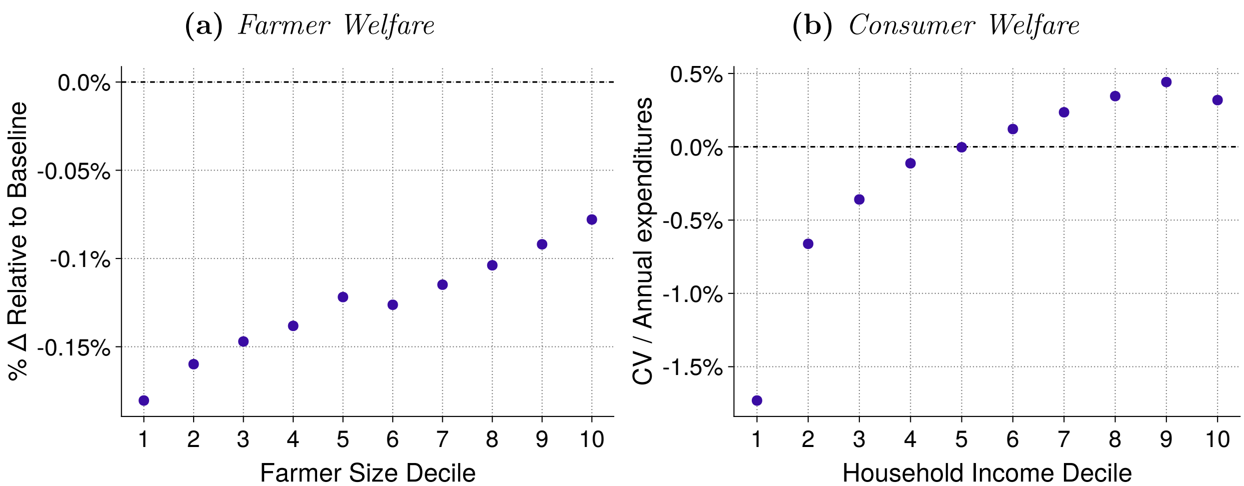

Notes. This figure displays the distributional welfare impacts of a counterfactual scenario where access to subsidized PDS grains is allocated randomly across households, rather than according to the baseline targeting criteria. Panel A plots the percentage change in farmer welfare (net utility) by farm size decile, which is driven by general equilibrium price effects. Panel B shows the change in consumer welfare, measured as compensating variation, expressed as a percentage of baseline annual household expenditures by household income decile.

This redistribution generates a negative demand shock in the private market because the marginal propensity to consume (MPC) staples varies inversely with income. Recall that in our demand specification, household utility is Cobb-Douglas and the parameter

$\alpha_{h,staple}$ governs the share of the budget allocated to staple crops. Empirically, low-income households have a much higher $\alpha_{h,staple}$ than high-income households. Since PDS transfers are inframarginal for almost all households, they function as fungible cash transfers.

The random allocation counterfactual effectively redistributes income from agents with a high MPC (the poor) to agents with a low MPC (the rich). Because higher-income households have lower marginal propensity to consume staples, their total consumption rises negligibly in response to the transfer. On the other hand, lower-income households respond to the loss of transfer value with a larger contraction in total consumption. The magnitude of the poor's contraction exceeds the rich's expansion, resulting in a net decline in aggregate demand that depresses private market prices and reduces the incomes of small farmers who rely on it.

Figure F.2: *Impact of Random Allocation of PDS Grains on Quantity of Rice and Wheat*

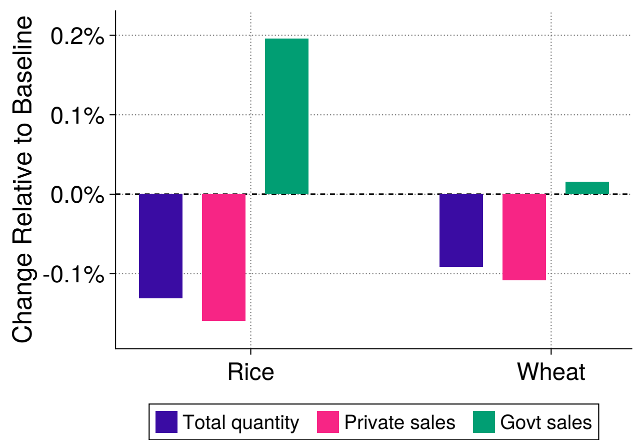

Notes. This figure shows the percentage change in aggregate quantities for rice and wheat under the random PDS allocation counterfactual relative to the baseline. For each crop, the bars depict the percentage change in total output, sales in private markets, and sales to government buyers. In this simulation, MSP levels and the probability of finding a government buyer are held fixed at baseline levels. Consequently, changes in government sales (procurement) are driven entirely by changes in equilibrium private prices: as the reallocation of PDS grains alters private demand and market prices, the fixed MSP becomes relatively more or less attractive to farmers, endogenously shifting the total quantity procured and subsequently distributed through the PDS.

F.3 Additional Counterfactual Figures

Figure F.3: *Impact on Prices and Output of Non-Staple Crops*

(a) *Price* **(b)** *Output*

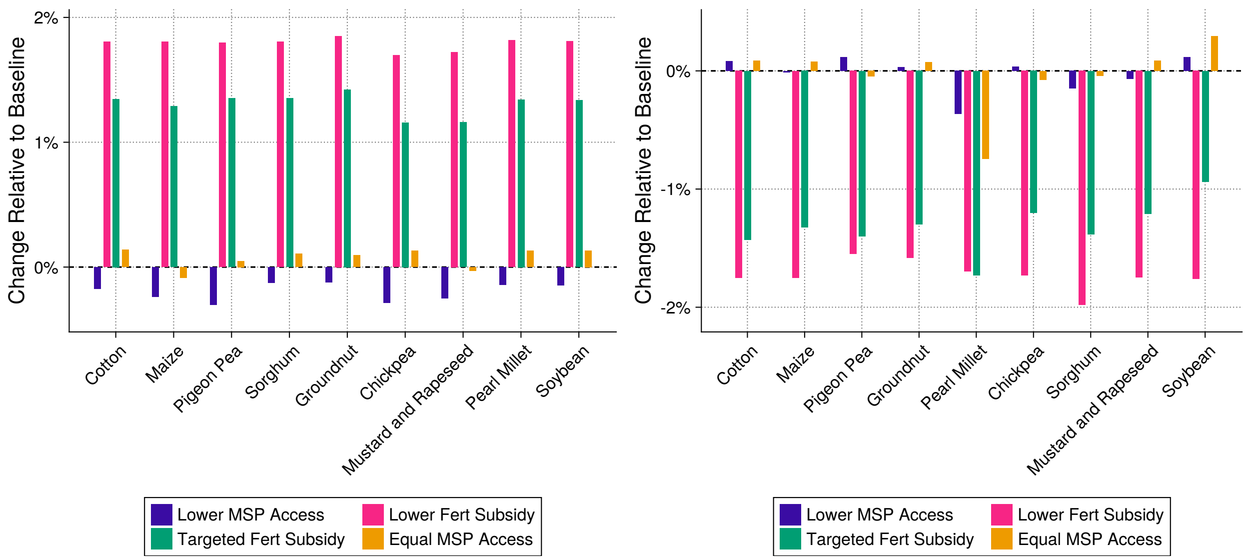

Notes. This figure shows the impact of counterfactual policy changes on prices (Panel A) and quantities (Panel B) of non-staple crops. The bars represent percent changes relative to the baseline for each crop under four policy scenarios: lower MSP access, lower fertilizer subsidy, targeted fertilizer subsidy, and equal MSP access across farmer sizes and locations in states with meaningful MSP procurement in the baseline.

Figure F.4: *Probability of Finding Government Buyer*

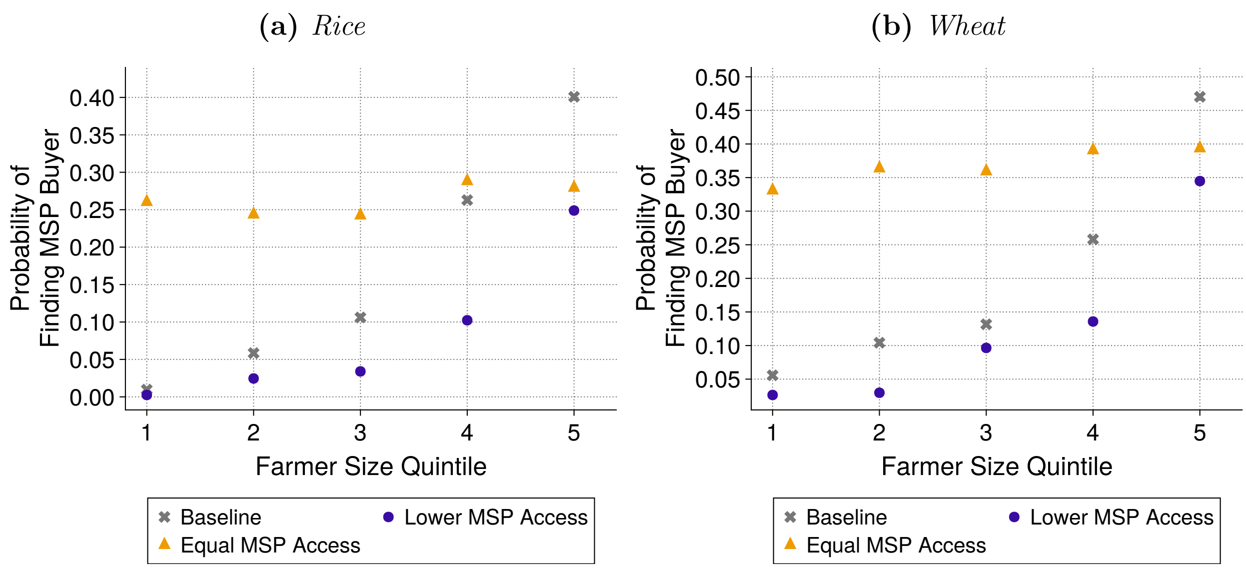

Notes. This figure plots the share of farmers with access to a government (MSP) buyer in our simulations by farm size quintile. The plot data are restricted to farmers in states with non-negligible government procurement at baseline. Panel A shows results for rice farmers, and Panel B shows results for wheat farmers. Each panel compares the simulated baseline scenario with the two MSP-related counterfactuals: lower MSP access and equal MSP access. Under the equal access scenario, all farmers in states with meaningful baseline procurement are assigned a common probability of finding a government buyer. Farmers are sorted into quintiles based on their total farm area using sampling weights from the 77th round of the NSS.

Figure F.5: *Impact on Price Risk by Farmer Size*

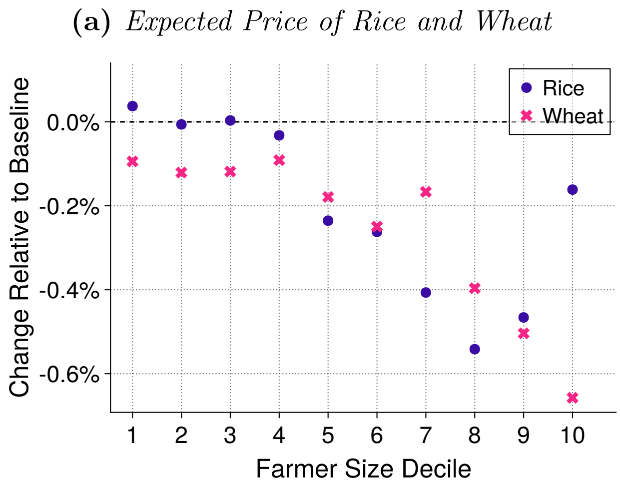

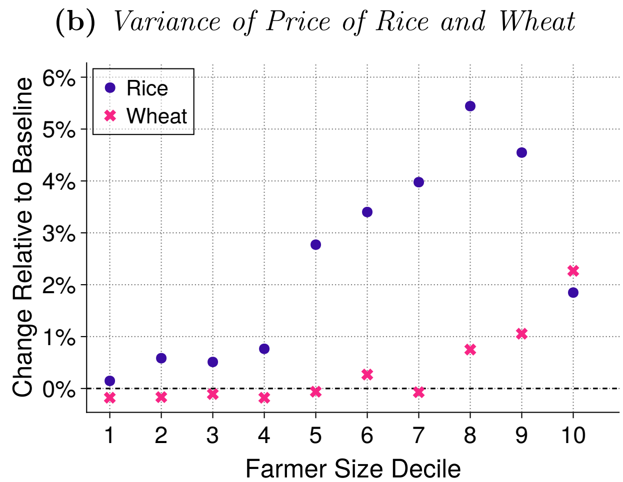

Notes. This figure shows the change in price risk faced by farmers under the "lower MSP access" counterfactual, by farm size decile. Price risk is characterized by the mean and variance of the ex-ante price distribution faced by farmers when making their planting decisions. Panel A plots the percentage change in the expected price. Panel B plots the percentage change in the price variance. Both panels show results for rice and wheat. Farmers are sorted into deciles based on their total farm area using sampling weights from the 77th round of the NSS.

Figure F.6: *Change in Share of Farmers Growing Rice and Wheat by Farmer Size Group*

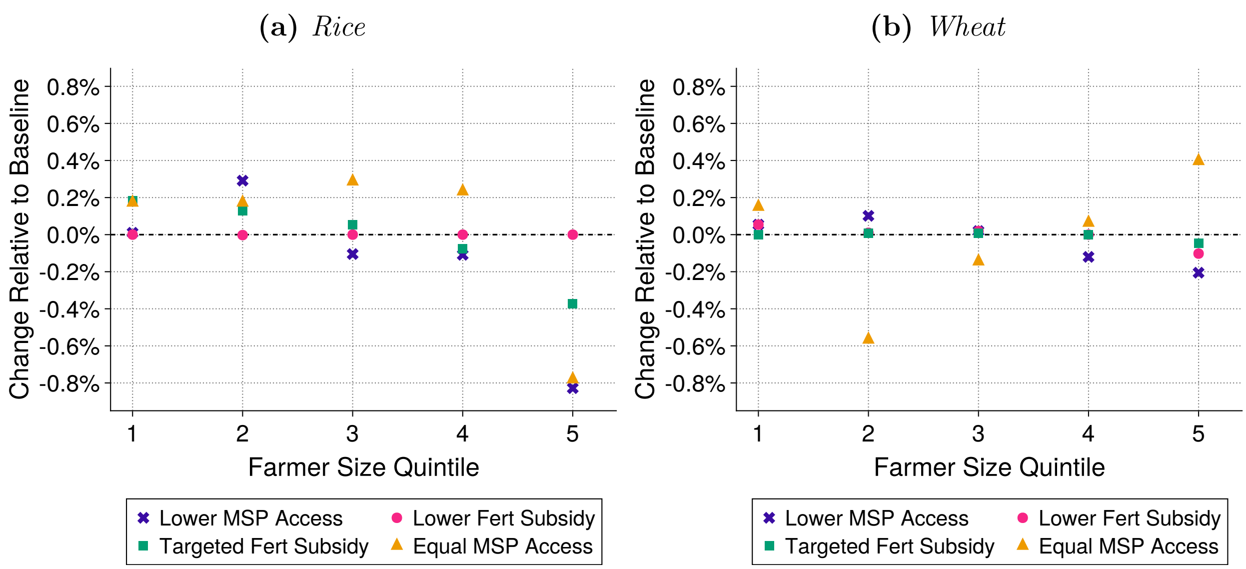

Notes. This figure shows the change in the share of farmers who include rice (Panel A) or wheat (Panel B) in their crop portfolio, by farm size quintile. The y-axis reports the percentage change in this share relative to the baseline for each of the four counterfactuals. Farmers are sorted into quintiles based on their total farm area using sampling weights from the 77th round of the NSS.

Figure F.7: *Change in Average Area Allocated to Rice and Wheat by Farmer Size Group*

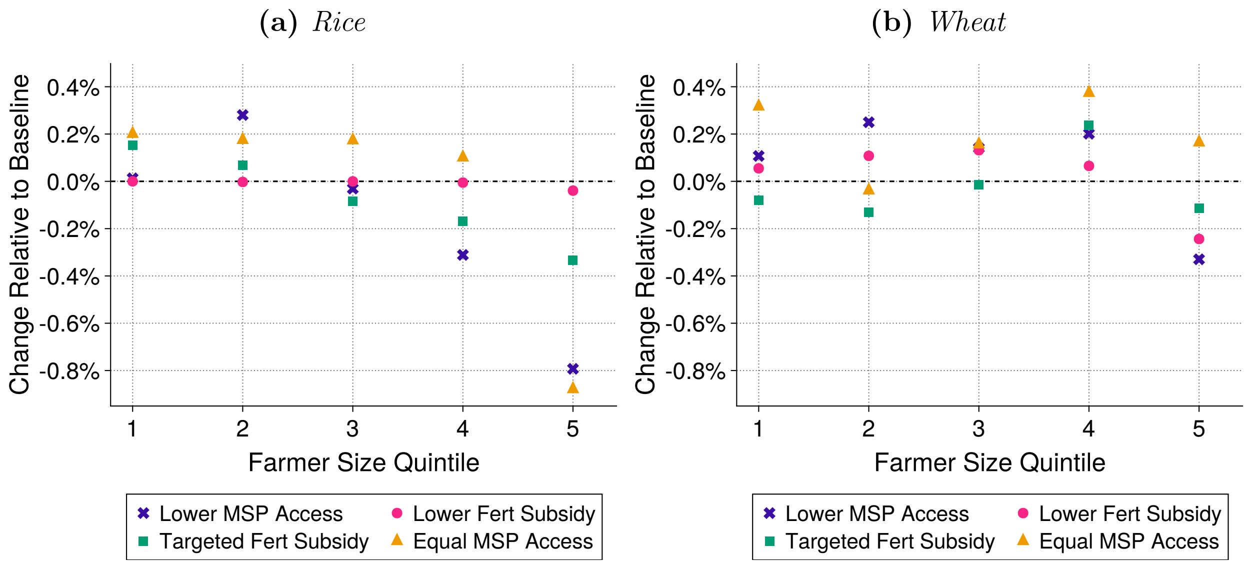

Notes. This figure shows the change in the average area allocated to rice (Panel A) or wheat (B) by farmers, broken down by farm size quintile. The y-axis reports the percentage change in average allocated area relative to the baseline for each of the four counterfactuals. Farmers are sorted into quintiles based on their total farm area using sampling weights from the 77th round of the NSS.

Figure F.8: *Impact on Input Use*

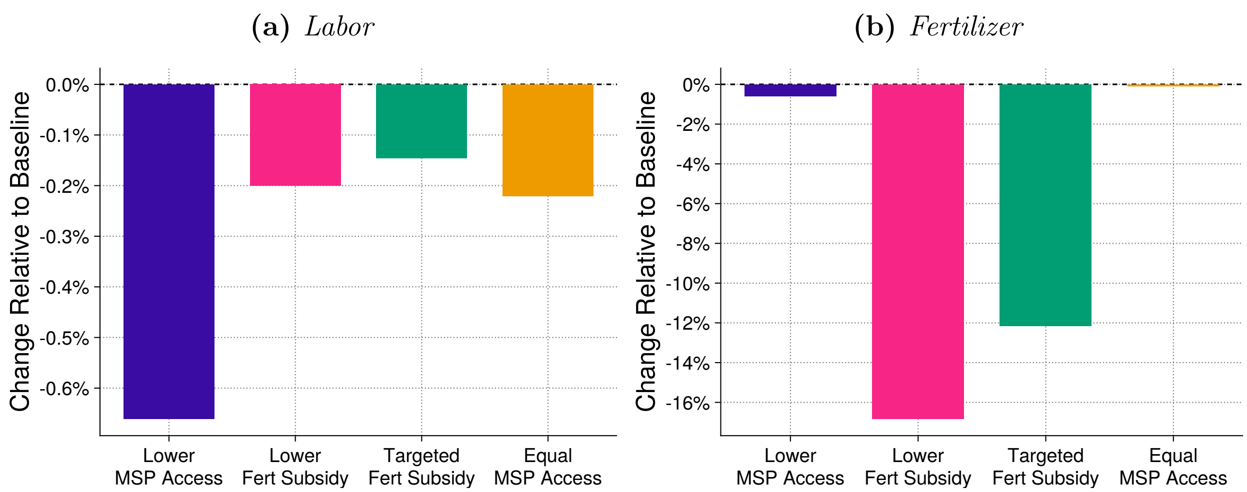

Notes. This figure shows the percentage change in the aggregate use of two of the three agricultural inputs for each counterfactual policy, relative to the baseline. Panel A shows the change in total labor demand (in hours). Panel B shows the change in total fertilizer use (in kg).

Figure F.9: *Change in Fertilizer Usage Intensity Across Farmer Size Groups*

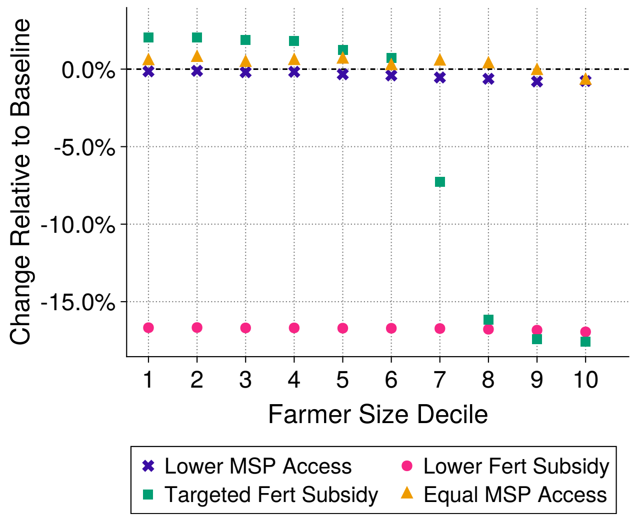

Notes. This figure shows the percentage change in average fertilizer application intensity (kg per hectare) for each farmer size decile, relative to the baseline. Farmers are sorted into deciles based on their total farm area using sampling weights from the 77th round of the NSS.

Figure F.10: *Average Farmer Size by Decile*

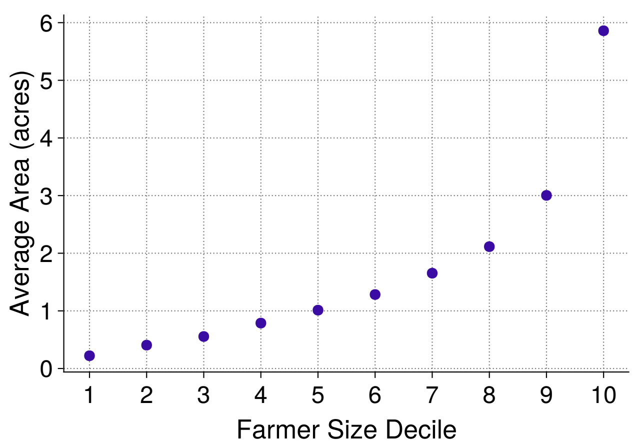

Notes. This figure plots the average farmer size in acres for farmers in each size decile. Farmers are sorted into deciles based on their total farm area using sampling weights from the 77th round of the NSS.