

Scaling Up Agricultural Insurance

Jeremy Majerovitz

Heitor S. Pellegrina

University of Notre Dame

University of Notre Dame*

June 2026

Abstract

We study the general equilibrium effects of scaling up rainfall insurance in low-income agricultural economies. Microeconomic experiments show that risk is an important constraint for farmers in developing countries and that rainfall insurance increases investment and shifts crop choice toward riskier, higher-return crops. We develop a model of farmer-households and insurance to quantify these effects. We derive a sufficient-statistics characterization of the optimal subsidy, which trades off three terms: a fiscal externality from subsidizing insurance below cost, a pecuniary externality that captures how crop-price changes reallocate consumption across states of the world, and a fiscal incidence term that captures how the subsidy transfers resources from those who do not buy insurance towards those who do. We calibrate the model to Ghana's agricultural economy, combining experimental moments from the literature with additional moments from household survey data. Subsidizing insurance leads to increased insurance take-up, fertilizer use, and agricultural output, and general equilibrium forces significantly amplify these effects. However, we find that the optimal insurance subsidy is close to zero.

Keywords: Agricultural Insurance, Investment Risk, Risk Sharing

*Acknowledgments: We are grateful to Joe Kaboski, and to seminar participants at the Y-RISE Annual Conference and the Carnegie-Rochester-NYU Conference, for thoughtful comments on the paper. We thank Florian Trouvain for an insightful discussion. We are grateful to Pace Phillips at Innovations for Poverty Action (IPA) for helping us better understand rainfall insurance in Ghana. We thank Zeina Shalaby for excellent research assistance on this project. We thank the University of Notre Dame's Building Inclusive Growth (BIG) Lab for support. Email: jmajerov@nd.edu, hpelleg3@nd.edu.

1 Introduction

In developing countries, a large share of the population works in agriculture and depends heavily on rainfall for their livelihoods. Being at the whim of the weather means that consumption is risky, especially where farmers lack access to agricultural insurance. Uninsured weather risk can deter investments in agricultural inputs even if their expected return is high, and push farmers to crops that have a lower expected return but greater resilience to rainfall shocks. Weather risk has also been proposed as a potential explanation for the apparent underutilization of inputs and inefficient crop choice in developing countries ([Donovan, 2021](#); [Allen and Atkin, 2022](#)).

Recognizing the challenges that uninsured weather shocks impose on farmers in developing countries, a large literature has studied consumption smoothing mechanisms and policy interventions that can help farmers cope with risk. A prominent strand of this literature examines index insurance tied to rainfall patterns (see [Carter et al. \(2017\)](#) for a review), with some papers finding significant effects of insurance provision on crop choices and investment ([Karlan et al., 2014](#)). More recently, there has been a growing interest in understanding how successful small-scale programs will perform when governments scale them up ([Mobarak, 2022](#); [Buera et al., 2023](#)). Yet the aggregate effects of scaling up agricultural insurance remain poorly understood.

Should governments scale up and subsidize rainfall insurance? To answer this question we build a general equilibrium model of agricultural households that make production and consumption decisions under rainfall risk, and sell a variety of crops to a common national market. These households live in different regions, each of which faces its own idiosyncratic weather shock. To insure against these shocks, households can buy rainfall insurance from a diversified national insurance firm. However, as is true in both developing and developed economies, the insurance is sold at above actuarially fair prices, because the insurance firm must cover its servicing costs. The government may choose to subsidize insurance, financed by lump-sum taxes.

Our main analytical result shows that the optimal subsidy balances three forces: a fiscal externality, a pecuniary externality, and a fiscal incidence term. The fiscal externality is standard. Increasing the generosity of the subsidy pushes farmers to take up more (subsidized) insurance, which drains the government's budget. This is balanced against a novel pecuniary externality, which reflects how changes in crop prices redistribute resources from the good state to the bad state of the world. Increasing the generosity of the subsidy encourages farmers to shift into riskier crops, which reduces the relative price of those crops in equilibrium. This change in crop prices shifts income from the good state to the bad state,

by reducing the relative importance of the risky crop in the farmer’s portfolio.

In our model, households face heterogeneous costs to participate in the insurance market, causing some households to not purchase any insurance. This introduces a third term to our sufficient-statistics characterization of the optimal subsidy, which we call the fiscal incidence term. When different households purchase different amounts of insurance, a subsidy funded by lump-sum taxes is a transfer from the households that buy less insurance to the households that buy more. In our calibrated model, we find that households that participate in the insurance market have a lower average marginal utility of consumption than those that do not participate. Thus, the subsidy redistributes the wrong way, amplifying differences in marginal utility, and the fiscal incidence term pushes the optimal subsidy lower.

In the absence of the fiscal incidence term (e.g. in a model where all households participate in the insurance market), we find that the optimal subsidy will be greater than zero, but never enough to implement full insurance. At zero subsidy, making the subsidy more generous imposes zero fiscal externality but has positive benefits through the pecuniary externality. Thus, a subsidy greater than zero is preferred. At full insurance, the fiscal externality becomes substantial but the pecuniary externality goes away: if farmers are fully insured, then there is no longer any benefit to redistributing resources across states. Thus, the subsidy must be intermediate: between zero and full insurance.

However, when we introduce an extensive margin to the insurance purchase decision, the fiscal incidence term can push the optimal subsidy downwards. If the pecuniary externality is small, it is even possible for the optimal subsidy to be negative: a tax on insurance rather than a subsidy.

Our analysis of optimal policy connects our model to the theory of the second-best. In a first-best equilibrium (i.e. full insurance), pecuniary externalities would be zero. Although insurance markets are incomplete in our model, subsidizing insurance does not fix the problem. The fundamental friction is that there are servicing costs associated with providing insurance, which the subsidy does not undo. Subsidies are only beneficial indirectly, through their effect on crop prices.

Guided by our sufficient-statistics result, we calibrate the model using experimental moments from [Karlan et al. \(2014\)](#), supplemental moments from a household survey (the Ghana Living Standards Survey), and additional parameters from the literature. [Karlan et al. \(2014\)](#) experimentally estimate the price elasticity of demand for rainfall insurance, which we use to pin down the coefficient of relative risk aversion in our model. Since the fiscal externality depends on how insurance demand responds to a subsidy, this establishes a tight link between the data and our estimate of the optimal subsidy. We calibrate the distribution of participation costs to match estimated insurance take-up at different prices, which allows us

to measure the fiscal incidence term. We use additional moments from the Ghana Living Standards Survey to calibrate the agricultural production function, including the relative risk and return of different crops. We also match the experimental estimate of how the allocation of land to risky crops responds to insurance subsidies, an essential mechanism behind the pecuniary externality. Along with the demand elasticity of substitution across crops, these parameters determine how changes in the insurance subsidy affect relative crop prices in general equilibrium. These are thus key inputs for measuring the pecuniary externality.

We begin by examining the effects of scaling up agricultural insurance, comparing the effects in partial and general equilibrium. Subsidizing insurance down to the actuarially fair price leads to higher take-up in partial equilibrium, as we move along the household's demand curve. In general equilibrium, there are two additional forces. The first is a crop-price channel, which is closely connected to the pecuniary externality. Insurance subsidies push farmers to grow riskier crops, which lowers the relative price of these crops. This reduces the riskiness of the farmer's portfolio and thus lowers insurance demand, although the effect is quantitatively small. The second force is connected with the fiscal externality. Subsidizing insurance below cost requires taxing households, making them poorer. Poorer households have higher (absolute) risk aversion, which raises their demand for insurance. Quantitatively, the fiscal channel is more important than the crop-price channel for take-up, and take-up is dramatically higher in general equilibrium than in partial equilibrium.

Due to cheaper insurance and the dramatic increase in take-up, farmers are more willing to take on risk, leading to greater fertilizer use and a shift towards riskier crops to maximize expected returns. The total value of agricultural production is 13% higher once insurance is subsidized down to the actuarially fair price. These changes are significantly amplified by general equilibrium forces.

However, the increase in agricultural output and insurance take-up does not necessarily imply that the government should scale up the program. We find that the optimal subsidy is near zero. Moreover, the fiscal incidence term dominates the pecuniary externality, and so the optimal subsidy is negative: the optimal policy is a tax on insurance equal to 0.09% of the cost of insurance.

Why is the optimal subsidy so small? One reason is that the elasticity of demand for insurance is large: [Karlán et al. \(2014\)](#) estimate an elasticity of demand of -4.1. This high elasticity means that subsidies have a strong effect on insurance demand, and so the fiscal externality grows large as soon as we reach even a modest subsidy. Another reason that the optimal subsidy is small is because the pecuniary externality is small. Although crop insurance is valuable to farmers, this does not provide a rationale for the government to subsidize insurance. Instead the pecuniary externality only includes the indirect benefits

of insurance, through changes in crop prices. Quantitatively, these indirect effects are too small to justify a large subsidy. Moreover, although the optimal subsidy would be small even without the fiscal incidence term, the fiscal incidence term pushes the optimal subsidy to be even smaller, and to be negative.

The rest of the paper is organized as follows. The remainder of this section discusses the related literature. Section 2 presents the model; Section 3 presents our sufficient-statistics results on the optimal insurance subsidy; Section 4 discusses our calibration of the model; Section 5 uses the calibrated model to compute counterfactuals and the optimal subsidy; and Section 6 concludes the paper.

Related Literature. Our paper relates to different strands of the literature. First, we build on an extensive literature on agricultural risk in developing countries, dating back to the seminal work of [Townsend \(1994\)](#), [Paxson \(1992\)](#), and [Udry \(1994\)](#). More recently, this literature has studied how removing agricultural risk affects labor market decisions and a host of investment choices ([Donovan, 2021](#); [Arteaga et al., 2025](#); [Brooks and Donovan, 2020](#)). Several papers have estimated well-identified effects of agricultural insurance on crop choices, investment, and take-up using randomized controlled trials ([Karlan et al., 2014](#); [Lane, 2024](#)). We complement this work by examining the impact of scaling up agricultural insurance and whether governments should subsidize it, using the elasticities estimated in these experiments to discipline our structural model.

By bringing a general equilibrium perspective to the results from micro-level studies, our paper relates to growing research in development economics employing micro-to-macro approaches (see review by [Buera et al., 2023](#)). Papers in this literature have studied, for example, the aggregate effects of fertilizer subsidies ([Bergquist et al., 2022](#)), of microcredit interventions ([Kaboski and Townsend, 2011](#)), and of rural-urban migration ([Lagakos et al., 2023](#)). Here, we add to this literature by studying agricultural insurance in general equilibrium. We characterize analytically how general equilibrium forces shape the optimal insurance subsidy, and find quantitatively that pecuniary externalities that emerge in general equilibrium are small.

Finally, our work relates to the literature on optimal policy and sufficient statistics. In this literature, the envelope theorem allows us to limit our analysis of optimal policy to a small set of sufficient statistics ([Chetty, 2009](#)). In an environment with no distortions other than the policy, this gives rise to the traditional fiscal externality. In environments with additional distortions, the welfare effect of a change in policy will also include additional effects through externalities in incomplete markets ([Kleven, 2021](#)). In our setting, this is what gives rise to the pecuniary externality and the fiscal incidence term. Recent work in

macroeconomics and development has begun to study policy interventions through a similar optimal policy lens; for example, see [Brooks and Donovan \(2026\)](#) for an analysis of optimal dynamic fertilizer subsidies.

2 Model

We begin by laying out our model. For the quantitative implementation of the model, we will use particular functional forms for household utility, crop demand, and crop production functions; we note those here in laying out the model. However, our sufficient-statistics results, which we discuss in [Section 3](#), will be consistent with more general functional forms.

2.1 Model Setup

2.1.1 Household's Problem

Environment and Utility. Regions are indexed by $i \in [0, 1]$. Each region is ex-ante identical, and is inhabited by a representative household. All agents are price takers and all markets are competitive, unless otherwise stated.

In each region, there are two states of the world, the high-income state, s_i^H , and the low-income state, s_i^L .¹ Shocks are i.i.d. across regions, so there is no aggregate uncertainty.

Household utility in region i is given by

$$U_i = \mathbb{E}[u(c_i(s_i))] \tag{1}$$

$$= P(s_i^H) \cdot u(c_i(s_i^H)) + P(s_i^L) \cdot u(c_i(s_i^L)) \tag{2}$$

where $u(c)$ is a utility function, $c_i(s_i)$ is the household's consumption as a function of their local state, s_i , and $P(s_i^H)$ and $P(s_i^L)$ are the probabilities of the good state and bad state, respectively. In the quantitative implementation of the model, we will use a constant relative risk aversion (CRRA) utility function, $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$.

Farm Income. The representative household in region i operates a farm to generate income. The farmer has a unit endowment of land, which she can allocate towards G different crops. She must allocate this land before knowing which state of the world her region will be in, although her productivity in each crop is state-dependent. Her land allocated to crop g is denoted $l_{g,i}$, and her output of crop g is denoted $y_{g,i}$. She also can make her land more

¹To simplify our exposition, we allow for only two states of the world, but it is possible to extend the model and all of our primary results to incorporate any discrete number of states of the world.

productive by purchasing fertilizer, which we denote f_i . For simplicity, fertilizer increases the productivity of all crops. Her output of crop g is a function of the land allocation, fertilizer choice, and the state of her region:

$$y_{g,i} = y_{g,i}(s_i, l_{g,i}, f_i) \quad (3)$$

The farmer chooses her land allocation subject to the resource constraint

$$\sum_{g=1}^G l_{g,i} = 1 \quad (4)$$

and to non-negativity constraints

$$f_i \geq 0 \text{ and } l_{g,i} \geq 0. \quad (5)$$

In the quantitative implementation of the model, the farmer's production function is

$$y_{g,i}(s_i, l_{g,i}, f_i) = z_g(s_i) \cdot l_{g,i}^{\theta_g} f_i^{\phi} \quad (6)$$

The curvature parameter θ_g reflects, in a reduced form way, the notion that different plots of land will have different comparative advantage. If the farmer only allocates one acre to maize, then she will choose an acre that is very productive at growing maize; if she allocates ten acres to maize, then the marginal acre may not be especially good for growing maize. The productivity $z_g(s_i)$ reflects the fact that productivity varies across crops and across states: different crops will have different expected productivity and different sensitivities to shocks.

The farmer sells her crops for crop-specific prices p_g . Crops are costlessly tradable across regions, so p_g does not vary across regions. Since there is no aggregate risk, p_g is not state-dependent. Her agricultural profits are thus given by:

$$\pi_i(s_i) = \sum_{g=1}^G p_g \cdot y_{g,i}(s_i, l_{g,i}, f_i) - f_i \quad (7)$$

$$= \sum_{g=1}^G p_g \cdot z_g(s_i) \cdot l_{g,i}^{\theta_g} f_i^{\phi} - f_i. \quad (8)$$

Agricultural Insurance. To reduce its risk, the household may buy agricultural insurance. Agricultural insurance is an asset that pays off in the region's bad state, s_i^L . For region i , we write x_i to denote the amount of insurance that the region's household has purchased.

We write $p_{x,i}$ to denote the price of this insurance. For each unit of insurance purchased, the insurance firm will pay the household one in the bad state, and zero in the good state.² For simplicity of exposition, we constrain $x_i \geq 0$, meaning that households cannot sell insurance.

Household Budget Constraint. The household must satisfy a budget constraint in either state of the world. The household gets income from agricultural profits, $\pi_i(s_i^H)$ and $\pi_i(s_i^L)$, and receives a dividend from an insurance firm, Π^I . It must pay out for agricultural insurance, but it gets back a payout on its insurance in the bad state. At the end of the period, when uncertainty is realized, the household uses whatever resources are left over to pay for consumption. Thus, the household's budget constraint, state by state, is given by:

$$\begin{aligned} c_i(s_i^H) &= \pi_i(s_i^H) + \Pi^I - p_{x,i} \cdot x_i \\ c_i(s_i^L) &= \pi_i(s_i^L) + \Pi^I - p_{x,i} \cdot x_i + x_i. \end{aligned} \quad (9)$$

Extensive Margin. In the data, a significant share of households do not take up insurance even when offered the actuarially fair price.³ To incorporate an extensive margin to insurance take-up, we add in a utility cost that the household must pay to participate in insurance markets. Before making any decisions, household i draws utility shocks ε_i^0 and ε_i^1 , which are i.i.d. across regions. Let V^1 be the household's unconstrained expected utility, not including utility shocks. Let V^0 be the household's expected utility, under the constraint that it buys zero insurance. If the household chooses to participate in the insurance market, then it may buy insurance if it chooses to do so, and receives $V^1 + \varepsilon_i^1$.⁴ If the household chooses not to participate in the insurance market, then it must buy zero insurance, and receives $V^0 + \bar{\varepsilon} + \varepsilon_i^0$. Thus, the net utility cost of participating in the insurance market is $\bar{\varepsilon} + \varepsilon_i^0 - \varepsilon_i^1$.

In the quantitative implementation of the model, we assume that the utility shocks ε_i^0 and ε_i^1 are drawn i.i.d. from the Type I extreme value distribution (also known as Gumbel), with inverse scale parameter χ . This gives us closed-form expressions for the probability of participating in the insurance market

$$P(V^1 + \varepsilon_i^1 > V^0 + \bar{\varepsilon} + \varepsilon_i^0) = \frac{\exp(\chi V^1)}{\exp(\chi V^1) + \exp(\chi(V^0 + \bar{\varepsilon}))}, \quad (10)$$

²In equilibrium, insurance will have an expected return that is either zero or negative. Thus, because risk is idiosyncratic, no household would wish to buy insurance on another region.

³Karlan et al. (2014) report that only 40-50% of farmers take up insurance at the actuarially fair price.

⁴A household that chooses to participate may still choose to buy zero insurance. If insurance is actuarially unfair, then with CRRA utility rich households may prefer to self-insure.

and also for the overall value function

$$V := \mathbb{E} \left[\max \{ V^0 + \bar{\varepsilon} + \varepsilon_i^0, V^1 + \varepsilon_i^1 \} \right] \quad (11)$$

$$= \frac{1}{\chi} \cdot \log \left(\exp(\chi V^1) + \exp(\chi(V^0 + \bar{\varepsilon})) \right) + \frac{C}{\chi}, \quad (12)$$

where C is Euler's constant.

Household Maximization Problem. The household takes crop prices $\{p_g\}_{g=1}^G$, the insurance price $p_{x,i}$, and insurance firm profits Π^I as given. Summarizing the above, the household solves the following maximization problem if it chooses to participate in the insurance market:

$$V^1 := \max_{\{l_{g,i}\}_{g=1}^G, f_i, x_i} \mathbb{E} \left[u \left(\pi_i(s_i) + \Pi^I - p_{x,i} \cdot x_i + 1(s_i = s_i^L) \cdot x_i \right) \right] \quad (13)$$

s.t.

$$\pi_i(s_i) = \sum_{g=1}^G p_g \cdot y_{g,i}(s_i, l_{g,i}, f_i) - f_i$$

$$\sum_{g=1}^G l_{g,i} = 1$$

$$x_i \geq 0, f_i \geq 0 \text{ and } l_{g,i} \geq 0.$$

where $1(s_i = s_i^L)$ is an indicator that is equal to one in the bad state. The non-participating household solves a similar problem, but with the constraint $x_i = 0$. This simplifies the problem:

$$V^0 := \max_{\{l_{g,i}\}_{g=1}^G, f_i} \mathbb{E} \left[u \left(\pi_i(s_i) + \Pi^I \right) \right] \quad (14)$$

s.t.

$$\pi_i(s_i) = \sum_{g=1}^G p_g \cdot y_{g,i}(s_i, l_{g,i}, f_i) - f_i$$

$$\sum_{g=1}^G l_{g,i} = 1$$

$$f_i \geq 0 \text{ and } l_{g,i} \geq 0.$$

Given the values of V^1 and V^0 , and the draws of $(\varepsilon_i^0, \varepsilon_i^1)$, the household chooses to participate if $V^1 + \varepsilon_i^1 > V^0 + \bar{\varepsilon} + \varepsilon_i^0$.

2.1.2 Closing the Model

Fertilizer Production. The farmer buys fertilizer from a competitive fertilizer sector that produces fertilizer with a linear technology.⁵ This can represent local production, or import from abroad, as long as the supply curve is horizontal.⁶ As a normalization, we say that units of the final good can be converted one-to-one into units of fertilizer, so the price of fertilizer is one.

Final Good. Agricultural crops are aggregated into the final good, which is numeraire ($P = 1$). Let Y denote the aggregate supply of the final good, and let Y_g denote the aggregate supply of crop g . We thus have:

$$Y_g = \int_0^1 y_{g,i}(s_i) di \quad \forall g, \quad (15)$$

and

$$Y = Y\left(\{Y_g\}_{g=1}^G\right), \quad (16)$$

where $Y(\cdot)$ is an aggregator with constant returns to scale (homogeneous of degree one). Note that since the aggregator has constant returns to scale, it does not matter whether this aggregation is done by a “final good producer” or by each household.

In our quantitative implementation of the model, the aggregation is CES, with elasticity of substitution σ . We have:

$$Y = \left(\sum_{g=1}^G Y_g^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}. \quad (17)$$

The price of each crop is determined in a competitive market, which gives us:

$$p_g = \frac{\partial Y\left(\{Y_g\}_{g=1}^G\right)}{\partial Y_g}, \quad (18)$$

which, under CES demand, becomes:

$$p_g = \left(\frac{Y_g}{Y}\right)^{-1/\sigma}. \quad (19)$$

⁵Since the fertilizer sector is competitive and uses a linear technology, it earns zero profits.

⁶In most developing countries, the majority of fertilizers are imported from abroad; see for example [Farrokhi and Pellegrina \(2023\)](#).

Insurance Firm. There is a representative insurance firm that offers agricultural insurance to each region. To cover the insurance, the insurer must pay out in the bad states of the world. Since risk is idiosyncratic and there is a continuum of regions, the insurance firm is fully diversified.

In addition to the direct cost of insurance, the insurer also pays a per-unit servicing cost, τ , and earns a per-unit markup, μ . The servicing cost represents true resource costs associated with providing insurance, while the markup represents economic profits.

The insurance firm thus offers the prices

$$p_{x,i} = P(s_i^L) + \tau + \mu \forall i \quad (20)$$

and makes profits

$$\begin{aligned} \Pi^I &= \int_0^1 (p_{x,i} - P(s_i^L) - \tau) \cdot x_i di \\ &= \mu \cdot \int_0^1 x_i di \end{aligned} \quad (21)$$

Note that since some households i do not participate in the insurance market, x_i equals zero for those households. The insurance firm is owned by households, so its profits end up rebated to households as a dividend. We assume that each household owns an equal share of the insurance firm.

Equilibrium. An equilibrium is a set of prices and allocations that satisfies the following conditions. Households take crop prices $\{p_g\}_{g=1}^G$, insurance prices $\{p_{x,i}\}_{i \in [0,1]}$, and insurance firm profits Π^I as given. Aggregate output of the final good, Y , is given by Equation 16. Crop prices are given by Equation 18, insurance prices are given by Equation 20, and the insurance firm's profits are given by Equation 21. Each household makes an insurance market participation decision, chooses its land allocation $\{l_{g,i}\}_{g=1}^G$, fertilizer f_i , and the amount of insurance purchased x_i (non-participants must choose $x_i = 0$), in order to maximize their expected utility given their draw of utility shocks $(\varepsilon_i^0, \varepsilon_i^1)$, subject to the household budget constraint in Equation 9 and to the technological, resource, and non-negativity constraints given in Equations 3, 4, and 5.

The aggregate resource constraint is

$$\int_0^1 c_i di = Y - \int_0^1 f_i di - \int_0^1 \tau \cdot x_i di \quad (22)$$

Note, however, that since all profits in the economy are rebated to households, the aggregate

resource constraint will hold automatically as long as the other conditions hold.⁷

2.2 Separation and Efficiency

Assuming that the household allocates some land to each crop (this will always be the case for the production function we use in the quantitative application), the household's production decisions are characterized by the following first-order conditions:

$$\mathbb{E} \left[\frac{u'(c_i(s_i))}{\mathbb{E}[u'(c_i)]} \cdot p_g \frac{\partial y_{g,i}}{\partial l_{g,i}} \right] = \lambda_i \forall g \quad (23)$$

The household will equalize the risk-adjusted return on each crop. The ratio $\frac{u'(c_i(s_i))}{\mathbb{E}[u'(c_i)]}$ is the stochastic discount factor, representing the fact that states of the world with low consumption have higher marginal utility, and thus the household values returns in the bad state more than returns in the good state.

If insurance demand is positive, $x_i > 0$, the household's insurance demand is characterized by the following first-order condition:

$$\mathbb{E}[u'(c_i)] \cdot p_{x,i} = u'(c_i(s_i^L)) \cdot P(s_i^L) \quad (24)$$

The household balances the price of insurance, which must be paid in all states of the world, against the marginal benefit of an extra unit of consumption that only pays off in the bad state.

Combining these first-order conditions, we get that if insurance demand is positive and some land is allocated to each crop, the risk-adjusted return on each crop will be given by:

$$\mathbb{E} \left[\frac{u'(c_i(s_i))}{\mathbb{E}[u'(c_i)]} \cdot p_g \frac{\partial y_{g,i}}{\partial l_{g,i}} \right] = p_{x,i} \cdot p_g \frac{\partial y_{g,i}(s_i^L)}{\partial l_{g,i}} + (1 - p_{x,i}) \cdot p_g \frac{\partial y_{g,i}(s_i^H)}{\partial l_{g,i}} \quad (25)$$

$$= \lambda_i \forall g \quad (26)$$

This result is a separation result. It tells us that, as long as the farmer buys some insurance, her preferences and other sources of income do not matter for her production decisions. We can solve for her production decisions knowing just her production technology, crop prices, and the price of insurance.⁸

⁷The profits of the insurance firm are rebated to households, the fertilizer sector makes zero profits, and each farmer keeps her own agricultural profits. Since pricing is competitive and the aggregator is homogeneous of degree one, the final good aggregator makes no profits by Euler's Theorem.

⁸Note that the separation result only applies to households that buy insurance. For households that buy zero insurance, we must solve the full household problem in order to pin down the stochastic discount factor

The result shows that the farmer behaves as if maximizing expected profits, but weighting profits in each state using the state-prices $p_{x,i}$ and $(1 - p_{x,i})$ rather than the true probabilities. When insurance is actuarially fair, $p_{x,i} = P(s_i^L)$, these state-prices coincide with the true probabilities and households maximize expected profits. When insurance is less than actuarially fair ($\tau + \mu > 0$), either because servicing costs make transferring consumption across states costly (τ) or because market imperfections prevent farmers from obtaining actuarially fair coverage (μ), consumption is no longer equalized across states, and aggregate output falls as households deviate from expected profit maximization. Hence, a distortion in the insurance market translates directly into a distortion in the productive efficiency of the economy.

3 Welfare and Optimal Subsidy

To scale up rainfall insurance, the government can provide a subsidy. We will consider per-unit subsidies to insurance that are financed by lump-sum taxes.

In our model, it is straightforward to see that subsidies are equivalent to (negative) markups. Markups raise the price of insurance, $p_{x,i}$, and the profits, Π^I , are rebated back to households. Equivalently, subsidies lower the price of insurance, with the cost coming out of the household's budget constraint. Since subsidies enter the model in exactly the same way as insurance, we will say that there is a baseline markup on insurance, μ^{Baseline} , and then the actual markup is equal to the baseline markup minus the subsidy:

$$\mu = \mu^{\text{Baseline}} - \text{Insurance Subsidy} \quad (27)$$

Subsidies are thus a tool that allows the government to choose the markup, μ . The government's problem is to select the optimal markup, μ^* , that maximizes household welfare.

We first analyze the special case without an extensive margin, which provides additional intuition. We then show how the results change when we incorporate the extensive margin.

3.1 The General Equilibrium Effect of Subsidies on Welfare

To determine the optimal subsidy we first study how subsidies affect welfare in general equilibrium. The household takes prices and insurer profits as exogenous. However, equilibrium prices and insurer profits will change in response to a change in the markup. We thus take the household's indirect utility function, $U(p_x, \Pi^I, \{p_g\}_{g=1}^G)$, and totally differentiate with

$\frac{u'(c_i(s_i))}{\mathbb{E}[u'(c_i)]}$.

respect to the markup. In what follows, we suppress the household subscript i . This yields:

$$\frac{dU}{d\mu} = \frac{\partial U}{\partial p_x} \cdot \frac{dp_x}{d\mu} + \frac{\partial U}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu} + \sum_{g=1}^G \frac{\partial U}{\partial p_g} \cdot \frac{dp_g}{d\mu} \quad (28)$$

$$= \mathbb{E} \left[u'(c) \cdot \left(\underbrace{\frac{\partial c}{\partial p_x} \cdot \frac{dp_x}{d\mu}}_{\text{Cost of Insurance}} + \underbrace{\frac{\partial c}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu}}_{\text{Insurer Profits}} + \underbrace{\sum_{g=1}^G \frac{\partial c}{\partial p_g} \cdot \frac{dp_g}{d\mu}}_{\text{Farmer Profits}} \right) \right] \quad (29)$$

where the second line uses the envelope theorem. Note that the total derivatives, $\frac{d}{d\mu}$, reflect the general equilibrium effect of a change in the markup on that variable.

An increase in the markup will reduce household welfare by raising the cost of insurance, but increases household welfare by increasing insurer profits, which are rebated back to the household. With some manipulations, we can derive the net effect of these two forces on the household's income:

$$\frac{\partial c}{\partial p_x} \cdot \frac{dp_x}{d\mu} + \frac{\partial c}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu} = \mu \cdot \frac{dx}{d\mu} \quad (30)$$

If the initial markup is zero, then the increase in insurer profits will exactly cancel out the increase in the cost of insurance. However, when the markup is non-zero, changing the markup will affect household income. Taking into account just these first two channels, household income is maximized when $\mu = 0$.⁹

Finishing our derivation by noting that $\frac{\partial c}{\partial p_g} = y_g$, we obtain our main expression for the marginal effect of the markup on welfare:

$$\frac{dU}{d\mu} = \mathbb{E} \left[u'(c) \cdot \left(\mu \cdot \frac{dx}{d\mu} + \sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right] \quad (31)$$

Note that since the subsidy decreases the markup one-for-one, the derivative of welfare with respect to the subsidy is simply the negative of this expression.

3.2 A Sufficient-Statistics Approach for the Optimal Subsidy

To find the optimal insurance subsidy, note that at the optimal markup, μ^* , the derivative of welfare with respect to the markup must be equal to zero. Setting $\frac{dU}{d\mu} = 0$ and rearranging terms, we have that, at the optimal subsidy:

⁹Since $\frac{dx}{d\mu}$ is negative (a higher markup will lower insurance demand), raising the markup will lower household income if μ is positive, and will raise household income if μ is negative.

$$\text{At Optimal Subsidy: } \underbrace{-\mu \cdot \frac{dx}{d\mu}}_{\text{Fiscal Externality}} = \underbrace{\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot \left(\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right]}_{\text{Pecuniary Externality/Risk-Sharing}} \quad (32)$$

The optimal subsidy balances two competing externalities. The fiscal externality represents the net effect of the subsidy on household income. If the baseline markup μ^{Baseline} is zero, then this can be thought of as a traditional fiscal externality. An increase in the generosity of the subsidy increases demand for insurance, which drains the government's budget if the subsidy is positive.

The pecuniary externality reflects how changes in crop prices transfer resources from the good state to the bad state of the world. An increase in the generosity of the subsidy induces farmers to produce more of the risky crop relative to the low-risk crop. This lowers the relative price of the risky crop and increases the relative price of the low-risk crop, transferring income from the good state to the bad state. Note that the pecuniary externality has no effect on expected income, but it is valuable to the household as long as consumption in the bad state is lower than consumption in the good state.¹⁰

When insurance is priced as if it were competitive, $\mu = \mu^{\text{Baseline}} - \text{Insurance Subsidy} = 0$, the fiscal externality will be zero. This reflects the notion that around the competitive optimum, the first-order cost of distorting the insurance market with markups is zero.

To better understand the pecuniary externality, we can also express it as follows, under the assumption that insurance demand is positive:¹¹

$$\text{If } x > 0: \mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot \left(\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right] = (\mu + \tau) \cdot \sum_{g=1}^G (y_g(s^L) - y_g(s^H)) \cdot \frac{dp_g}{d\mu} \quad (33)$$

When insurance is actuarially fair, $\mu + \tau = 0$, the pecuniary externality will be zero.¹²

¹⁰We can use competitive crop markets and constant-returns-to-scale aggregation to prove that the pecuniary externality has no effect on average income: that is, $\mathbb{E} \left[\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right] = 0$. First, note that since we have competitive prices and $P = 1$, $p_g = \frac{dY}{dy_g}$. Since the aggregator has constant returns to scale, Euler's Theorem then gives us $\mathbb{E} \left[\sum_{g=1}^G p_g \cdot y_g \right] = Y$. Then, by the product rule, $\frac{dY}{d\mu} = \mathbb{E} \left[\sum_{g=1}^G \left(p_g \cdot \frac{dy_g}{d\mu} + y_g \cdot \frac{dp_g}{d\mu} \right) \right]$. But, by the chain rule, $\frac{dY}{d\mu} = \mathbb{E} \left[\sum_{g=1}^G \frac{dY}{dy_g} \cdot \frac{dy_g}{d\mu} \right] = \mathbb{E} \left[\sum_{g=1}^G p_g \cdot \frac{dy_g}{d\mu} \right]$. Subtracting our chain rule expression from our product rule expression yields $\mathbb{E} \left[\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right] = 0$.

¹¹To derive this, we use the first-order condition for insurance demand, $p_x = P(s^L) \frac{u'(c)}{\mathbb{E}[u'(c)]}$, combined with the fact that the pecuniary externality has no effect on expected income, $\mathbb{E} \left[\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right] = 0$, and the insurance supply equation, $p_x = P(s^L) + \mu + \tau$.

¹²Recall that insurance is actuarially fair if the expected payout to a policy holder is equal to the price of

The pecuniary externality does not increase average income, it only transfers income across states. If insurance is actuarially fair then farmers will already have purchased enough insurance to equalize their consumption in the good state and the bad state, so there will be no benefit to transferring income across states.

3.3 The Optimal Subsidy Provides Insurance Below Cost, But Does Not Implement Full Insurance

From the analysis above, we can see that the optimal markup will be negative but will not implement full insurance: $-\tau < \mu^* < 0$. To see why this is true, we can imagine what would happen if the markup were outside of the relevant range. At $\mu = 0$, the fiscal externality is zero, but the pecuniary externality will still make subsidy beneficial. Thus, the optimal markup is negative, $\mu^* < 0$. At $\mu + \tau = 0$, the pecuniary externality is zero, but the fiscal externality will make a subsidy costly. Thus, the optimal markup is not negative enough to implement full insurance, $\mu^* > -\tau$.

What does this tell us about the optimal subsidy? The optimal subsidy will implement a markup that is less than zero but will not go all the way to full insurance. If there is a markup at baseline, the optimal subsidy will first undo that markup. Then, the optimal subsidy will go somewhat further, exploiting the pecuniary externality until it is balanced by the fiscal externality.

This analysis connects to more general principles about efficiency in incomplete markets. When markets are competitive ($\mu = 0$) and complete ($\tau = 0$), both the fiscal and pecuniary externalities are zero, and the first welfare theorem guarantees that the efficient allocation is achieved. However, away from the first-best, pecuniary externalities need no longer cancel out. In our setting, the market incompleteness created by $\tau > 0$ is what allows for non-zero pecuniary externalities even in the competitive allocation ($\mu = 0$), and thus calls for a subsidy that goes beyond the traditional result of simply undoing markups.

3.4 Optimal Subsidies with an Extensive Margin

We now return to the full model, which incorporates an extensive margin.

We are interested in the derivative of the value function, V , with respect to the markup, the insurance. In our model, this requires $p_{x,i} = P(s^L)$. Since Equation 20 gives us that $p_{x,i} = P(s^L) + \mu + \tau$, this is equivalent to $\mu + \tau = 0$.

μ . We have:

$$\begin{aligned}\frac{dV}{d\mu} &= \frac{d}{d\mu} \mathbb{E} [\max \{V^0 + \bar{\varepsilon} + \varepsilon^0, V^1 + \varepsilon^1\}] \\ &= P(\text{Participates}) \cdot \frac{dV^1}{d\mu} + (1 - P(\text{Participates})) \cdot \frac{dV^0}{d\mu}\end{aligned}\tag{34}$$

Note that there is no term reflecting how a change in markups affects the participation decision: this is because households at the margin of participating get the same utility from participation and non-participation.

Totally differentiating and applying the Envelope Theorem, we have:

$$\frac{dV^0}{d\mu} = \mathbb{E} \left[u'(c) \cdot \left(\mu \cdot \frac{d\bar{x}}{d\mu} + (\bar{x} - x) + \sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \middle| \text{Does Not Participate} \right]\tag{35}$$

$$\frac{dV^1}{d\mu} = \mathbb{E} \left[u'(c) \cdot \left(\mu \cdot \frac{d\bar{x}}{d\mu} + (\bar{x} - x) + \sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \middle| \text{Participates} \right]\tag{36}$$

$$\frac{dV}{d\mu} = \mathbb{E} \left[u'(c) \cdot \left(\mu \cdot \frac{d\bar{x}}{d\mu} + (\bar{x} - x) + \sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right]\tag{37}$$

where $\bar{x} := \int_0^1 x_i di$ is the average insurance demand for all households, as opposed to x which represents the household's own insurance demand (which is zero for non-participants). To study the optimal subsidy, we set $\frac{dV}{d\mu} = 0$. This yields:

$$\text{At Optimal Subsidy: } \underbrace{-\mu \cdot \frac{d\bar{x}}{d\mu}}_{\text{Fiscal Externality}} = \underbrace{\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot (\bar{x} - x) \right]}_{\text{Fiscal Incidence Term}} + \underbrace{\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot \left(\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right]}_{\text{Pecuniary Externality/Risk-Sharing}}$$

This characterization of the optimal subsidy is similar to Equation 32 from the special case without an extensive margin, but with two differences. First, we now have a fiscal incidence term, $\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot (\bar{x} - x) \right]$, which reflects how the subsidy transfers resources across households with different demand for insurance. Since the subsidy is financed through lump-sum taxation of all households, but only participants benefit from lower insurance prices, the subsidy transfers resources from non-participants to participants. If non-participants have higher marginal utility on average, then this will be a negative effect of the subsidy.

Second, there is a subtle difference in the interpretation of the pecuniary externality term. In the baseline model, changes in relative crop prices transferred resources from the good state, s^H , to the bad state, s^L . Adding an extensive margin into the model adds another

dimension to the state space: whether the household participates in the insurance market or not. The pecuniary externality now reflects transfers between four different states, instead of the original two. If the marginal utility of consumption differs between participants and non-participants, and if crop production decisions differ between these two groups, then the pecuniary externality will incorporate transfers between participants and non-participants.

Adding an extensive margin to the model can also overturn our earlier result that the optimal subsidy must be positive. In the baseline model, when the markup is zero, the fiscal externality is zero, so a subsidy must have a beneficial effect through the pecuniary externality. With an extensive margin, the fiscal incidence term can counterbalance the pecuniary externality, even when the fiscal externality is zero. This means that the optimal subsidy could be zero, or even negative.

4 Calibration

We calibrate our model to the agricultural economy in Ghana. We use three sources for calibration: experimental moments from [Karlan et al. \(2014\)](#), supplemental moments from the Ghana Living Standards Survey, and external calibration to parameters from [Sotelo \(2020\)](#) and [Donovan \(2021\)](#). The model is just identified and so we search for parameters that exactly hit the target moments. Although in principle the model parameters are jointly identified by all moments, our strategy enables us to draw a close connection between parameters and the particular moments that pin them down.

Our quantification proceeds in three steps. First, without solving the model, we calibrate some parameters externally: the elasticity of substitution σ , the elasticity of output with respect to fertilizer ϕ , the baseline markup μ^{Baseline} , the servicing cost τ , and the probability of the bad state $P(s^L)$. Second, we use moments from the Ghana Living Standards Survey to compute average land shares, and state-dependent yields for each crop. We show that these moments identify the state-dependent productivities $\{z_g(s_i)\}_{g=1}^G$ and the curvature parameters $\{\theta_g\}_{g=1}^G$, up to a constant θ^* . Finally, to identify the remaining parameters, we simulate our model for a given guess of parameters, and simulate what would happen if we randomly offered actuarially fair insurance to some households ($p_{x,i} = P(s^L)$), but kept other equilibrium objects (crop prices $\{p_g\}_{g=1}^G$ and insurer profits Π^I) constant. We match the results of this partial-equilibrium experiment to experimental moments from [Karlan et al. \(2014\)](#), and use this to pin down the coefficient of relative risk aversion γ , the insurance take-up parameters χ and $\bar{\varepsilon}$, and the crop curvature parameter θ^* . Pinning down θ^* pins down the full set of crop curvature parameters and the state-dependent productivities, so the model's parameters are fully identified.

Table 1 offers a summary of our calibration values. The table notes each of the parameters in our model, the target moment that identifies that parameter, and our calibrated value. In the remainder of this section, we discuss the calibration strategy for each of these parameters.

Table 1: Calibration Table

Parameter	Description	Value	Calibration Target
σ	Elasticity of Substitution Across Crops	2.4	Sotelo (2020)
ϕ	Elasticity of Output w.r.t. Fertilizer	0.4	Donovan (2021)
τ	Servicing Cost of Insurance	$0.4 \times P(s^L)$	Computed in Karlan et al. (2014)
$P(s^L)$	Probability of Low Rainfall	0.68	Rainfall Data
$\{\theta_g\}_{g=1}^G$	Curvature Parameter by Crop	See Table 2	Relative Yields Across Crops and Effect of Insurance on Crop Shares
$\{z_g(s_i)\}_{g=1}^G$	State-Dependent Productivity by Crop	See Table 2	Yield Regressions and Land Shares
γ	Coefficient of Relative Risk Aversion	4.4	Elasticity of Demand for Insurance
$(\chi, \bar{\varepsilon})$	Participation Cost Parameters	(0.16, 8.7)	Insurance Take-Up Probabilities
μ^{Baseline}	Markup on Insurance	0	By Assumption

4.1 Step 1: Externally Calibrated Parameters

We begin with parameters that we can calibrate externally, without solving the model.

Elasticity of Substitution Across Crops (σ). For the elasticity of substitution across crops, σ , we pick $\sigma = 2.4$ based on Sotelo (2020).¹³ The elasticity of substitution across crops plays an important role in our analysis because it is an input into the pecuniary externality. An increase in the generosity of the insurance subsidy will increase the quantity of risky crops that households produce. How that quantity effect translates into an equilibrium price effect, $\frac{dp_g}{d\mu}$, will depend on σ . Lower values of σ will translate into larger price effects. By using a fairly low estimate of σ , we give the pecuniary externality the best chance to matter.

Elasticity of Output with Respect to Fertilizer (ϕ). For the elasticity of output with respect to fertilizer, ϕ , we follow Donovan (2021) and set $\phi = 0.4$. Donovan bases his estimate on the cost share of intermediate inputs in U.S. agriculture. Note that although we refer to it as fertilizer, we interpret this as a stand-in for various intermediate inputs that are used in agricultural production.

¹³Using Peruvian data, he estimates an elasticity of $\sigma = 2.4$. This estimate is in broadly the same range as the rest of the literature; for example, Ghose et al. (2023) use the same IV strategy as Sotelo to estimate an elasticity of substitution across crops of $\sigma = 1.7$ in Sri Lanka.

Insurance Servicing Costs (τ) and Baseline Markup (μ^{Baseline}). On the supply side, the price of insurance will be the sum of three terms: the probability of the bad state $P(s^L)$, the servicing cost τ , and the markup μ . We calibrate $\tau + \mu$ using the actuarially fair price and the estimated market price of the Year 2 insurance product in Karlan et al. (2014). For the Year 2 insurance, Karlan et al. use a market price of 14 Ghanaian cedi (GHC) and an actuarially fair price of 10 GHC. Note that:

$$\frac{\text{Market Price}}{\text{Actuarially Fair Price}} = \frac{P(s^L) + \mu^{\text{Baseline}} + \tau}{P(s^L)} = \frac{14}{10} \quad (38)$$

which implies that $\mu^{\text{Baseline}} + \tau = 0.4 \times P(s^L)$.

To focus on the competitive benchmark, we calibrate $\mu^{\text{Baseline}} = 0$. This allows us to isolate the pecuniary externality as the rationale for insurance subsidy, which is the novel contribution of our paper. If there were a pure markup on insurance at baseline, then the optimal subsidy would undo this markup, as is already well-understood in other settings. However, we view it as plausible that servicing costs make up the majority of the gap between the actuarially fair price and the market price. Marketing insurance is quite expensive, especially to remote farmers in a poor country like Ghana. The insurer must employ a salesman and rent a vehicle to travel from household to household, selling policies that only amount to a few dollars; the insurer must then pay additional costs to run the firm, and to distribute payments if there is a payout.¹⁴

Calibrating the baseline markup μ^{Baseline} to zero, we obtain servicing costs equal to $\tau = 0.4 \times P(s^L)$.

4.2 Step 2: Estimating Land Shares and Yields

To estimate yields and land shares, we use moments from the Ghana Living Standards Survey (GLSS). These data provide us with households' revenue and land allocation by crop. The revenue data is available from rounds 4 to 7 (1998-1999, 2005-2006, 2012-2013, and 2016-2017).

Measuring Land Shares. The GLSS only asks households about the area harvested for certain crops.¹⁵ To back out land shares across crops, we use an imputation approach described in Appendix A. To calibrate our model, we focus on maize, sorghum, groundnut

¹⁴We are grateful to Pace Phillips at Innovations for Poverty Action (IPA) for helping us better understand these servicing costs.

¹⁵The GLSS divides crops into two groups based on typical harvest frequency. For frequently harvested crops, they do not ask the household for those crops.

(peanut), beans, rice, cassava, plantain, yam, and cocoyam (taro). These are the crops with the largest area harvested, excluding cocoa.¹⁶ We list our estimated land shares in Column 1 of Table 2.

Table 2: Crop-Specific Estimates

Crop	$\mathbb{E}[l_{g,i}]$ (1)	δ_g (2)	β_g (3)	θ_g (4)	$z_g(s^H)$ (5)	$z_g(s^L)$ (6)
Beans	0.07	0.00	-0.07	0.78	0.41	0.39
Cassava	0.13	0.51	-0.66	0.60	1.50	0.77
Cocoyam	0.07	0.06	-0.65	0.90	0.80	0.42
Groundnut	0.10	0.56	-0.76	0.56	1.39	0.65
Maize	0.28	0.37	-0.25	0.57	2.42	1.89
Plantain	0.08	0.55	-0.59	0.54	0.81	0.45
Rice	0.05	0.69	-0.75	0.49	0.65	0.31
Sorghum	0.14	-0.04	-0.62	0.90	1.27	0.68
Yam	0.07	1.66	-0.53	0.17	2.01	1.18

Notes: This table reports crop-specific estimates, reported to the second decimal place. Column 1 reports estimated land shares. Columns 2 and 3 report our estimates of Equations 39 and 40. We use these estimates to recover the crop-specific curvature parameters $\{\theta_g\}_{g=1}^G$, which are reported in Column 4. Finally, we use the curvature parameters, land shares, and yields to back out the productivity parameters $\{z_g(s^H)\}_{g=1}^G$ and $\{z_g(s^L)\}_{g=1}^G$, which we report in Columns 5 and 6.

Measuring Average Yields. To back out average yields, we focus on the round 7 data. Using data on farming households, we run the following regression:

$$\log(\text{Yield}_{irg}) = \alpha_r + \delta_g + \varepsilon_{irg} \quad (39)$$

where $\log(\text{Yield}_{irg})$ is the log yield (harvest value divided by area harvested) for household i , in region r , for crop g . To compute yields, we trim harvest values by crop-year at the 1st and 99th percentile, compute yields using either actual or imputed acreage, and then trim the resulting yields at the 1st and 99th percentile. The crop fixed effects are reported in Column 2 of Table 2. The omitted crop fixed effect is for beans, so each fixed effect is relative to beans.

Rainfall Data and Regressions. We supplement the GLSS data with daily rainfall data from Google Earth Engine. We use the CHIRPS data set from UC Santa Barbara. This data set provides data from 1981-2025 with 0.05 degree (≈ 5.5 km) pixel resolution. For each pixel, we compute a “bad weather” dummy for each year, following the exact definition used

¹⁶We drop cocoa because we believe that cocoa farming is distinct from farming the other foodstuffs in our model.

for payouts from the Year 2 insurance product in [Karlan et al. \(2014\)](#).¹⁷ Taking the average across pixel-years, we compute that there was a 68% probability of bad weather. We use this to calibrate the probability of the bad state in our model, $\Pr(s^L) = 0.68$.

To estimate how crop-specific productivity depends on the weather, we use the rounds 4-7 data to run the following regressions, crop by crop. For crops where we can observe yields directly (beans, groundnut, maize, rice, and sorghum), we regress log yields on a bad weather variable and on region and time fixed effects. We run the regression:

$$\log(\text{Yield}_{irgt}) = \beta_g \cdot \text{Bad Weather}_{rt} + \gamma_{rg} + \theta_{gt} + \varepsilon_{irgt} \quad (40)$$

where $\log(\text{Yield}_{irgt})$ is the log yield for household i , in region r , for crop g , in round t , and Bad Weather_{rt} is the average of the bad weather dummy across all pixels in region r at time t . For crops where we cannot observe yields, we instead use the harvest value as the outcome variable:

$$\log(\text{Harvest Value}_{irgt}) = \beta_g \cdot \text{Bad Weather}_{rt} + \gamma_{rg} + \theta_{gt} + \varepsilon_{irgt} \quad (41)$$

Although using the yield is preferable because it leads to more precisely estimated coefficients, both approaches are model-consistent, since the land allocation is chosen before the productivity is realized in the model.

We use the same trimming procedure for these regressions as we use for the average yield regressions. The resulting β_g estimates are reported in the Column 3 of Table 2.

Calibrating Parameters Using the GLSS Moments. With these estimates, we are able to calibrate the model to match the estimates. To map our regression estimates to the model, we assume that:

$$\frac{p_g \cdot \mathbb{E}[y_{g,i}]}{l_g} = \exp(\delta_g) \quad (42)$$

$$\frac{p_g \cdot y_{g,i}(s_i^L)}{l_g} = \exp(\beta_g) \cdot \frac{p_g \cdot y_{g,i}(s_i^H)}{l_g} \quad (43)$$

and that the land shares correspond to the land shares we computed using the round 7 data.

Identifying Production Parameters with the GLSS. For each crop, we need to identify three production function parameters: the curvature parameter θ_g , productivity in the good state $z_g(s^H)$, and productivity in the bad state $z_g(s^L)$. We can largely identify these using

¹⁷The insurance product defined the growing season from June to September. A day was considered “dry” if there was less than or equal to 1 millimeter of rain, and “wet” otherwise. The insurance product would trigger a payout if there were 12 or more consecutive dry days, and/or 7 or more consecutive wet days.

data on land use and yields.

How do we identify the curvature parameter, θ_g ? Recall the farmer's first-order condition for land: farmers equalize the risk-adjusted marginal revenue product of land across crops, $\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} p_g \cdot \frac{dy_g}{dl_g} \right]$. As long as insurance demand is positive, we can use the separation result from Equation 25 to rewrite this as a function of the price of insurance, p_x .¹⁸ Plugging in our specific functional form for the production function, we then have:

$$\lambda = \mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} p_g \cdot \frac{dy_g}{dl_g} \right] \quad (44)$$

$$= \theta_g \cdot \frac{p_g \cdot (p_x \cdot y_g(s^L) + (1 - p_x) \cdot y_g(s^H))}{l_g} \quad (45)$$

$$= \theta_g \cdot \frac{p_g \cdot (\mathbb{E}[y_g] + (\mu + \tau) \cdot (y_g(s^L) - y_g(s^H)))}{l_g} \quad (46)$$

Since we have estimated average yields in each state from the GLSS, and we have already externally calibrated $\mu + \tau$, this formula identifies the crop-specific curvature parameters θ_g up to scale. To fully pin down the curvature parameters, we will supplement these equations with an additional experimental moment from Karlan et al., which we describe in the next subsection.

Next, once we know θ_g , we can back out the productivity:

$$z_g(s_i) \cdot f_i^\phi = y_{g,i}(s_i) / l_{g,i}^{\theta_g} \quad (47)$$

Note that the fertilizer term, f_i^ϕ , is constant across crops and across states, and so it falls out as a normalization. Thus, to back out state-specific productivity, we can use the land shares, l_g , and state-specific output, $y_g(s)$, for each crop. We observe these objects in the GLSS, with the caveat that in practice we observe revenue rather than output. However, since we already know the elasticity of substitution σ , we are able to back out the crop price p_g from the revenue share, $\mathbb{E}[p_g \cdot y_g]$, and thus back out the model-implied quantity y_g .¹⁹

4.3 Step 3: Matching Experimental Moments

We complete our calibration by matching experimental moments from Karlan et al. (2014).

¹⁸In the extended model, the separation result will no longer hold for all households because some households will not purchase insurance. However, to simplify the calibration, we will continue to use the separation-based formula.

¹⁹Given that our demand system does not feature crop-specific demand multipliers, this step is partly a normalization.

Simulating the Partial-Equilibrium Experiment. Karlan et al. (2014) randomized the price of insurance offered to farmers. To match their experimental estimates, we replicate the Karlan et al. experiment within our model. To do this, we first solve the model in general equilibrium for a given parameter guess. Then, we vary the household-specific price of insurance, $p_{x,i}$, while holding fixed the other equilibrium quantities: insurer profits Π^I , and crop prices $\{p_g\}_{g=1}^G$. We compare the behavior of two groups of households: a “control” group that is offered the market price $p_{x,i} = P(s^L) + \mu^{\text{Baseline}} + \tau$, and a “treated” group that is offered the actuarially fair price $p_{x,i} = P(s^L)$. Note that the experiment is in partial equilibrium, and only requires us to solve the household’s problem at a new insurance price, whereas in general equilibrium there would also be changes in crop prices and insurer profits.

Identifying Risk Aversion (γ) from the Insurance Demand Elasticity. To identify the coefficient of relative risk aversion, γ , we match the partial-equilibrium price elasticity of demand for rainfall insurance to the experimentally estimated elasticity in Karlan et al. (2014). They estimate an elasticity of demand of -4.1 between the market price and the actuarially fair price.²⁰ To match this partial equilibrium elasticity, we compute the arc price elasticity of demand for insurance between the baseline market price and the actuarially fair price using the following formula:

$$\begin{aligned} & \text{Arc Price Elasticity of Demand} && (48) \\ & = \frac{(\mathbb{E}[x_i \mid \text{Treated}] - \mathbb{E}[x_i \mid \text{Control}]) / (\mathbb{E}[x_i \mid \text{Treated}] + \mathbb{E}[x_i \mid \text{Control}])}{(\text{Actuarially Fair Price} - \text{Market Price}) / (\text{Actuarially Fair Price} + \text{Market Price})} \end{aligned}$$

where $\mathbb{E}[x_i \mid \text{Treated}]$ is the average insurance demand for a treated household and $\mathbb{E}[x_i \mid \text{Control}]$ is the equivalent for a control household.

The elasticity of demand for insurance is informative about the coefficient of relative risk aversion, γ . To see why, note that at actuarially fair prices all households that participate in the insurance market will demand full insurance, regardless of their risk aversion. Above the actuarially fair price, the participating household’s demand for insurance will be decreasing in the level of risk aversion. Thus, the (negative) elasticity between the actuarially fair price and the market price will be decreasing in risk aversion, γ .²¹ If the farmer is nearly risk-neutral, demand will quickly drop as we move away from the actuarially fair price, implying

²⁰We obtain this estimate from Figure 3 of their paper, comparing insurance demand at 9.5 Ghanaian cedis to demand at 14 cedis. This is an elasticity of total demand, including zeros.

²¹This argument is exact in the absence of participation costs. With participation costs, things are somewhat more complicated: the elasticity of demand for participating households is decreasing in γ , but an increase in risk aversion can also raise insurance take-up at the actuarially fair price, and so risk aversion has ambiguous effects on the extensive margin of demand. We use additional moments to pin down the extensive margin of demand.

a large elasticity. If the farmer is very risk-averse then demand will stay high even at higher prices, so the elasticity will be smaller.

Identifying the Extensive Margin of Insurance Demand (χ and $\bar{\varepsilon}$). For the extended model with an extensive margin, we must also calibrate χ and $\bar{\varepsilon}$ in order to match the extensive margin of insurance take-up. To do this, we run the same partial-equilibrium experiment as before, but add in two additional moments from [Karlan et al.](#): the probability of insurance take-up at the market price, and the probability of take-up at the actuarially fair price. [Karlan et al.](#) report that take-up ranges from 40-50% at actuarially fair prices, and from 10-20% at market prices.²² We target a take-up rate of 40% at the actuarially fair price and 20% at the market price. We search for parameters that can hit these target moments while still obtaining an arc elasticity of demand of -4.1. We estimate a coefficient of relative risk aversion of $\gamma = 4.4$, and utility cost parameters $\chi = 0.16$ and $\bar{\varepsilon} = 8.7$.

Identifying the Curvature (θ^*) from experimental evidence. Our moments from the Ghana Living Standards Survey identify the curvature parameters θ_g across crops, but only up to scale. That is, for each crop we know that $\theta_g = C_g \cdot \theta^*$, where C_g is identified from the GLSS moments and θ^* is common across crops. To fully pin down the system of equations and identify θ^* , we incorporate another experimental moment from [Karlan et al. \(2014\)](#).

[Karlan et al.](#) use the randomized offer of different insurance prices as an instrument for insurance take-up. They use this IV strategy to estimate an effect of insurance on various outcomes, including the share of land devoted to the cultivation of the risky crop. In the particular subregion of Ghana that they study, maize is the risky crop. They find that insurance raises the share of land devoted to maize by 9 percentage points, from a baseline of 31%.

To replicate this experimental estimate in our model, we replicate the instrumental variables regression in [Karlan et al. \(2014\)](#). This can be easily computed with the Wald estimator. The numerator (reduced form) is the change in the share of land devoted to risky crops, and the denominator (first stage) is the change in the probability of insurance take-up. Our IV estimate is thus given by:

$$\beta^{IV} = \frac{\mathbb{E} \left[\sum_{g \in \text{Risky Crops}} l_{g,i} \mid \text{Treated} \right] - \mathbb{E} \left[\sum_{g \in \text{Risky Crops}} l_{g,i} \mid \text{Control} \right]}{P(x_i > 0 \mid \text{Treated}) - P(x_i > 0 \mid \text{Control})} \quad (49)$$

where $\sum_{g \in \text{Risky Crops}} l_{g,i}$ is the share of land that household i devotes to risky crops, and

²²[Karlan et al.](#) note that their take-up estimates are very similar to those found by [Mobarak and Rosenzweig \(2012\)](#).

$x_i > 0$ indicates that household i purchased insurance.

A few things are worth noting here. First, it is only possible to compute this IV in the context of the extended model with an extensive margin: in the baseline model without participation costs, the probability of taking up insurance is either zero or one, and so the first stage will typically be zero. Second, the model makes clear that this IV estimator does not quite satisfy an exclusion restriction. A farmer who takes up insurance will take on more risk if the insurance is actuarially fair than she would at market prices, so there is a direct effect of the instrument on the outcome. We sidestep this issue by treating the IV estimate as a moment to be matched rather than attributing it a precise causal interpretation. Finally, [Karlan et al.](#) offer insurance at many different prices and change the insurance product from year to year; they use all of these randomized prices as instruments in their IV regression. We simplify this by simulating just an offer at the market price and an offer at the actuarially fair price.

We divide the crops into risky and non-risky groups based on the calibrated variance of their log productivity. We place groundnut, rice, and cassava into the risky group; these crops make up 28.34% of land at baseline. We select θ^* to match a one-third increase in the share of land devoted to risky crops.²³ We report the resulting estimates of curvature parameters and state-dependent productivities in Columns 4-6 of Table 2.

Although in principle all moments jointly identify the parameters, this moment will be important for identifying θ^* because the curvature parameter determines how much land quantities respond to differences in risk-adjusted productivity. Since actuarially fair insurance makes risky crops more productive in risk-adjusted terms, the curvature parameter will be critical for determining how land allocation responds to insurance.

5 Counterfactuals

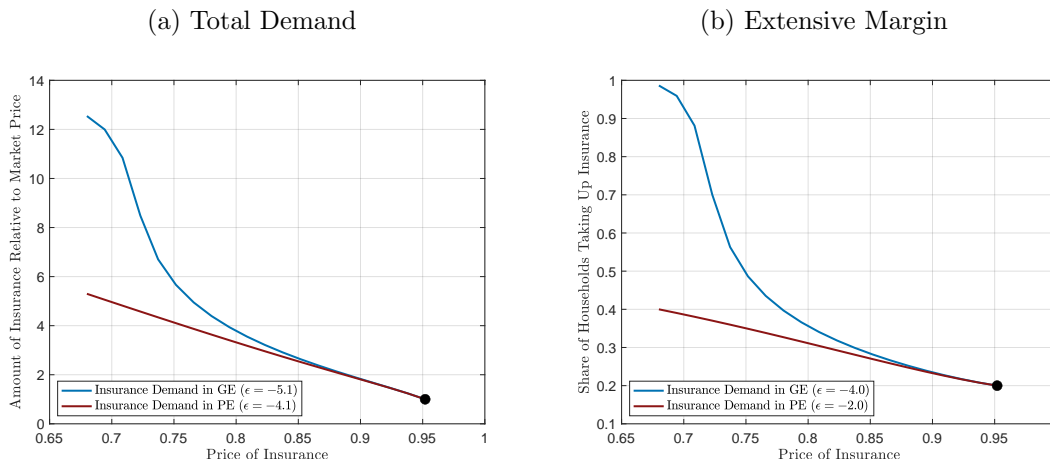
Using our calibrated model we now move on to estimate counterfactuals. For each counterfactual we will consider what happens as the government uses the insurance subsidy to change the markup, μ . Throughout, we compare the partial-equilibrium effect, in which other prices and insurer profits are held fixed, to the general-equilibrium effect, where equilibrium outcomes respond endogenously to the subsidy. We interpret the former as the case in which the government subsidizes the program in a single region, and the latter as the case in which the government expands the program to all regions.

We start with a detailed quantitative analysis of the impact of subsidies on the demand for insurance. We then move to the impact on various other outcomes: crop prices, land

²³We cap each θ_g at 0.9, so $\theta_g = \min(0.9, C_g \cdot \theta^*)$.

allocation, agricultural output, and fertilizer use. To understand mechanisms, we also discuss how the marginal revenue product of land changes with the subsidy. Finally, we conclude the section by analyzing the optimal subsidy, and studying how the optimal subsidy changes with different levels of risk aversion.

Figure 1: Insurance Demand in Partial vs. General Equilibrium



Notes: Panel (a) shows the change in the total amount of insurance purchased by households, including zeros, when we move the price from the actuarially fair price to the market price. Panel (b) shows the change in the share of households taking up insurance when we move the price from the actuarially fair price to the market price. We calibrate γ to match the arc elasticity of insurance demand with respect to price, and $\bar{\epsilon}$ and χ to match a 20% take-up rate at market prices, and to match the (partial-equilibrium) change in take-up rates between the market price and the actuarially fair price. The PE curve holds crop prices and insurance firm profits fixed, while the GE curve allows all crop prices and profits in the economy to adjust in response to the subsidy.

5.1 Impact of Insurance Subsidies on Demand for Insurance

Figure 1 starts by showing how insurance demand changes in response to the subsidy in partial equilibrium and in general equilibrium. Panel 1a shows the total demand for insurance, while Panel 1b shows the extensive margin: the share of households buying any insurance. In both panels, the black dot shows the price of insurance and insurance demand in the baseline equilibrium where the subsidy is zero.

Each panel has two curves: a partial-equilibrium curve and a general-equilibrium curve. Both curves trace out insurance demand as we move the price of insurance from the market price to the actuarially fair price. What distinguishes these two lines is what is happening in the background to other prices and to insurer profits. The partial equilibrium curve holds insurer profits Π^I and crop prices $\{p_g\}_{g=1}^G$ fixed at the same level as in the baseline equilibrium, and simply traces out household demand for insurance as a function of the price of insurance, $p_{x,i}$. The general equilibrium curve solves for equilibrium profits and prices at

each subsidy and reports the equilibrium insurance demand with all prices changing.

Although in principle the response of insurance demand in general equilibrium is distinct from the response in partial equilibrium, in practice the two curves are fairly similar. We calibrate to a partial-equilibrium arc elasticity of insurance demand of -4.1; the general-equilibrium elasticity is -5.1.²⁴ There is somewhat more divergence for the extensive margin: the partial-equilibrium take-up elasticity is -2.0, while the general equilibrium elasticity is -4.0.

The general-equilibrium response of insurance demand to the subsidy is the essential input into the fiscal externality. The fact that the general-equilibrium elasticity is similar to the partial-equilibrium elasticity confirms that there is a close connection between the sufficient statistics we derive from our macroeconomic model and the estimated demand elasticity that we get directly from the [Karlan et al. \(2014\)](#) experiment.

5.2 Impact of Insurance Subsidy on Agricultural Activity

Table 3 shows the impact of moving from the market price to the actuarially fair price of insurance across five sets of outcomes: insurance take-up, crop prices, land allocation, agricultural output, and fertilizer use. We compare the partial-equilibrium effects with the general-equilibrium ones, which are triggered by scaling up agricultural insurance.

Insurance. Panel A measures how the subsidy affects insurance participation. The share of farmers purchasing insurance rises from 20 percent at the market price to 40 percent in partial equilibrium and 99 percent in general equilibrium. On the intensive margin, the amount of coverage purchased, conditional on take-up, rises dramatically: by 165% in partial equilibrium and 154% in general equilibrium.

The large gap between partial- and general-equilibrium take-up reflects two forces, which we decompose quantitatively.²⁵ The first is the crop-price channel: as more farmers insure and shift land toward risky crops, the equilibrium price of those crops falls, making them less attractive and mildly reducing the demand for insurance. This force is quantitatively small, lowering take-up by 2 percentage points relative to partial equilibrium. The second, and dominant, force is a fiscal channel. Moving to the actuarially fair price requires the insurer to operate at a loss, and those losses are paid for by households through taxes (equivalently, the dividend Π^I is negative). This makes the shocks larger relative to the farmers average

²⁴Both of these elasticities are computed going from the market price to the actuarially fair price.

²⁵Specifically, we compute the share of farmers taking up insurance using Equation 10 under different price assumptions. To measure the crop-price channel, we compute take-up using crop prices from general equilibrium while holding Π^I fixed. To measure the fiscal channel, we compute take up using Π^I from general equilibrium while holding crop prices fixed at their baseline levels.

income, increasing her demand for insurance.²⁶ Quantitatively, this fiscal channel raises take-up by 59 percentage points relative to partial equilibrium, accounting for essentially the entire gap between partial and general equilibrium.

Crop Prices. Panel B shows how equilibrium crop prices respond to the subsidy. By construction, partial equilibrium holds crop prices fixed, so the partial equilibrium column replicates the baseline. In general equilibrium, the reallocation of land toward riskier crops increases their supply, pushing down their relative price. The price of high-risk crops falls by 2.7%, while the price of low-risk crops rises by 1.3%.²⁷ This general-equilibrium price adjustment will be a key input into the pecuniary externality.

Land Allocation. Panel C documents the reallocation of land across crop types. At the market price, farmers face uninsured weather risk and tilt their portfolios toward safer crops with lower expected returns on the margin. When insurance is subsidized to the actuarially fair price, this misallocation is corrected. In partial equilibrium, the share of land devoted to high-risk crops rises from 29 percent to 30 percent. In general equilibrium the share devoted to high-risk crops rises even further to 32 percent, as a much larger share of farmers take up insurance when the program is scaled up.

Fertilizer. Panel D shows that the subsidy also stimulates higher input use. Aggregate fertilizer rises by 14.1% in partial equilibrium and 35.3% in general equilibrium. This reflects the complementarity between insurance and intensive input use, as in [Donovan \(2021\)](#). Fertilizer is a risky investment, so insuring farmers against risk leads them to apply more fertilizer.

Agricultural Output. Finally, Panel E reports results for agricultural output. The total value of agricultural production rises by 5.4% in partial equilibrium and 13.2% in general equilibrium, reflecting the efficiency gains from land reallocation and increased fertilizer use. The gain is larger in high-risk crops, whose output value rises by 9.0% in partial equilibrium and 17.7% in general equilibrium, while the output value of low-risk crops rises by 3.9% in partial equilibrium and 11.4% in general equilibrium. Note, however, that these output gains

²⁶Formally, this is because the farmer has decreasing absolute risk aversion (DARA): as she gets wealthier, her absolute risk aversion declines. CRRA utility exhibits decreasing absolute risk aversion, as do most utility functions that are used in practice.

²⁷We define high-risk crops in the same way as we did earlier in the calibration to match the experimental moments. Low-risk crops are all the crops that are not high-risk crops. We compute the weighted average percent change in prices, weighting by baseline revenue, as in the Laspeyres index.

do not necessarily translate into an increase in welfare, since providing insurance is costly. We return to this point when we study the optimal subsidy.

Table 3: Impact of Subsidy With and Without General Equilibrium Effects

	Levels			% Change	
	Baseline	PE	GE	Δ PE	Δ GE
	(1)	(2)	(3)	(4)	(5)
<i>A. Insurance Demand</i>					
Extensive Margin	0.20	0.40	0.99	+100.00	+393.27
Intensive Margin	–	–	–	+164.94	+154.30
<i>B. Crop Prices</i>					
High-Risk Crops	–	–	–	+0.00	-2.73
Low-Risk Crops	–	–	–	+0.00	+1.29
<i>C. Land Allocation</i>					
High-Risk Crops	0.29	0.30	0.32	+6.29	+13.06
Low-Risk Crops	0.71	0.70	0.68	-2.52	-5.23
<i>D. Inputs</i>					
Fertilizer	–	–	–	+14.11	+35.28
<i>E. Output Value</i>					
High-Risk Crops	–	–	–	+9.02	+17.66
Low-Risk Crops	–	–	–	+3.88	+11.37
Total	–	–	–	+5.35	+13.17

Notes: This table shows the impact of subsidizing insurance by reducing its price from the market price to the actuarially fair price on various outcomes. Column 1 reports the value of each variable in the baseline simulation, where insurance is priced at the market price. Column 2 reports outcomes when the price of insurance is subsidized down to the actuarially fair price in partial equilibrium, holding crop prices and insurer profits fixed. Column 3 reports outcomes in general equilibrium, when crop prices and insurer profits are allowed to adjust endogenously. For Columns 1-3, we omit entries for outcomes that do not have well-defined units. Columns 4 and 5 express the same information as percent changes relative to the baseline. Panel A reports insurance demand: the extensive margin reports the probability of taking up insurance, and the intensive margin reports the insurance demand conditional on take-up. Panel B reports effects on crop prices. We use the same definition of high-risk crops as earlier, and define low-risk crops as all crops that are not high risk. We compute price changes as baseline-revenue-weighted percent change in prices, as in a Laspeyres index. Panel C reports the land allocation to low- and high-risk crops as a share of all land. Panel D reports fertilizer use, and Panel E reports the value of agricultural output.

5.3 Impact of Subsidies on Marginal Revenue Products

We next study how incomplete insurance generates gaps in the return to different crops. In Figure 2 we show the expected marginal revenue product of land for rice, plantains, and beans, for farmers who participate in the insurance market, as we vary the subsidy from zero (the vertical dashed line on the left) to full insurance (the vertical dashed line on the right). In our calibrated model, rice and plantains, represented by the dark and light blue lines

respectively, are relatively risky crops, while beans, represented by the red line, is a safer crop. This difference in the relative riskiness of each crop creates a gap in expected returns at baseline. The expected marginal return on rice is substantially higher than the return on beans, reflecting compensation for risk. Lack of insurance thus creates misallocation on the production side of the economy; relative to the first-best, farmers allocate too much land to beans and not enough to rice and plantains.

Farmers do not equalize the expected marginal revenue product of land across crops; instead they equalize the risk-adjusted return, $\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot MRPL \right]$. We plot the risk-adjusted return in black. Because all of the crops in our calibration are risky (they return more in the good state than the bad state), the risk-adjusted return is below the expected return of any of the crops.

As the insurance subsidy increases, farmers who buy insurance reallocate their land towards riskier crops. This shrinks the gap in expected returns across crops, as we can see in the figure. The risk-adjusted return rises as the crop allocation becomes more efficient and as the risk premium shrinks. At the actuarially fair price, the farmer is fully insured and all of the returns converge to the same value. If all households are fully insured, farmers implement the same allocation of land as they would in the first-best.

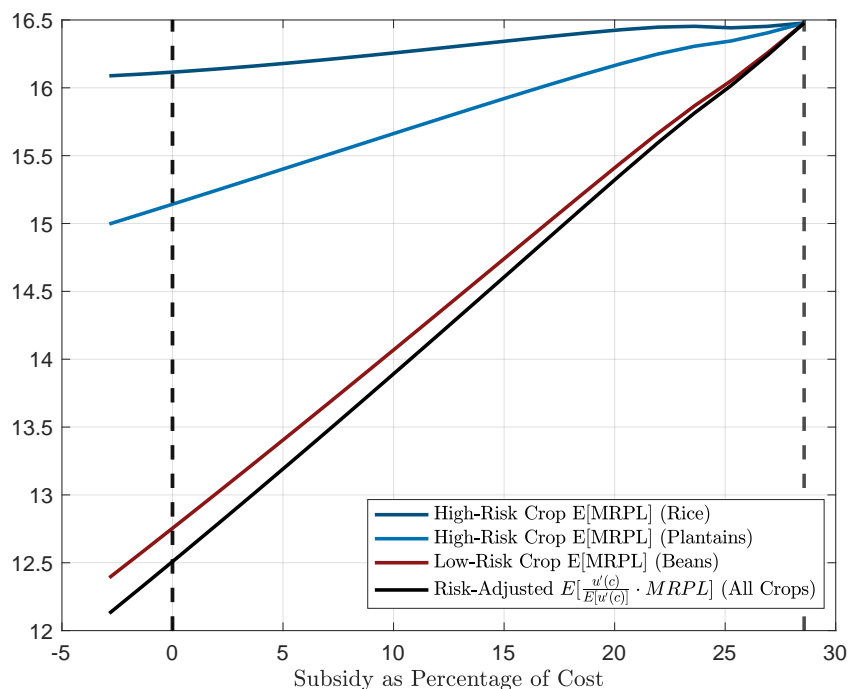
5.4 Optimal Subsidy

We now use our calibrated model to quantify the optimal subsidy, as well as the fiscal externality, pecuniary externality, and fiscal incidence term. Figure 3 plots these three terms, as well as household welfare as a function of the subsidy. The fiscal externality, $-\mu \frac{dx}{d\mu}$, is plotted in red. The fiscal externality is zero at zero subsidy, but grows quickly as the subsidy increases. The fiscal externality tells us that positive subsidies are harmful: an increase in the subsidy induces more demand, which drains the government budget. Note, however, that the sign of the fiscal externality flips if the subsidy becomes negative; the fiscal externality is a force that pushes towards zero subsidy.

The pecuniary externality, $\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot \left(\sum_{g=1}^G y_g \cdot \frac{dp_g}{d\mu} \right) \right]$, is plotted as a dashed blue curve. This externality measures how a change in the markup affects welfare by changing equilibrium crop prices, transferring resources across states. Note that the externality is negative, because it is expressed in terms of a change in markup; a negative pecuniary externality means a positive effect of subsidy. In contrast to the fiscal externality, the pecuniary externality is fairly small.


We plot the sum of the pecuniary externality and the fiscal incidence term as a solid blue curve. Whereas the pecuniary externality is always negative, the fiscal incidence term

Figure 2: Marginal Revenue Product of Land as a Function of Subsidy (Insurance Market Participants)



Notes: This figure shows the expected marginal revenue product of land for various crops, for households that participate in the insurance market. The expected marginal revenue product of land is defined as the derivative of expected revenue from sales of that crop, with respect to the land allocated to that crop. Blue and red curves show the expected marginal revenue product of land for rice, plantains, and beans. The black line shows the risk-adjusted marginal revenue product of land, which is defined as $\mathbb{E} \left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot MRPL \right]$. The risk-adjusted marginal revenue product of land is the same for all crops. The figure shows how the marginal revenue product of land varies, in general equilibrium, as the subsidy moves the price of insurance from the market price to the actuarially fair price.

Figure 3: Optimal Subsidy (Calibrated Model with Extensive Margin)



Notes: This figure shows the optimal subsidy under our calibrated model. The fiscal externality is plotted in red, the pecuniary externality is plotted in dashed blue, and the sum of the pecuniary externality and the fiscal incidence term is plotted as a solid blue curve. Equilibrium welfare is plotted in black. The optimal subsidy is where the fiscal externality intersects the sum of the pecuniary externality and fiscal incidence term, and is marked with a vertical dashed line. The optimal subsidy is -0.09% of the cost of insurance.

is positive, shifting the solid blue curve up relative to the dashed blue curve. The fiscal incidence term reflects the fact that the subsidy redistributes from households that do not buy insurance to households that do buy insurance, and this redistribution reduces expected welfare.

Household welfare is plotted on the figure in black. Welfare is maximized at the subsidy level where the fiscal externality (red curve) crosses the sum of the pecuniary externality and the fiscal incidence term (solid blue curve). The dashed vertical line marks the optimal subsidy.

In our calibrated model, the optimal subsidy is small, at only -0.09% of the cost of insurance.²⁸ The small subsidy reflects the fact that insurance demand is elastic, and so the fiscal externality grows quickly, restraining any subsidy from deviating too far from zero. Moreover, the fiscal incidence term counteracts the pecuniary externality, and in fact we end up with an optimal subsidy that is negative. Somewhat surprisingly, the optimal policy in our model is a small tax on rainfall insurance.

5.5 How Does the Optimal Subsidy Vary with Risk Aversion?

Risk aversion is critical for understanding the value of insurance. To better understand the role played by risk aversion in determining the optimal subsidy, we study how the optimal subsidy varies as we change the coefficient of relative risk aversion, γ .

5.5.1 Model Without Extensive Margin


To begin our analysis, we study the simpler version of our model, with no extensive margin. For this version of the model, we keep the same parameters as before, but we shut down the extensive margin (all households now participate in the insurance market), and we recalibrate γ to match the partial-equilibrium elasticity of demand for insurance. This yields an estimate of $\gamma = 3.6$.

In Figure 4, we plot the fiscal and pecuniary externalities for this version of the model, under the calibrated risk aversion of $\gamma = 3.6$, and for a dramatically increased level of risk aversion, $\gamma = 36$. The pecuniary externality, plotted in blue, is unaffected by the level of risk aversion. As shown in Equation 33, the separation result implies that as long as insurance demand is positive for all households, the pecuniary externality depends only on production parameters and insurance prices.

An increase in the household's risk aversion reduces the elasticity of demand for insurance, which makes the fiscal externality smaller. The fiscal externality under the calibrated γ is

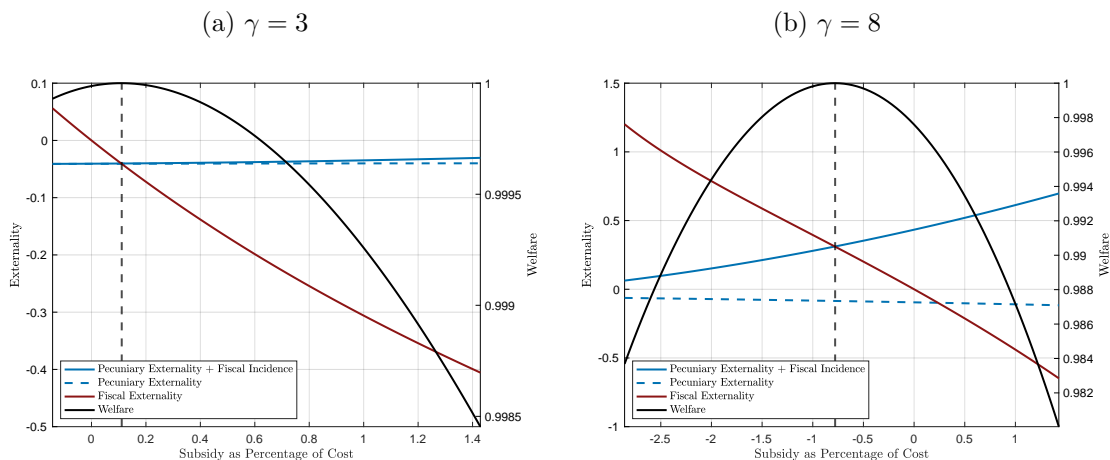
²⁸We express the subsidy as a percentage of the cost of insurance. So, $\frac{\text{Optimal Subsidy}}{P(s^L)+\tau} = -0.09\%$.

Figure 4: The Optimal Subsidy Increases with Risk Aversion (Model Without Extensive Margin)



Notes: This figure shows the optimal subsidy under a version of the model where the extensive margin of insurance demand is turned off, and all households participate in the insurance market. The coefficient of relative risk aversion is recalibrated to match the partial-equilibrium demand for insurance. The fiscal externality under the recalibrated model is plotted in solid red, and the fiscal externality under a model with elevated risk aversion is plotted in dashed red. The pecuniary externality is plotted in blue, and is the same for both levels of risk aversion. The optimal subsidies occur where the fiscal externality curves intersect the pecuniary externality, and are marked with vertical dashed lines.

Figure 5: Optimal Subsidy Under Alternative Risk Aversion (Model with Extensive Margin)



Notes: This figure shows the optimal subsidy under alternative levels of risk aversion. Panel (a) shows results for a lowered level of risk aversion, while Panel (b) shows results for an elevated level of risk aversion. The fiscal externality is plotted in red, the pecuniary externality is plotted in dashed blue, and the sum of the pecuniary externality and the fiscal incidence term is plotted as a solid blue curve. Equilibrium welfare is plotted in black. The optimal subsidy is where the fiscal externality intersects the sum of the pecuniary externality and fiscal incidence term, and is marked with a vertical dashed line.

plotted with a solid red line, while the fiscal externality under the elevated γ is plotted with a dashed red line. With this smaller fiscal externality, the crossing point between the fiscal and pecuniary externalities shifts out. The optimal subsidy rises from 0.26% of cost to 1.5% of cost.

This comparative static highlights how the large fiscal externality drives the low optimal subsidy. However, even with an extremely high risk aversion the optimal subsidy is small. Moreover, this risk aversion implies a counterfactually small elasticity of insurance demand. The insurance demand elasticity that we use is obtained directly from the [Karlan et al. \(2014\)](#) experiment. We interpret this as suggesting that the result of a small optimal subsidy is robust: a large subsidy would require a much smaller fiscal externality, which would strongly contradict the experimental data.

5.5.2 Model with an Extensive Margin

We now move to the more complicated case, where the model features an extensive margin. In Figure 5, we show how the optimal subsidy changes with different levels of risk aversion. Since our calibrated model featured a risk aversion of $\gamma = 4.4$, we examine how the results change if risk aversion is decreased ($\gamma = 3$) or increased ($\gamma = 8$).

In the model with an extensive margin, the fiscal externality, pecuniary externality, and fiscal incidence term can all change in response to a change in risk aversion. As before,

the fiscal externality term changes because risk aversion affects the demand for insurance. Unlike before, the pecuniary externality term now can depend on the level of risk aversion: this is because some households do not purchase insurance, and so the separation result no longer pins down the pecuniary externality. The fiscal incidence term will be affected by risk aversion as well, because changes in risk aversion will change both the probability of participating in the insurance market and the amount of insurance purchased by participating households.

As a result, in the model with an extensive margin, it is possible for the optimal subsidy to either increase or decrease with the level of risk aversion.²⁹ In our setting, we find that the optimal subsidy decreases with γ : we find an optimal subsidy of 0.11% under $\gamma = 3$, an optimal subsidy of -0.09% under the calibrated $\gamma = 4.4$, and an optimal subsidy of -0.78% under $\gamma = 8$. However, what all of these optimal subsidies have in common is that they are small. This is a robust finding from our analysis: plausible levels of risk aversion all yield small optimal subsidies.

5.6 Why is the Optimal Subsidy Small?

Across calibrations, we find that the optimal subsidy is consistently small. Looking to our sufficient statistics result, there are three reasons for this. First, the experimental evidence shows that insurance demand is elastic, and so the fiscal externality grows quickly as the subsidy moves away from zero. This makes the fiscal externality a powerful force pulling the optimal subsidy back toward zero. Second, subsidizing insurance only affects crop prices indirectly, and so the pecuniary externality is small relative to the fiscal externality. Finally, the fiscal incidence term pushes against the pecuniary externality and overpowers it in our calibrated model.

The optimal subsidy is small because it is the wrong tool to address the market failure. Traditionally, the targeting principle tells us that the optimal policy is to address a market failure at its source (Bhagwati and Ramaswami, 1963). In our setting, the source of the market failure is the servicing cost τ , which makes providing insurance costly. If the servicing cost could be reduced to zero, then the market outcome would achieve the first-best. In practice, this is typically not a tool available to the policymaker.

Yet even under this constraint, the policymaker can still do better than the laissez-faire equilibrium by changing relative crop prices. But the targeted tool for that job would be taxes and subsidies on crops, not an insurance subsidy. This is somewhat surprising, given that the origin of the distortion is in the insurance market. But insurance subsidies do not

²⁹In fact, the optimal subsidy can even be non-monotone in γ .

affect the servicing cost, and they only indirectly affect crop prices. Moreover, subsidizing insurance creates an additional fiscal externality in the insurance market. Insurance subsidies are the wrong tool for the job.

6 Conclusion

In this paper we study the general equilibrium effects of scaling up agricultural insurance. We derive a sufficient-statistics characterization of the optimal subsidy, which balances three forces: a fiscal externality, which reflects the effect of subsidy on the government budget, a pecuniary externality, which reflects how changes in crop prices transfer resources from the good state to the bad state, and a fiscal incidence term, which reflects how subsidies transfer resources from those who do not buy insurance to those who do. We calibrate the model to the agricultural economy in Ghana using experimental data from [Karlán et al. \(2014\)](#), as well as additional moments from the Ghana Living Standards Survey and parameters from the literature.

We find that the optimal subsidy is small, and negative at -0.09% of the cost of insurance. That is, the optimal policy is a very modest tax on insurance. The optimal subsidy is small because the pecuniary externality we measure in the data is modest, while the fiscal externality grows rapidly because insurance demand is fairly elastic. The optimal subsidy is negative because the fiscal incidence term counteracts the pecuniary externality, and pushes the optimal subsidy to be negative. The fact that the optimal subsidy is small in magnitude is a robust result of our analysis: generating even a modest subsidy would require a counterfactually low elasticity of insurance demand.

Our analysis highlights that even if insurance is valuable to farmers, and even if it significantly raises agricultural output, this alone is not a rationale for subsidy. Farmers already internalize the benefits of insurance from reduction in risk and from the ability to pursue riskier, more profitable agricultural investments. The rationale for a subsidy only comes through the pecuniary externality: the uninternalized benefits coming through equilibrium changes in crop prices. This rationale is too weak to rationalize any meaningful government intervention in the rainfall insurance market. Even in alternative models with additional externalities, the fiscal externality will be a powerful force pushing the optimal subsidy towards zero in any model that matches the elastic demand for insurance.

More broadly, our paper highlights the importance of general equilibrium analysis for understanding optimal policy under incomplete markets. Quantifying equilibrium effects on prices is essential for measuring the pecuniary externality, and balancing it against the fiscal externality and fiscal incidence of the policy.

References

- Allen, T. and D. Atkin (2022). Volatility and the gains from trade. *Econometrica* 90(5), 2053–2092.
- Arteaga, J., N. de Roux, M. Gáfaró, A. Ibanez, and H. S. Pellegrina (2025). Farm size distribution, weather shocks, and agricultural productivity. *Documento CEDE* (2025-5).
- Bergquist, L. F., B. Faber, T. Fally, M. Hoelzlein, E. Miguel, and A. Rodríguez-Clare (2022). Scaling agricultural policy interventions. Technical report, National Bureau of Economic Research.
- Bhagwati, J. and V. K. Ramaswami (1963, February). Domestic distortions, tariffs and the theory of optimum subsidy. *Journal of Political Economy* 71(1), 44–50.
- Brooks, W. and K. Donovan (2020). Eliminating uncertainty in market access: The impact of new bridges in rural nicaragua. *Econometrica* 88(5), 1965–1997.
- Brooks, W. and K. Donovan (2026). Industrial policy with development characteristics: Fertilizer policy in times of crisis. Technical report.
- Buera, F. J., J. P. Kaboski, and R. M. Townsend (2023). From micro to macro development. *Journal of Economic Literature* 61(2), 471–503.
- Carter, M., A. De Janvry, E. Sadoulet, and A. Sarris (2017). Index insurance for developing country agriculture: a reassessment. *Annual Review of Resource Economics* 9, 421–438.
- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics* 1(1), 451–488.
- Donovan, K. (2021). The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity. *The Review of Economic Studies* 88(5), 2275–2307.
- Farrokhi, F. and H. S. Pellegrina (2023). Trade, technology, and agricultural productivity. *Journal of Political Economy* 131(9), 2509–2555.
- Ghose, D., E. Fraga, and A. Fernandes (2023). Fertilizer import bans, agricultural exports, and welfare: Evidence from sri lanka. *Policy Research Working Paper 10642*, 793–831.
- Kaboski, J. P. and R. M. Townsend (2011). A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica* 79(5), 1357–1406.

- Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics* 129(2), 597–652.
- Kleven, H. J. (2021, August). Sufficient statistics revisited. *Annual Review of Economics* 13(1), 515–538.
- Lagakos, D., A. M. Mobarak, and M. E. Waugh (2023). The welfare effects of encouraging rural–urban migration. *Econometrica* 91(3), 803–837.
- Lane, G. (2024). Adapting to climate risk with guaranteed credit: Evidence from bangladesh. *Econometrica* 92(2), 355–386.
- Mobarak, A. M. (2022). Assessing social aid: the scale-up process needs evidence, too. *Nature* 609(7929), 892–894.
- Mobarak, A. M. and M. R. Rosenzweig (2012). Selling formal insurance to the informally insured.
- Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in thailand. *The American Economic Review*, 15–33.
- Sotelo, S. (2020). Domestic trade frictions and agriculture. *Journal of Political Economy* 128(7), 2690–2738.
- Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica: journal of the Econometric Society*, 539–591.
- Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern nigeria. *The Review of Economic Studies* 61(3), 495–526.

A Data

Imputing Land Shares. The GLSS only collects area harvested at the household level for a subset of crops. Fortunately, we are able to estimate the land shares devoted to each crop using an imputation procedure. In round 7, the GLSS collects area information at the plot level, as well as a list of what crops are harvested on that plot. We impute the land shares by crop in three steps. First, we trim the harvest value by crop across plots at the 1st and 99th percentile. Second, for each region-crop, we compute the yield among single-cropped plots as the total harvest value for that region-crop divided by the total area for that region-crop. Third, we return to the full sample of households and divide the harvest value of each crop by the yield we computed for that region-crop in order to get the imputed acreage.

The results are in Table 4. For beans, groundnut, maize, rice, and sorghum, we can compare the actual total acreage harvested with the imputed acreage. This comparison suggests that the imputation procedure is fairly accurate, at least for the purposes of an aggregate moment like land shares.

Table 4: GLSS Estimates

Crop	Acreage (1)	Imputed Acreage (2)
Beans	2,815.74	3,093.89
Cassava		4,907.58
Cocoyam		2,557.52
Groundnut	4,031.86	4,916.80
Maize	11,053.42	11,691.24
Plantain		2,960.37
Rice	2,061.25	3,037.53
Sorghum	5,542.16	4,184.56
Yam		2,889.68