Scaling up Agricultural Insurance

Jeremy Majerovitz
University of Notre Dame

Heitor S. Pellegrina
University of Notre Dame*

March 2026
(Preliminary)

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 result for the optimal subsidy, which trades off a fiscal externality from subsidizing insurance below cost against a pecuniary externality that captures how crop-price changes reallocate consumption across states of the world. We calibrate the model to Ghana's agricultural economy, combining experimental moments from Karlan et al. (2014) with additional moments from the Ghana Living Standards Survey.

Keywords: Agricultural Insurance, Investment Risk, Risk Sharing

*Acknowledgments: We are grateful to Joe Kaboski, and to seminar participants at the Y-RISE Annual Conference, for thoughtful comments on the paper. 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, which is especially costly for farmers operating near subsistence. Moreover, weather risk can reduce investment even if the expected return is high, and push farmers to crops that have a lower expected return but are more resilient to rainfall shocks. Weather risk is often pointed to as a potential explanation for the apparent underutilization of inputs and inefficient crop choice in developing countries (Donovan, 2021).

Recognizing the challenges that weather shocks impose on farmers in developing countries, a large literature has studied insurance and consumption smoothing mechanisms that can help farmers cope with agricultural risk. A prominent strand of this literature studies 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 our understanding of the aggregate effects of scaling up agricultural insurance remains limited.

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 ensure against these shocks the 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 two forces: a fiscal externality and a pecuniary externality. 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.

Our analysis finds that the optimal subsidy is greater than zero, but never enough to implement full insurance. At zero subsidy, making the subsidy more generous imposes zero fiscal externality but 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.

This also 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 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 use additional moments from the Ghana Living Standards Survey to calibrate the agricultural production function, including the relative risk and return of different crops. Moreover, to isolate the role of pecuniary externalities in our quantitative analysis, we assume that insurance firms earn no markup. 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. This is thus a key input for measuring the pecuniary externality.

We then use our calibrated model to run counterfactuals and to find the optimal subsidy. We find that the optimal subsidy is small, at just 0.1% of the cost of insurance. Why is the optimal subsidy so small? One reason is that the elasticity of demand for insurance is large: Karlan et al. (2014) estimate an elasticity of demand of -2.0. 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. The other 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.

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.

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 investment choices (Donovan, 2021; Brooks and Donovan, 2020). In particular, 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 whether governments should scale up and subsidize agricultural insurance, using the elasticities estimated in these experiments to discipline our structural model.

We also contribute to growing research examining the macro-economic implications of the micro-level decisions made by heterogeneous agents in the context of developing countries—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 contribute 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.

2 Model

We begin by laying out our baseline static model, and then we discuss extensions. 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 will be more general, and will not depend on functional form.

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. Shocks are i.i.d. across regions, so there is no aggregate uncertainty.

Household utility in region i is given by

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

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

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

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 $y_{g,i}$. Her output of crop g is a function of the allocation of land and of the state of her region:

$$y_{g,i} = y_{g,i}\left(s_i, l_{g,i}\right) z_g\left(s_i\right) \cdot l_{g,i}^{\theta_g} \tag{3}$$

and she chooses her land allocation subject to the resource constraint:

$$\sum_{g=1}^{G} l_{g,i} = 1 \tag{4}$$

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

$$y_{g,i}\left(s_i, l_{g,i}\right) = z_g\left(s_i\right) \cdot l_{g,i}^{\theta_g} \tag{5}$$

The curvature parameter θ_g reflects the notion that different plots of land will have different comparative advantages. 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 $z_g\left(s_i\right)$ 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:

$$\pi_i(s_i) = \sum_{g=1}^{G} p_g \cdot y_{g,i}(s_i, l_{g,i}) \tag{6}$$

$$= \sum_{g=1}^{G} p_g \cdot z_g(s_i) \cdot l_{g,i}^{\theta_g} \tag{7}$$

Final Good (Aggregation Across Crops) 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_j denote the aggregate supply of crop g. We thus have:

$$Y_g = \int_0^1 y_{g,i}(s_i)\, di \,\forall g \tag{8}$$

$$Y = Y\left(\{Y_g\}_{g=1}^{G}\right) \tag{9}$$

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: the two are mathematically equivalent.

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

$$Y = \left(\sum_{g=1}^{G} Y_g^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \tag{10}$$

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

$$p_g = \left(\frac{Y_g}{Y}\right)^{-1/\sigma} \tag{11}$$

Agricultural Insurance There is a representative insurance firm that offers agricultural insurance to each region. 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 representative household has purchased: the insurance firm will pay the household x_i in the bad state, and zero in the good state.[1] We write p_i^x to denote the price of this insurance. For

[1] In 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.

simplicity of exposition, we constrain $x_i \geq 0$, meaning that households cannot sell insurance.

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

$$p_{x,i} = P\left(s_i^L\right) + \tau + \mu \,\forall i \tag{12}$$

and makes profits

$$\begin{aligned} \Pi^I &= \int_0^1 \left(p_{x,i} - P\left(s_i^L\right) - \tau\right) \cdot x_i di \\ &= \mu \cdot \int_0^1 x_i di \end{aligned} \tag{13}$$

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.

Household Budget Constraint The household must satisfy a budget constraint in either state of the world. The household gets income from agricultural profits and from its dividend from the insurance firm. 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:

$$c_i\left(s_i^H\right) = \pi_i\left(s_i^H\right) + \Pi^I - p_{x,i} \cdot x_i \tag{14}$$

$$c_i\left(s_i^L\right) = \pi_i\left(s_i^L\right) + \Pi^I - p_{x,i} \cdot x_i + x_i \tag{15}$$

Equilibrium An equilibrium is defined as follows. Households take prices and insurance firm profits as given. Crop prices are given by Equation 11, insurance prices are given by Equation 12, and the insurance firm's profits are given by Equation 13. Each household chooses its land allocation, $\{l_{g,i}\}_{g=1}^G$, and the amount of insurance purchased, x_i, in order to maximize their expected utility. The aggregate resource constraint is

$$\int_0^1 c_i di = Y - \int_0^1 \tau \cdot x_i di \tag{16}$$

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.[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:

$$\mathbb{E}\left[\frac{u'\left(c_i\left(s_i\right)\right)}{\mathbb{E}\left[u'\left(c_i\right)\right]} \cdot p_g \frac{\partial y_{g,i}}{\partial l_{g,i}}\right] = \lambda_i \, \forall g \tag{17}$$

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.

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

$$\mathbb{E}\left[u'\left(c_i\right)\right] \cdot p_{x,i} = u'\left(c_i\left(s_i^L\right)\right) \cdot P\left(s_i^L\right) \tag{18}$$

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:

$$\mathbb{E}\left[\frac{u'\left(c_i\left(s_i\right)\right)}{\mathbb{E}\left[u'\left(c_i\right)\right]} \cdot p_g \frac{\partial y_{g,i}}{\partial l_{g,i}}\right] = p_{x,i} \cdot p_g \frac{\partial y_{g,i}\left(s_i^L\right)}{\partial l_{g,i}} + \left(1 - p_{x,i}\right) \cdot p_g \frac{\partial y_{g,i}\left(s_i^H\right)}{\partial l_{g,i}} \tag{19}$$

$$= \lambda_i \, \forall g \tag{20}$$

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

The result shows us that the farmer will behave as if she is maximizing expected profits, but weighting profits in each state using the "state-prices" $p_{x,i}$ and $(1 - p_{x,i})$ instead of the actual probabilities. If insurance is actuarially fair, $p_{x,i} = P\left(s_i^L\right)$, then these state-prices will be the true probabilities, and so households will maximize true expected profits. The

[2]The profits of the insurance firm are rebated to households, 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.

equilibrium with actuarially fair insurance will maximize aggregate output. If insurance is less than actuarially fair ($\tau + \mu > 0$), then aggregate output will be lowered because the household is not maximizing true expected profits. Thus, a distortion in the insurance market will distort the production efficiency of the economy.

2.1 Extensions

Dynamics. We are working on an extension of the model that allows for dynamic investment decisions along the lines of Donovan (2021). Specifically, the economy operates over discrete time. Each period is broken down into two stages, a cultivation and a harvest stage, separated by the realization of the weather regime. In the cultivation stage, households choose how much land to allocate to each crop and how much fertilizer to buy using their savings before the realization of the weather. In the harvest stage, production is realized and households choose how much to consume and how much to save for the following period.

Extensive vs. Intensive Margin. We are also planning to incorporate both extensive and intensive margins of insurance take-up. In our current formulation, all households purchase some insurance whenever the price is sufficiently low, yielding only adjustments on the intensive margin. We plan to introduce an extensive margin by allowing for idiosyncratic preferences for insurance within each region, so that only a fraction of households purchase coverage in equilibrium. Under suitable assumptions on the distribution of these preferences, we expect the model to deliver tractable closed-form expressions for regional take-up rates.

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:

$$\mu = \mu^{\text{Baseline}} - \text{Insurance Subsidy} \tag{21}$$

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.

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, $U\left(p_x, \Pi^I, \{p_g\}_{g=1}^G\right)$, and totally differentiate with respect to the markup.[3] This yields:

$$\frac{dU}{d\mu} = \frac{\partial U}{\partial p_x} \cdot \frac{dp_x}{d\mu} + \frac{\partial U}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu} + \sum_{g=1}^{G} \frac{\partial U}{\partial p_g} \cdot \frac{dp_g}{d\mu} \tag{22}$$

$$= \mathbb{E}\left[u'(c) \cdot \left(\underbrace{\frac{\partial c}{\partial p_x} \cdot \frac{dp_x}{d\mu}}_{\text{Cost of Insurance}} + \underbrace{\frac{\partial c}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu}}_{\text{Insurer Profits}} + \underbrace{\sum_{g=1}^{G} \frac{\partial c}{\partial p_g} \cdot \frac{dp_g}{d\mu}}_{\text{Farmer Profits}} \right)\right] \tag{23}$$

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:

$$\frac{\partial c}{\partial p_x} \cdot \frac{dp_x}{d\mu} + \frac{\partial c}{\partial \Pi^I} \cdot \frac{d\Pi^I}{d\mu} = \mu \cdot \frac{dx}{d\mu} \tag{24}$$

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 $\mu = 0$.[4]

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:

[3] For this section, we suppress the household subscript i.

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

$$\frac{dU}{d\mu} = \mathbb{E}\left[u'(c) \cdot \left(\mu \cdot \frac{dx}{d\mu} + \sum_{g=1}^{G} y_g \cdot \frac{dp_g}{d\mu}\right)\right] \tag{25}$$

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:

$$\text{At Optimal Subsidy:} \quad \underbrace{-\mu \cdot \frac{dx}{d\mu}}_{\text{Fiscal Externality}} = \underbrace{\mathbb{E}\left[\frac{u'(c)}{\mathbb{E}\left[u'(c)\right]} \cdot \left(\sum_{g=1}^{G} y_g \cdot \frac{dp_g}{d\mu}\right)\right]}_{\text{Pecuniary Externality/Risk-Sharing}} \tag{26}$$

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.[5]

[5] We 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 $\mathbb{E}\left[\sum_{g=1}^{G} y_g \cdot \frac{dp_g}{d\mu}\right] = 0$.

When the insurance market is competitive, $\mu = 0$, the fiscal externality will be zero. This reflects the notion that around the competitive optimum, the cost of distorting the insurance market with markups is second order.

To better understand the pecuniary externality, we can also express it as follows, under the assumption that insurance demand is positive:[6]

$$\text{If } x > 0\text{: } \mathbb{E}\left[\frac{u'(c)}{\mathbb{E}[u'(c)]} \cdot \left(\sum_{g=1}^{G} y_g \cdot \frac{dp_g}{d\mu}\right)\right] = (\mu + \tau) \cdot \sum_{g=1}^{G} \left(y_g\left(s^L\right) - y_g\left(s^H\right)\right) \cdot \frac{dp_g}{d\mu} \tag{27}$$

When insurance is actuarially fair, $\mu + \tau = 0$, the pecuniary externality will be zero. The pecuniary externality does not increase average income, it only transfers income across states. If insurance is actually 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

[6] To derive this, we use the first-order condition for insurance demand, $p_x = P\left(s^L\right)\frac{u'(c)}{\mathbb{E}[u'(c)]}$, 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, $p_x = P\left(s^L\right) + \mu + \tau$.

Parameter	Description	Value	Calibration Target
σ	Elasticity of Substitution Across Crops	2.4	Sotelo (2020)
γ	Coefficient of Relative Risk Aversion	7.4	Elasticity of Demand for Insurance
$P\left(s^L\right)$	Probability of Low Rainfall	0.68	Rainfall Data
$\{\theta_g\}_{g=1}^G$	Curvature Parameter by Crop	—	Relative Yields Across Crops
$\{z_g\left(s_i\right)\}_{g=1}^G$	State-Dependent Productivity by Crop	—	Yield Regressions and Land Shares
τ	Servicing Cost of Insurance	$0.4 \times P\left(s^L\right)$	Computed in Karlan et al. (2014)
μ^{Baseline}	Markup on Insurance	0	By Assumption

Table 1: Calibration Table

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.

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). 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.

We summarize our calibration strategy in Table 1. 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.

4.1 Using the Karlan et al. Experiment to Identify Risk Aversion (γ)

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). As part of their experiment, Karlan et al. randomized the price of insurance offered to farmers. They estimate an elasticity of demand of -2.0, between the market price and the actuarially fair price.

To match this partial equilibrium elasticity, 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

compute the price elasticity of demand for insurance between the baseline market price and the actuarially fair price and vary risk aversion γ until we hit the partial-equilibrium elasticity of demand exactly.

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 will demand full insurance, regardless of their risk aversion. Above the actuarially fair price, the 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. If the farmer is nearly risk-neutral, demand will quickly drop away from the actuarially fair price, implying 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.

We estimate a coefficient of relative risk aversion of $\gamma = 7.4$, which allows us to exactly match the insurance demand elasticity of -2.0.

4.2 Identifying Production Parameters

For each crop, we need to identify three production function parameters: the curvature parameter θ_g, productivity in the good state $z_g\left(s^H\right)$, and productivity in the bad state $z_g\left(s^L\right)$. We identify these using data on land use and yields. First, note that if we know θ_g, then we can back out the productivity:

$$z_g\left(s_i\right) = y_{g,i}\left(s_i\right) / l_{g,i}^{\theta_g} \tag{28}$$

Thus, to back out state-specific productivity, we need to be able to measure land shares, l_g, and state-specific output, $y_g\left(s\right)$, for each crop. We will be able to observe both of these in the data. Note that in practice we will observe revenue, rather than output. However, we are able to back out the crop price p_g from the revenue share, $\mathbb{E}\left[p_g \cdot y_g\right]$, and thus back out the model-implied quantity y_g.[7]

How do we identify the curvature parameter, θ_g? Recall the farmer's first-order condition for the land allocation: farmers equalize the risk-adjusted marginal revenue product of land across crops, $\mathbb{E}\left[\frac{u'(c)}{\mathbb{E}[u'(c)]} p_g \cdot \frac{dy_g}{dl_g}\right]$. Since insurance demand is positive, we can use the separation result from Equation 19 to rewrite this as a function of the price of insurance, p_x. Plugging

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

Crop	Acreage	Imputed Acreage	δ_g	β_g
Beans	2,815.74	3,093.89	0.00	0.23
Cassava		4,907.58	0.51	-0.66
Cocoyam		2,557.52	0.06	-0.65
Groundnut	4,031.86	4,916.80	0.56	-0.51
Maize	11,053.42	11,691.24	0.37	-0.21
Plantain		2,960.37	0.55	-0.59
Rice	2,061.25	3,037.53	0.69	-0.37
Sorghum	5,542.16	4,184.56	-0.04	-0.21
Yam		2,889.68	1.66	-0.53

Table 2: GLSS Estimates

in our specific functional form for the production function, we then have:

$$\lambda = \mathbb{E}\left[\frac{u'(c)}{\mathbb{E}[u'(c)]} p_g \cdot \frac{dy_g}{dl_g}\right] \tag{29}$$

$$= \theta_g \cdot \frac{p_g \cdot \left(p_x \cdot y_g\left(s^L\right) + (1-p_x) \cdot y_g\left(s^H\right)\right)}{l_g} \tag{30}$$

$$= \theta_g \cdot \frac{p_g \cdot \left(\mathbb{E}[y_g] + (\mu+\tau) \cdot \left(y_g\left(s^L\right) - y_g\left(s^H\right)\right)\right)}{l_g} \tag{31}$$

If insurance were actuarially fair, the ratio of average yields across crops would pin down the ratio of curvature parameters, θ_g. When insurance is not actuarially fair, the average yield is instead replaced by the risk-adjusted yield. Either way, this will be something we can measure with data on land shares and on output in the good and bad states.

4.2.1 Ghana Living Standards Survey (GLSS) Data

To pin down production parameters, we use moments from the Ghana Living Standards Survey (GLSS). These data provide us with household's 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). However, the GLSS only asks households about the area harvested for certain crops.[8]

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

Imputing Land Shares 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 the first two columns of Table 2. To calibrate our model, we focus on maize, sorghum, groundnut (peanut), beans, rice, cassava, plantain, yam, and cocoyam (taro). These are the crops with the largest area harvested, excluding cocoa.[9] 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.

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

$$\log\left(\text{Yield}_{irg}\right) = \alpha_r + \delta_g + \varepsilon_{irg} \tag{32}$$

where $\log\left(\text{Yield}_{irg}\right)$ 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 the third column 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 dataset 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 for payouts from the Year 2 insurance product in Karlan et al. (2014).[10] Taking the average

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

[10] The insurance product defined the growing season from June to September. A day was considered "dry" if there was less than 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.

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, $P\left(s^L\right) = 0.68$.

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:

$$\log\left(\text{Yield}_{irgt}\right) = \beta_g \cdot \text{Bad Weather}_{rt} + \gamma_{rg} + \theta_{gt} + \varepsilon_{irgt} \tag{33}$$

where $\log\left(\text{Yield}_{irgt}\right)$ 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:

$$\log\left(\text{Harvest Value}_{irgt}\right) = \beta_g \cdot \text{Bad Weather}_{rt} + \gamma_{rg} + \theta_{gt} + \varepsilon_{irgt} \tag{34}$$

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 final column 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 data, we assume that:

$$\begin{aligned}
\frac{p_g \cdot \mathbb{E}\left[y_{g,i}\right]}{l_g} &= \exp\left(\delta_g\right) \\
\frac{p_g \cdot y_{g,i}\left(s_i^L\right)}{l_g} &= \exp\left(\beta_g\right) \cdot \frac{p_g \cdot y_{g,i}\left(s_i^H\right)}{l_g}
\end{aligned}$$

and that the land shares correspond to the land shares we computed using the round 7 data. Thanks to the separation hypothesis, we are able to use δ_g and β_g to identify the crop-specific curvature parameters θ_g up to a constant. To pin them down fully, we assume $\theta_{\text{Maize}} = 0.4$, drawing on the land share for maize estimated in Sotelo (2020). Then, once we know σ we can solve for p_g, and we obtain closed-form expressions for the state-specific productivities. Thus, the calibration of the production parameters can be done without solving the full model, and does not depend on the risk aversion parameter γ.

4.3 Elasticity of Substitution Across Crops (σ)

For the elasticity of substitution across crops, σ, we use an estimate from Sotelo (2020). Using 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.

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.

4.4 Insurance Servicing Costs (τ) and Markup (μ)

On the supply side, the price of insurance will be the sum of three terms: the probability of the bad state $P\left(s^L\right)$, 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:

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

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

To focus on the competitive benchmark, we calibrate $\mu^{\text{Baseline}} = 0$. This allows us to focus on the pecuniary externality as the rationale for insurance subsidy, which is the novel contribution of our paper. Of course, 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 vast 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.[11]

[11] We are grateful to Pace Phillips at Innovation for Poverty Action (IPA) for helping us better understand these servicing costs.

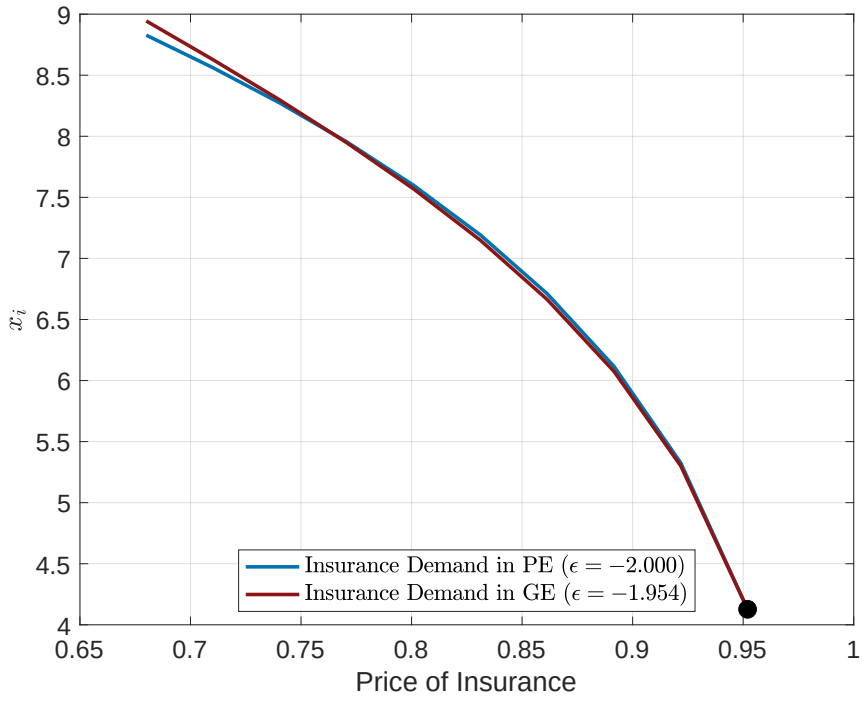


Figure 1: Insurance Demand in Partial vs. General Equilibrium

Calibrating the baseline markup μ^{Baseline} to zero, we obtain servicing costs equal to $\tau = 0.4 \times P\left(s^L\right) = 0.272$.

5 Results and 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.

We begin by showing how insurance demand changes in response to the subsidy in general equilibrium. The results are shown in Figure 1. The black dot shows the price of insurance and insurance demand in the baseline equilibrium where the subsidy is zero. The red line traces out what happens to the price of insurance and the quantity demanded as the subsidy rises, until it reaches the actuarially fair price. For comparison the blue line shows the partial equilibrium demand curve for insurance, in the baseline equilibrium. What distinguishes these two lines is what is happening in the background to other prices and to insurer profits. The partial equilibrium curve holds fixed insurer profits, Π^I, and crop prices, $\{p_g\}_{g=1}^G$, 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

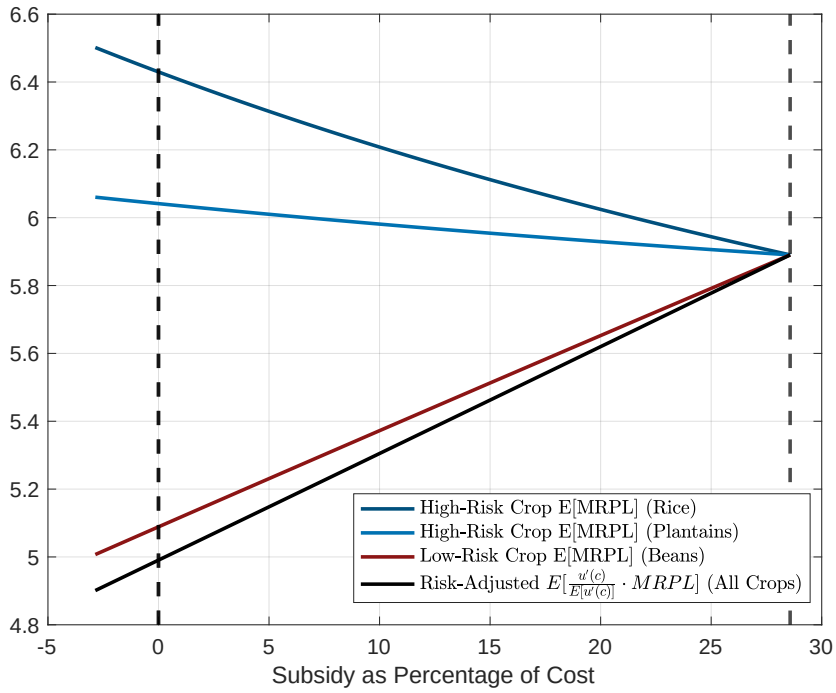


Figure 2: Marginal Revenue Product of Land as a Function of Subsidy

and reports the equilibrium insurance demand with all prices changing.

Although in principle the response of insurance demand in general equilibrium versus partial equilibrium are distinct objects, in practice the two curves are extremely similar. We calibrate to a partial-equilibrium elasticity of insurance demand of -2.000; the general-equilibrium elasticity is -1.954.[12] We view this as reassuring. 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 almost identical to the partial-equilibrium elasticity confirms that there is a tight 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.

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 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

[12]Both of these elasticities are computed going from the market price to the actuarially fair price.

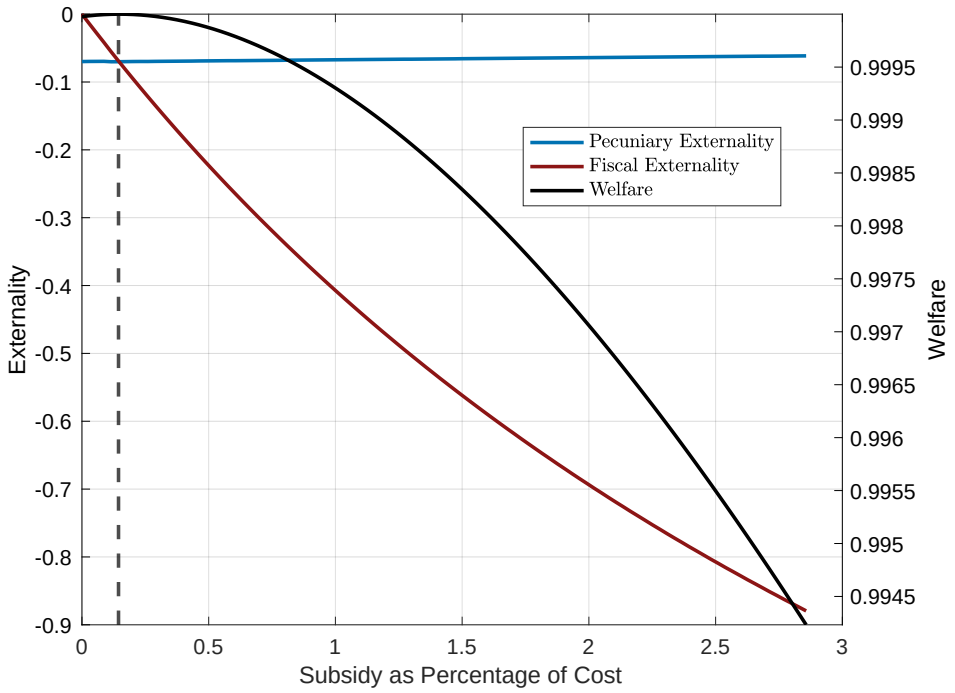


Figure 3: Fiscal vs. Pecuniary Externalities and the Optimal Subsidy

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 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. Under full insurance, farmers implement the same allocation of land as they would in the first-best.

We now use our calibrated model to quantify the fiscal and pecuniary externalities and to compute the optimal subsidy. Figure 3 plots these externalities, 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

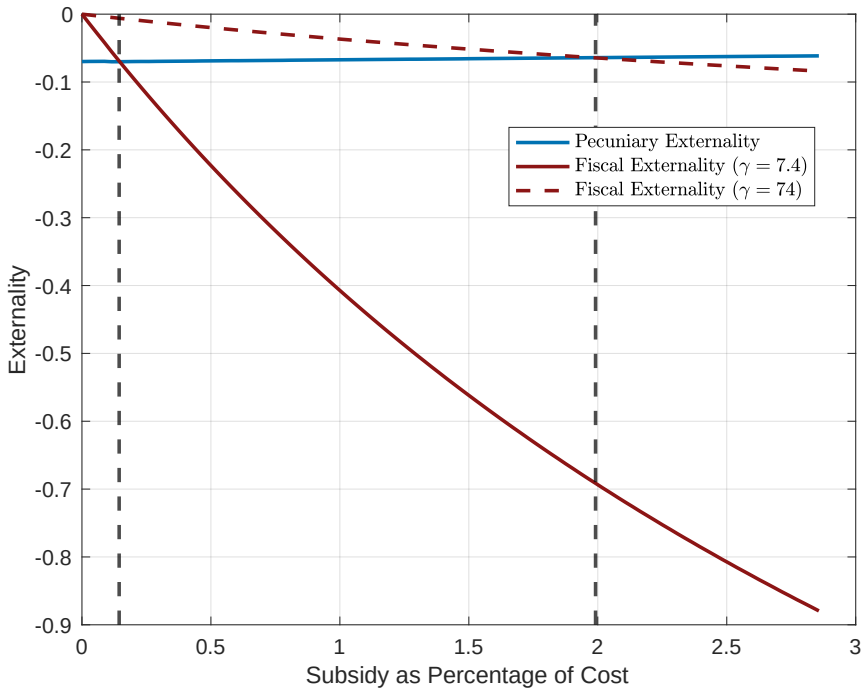


Figure 4: The Optimal Subsidy Increases with Risk Aversion

fiscal externality is negative: an increase in the subsidy induces more demand, which drains the government budget.

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 in blue. 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.

Household welfare is plotted on the figure in black. Welfare is maximized at the subsidy level where the fiscal and pecuniary externalities cross. The dashed vertical line marks the optimal subsidy.

In our calibrated model, the optimal subsidy is small, at only 0.1% of the cost of insurance.[13] The small subsidy reflects the fact that the pecuniary externality is small, while the fiscal externality quickly becomes substantial. Thus, the subsidy at which the two externalities are equalized is small.

To better understand what drives the small optimal subsidy, we can examine what hap-

[13]We express the subsidy as a percentage of the cost of insurance. So, $\frac{\text{Optimal Subsidy}}{P(s^L)+\tau} = 0.1\%$.

pens as we vary the household's coefficient of relative risk aversion, γ. In Figure 4, we plot the fiscal and pecuniary externalities for our baseline calibration, $\gamma = 7.4$, and for a dramatically increased level of risk aversion, $\gamma = 74$. The pecuniary externality, plotted in blue, is unaffected by the level of risk aversion. As shown in Equation 27, the separation result implies that as long as insurance demand is positive, 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 baseline calibration is plotted with a solid red line, while the fiscal externality under the $\gamma = 74$ calibration 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 to 2% 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 only 2%. 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.

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 two forces: a fiscal externality, which reflects the effect of subsidy on the government budget, and a pecuniary externality, which reflects how changes in crop prices transfer resources from the good state to the bad state. This result implies that the optimal subsidy is positive, but that it will never be large enough to implement full insurance. We calibrate the model to the agricultural economy in Ghana using experimental data from Karlan 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, at only 0.1% of the cost of insurance. The small optimal subsidy reflects the fact that the pecuniary externality we measure in the data is modest, while the fiscal externality grows rapidly because insurance demand is fairly elastic. A small optimal subsidy 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, 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. 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 of the policy.

References

Bergquist, L. F., Faber, B., Fally, T., Hoelzlein, M., Miguel, E., and Rodríguez-Clare, A. (2022). Scaling agricultural policy interventions. Technical report, National Bureau of Economic Research.

Brooks, W. and Donovan, K. (2020). Eliminating uncertainty in market access: The impact of new bridges in rural nicaragua. *Econometrica*, 88(5):1965–1997.

Buera, F. J., Kaboski, J. P., and Townsend, R. M. (2023). From micro to macro development. *Journal of Economic Literature*, 61(2):471–503.

Carter, M., De Janvry, A., Sadoulet, E., and Sarris, A. (2017). Index insurance for developing country agriculture: a reassessment. *Annual Review of Resource Economics*, 9:421–438.

Donovan, K. (2021). The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity. *The Review of Economic Studies*, 88(5):2275–2307.

Ghose, D., Fraga, E., and Fernandes, A. (2023). Fertilizer import bans, agricultural exports, and welfare: Evidence from sri lanka. *Policy Research Working Paper*, 10642:793–831.

Kaboski, J. P. and Townsend, R. M. (2011). A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica*, 79(5):1357–1406.

Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics*, 129(2):597–652.

Lagakos, D., Mobarak, A. M., and Waugh, M. E. (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.

Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in thailand. *The American Economic Review*, pages 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*, pages 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.