Perceptions and Participation: Research Design with Low Program Enrollment and Heterogeneous Impacts in Development.

Conner Mullally
PhD Candidate
Department of Agricultural and Resource Economics
University of California, Davis
mullally@primal.ucdavis.edu

Abstract: Recent emphasis in applied development economics has been on evaluating complex financial market interventions, such as microfinance programs, savings mechanisms, and innovative insurance products. However, these programs are often plagued by low participation, making the task of impact evaluation more difficult due to imprecision of estimates. Furthermore, the program impacts identified by different designs will vary when the effects of participation are heterogeneous. These two factors can combine to leave researchers with difficult choices. This paper is a methodological exploration of research design in the presence of low participation and program effect heterogeneity. To put this issue in context, I use the example of index insurance, an innovative financial tool characterized by low participation rates. I focus on the choice between a research design based on randomized eligibility and a randomized encouragement design. Randomized encouragement designs offer a stronger incentive for program participation to a randomly chosen subpopulation, and then use the incentive as an instrumental variable in econometric impact evaluation. When impacts are heterogeneous, the effect estimated by a randomized encouragement design will be biased relative to the impact of a program on participants. However, the possibility of greatly enhanced precision in estimation means that randomized encouragement designs may yield estimates closer to the truth than a randomization of eligibility in a given sample. In addition, greater unobserved heterogeneity will not necessarily increase this bias, contrary to intuition. These conclusions depend on the nature of the program and outcomes being studies, and ought to be considered carefully by researchers weighing alternative research designs.

1 Introduction

Direct randomization into treatment and control groups is a powerful tool for the evaluation of development programs, as it will identify the distribution of treated and untreated outcomes under relatively weak assumptions (Duflo, Glennerster and Kremer 2006). However, randomizing *treatments* in this way is ruled out in many types of programs where participation must be voluntary. For example, in the context of financial market interventions in developing countries, households usually cannot be compelled to take out a loan from a micro-lender or be forced to purchase insurance.

Two alternative research designs are a randomization of *eligibility* and a randomized *encouragement*. The former entails randomly assigning program eligibility but allowing for voluntary participation, while the latter consists of allowing all households to be eligible but offering a randomly chosen group an extra incentive for participation that raises enrollment; in both cases, it is assumed that randomization has no direct effect on the outcome. A randomization of eligibility will identify the average impact of a program on participants, while the randomized encouragement will identify the average impact on participants induced to join the program by receiving the extra incentive (Imbens and Angrist 1994).

When treatment effects are heterogeneous, these two effects will not usually coincide (Heckman and Vytlacil 2007). This would seem to suggest that randomizing eligibility is the superior strategy. But when participation rates are low among eligible households not offered any extra incentive for enrollment, the identifying power of the randomization of eligibility can become very weak, and it will no longer be clear which strategy will yield the estimated effect that is closer to the truth on average; this average distance from the true effect can be captured by the Mean Square Error of each estimator.

This essay is a methodological study of the tradeoffs faced by development economists when choosing a research design for impact evaluation given low program participation rates and heterogeneous effects, two characteristics that often describe interventions in financial markets in poor countries (McKenzie 2009). Poorly functioning financial markets are thought to be an important determinant of poverty, as market failures drive households to use costly alternative strategies to manage risk and maintain smooth consumption levels (Rosenzweig and Binswanger 1993; Carter, et al. 2007; Dercon 2004). Identifying effective financial market interventions is thus a worthy policy goal, but can be made exceedingly difficult by low program enrollment rates.

Low participation in financial programs might be accepted as a fact of life if it were the result of fully informed households optimizing their behavior. Researchers in this case would have to use extremely large sample sizes for econometric analysis. But the literature on financial market participation in the rich world provides evidence for the influence of factors such as trust in institutions and financial education on household decision making in this context (Guiso, Sapienza and Zingales 2008; Guiso and Jappelli 2009). These variables are likely to be even more prominent in less developed areas where households have little experience with formal finance and low levels of education, making it far from clear that a lack of program benefits is what is driving low financial program enrollment.

This suggests that it would be a worthwhile endeavor to increase participation rates in a way that allows for sound econometric analysis of program benefits. This essay uses the example of index insurance to examine the potential of the randomized encouragement design strategy as a means of doing just that. Index insurance is a financial market innovation which has been characterized by low participation rates in developing countries (Hellmuth et al., 2009). In Section 2, I briefly summarize the literature on demand for index insurance. In Section 3, I build a theoretical model in which households are faced with the choice of engaging in a safe activity, a risky high-return activity, or pairing the risky high-return activity with index insurance. To place the model in context, I parameterize and simulate it using data on cotton production in the southern coast of Peru, which has been the site of an index insurance pilot project for the past few years. I then demonstrate what happens when farmers lack trust in the insurer, and believe they will be cheated out of a portion of the indemnity. As a result of this lack of trust, participation falls, and households that could benefit from index insurance do not do so.

In Section 4, I link the theoretical model to econometric program evaluation, and examine the problem of measuring the impacts of index insurance in this model economy from the perspective of an analyst who cannot observe household-level heterogeneity that is driving both activity choice and insurance demand. A randomization of eligibility and a randomized encouragement design are compared in this context, bringing into sharp relief the strengths and weaknesses of each approach. If researchers focus solely on unbiasedness, then randomizing eligibility will be preferred. But when participation rates are low, randomized encouragements may be preferable to randomized eligibility when the choice of research design is based on traditional model selection criteria such as Mean Square Error. I also explore the effects of greater unobserved heterogeneity that is correlated with program uptake and outcomes of interest on the program effects identified by each strategy. While intuition might suggest that greater heterogeneity of this sort would result in larger differences between the program effects identified via each strategy, I demonstrate that this is not necessarily the case. Depending on the characteristics of the program and outcome variables being studied, more heterogeneity of this sort can actually bring the program effect closer to one another in expectation. Section 5 demonstrates these points via a simulation, and Section 6 concludes.

2 Index insurance: an example of an innovative financial market intervention

2.1 Definition of index insurance

As summarized above, households in developing countries may adopt costly risk management strategies. This is of particular concern among rural households, given the risks associated with agriculture. Crop insurance seems like a natural response to this problem, but traditional crop insurance based on compensating farmers for losses relative to a farm-specific historical level of output or income has proven extremely expensive, due to the incentive problems introduced by covering household-specific risks without the ability to perfectly monitor behaviors (Hazell 1992).

High cost makes traditional insurance a less than desirable policy option for governments in developing countries. Index insurance is an alternative to traditional crop insurance that bases payouts on an index that is highly correlated with crop yields but beyond the control of any single household. For example, and index product that bases payouts on rainfall levels would pay households if precipitation levels were to fall below a “strike point,” with the payout increasing in the size of the shortfall. Since it does not cover all risks, index insurance will not offer protection to the same extent as traditional crop insurance, but it avoids the incentive problems which have made the latter costly and fiscally unsustainable.

2.2 Demand for index insurance

The empirical literature to date on index insurance generally reports low levels of demand. While variables such as price and access to liquidity are found to affect demand, other factors outside the scope of a model of perfectly informed optimizing households also influence uptake. Giné, Townsend, and Vickery (2007) found that being a past credit client of the institution offering a rainfall insurance product in the Indian village of Andhra Pradesh has a positive and significant effect on demand; financial sophistication and trust in the institution could both explain this finding. Focusing on the same insurance product studied by Giné, Townsend, and Vickery, Cole et al. (2010) randomized factors that might influence demand for rainfall insurance and measured their effects using a sample drawn from a larger number of villages. Varying the price of the insurance and providing household with enough cash to buy a policy had significant effects on demand in the expected direction, but a household sales visit also had a large and positive effect.

Cai et al. (2009) measured the impact of an insurance product for sows on investment by farmers in China. In order to generate exogenous variation in demand, the authors randomly varied the incentives faced by marketing agents charged with selling the insurance policies. Living in a village served by a marketing agent facing stronger incentives had a strongly positive and significant impact on demand. The authors also found that households which were already receiving government subsidies from other programs were significantly more likely to buy insurance, and argue that receipt of subsidies is a valid proxy for trust in government programs.

Giné and Yang (2009) studied the effect of bundling in-kind loans of fertilizer and high yield maize and groundnut seeds to farmers in Malawi with a rainfall insurance product. After measuring an unexpectedly negative impact of being offered the bundled contract on credit uptake, Giné and Yang separately examined factors affecting credit demand for the subsamples offered the bundled and unbundled loans. Acceptance of the bundled loan was positively correlated with education levels, whereas this was not the case for the unbundled contract, suggesting that the cognitive burden of evaluating the complex insurance contract depressed demand.

The literature on demand for index insurance reinforces the conclusions of studies of participation in financial markets in the rich world: factors such as trust that might not be part of a typical model of the program participation decision are important determinants of enrollment. In the next section, I model the decision to participate in an index insurance program, first assuming that households are perfectly informed and then assuming that households believe the insurance company will systematically underreport shortfalls of the insured risk from the strike point.

3 A model of demand for index insurance and beliefs about the insured risk

3.1 Model structure: putting the problem in context

The model consists of farmers choosing between two activities: planting cotton on a single hectare of land, or planting a subsistence crop on the same single hectare. The subsistence crop guarantees a risk free return of w. Cotton is a risky activity, as output in each period is vulnerable to a covariate shock, ε^c, that is common to all households, and a households-specific shock, ε^s. In this model, I define the common shock at time t, ε_t^c, as the deviation of average output per hectare, or area-yields, from its mean:

$$\varepsilon_t^c = \mu - \overline{q}_t \tag{1}$$

where $\overline{q}_t$ is area-yields at time *t* and μ is the mean of $\overline{q}_t$. The price of cotton is fixed at unity, making output and revenue identical. The cotton yield farmer *i* at time *t* is:

$$q_{it} = \mu_i - \beta_i\left(\mu - \overline{q}_t\right) - \varepsilon_t^s = \mu_i - \beta_i\varepsilon_t^c - \varepsilon_t^s \qquad (2)$$

This production function is adapted from Miranda (1991). It states that in each period a farmer planting cotton receives his or her mean yield, μ, net of any covariate or idiosyncratic shocks, i.e., ε_t^c and ε_t^s, respectively . Note that the shocks can be harmful or beneficial, depending on the signs of ε_t^c and ε_t^s. The shocks ε_t^c and ε_t^s are assumed to be continuous, jointly independent, and symmetrically distributed about a mean of zero. Both terms are identically distributed across farmers, although for the present purpose this need not be the case for ε^s.

The effect of the covariate shock is multiplied by β_i, a parameter whose value varies across farmers. Thus while the magnitude of ε_t^c is the same for all households in a given period *t*, its impact will depend on the value of β_i for a given farm. It is assumed that β_i is weakly greater than zero for all farmers. Taken together, the individual β_i parameters form the population level distribution of β, which is symmetric and centered at unity with a variance of σ_β^2, but truncated at zero due to the fact that all β_i are positive.

To make sense of the β_i parameter, one could imagine that β_i is related to physical characteristics of the farm that would make it more or less vulnerable to covariate shocks. For example, in an area where agriculture is dependent upon irrigation, β_i might capture distance from the farms to the primary irrigation canal. Farms located close to the canal would always have sufficient water, and therefore might have low β_i values. Farms located further from the canal would be more sensitive to water availability, and have higher β_i values as a result.[1]

For simplicity, it is assumed that the mean level of output μ_i is the same across households, so that $\mu_i = \mu$ for all *i*. This is not a realistic assumption, but adding additional heterogeneity would distract from β_i. Differences in β_i across farmers will drive variation in the potential benefits to insuring against the covariate shock ε^c, and this is the focus of the model.

From equation (2), the variance of yield for a single farmer can be written as:

$$\sigma_{q_i}^2 = \beta_i^2 \sigma_c^2 + \sigma_s^2 \tag{3}$$

Yield variance is thus the sum of a component due to variation in area-yields, $\beta_i^2 \sigma_c^2$, and another due to all other sources of risk, σ_s^2. Differences in the variance of yield across farmers are driven by heterogeneity in sensitivity to the covariate shock as captured by β_i.

3.2 The planting decision without index insurance

A farmer's planting decision will be driven his ex-ante evaluation of the benefits of planting cotton versus those of the subsistence crop. Denote by $plant_i = 1$ the decision to grown cotton by farmer *i*, and $plant_i = 1$ planting the subsistence crop. The expected utility of planting cotton is:

$$EU_{plant_i=1} = \mu^2 - \gamma\left(\beta_i^2 \sigma_c^2 + \sigma_s^2\right) \tag{4}$$

and for the subsistence crop:

---

[1] In the long run, the value of β_i for each farmer would likely be a choice variable, as pointed out by Chambers and Quiggin (2002). I assume this parameter is fixed.

$$EU_{plant_i=0} = w^2 \tag{5}$$

This mean-variance expected utility function is taken from Nelson and Escalante (2004). Squaring the mean implies that farmer preferences are characterized by constant relative risk aversion, where the coefficient of relative risk aversion is given by γ .

Farmers will plant cotton if $EU_{plant_i=1} - EU_{plant_i=0} > 0$. Assuming that $\mu > w$, differences in planting decisions across farmers will be determined by heterogeneity in β_i . Specifically, farmers with high values of β_i will deem cotton production as too risky, as their output is highly sensitive to covariate risk. Setting the difference between (4) and (5) to zero and rearranging yields a critical value on β_i , beyond which farmers plant the subsistence crop:

$$EU_{plant_i=1} - EU_{plant_i=0} > 0 \leftrightarrow \beta_i < \sqrt{\frac{\mu^2 - \gamma\sigma_s^2 - w^2}{\gamma\sigma_c^2}} \equiv \beta_0 \tag{6}$$

3.3 The area-yield insurance contract

Other things being equal, farmers would be more inclined to move out of subsistence crop production if the risk associated with growing cotton could be reduced, and any such shift would increase expected output in the economy. Introducing area-yield insurance (AYI) is one way this might be accomplished. I assume that the AYI contract has an indemnity function, I_t , which takes the following form:

$$I_t = \max\left[0, \mu - \overline{q}_t\right] = \max\left[0, \varepsilon_t^c\right] \tag{7}$$

The indemnity pays farmers the difference between the mean of area-yields, μ , and area-yields at time $t, \overline{q}_t$, when this difference is positive, and nothing otherwise. To purchase AYI, farmers must pay the premium, τ , which is equal to the expected indemnity, *r*, plus a “loading” term, *L*:

$$\tau = E(I) + L = E\left(\varepsilon^c \mid \varepsilon^c > 0\right)P\left(\varepsilon^c > 0\right) + L = r + L \tag{8}$$

Since it includes a loading term, the AYI contract is not actuarially fair.[2]

The expected utility of insurance purchasers will be affected through the expected value of the indemnity, r, as well as the variance of the indemnity, σ_I^2, and its covariance with the common shock, $\sigma_{c,I}$. The variance of the indemnity is:

$$\sigma_I^2 = \left(\sigma_c^2 \mid \varepsilon^c > 0\right)P\left(\varepsilon^c > 0\right) + r^2 \tag{9}$$

The covariance of the indemnity with ε^c is:

$$\sigma_{c,I}^2 = \left[\left(\sigma_c^2 \mid \varepsilon^c > 0\right) + E\left(\varepsilon^c \mid \varepsilon^c > 0\right)\right]P\left(\varepsilon^c > 0\right) \tag{10}$$

Equations (9) and (10) are derived in the appendix. Along with the value of β_i, the covariance term determines the risk reduction potential of AYI for each individual farmer.

3.4 The decision to buy area-yield insurance and its effect on activity choice

Letting $AYI_i = 1$ denote the decision to buy AYI and $AYI_i = 0$ opting not to do so, expected utility conditional on planting cotton and buying AYI is:

$$EU_{plant_i=1, AYI_i=1} = (\mu - L)^2 - \gamma\left(\beta_i^2\sigma_c^2 + \sigma_s^2 + \sigma_I^2 - 2\beta_i\sigma_{c,I}\right) \tag{11}$$

Taking the difference between the expected utility of cotton with AYI and the expected utility of uninsured cotton yields an additional critical value on β:

$$EU_{plant_i=1, AYI_i=1} - EU_{plant_i=1, AYI_i=0} \leftrightarrow \beta_i > \frac{(2\mu - L)L + \gamma\sigma_I^2}{2\gamma\sigma_{c,I}} \equiv \beta_1 \tag{12}$$

---

[2] "Loading" is what is added by insurance companies to cover costs of offering insurance in addition to making indemnity payments.

Farmers that always plant cotton will purchase AYI if their β_i values are greater than or equal to β_1. Smaller values of β_i imply that yields are not sufficiently sensitive to the common shock to make purchasing AYI worthwhile.

Introducing AYI into the economy may induce some farmers to switch from planting the subsistence crop to farming cotton with insurance. Comparing the difference between the expected utility of cotton with AYI and the subsistence crop give us two new critical values on β_i :

$$EU_{plant_i=1,AYI_i=1} - EU_{plant_i=0} > 0 \leftrightarrow \beta_2 \equiv \frac{\sigma_{c,I} - x}{\sigma_c^2} < \beta_i < \frac{\sigma_{c,I} + x}{\sigma_c^2} \equiv \beta_3$$
$$\text{where } x = \sqrt{\gamma\left(\sigma_{c,I}\right)^2 + \sigma_c^2\left[\left(\mu - L\right)^2 - w^2 - \gamma\left(\sigma_s^2 + \sigma_I^2\right)\right]} \tag{13}$$

Purchasing AYI results in disutility from the variance of the indemnity, σ_I^2 , and the decrease in expected returns caused by *L*. In order to switch from the subsistence crop to cotton with AYI, farmers must have values of β_i at least as large as β_2 , as this will guarantee a sufficient gain in expected utility via risk reduction to offset these other changes. Values of β_i greater than β_3 result in sensitivity to the covariate risk that is too large to be offset through the purchase AYI; farmers with very high values of the β_i parameter will therefore continue to plant the subsistence crop. These relationships between the insurance decision, activity choice, and the value of β_i for each farmer are depicted in Figure 1, which graphs the expected utility of the options discussed above as function of β_i under three different values of the loading parameter, *L*:

Figure 1: Expected Utility of Subsistence Farming, Cotton, and AYI

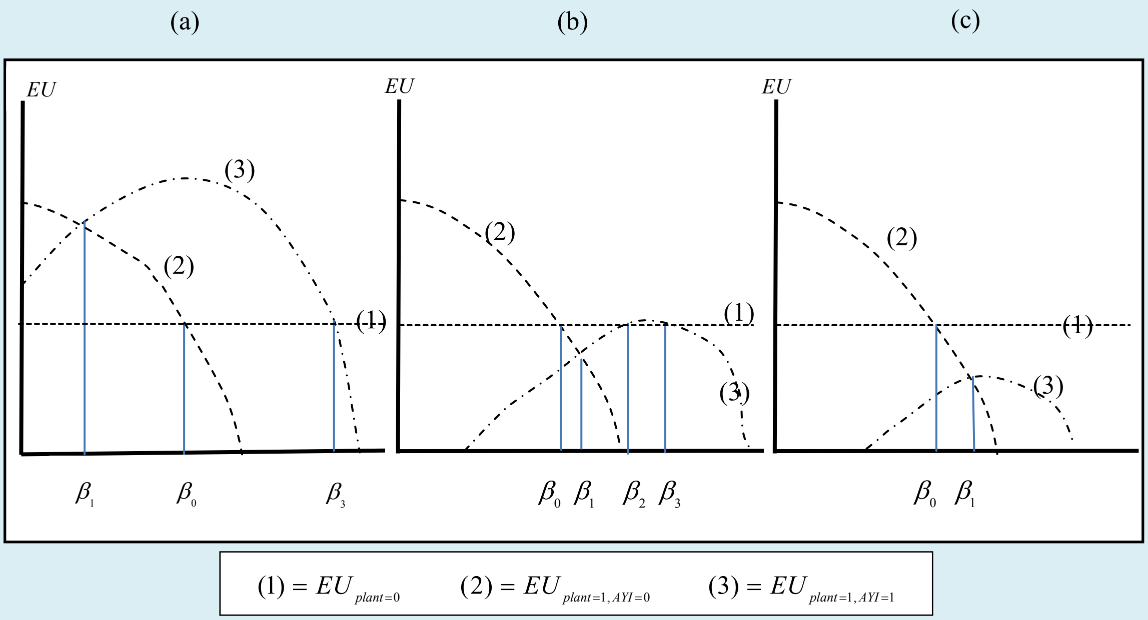

In all three panels, farmers with β_i values to the left of β_0 plant cotton prior to and following the introduction of index insurance, whereas those to the right of β_0 plant the subsistence crop in the absence of AYI. What happens following the introduction of insurance depends on the position of $EU_{plant_i=1,AYI_i=1}$, which is determined by the parameters of the AYI contract; it is only in panels (a) and (b) that introducing AYI has an impact on the proportion of farmers planting cotton.

In panel (a), there is a positive level of demand for AYI, as $EU_{plant_i=1,AYI_i=1}$ is greater than the expected utility of the other alternatives for farmers with β_i values between β_1 and β_3. Those purchasing insurance are a mix of farmers who plant cotton with or without AYI, i.e., those with β_i values between β_0 and β_1, and farmers who are induced to plant cotton by purchasing AYI, i.e., those between β_0 and β_3. The impact of purchasing AYI on activity choice, as measured by the change in the proportion of farmers planting cotton among those that buy AYI, will be somewhat diluted by the fact that some AYI buyers plant cotton no matter what.

In panel (b), L is larger, shifting $EU_{plant_i=1, AYI_i=1}$ down. Now farmers between β_2 and β_3 purchase AYI, and all of these farmers were previously engaged in the subsistence activity. The impact in this case is a complete shift in activity choice among AYI purchasers, as all are induced to shift from the safe activity to the risky activity of cotton farming. In panel (c), the loading term L is large enough such that no one purchases AYI. The expected utility of cotton with insurance is greater than the expected utility of uninsured cotton within a certain range of very high risk β_i values, but farmers within that range are better off planting the subsistence crop.

Stated more formally, farmers purchasing AYI must have values of β_i satisfying:

$$AYI_i = 1 \leftrightarrow \beta_3 \text{ exists and} \max\left[\beta_1, \beta_2\right] < \beta_i < \beta_3 \tag{14}$$

Panel (c) corresponds to the case in which β_2 and β_3 do not exist because the expected utility of cotton with AYI is everywhere below that of the subsistence crop, while β_2 does not exist in panel (a) because of the truncation of the β distribution. The β_i values of farmers planting cotton following the introduction of AYI must obey:

$$\beta_i < \beta_0 \text{ or} \max\left[\beta_1, \beta_2\right] < \beta_i < \beta_3 \tag{15}$$

The proportion of farmers falling within these different bounds on β_i will depend on the distribution of β within the population. Figure 2 below depicts this distribution and the critical values of β_i for the scenario depicted in panel (a) of Figure 1; i.e., non-zero demand for AYI, with a mix of farmers who always plant cotton and farmers induced to adopt cotton by AYI electing to buy index insurance:

Figure 2: Bounds on β_i Before and After AYI

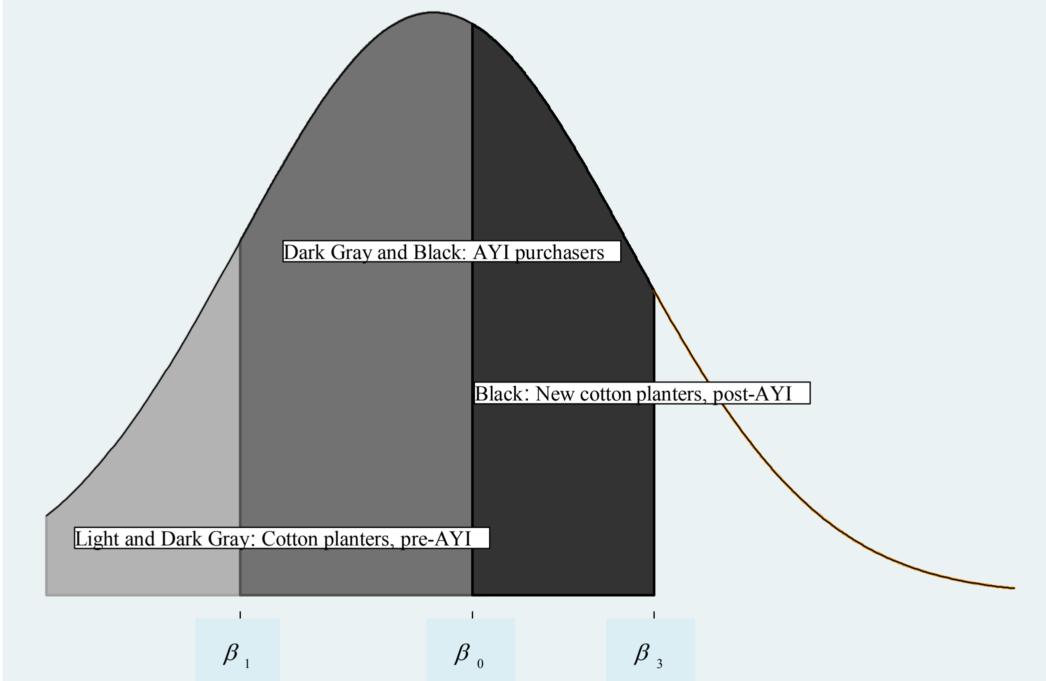

Figure 2 depicts β as a normally distributed random variable with a mean of 1 and a variance of 0.25, truncated at zero. Note that β_2 is negative in this case, and as a result is not depicted. Farmers planting cotton prior to the introduction of AYI have β_i values to the left of β_0. Once AYI is introduced, this upper bound on β_i for cotton planters is expanded to include farmers with β_i values between β_0 and β_3. Insurance purchasers whose activity choices are unaffected have β_i values between β_1 and β_0. Given the positions of the different critical values of β_i, the spread of the β distribution, σ_β^2, will determine the proportion of farmers who purchase AYI and the share of farmers that are induced to switch to cotton farming following the introduction of index insurance.

3.5 Lack of trust in the insurer and demand for area-yield insurance

If farmers know the different mean, variance, and covariance terms detailed above, then any farmer who is better off with AYI will purchase it. Suppose, however, that farmers believe the insurance company will "cheat" them in the sense that they will always underreport the size of the shortfall of area-yields from μ by some fixed proportion of μ, which we will label *g*. It is assumed for simplicity that *g* is identical across all farmers. For example, the insurance company would likely carry out a survey of yields in the region where the AYI product is sold at the end of each harvest, in order to determine whether a payout has been triggered. Rather than taking them at their word, farmers believe that the insurer always adds $g\mu$ to the measured area-yields; this is the most the insurers are believed to be able to "get away with" without drawing attention from regulatory bodies.

The result of this perception is that farmers will no longer view the indemnity function as being based on the distribution of $\overline{q}_t$, but on $\tilde{\overline{q}}_t$, where the latter is equal to:

$$\tilde{\overline{q}}_t = \overline{q}_t + g\mu \tag{16}$$

Farmer perceptions, whether correct or not, cause a location shift in the distribution of area-yields. The expected value of this subjective distribution of area-yields is:

$$E\left[\tilde{\overline{q}}_t\right] = E\left[\overline{q}_t + g\mu\right] = \tilde{\mu} \tag{17}$$

Note that this change in perception does not alter the distribution of the common shock, ε^c:

$$E\left[\tilde{\varepsilon}^c\right] = E\left[\tilde{\mu} - \tilde{\overline{q}}_t\right] = 0 \tag{18}$$

and

$$E\left[\left(\tilde{\varepsilon}^c\right)^2\right] = E\left[\left(\left(\mu + g\mu\right) - \left(\tilde{\overline{q}}_t + g\mu\right)\right)^2\right] = \sigma_c^2 \tag{19}$$

Since the common shock under the perception of a dishonest insurer is equal to the shock under the scenario of no perceived cheating, we can use the latter going forward rather than having to work with a new distribution.

What the perception of cheating does do is alter the definition of the indemnity function as viewed by farmers. The subjective indemnity function is:

$$\tilde{I}_t = \max\left[0, \varepsilon_t^c - g\mu\right] \tag{20}$$

The strike point under this new indemnity function is closer to the left tail of the ε^c distribution than under the true distribution, lowering the subjective probability of receiving an indemnity payout. This in turn leads farmers to perceive a larger difference between the premium and the expected indemnity; in effect, farmers view the contract as having a higher loading term. These impacts of the perception of cheating are summarized graphically in Figure 3:

Figure 3: Perception of Cheating, Area-yields, and the Indemnity

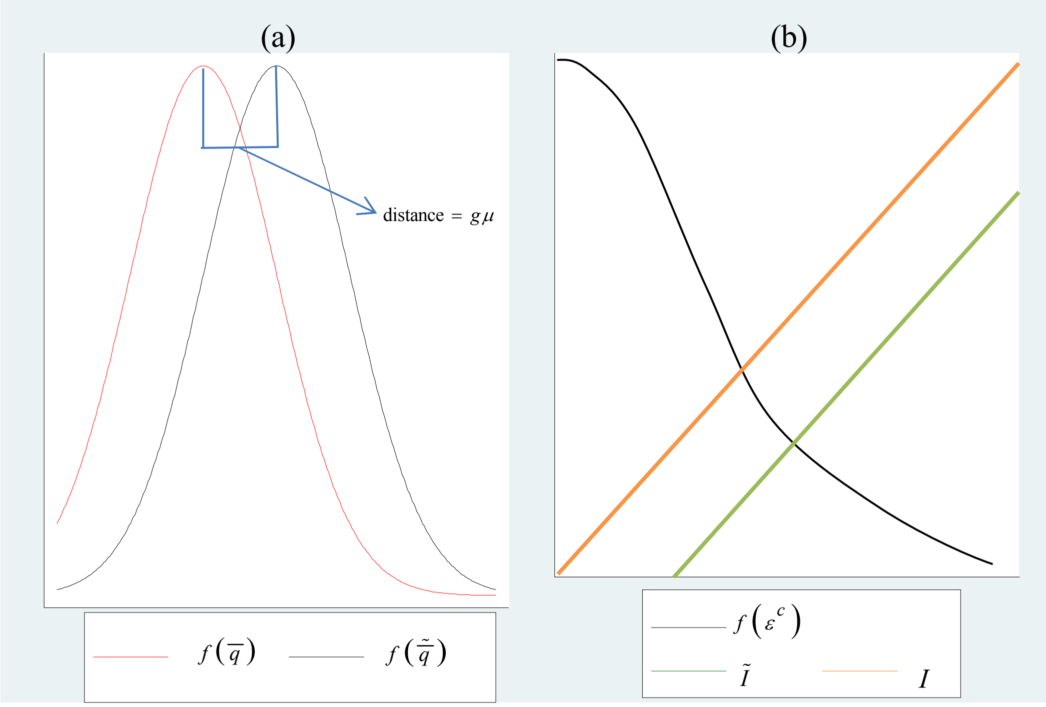

Panel (a) depicts the location shift of area-yields caused by the shared perception among farmers that in the insurer systematically exaggerated area-yields in each time period by an amount $g\mu$. Panel (b) shows the effect on the expected indemnity. Both I_t and $\tilde{I}_t$ are evaluated over the distribution of ε^c, but the latter is shifted downward by the amount $g\mu$ at every level of the common shock that generates a payout.

As a result, if the AYI contract were to be actuarially fair in the eyes of farmers, it would have to have a premium lower than the actuarially fair premium under the perception of no cheating, τ, which was defined above in equation (8). Thus the impact of the perception of cheating on the AYI contract can be interpreted as a perceived increase in the loading term, L, since it increases the perceived difference between the actuarially fair premium and what is charged by the insurer.

While the impact of g on the expected returns to the AYI contract is clear, the perception of cheating does not unequivocally reduce the risk-reduction potential of AYI. As shown in the appendix, this ambiguity stems from the fact that while the higher strike point reduces the covariance of the indemnity and the covariate shock, it may also reduce the variance of indemnity. The net effect of the decrease in expected returns and the change in total variance due to purchasing AYI given the subjective mean of area-yields cannot be signed without choosing explicit parameter values. To clear up this ambiguity and set the stage for linking the analytical model to the discussion of econometric methodology, this model is simulated in the following section using data on cotton farming in the southern coast of Peru. The results support the intuition: a higher g has a strongly negative effect on demand for AYI.

3.6 Assigning parameter values to the model

The model is parameterized using data from the Pisco Valley in Peru, which has been the site of an AYI pilot project for the past several years.[3] The Pisco Valley is home to around 5,000 farmers, approximately half to three quarters of which farm cotton in any given year with maize and beans comprising the other major crops. Farms are small, with an average size of around 5 hectares, and about 20 percent of farmers have access to credit from formal financial institutions. Land is productive in Pisco, but agriculture is subject to idiosyncratic risk in the form of localized flooding from breakdown of irrigation infrastructure, and systemic shocks from El Niño, which can cause flooding and pest outbreaks that can devastate the usually robust cotton plant.

Focus group discussions with farmers in Pisco suggest a strong disconnect between the distribution of area-yields as estimated using data from the Ministry of Agriculture in Peru, which were used to price the contract, and farmer perceptions of area-yields. The causes of this dissonance between the data and farmer beliefs are more heterogeneous than what has been modeled in this essay. For example, some farmers have a low opinion of the Ministry of Agriculture and thus may data see the data as unreliable, while others (particularly highly productive farmers) have a difficult time believing that area-yields will ever drop below the strike point. But in one sense the consequences are the same, as farmers see AYI as too expensive given the likelihood of the insurer reporting a level of $\overline{q}_t$ that would trigger a substantial payout.

---

[3] For additional details on this project, see Boucher and Mullally (2010).

Parameters values used in the simulations are given in Table 1 below:

Table 1: Parameter Values for Economic Model.		
$\mu = 1{,}876$ kg	$\gamma = 3.5$	$\sigma_\beta^2 = 0.25$
$\sigma_c^2 = 147{,}516$ kg^2	$\sigma_I^2 = 50{,}276$ kg^2	$r = 153$ kg
$\sigma_s^2 = 147{,}516$ kg^2	$\sigma_{c,I} = 73{,}576$ kg^2	$L = 23$ kg
$w = 1{,}565$ kg		

The values of $\mu, \sigma_c^2, \sigma_I^2, \sigma_{c,I}, r$ and L were taken from a 20 year time series of cotton yields in the Pisco Valley, provided by the Ministry of Agriculture. The data were de-trended using a linear model. Remaining parameters were chosen by the author with the purpose of generating a pre-AYI level of cotton participation similar to what is typically observed in Pisco.

3.7 Simulation results

Simulation results depicting how perceived cheating by the insurer affects AYI demand and activity choice are shown in Figure 4.

Figure 4: Perceived Cheating and Demand for AYI

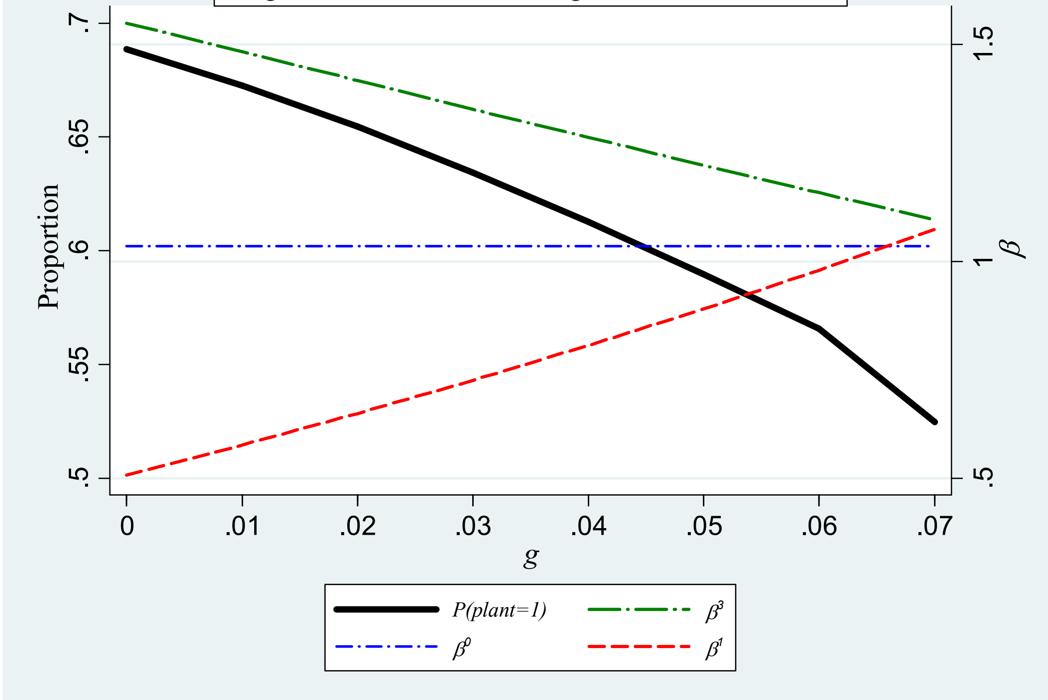

Along the horizontal axis is g, corresponding to the size of the amount insurers are believed to exaggerate area-yields as a proportion of μ. The middle dashed line show the critical value for cotton planters prior to the introduction of AYI, β_0, which is unaffected by perceptions of cheating on the part of the insurer. The uppermost dashed line gives the new upper bound on β_i for cotton planters once AYI is available, β_3, and lower dashed line is the minimum level of β_i for cotton planters, β_1. Farmers induced to plant cotton by AYI will have β_i values between β_0 and β_3, while insurance purchasers whose activity choices are unaffected will fall between β_1 and β_0. The effect of increasing g is to narrow the range of β_i values which fall between both sets of bounds, lowering demand for AYI and the increase in the proportion of farmers who plant cotton following the introduction of AYI. Beyond a value of g equal to 6 percent of μ, demand for AYI falls to zero.

3.8 Model implications

If the insurer is indeed cheating the farmers who purchase AYI in this model economy, then farmers who are dissuaded from buying it under this scenario but would do so if the insurer were honest are behaving optimally. But as described in Section 2, the literature on index insurance and other areas of finance shows that variation in the degree of trust in presumably honest institutions can have a strong effect on demand. This variation is all about perception and may have little to do with the benefits a household can expect to receive from participating in index insurance or any other financial market innovation.

At the same time, little is known empirically about the nature and magnitude of the impacts of many different financial innovations in developing countries. The ideal scenario would therefore appear to be to offer farmers the incentives necessary to experiment with new financial market interventions, allowing them to learn for themselves the benefits of these program, and do so in a way that allows for the measurement of impacts, i.e., the incentive must not directly affect outcomes of interest. Whether households can be convinced to try new programs is not a foregone conclusion. However, the studies summarized in Section 2 have had some success identifying variables that can significantly raise demand for index insurance. It therefore does appear possibly to raise uptake even in programs that have been strongly affected by low participation rates.

Introducing extra incentives for participation will alter the program effect that is ultimately estimated, leaving researchers with difficult choices. Assume that households must make a binary program participation decision. If program impacts vary across households, and the incentive shifts program participation from some positive percentage of households under no extra incentive to a higher proportion, then the resulting estimate will be a Local Average

Treatment Effect (LATE). The LATE will capture the average impact of the program on those that were induced to enroll by the extra incentive (Imbens and Angrist 1994).

4 Increasing demand for index insurance in the context of program evaluation

4.1 Tying the model to econometrics

Suppose we would like to estimate the average impact of AYI on activity choice among insurance purchasers in the economy described in Section 3, and assume that the individual β_i parameters are not observed. In other words, β_i is an unobserved source of heterogeneity across farmers that will affect both the decision to participate in index insurance and whether to engage in the risky activity of cotton farming. A simple comparison of average outcomes among AYI purchasers and non-purchasers will not yield an unbiased estimate of average program impacts on participants, because the two groups may systematically vary according to their respective distributions of β. Estimation of impacts will require an additional, observed source of variation in AYI demand that is uncorrelated with activity choice.

4.2 A randomization of eligibility and a randomized encouragement design

A method that can yield unbiased estimates of average treatment effects is directly randomizing individuals into and out of the index insurance program. This approach may be problematic in the present context for reasons discussed in the introduction to this essay. Alternatively, one could randomize a variable that affects demand for index insurance without affecting the variable to be used as the outcome. Comparisons of the groups formed by this randomization can yield unbiased estimates of average treatment effects. However, the average treatment effect that is estimated may vary depending on the randomization strategy used. This is because the randomization strategy chosen will affect the composition of the comparison groups. Unless impacts are the same for all individuals or vary randomly in a way that is not correlated with the decision to buy insurance, the average outcome in each group will depend on group composition (Heckman and Vytlacil, 2007).

For example, consider a randomization of eligibility to purchase index insurance. Conceptually, this is perhaps the simplest randomization of this sort, although it may be one of the most difficult to implement in practice. At the individual level, randomizing eligibility would require some sort of mechanism by which randomly chosen persons are not allowed to purchase index insurance, while others are left to purchase it as they please. This is what was done by Giné and Yang (2009) in their study of demand for loans bundled with rainfall insurance in Malawi.

Denote by $z_i = 1$ random assignment of farmer i to the eligible group, i.e., those who can purchase AYI at the market price $\tau = (r + L)$, and $z_i = 0$ random assignment to the ineligible group. Suppose that a proportion $\rho_{z_i=1}$ of the eligible group would buy index insurance. Using activity choice $plant_i$ as the outcome, the following average treatment effect could be estimated, using data on activity choice and demand for insurance among the eligible and ineligible groups:

$$\frac{P(plant_i = 1 \mid z_i = 1) - P(plant_i = 1 \mid z_i = 0)}{P(AYI_i = 1 \mid z_i = 1) - P(AYI_i = 1 \mid z_i = 0)} = \frac{P_{z_i=1} - P_{z_i=0}}{\rho_{z_i=1}} \tag{21}$$

$P(AYI_i = 1 \mid z_i = 1) = P_{z_i=1}$ is the proportion of eligible farmers planting cotton, also equal to the expected value of $plant_i$ in the AYI-eligible subpopulation.

Exploiting the fact that z_i is randomly assigned and that it does not directly affect the outcome $plant_i$, it can be shown that the expression given in (21) is equivalent to:

$$P\left(plant_{i,AYI_i=1}=1\mid AYI_i=1\right)-P\left(plant_{i,AYI_i=0}=1\mid AYI_i=1\right) \tag{22}$$

Note that the potential outcomes $plant_{i,AYI_i=1}$ and $plant_{i,AYI_i=0}$ are used in (22), i.e., what farmer *I* would do if he were to purchase AYI, rather than the observed outcome $plant_i$. Correspondingly, $P\left(plant_{i,AYI_i=1}=1\mid AYI_i=1\right)$ is the proportion of farmers who would plant cotton if they were to buy index insurance, conditional on being among the farmers that elect to buy index insurance, and $P\left(plant_{i,AYI_i=0}=1\mid AYI_i=1\right)$ is the proportion of farmers in this same group who would plant cotton without index insurance.

The expression in (22) is the change in the share of farmers planting cotton among the group that elects to buy area-yield insurance when it is made available at the premium τ. In the vocabulary of program evaluation, this is an example of the "Average Treatment on the Treated," or ATT. If index insurance is to be made available to all farmers at this same price, then this particular average treatment effect may be of greatest relevance to policymaking. This is because the group of insurance purchasers under the randomization of eligibility ought to strongly resemble the group of farmers that will buy index insurance when it is made widely available. For this reason this average treatment effect will be referred to in what follows as the "Policy Relevant Treatment Effect," or PRTE, using the terminology of Heckman and Vytlacil (2007).

Now consider a "randomized encouragement design" or allowing all households to purchase AYI but encouraging a randomly selected group to do so by offering them an extra incentive for enrollment. An example of a randomized encouragement design would be a voucher program that reduces the cost of participation in a social program.[4] Here I will use a randomized encouragement design that gives a randomly chosen subset of farmers a discount "coupon" enabling each to pay a lower price for index insurance. Denote by $c_i = 1$ assignment of farmer *i* to the encouraged group, and $c_i = 0$ if farmer *i* is not picked to receive a coupon. Suppose further that a proportion $\rho_{c_i=1}$ of the coupon group participates, while a share $\rho_{c_i=0}$ of farmers without coupons buys index insurance. This randomized encouragement yields the following estimator:

$$\frac{P\left(plant_i = 1 \mid c_i = 1\right) - P\left(plant_i = 1 \mid c_i = 0\right)}{P\left(AYI_i = 1 \mid c_i = 1\right) - P\left(AYI_i = 1 \mid c_i = 0\right)} = \frac{P_{c_i=1} - P_{c_i=0}}{\rho_{c_i=1} - \rho_{c_i=0}} \tag{23}$$

This is a "Local Average Treatment Effect" (LATE).[5] It can be shown that this expression is equivalent to:

$$P\left(plant_{i,AYI_i=1} = 1 \mid AYI_i = 1 \leftrightarrow c_i = 1\right) - P\left(plant_{i,AYI_i=0} = 1 \mid AYI = 1 \leftrightarrow c_i = 1\right) \tag{24}$$

where the expression $\left(AYI_i = 1 \leftrightarrow c_i = 1\right)$ should be read as "buys area-yield insurance if and only if given a coupon." In other words, equation (24) is the average change in activity choice due to having AYI among the group of farmers that would purchase AYI if they were to receive a coupon, but would otherwise not purchase it. This group is known as the "compliers" in the program evaluation literature.

---

[4] McKenzie (2009) offers some examples of encouragement designs in development.

[5] For this to represent an LATE, the encouragement must satisfy the "monotonicity" assumption (Imbens and Angrist 1994). In the case of the coupon, this assumption would require that the lower premium either encourages or has no effect on the AYI purchase decision; it cannot persuade some farmers to buy insurance and dissuade others. One could imagine some encouragement designs where satisfaction of this assumption would not be obvious.

4.3 Choosing between competing research designs: Mean Square Error

The fact that the LATE only captures average impacts on individuals induced to participate by the encouragement is a limitation; if program effects are heterogeneous, then each different possible encouragement could yield a new LATE, making interpretation of these effects difficult (Heckman and Vytlacil 2007). A corollary to this is that the effect estimated by the randomized encouragement design and that captured using the randomization of eligibility will not in general coincide. A randomized encouragement design can generate an unbiased estimate of the average impact of AYI on the group whose insurance purchase decision is determined by the encouragement, but it is a biased estimate of the PRTE when program impacts are heterogeneous.

Two possible responses to this criticism come to mind. The first is that in the context of new programs, changing the pool of participants via a randomized encouragement will allow households to update information with respect to the gains from participation. It seems reasonable to assume that the most effective means of allowing households to decide if they can benefit from financial market interventions is to let them experiment for themselves, or learn from the experiences of others. Randomized encouragement designs, carried out over multiple years, might accomplish this while enabling researchers to determine if the intervention has a beneficial effect on the subpopulation of compliers. This dynamic aspect of randomized encouragements is beyond the scope of this essay but a possible area of future research.

A second argument is based on the Mean Square Error of the estimator based on the randomization of eligibility versus that of the estimator based on a randomized encouragement given very low enrollment rates, where the "error" is the distance of the estimated effect from the true PRTE. The estimator generated by a randomized encouragement design will likely be a biased with respect to the PRTE. But rather than concentrating solely on bias, we might focus on the Mean Square Error of each estimator with respect to the PRTE. Low participation among eligible farmers who have not been given extra incentive to purchase AYI may cause estimates under a randomization of eligibility to be so imprecise that the randomized encouragement design yields the preferred estimator.

Suppose we draw a sample of n farmers, a proportion π of which are assigned to the eligible group. The MSE for the estimate of the estimator based on randomized eligibility is:

$$\frac{\left(1-\pi\right)P_{z_i=1}\left(1-P_{z_i=1}\right)+\pi P_{z_i=1}\left(1-P_{z_i=1}\right)}{n\pi\left(1-\pi\right)}\frac{1}{{\rho_{z_i=1}}^{2}} \tag{25}$$

Since the estimator based on the randomization of eligibility is unbiased with respect to the PRTE, the MSE given in (25) is equal to the variance of the estimator; the formula for the variance follows from the binary nature of $plant_i$.[6] Low program participation affects the MSE through its direct effect on participation rates, and via its effect on the variance of the outcome $plant_i$ among eligible farmers. Consider the derivative of (25) with respect to g:

$$\left[\frac{\partial P_{z_i=1}}{\partial g}\frac{\left(1-\pi\right)\left(1-2P_{z_i=1}\right)\rho_{z_i=1}}{n\pi\left(1-\pi\right){\rho_{z_i=1}}^{3}}\right]-\left[\frac{\partial \rho_{z_i=1}}{\partial g}\frac{2\left[\left(1-\pi\right)P_{z_i=1}\left(1-P_{z_i=1}\right)+\pi P_{z_i=0}\left(1-P_{z_i=0}\right)\right]}{n\pi\left(1-\pi\right){\rho_{z_i=1}}^{3}}\right] \tag{26}$$

Assuming that demand for AYI falls as g increases, the second term of this expression is always positive while the sign of the first term is ambiguous; the latter depends on the proportion of farmers eligible to purchase AYI that elect to plant cotton, i.e., $P_{z_i=1}$. This ambiguity makes it

[6] Although I assume away the possibility here, the estimate of the PRTE based on the randomization of eligibility can exhibit finite sample bias if participation among eligible farmers is sufficiently low (Bound, Jaeger and Baker 1995).

impossible to sign (26). What can be said is that the MSE of the estimator based on a randomization of eligibility is an increasing function of g for all values of $P_{z_i=1}$ greater than 0.25.[7]

Now consider the MSE for the estimator relying on a randomized encouragement design. This will be equal to the square of its bias (the difference between the LATE and the PRTE) and its variance. Note that the proportion of farmers without coupons in the encouragement design who elect to buy AYI is equal in expectation to the share of eligible farmers who purchase AYI in the randomization of eligibility. As a result, the fraction of farmers planting cotton when not given a coupon in the encouragement design is equal in expectation to the proportion of eligible farmers planting cotton in the randomization of eligibility. Using these facts and rearranging terms, the MSE for the estimator based on the randomized encouragement design can be expressed as:

$$\left[\frac{\rho_{z_i=1}\left(P_{c_i=1}-P_{z_i=0}\right)-\rho_{c_i=1}\left(P_{z_i=1}-P_{z_i=0}\right)}{\rho_{z_i=1}\left(\rho_{c_i=1}-\rho_{z_i=1}\right)}\right]^2+ \\ \frac{(1-\omega)P_{c_i=1}\left(1-P_{c_i=1}\right)+\omega P_{z_i=1}\left(1-P_{z_i=1}\right)}{n\omega(1-\omega)}\frac{1}{\left(\rho_{c_i=1}-\rho_{z_i=1}\right)^2} \tag{27}$$

The ω parameter is the proportion of farmers in the sample assigned coupons. The first line is the square of (LATE-PRTE). Here we see that the bias grows as the difference in the proportion of

[7] The derivative is positive if the second term is larger in absolute value than the first. $\partial\rho_{z_i=1}/\partial g$ is weakly greater in absolute value than $\partial P_{z_i=1}/\partial g$; at most 100 percent of the cotton farmers who switch from purchasing AYI to not doing so due to an increase in g also switch from cotton to the subsistence crop. Since $\rho_{z_i=1}$ is less than 1 and $\pi P_{z_i=0}\left(1-P_{z_i=0}\right)$ is positive, $2P_{z_i=1}\left(1-P_{z_i=1}\right)\geq 1-2P_{z_i=1}$ is sufficient for the second term to be greater than the first. This is true for all $P_{z_i=1}$ greater than 0.25.

coupon holders planting cotton increases relative to the proportion of AYI-eligible farmers planting cotton in the randomization of eligibility. This difference is driven by the impact of the incentive offered by the coupon on demand for AYI. The stronger the incentive offered by the encouragement design, the greater the change in the pool of insurance purchasers induced by the encouragement relative to the mix of farmers who would purchase AYI under a randomization of eligibility, and the bigger the absolute difference between the LATE and the PRTE.

The fact that the bias changes with the size of the coupon can be made explicit by examining the derivative of (LATE – PRTE) with respect to the size of the coupon:

$$\frac{\partial P_{c_i=1}}{\partial coupon}\frac{1}{\left(\rho_{c_i=1}-\rho_{z_i=1}\right)}-\frac{\partial \rho_{c_i=1}}{\partial coupon}\frac{P_{c_i=1}-P_{c_i=0}}{\left(\rho_{c_i=1}-\rho_{z_i=1}\right)^2} \tag{28}$$

This is equal to zero if:

$$\frac{\partial P_{c_i=1}}{\partial \rho_{c_i=1}}=\frac{P_{c_i=1}-P_{c_i=0}}{\left(\rho_{c_i=1}-\rho_{z_i=1}\right)} \tag{29}$$

The left-hand side is the marginal change in the share of farmers planting cotton due to a small change in the proportion of farmers with AYI. The right-hand side is the average change in cotton farming due to a change in AYI demand, where the average is evaluated over the range of AYI participation rates falling between $\rho_{c_i=1}$ and $\rho_{z_i=1}$. The two sides of (29) will only coincide if the average impact of purchasing AYI on the outcome is constant over some range of participation rates in the insurance program. In general this will not be the case when impacts of AYI are heterogeneous, and the absolute value of the bias of the LATE as an estimator of the PRTE will grow as the value of the coupon increases.

The second term of the MSE given in (27) is the variance of the estimator based on the randomized encouragement design. The impact of the coupon on the variance of this estimator is the mirror image of the impact of g on the precision of the estimator based on a randomization of eligibility. The derivative of the second term in (27) with respect to the size of the coupon is:

$$\frac{\partial P_{c_i=1}}{\partial coupon}\left[\frac{\omega\left(1-2P_{c_i=1}\right)\left(\rho_{c_i=1}-\rho_{z_i=1}\right)}{n\omega(1-\omega)\left(\rho_{c_i=1}-\rho_{z_i=1}\right)^3}\right]- \frac{\partial \rho_{c_i=1}}{\partial coupon}\left[\frac{2\left[\omega P_{c_i=1}\left(1-P_{c_i=1}\right)+(1-\omega)P_{c_i=0}\left(1-P_{c_i=0}\right)\right]}{n\omega(1-\omega)\left(\rho_{c_i=1}-\rho_{z_i=1}\right)^3}\right] \tag{30}$$

Using reasoning similar to that employed above with respect to the change in the variance of the estimator based on a randomization of eligibility, it can be shown that (30) is positive if $P_{c_i=1}$ is greater than 0.25. A necessary and sufficient condition for this is that at least a quarter of farmers eligible to buy AYI but not given coupons elect to plant cotton, i.e., $P_{c_i=0}=P_{z_i=1}\geq 0.25,$ as farmers with coupons will be at least as likely to purchase AYI and plant cotton as farmers paying the market price for index insurance. The effect of this greater precision on the MSE of the estimator will be tempered by an increase in bias, which will be non-zero when treatment effects are heterogeneous.

Note that the ambiguity of the change in the MSE of each estimator with respect to varying the incentives for participation in the AYI program stems from the fact that the variance of $plant_i$ depends on z_i. This will always be the case with binary outcomes when the instrument has identifying power. If the variance of the outcome of interest did not depend on z_i, which might be the case with many continuous outcome variables, the MSE of the estimator based on randomized eligibility would be strictly increasing with respect to g and the variance of the estimator based on the randomized encouragement design would always decrease with the size of the coupon; changes in precision would only occur through the $\rho_{c_i=1}$ and $\rho_{z_i=1}$ terms. The tradeoff between increased bias and greater precision captured by the MSE of the randomized encouragement design estimator would still be present, however. This suggests that in many applications, the bias-precision tradeoff offered by a randomized encouragement design would exist at all values of the outcome of interest.

5 A randomization of eligibility and a randomized encouragement design: simulating the tradeoffs

5.1 A randomization of eligibility

Using the parameter values given in Table 1, the PRTE and its 95 percent confidence interval as estimated using randomized eligibility are drawn as functions of *g* in Figure 5:

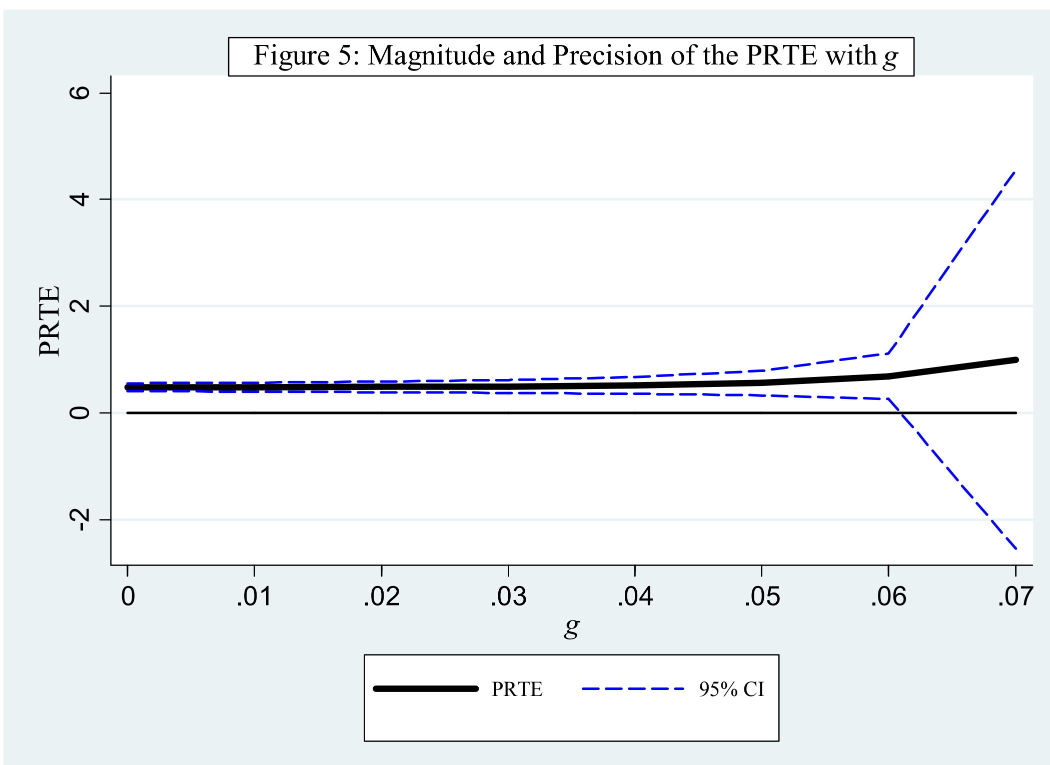

Figure 5 was drawn assuming a sample size of 1,000 farmers, half of which were assigned to the eligible group. Once g passes 7 percent of μ, the estimate of the PRTE becomes statistically insignificant. This is despite the fact that the PRTE is increasing monotonically with *g*. The impact of *g* on precision of estimation limits what can be stated about the effects of AYI on activity choice among insurance purchasers in this case.

The sharp increase in the confidence bounds around the PRTE reflects the fact that the impact of changes in participation rates on the precision of the estimator is nonlinear, and that as participation falls, the variance of the estimator increases at an increasing rate. This is in contrast to the effect of greater sample size. The inverse of the participation rate is squared in the formula for the variance of the estimator given in (25), whereas the inverse of the sample size enters as a linear term.

The PRTE is increasing with respect to *g* in Figure 5, but this is not true in general. The derivative of the estimator of the PRTE with respect to g is:

$$\frac{1}{\rho_{z_i=1}}\left[\frac{\partial P_{z_i=1}}{\partial g} - \text{PRTE}\frac{\partial \rho_{z_i=1}}{\partial g}\right] \tag{31}$$

While the second derivative in brackets will be greater than the first in absolute value, the PRTE is weakly less than one, and the expression cannot be signed. Increasing *g* tightens the upper and lower bounds on β_i for insurance purchasers as shown in Figure 4. Tightening the lower bound increases the PRTE, as this pushes farmers out of the insurance market who would plant cotton regardless of having purchased AYI, whereas tightening the upper bound has the opposite effect. The number of farmers of each type pushed away from buying insurance by a larger *g* will depend on the position of these upper and lower bounds in the β distribution, which in turn will be determined by the values of the parameters in the model.

5.2 A randomized encouragement design

Now consider a randomized coupon scheme. Figure 6 graphs the LATE as calculated using equation (23), its 95 percent confidence interval, and the difference between the PRTE and the LATE as a function of the size of the coupon. The cheating parameter g is fixed at 7 percent of μ:

Figure 6: Precision and Bias of the LATE

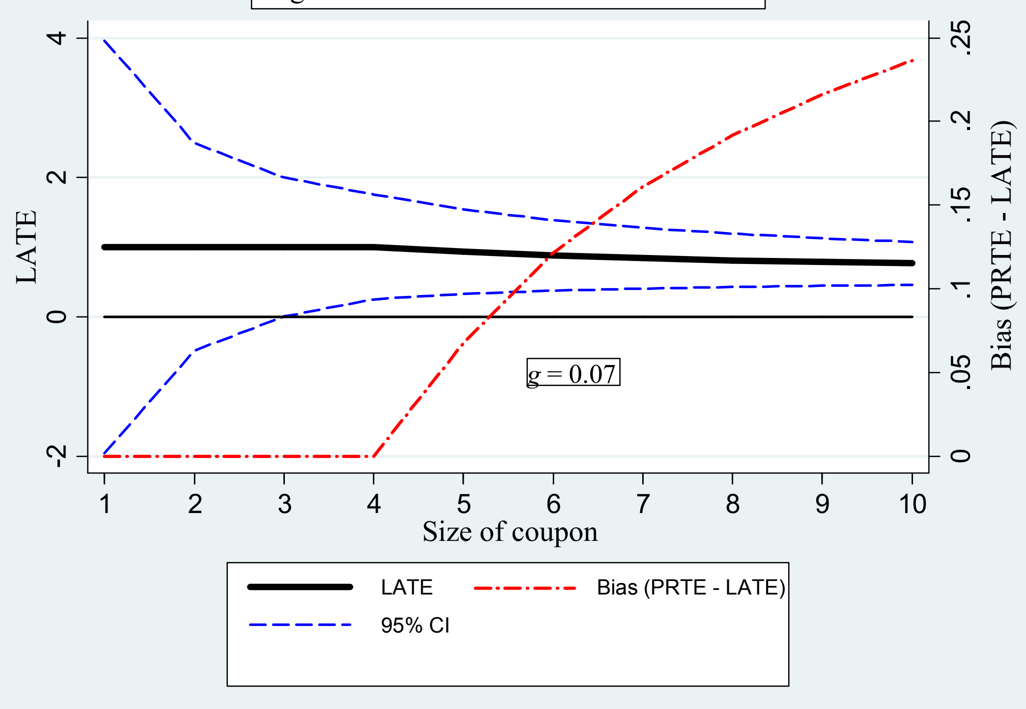

When the coupon reaches a value of 3, or approximately 2 percent of the total premium of 176 (including loading L), the estimated LATE becomes statistically significant. In reality we would expect that a substantially larger coupon would be needed to drive up participation rates, but the mechanics would be similar to those depicted in Figure 6.

Although it is difficult to see, the LATE is decreasing monotonically with the size of the coupon. The mechanics behind this are the mirror image of the impact of the g on the magnitude of the PRTE. Recall that the LATE measures the impact of AYI on compliers, i.e., farmers who are induced to purchase AYI by receiving a coupon. Farmers who switch from not purchasing AYI to doing so by an increase in the size of the coupon all belong to the group of compliers at the new coupon value. If a larger coupon brings in relatively more farmers whose activity choices are unaffected by having insurance, then the LATE will fall. This is what we observe in Figure 6.

Larger coupons also result in a greater difference between the LATE and the PRTE. When g is at 7 percent of μ, the PRTE is 1; increasing the size of g has resulted in the scenario depicted in panel (c) of Figure 1, and all insurance purchasers only plant cotton when purchasing AYI. The LATE shrinks from 1 at a coupon of 1 to 0.763 when the coupon is equal to 10. The bias is initially zero, as additional farmers brought into the insurance market by the coupon only plant cotton when purchasing insurance, but increases quickly with larger coupons.

5.3 Simulating the Mean Square Error

As was stated earlier, a more complete picture of how close one can expect the estimators generated by these two competing research designs to come to the truth is given by the MSE of each. These are depicted below as a function of coupon size in Figure 7:

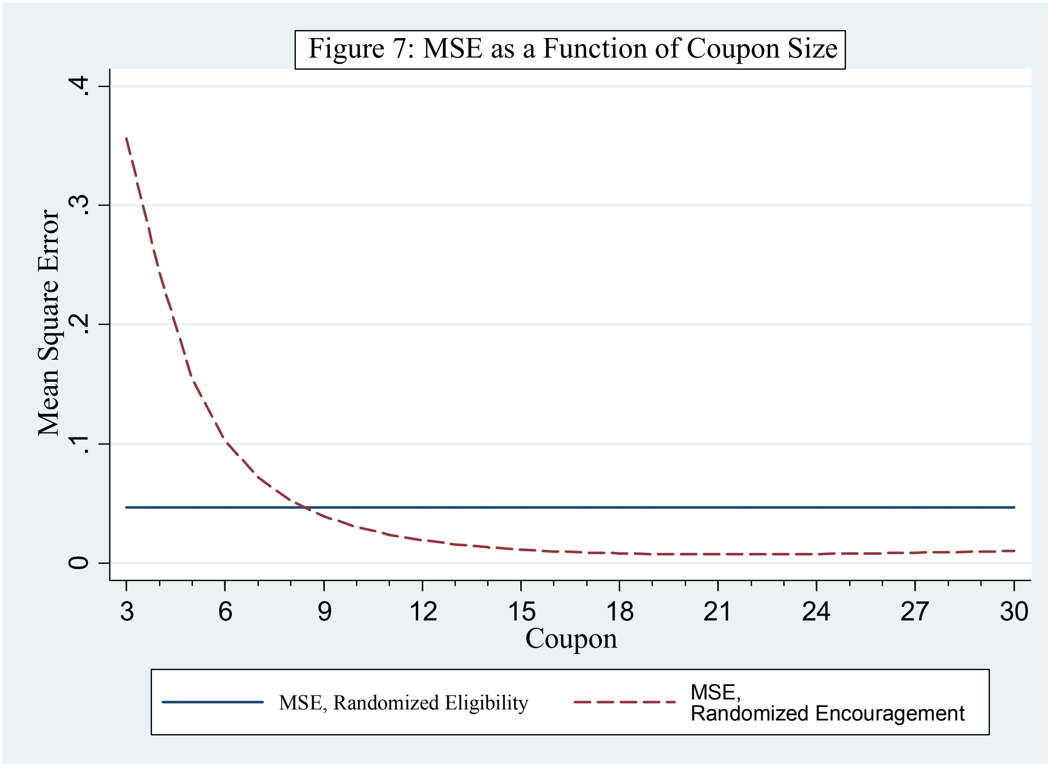

Figure 7: MSE as a Function of Coupon Size

Here g has been fixed at 6 percent of μ, rather than the 7 percent level used in Figure 6. The increase in the variance of the randomization of eligibility estimator was so sharp at $g = 0.07$ that comparing the MSE of the two estimators at $g = 0.06$ seemed more reasonable. The MSE of the estimator based on a randomization of eligibility does not change with the value of the coupon, and it is shown by the flat line between zero and 0.1. At small coupon values, the impact of the encouragement on demand for AYI is very small, and as a result the MSE of the estimator for the encouragement design as depicted by the dashed line in Figure 7 is quite high. Once the coupon reaches a value of 9, the MSE of the encouragement design estimator drops below that of the randomization of eligibility estimator, and continues to fall until the coupon amount reaches 21; this is the size of the encouragement that minimizes the MSE in this case. Beyond this coupon value, the MSE of the estimator based on the randomized encouragement begins to increase, as gains in precision diminish while bias with respect to the PRTE continues to grow.

When the MSE is used as the model choice criterion, a strong enough encouragement can make the randomized encouragement explored here preferable to a randomization of eligibility. All of the above simulation results were generated by holding unobserved heterogeneity as represented by the spread of the β distribution, σ_{β}^{2}, constant at 0.5. A question one might ask is to what extent the degree of unobserved heterogeneity present in the population will determine the potential of a randomized encouragement design to improve upon a randomization of eligibility with respect to MSE.

At first glance, one might expect that increasing σ_{β}^{2} would increase the MSE of the encouragement design relative to that of the randomization of eligibility. If there is more heterogeneity with respect to β, then this could translate into bigger differences between the average treatment effects identified by competing research designs. But whether this is the case is not obvious, as taking derivatives of the bias term in the MSE of the encouragement design estimator with respect to σ_{β}^{2} yields no definitive answer. Intuition suggests that greater heterogeneity could just as easily lead to a reduction in bias. The identifying power of the coupon scheme comes from the expansion of the lower and upper bounds on β_i for insurance purchasers generated by the lower price for AYI. The size of the movement of these bounds does not depend on σ_{β}^{2}, and a narrower spread of the β distribution would result in a coupon of a given size capturing more mass of the β distribution via the expansion of the bounds on β_i for insurance purchasers. Values of β_i among compliers would then be closer to the tails of the distribution of β, making them less like their counterparts who would purchase AYI at the market price. This suggests that less heterogeneity as represented by a smaller value of σ_β^2 could actually increase the bias of the randomized encouragement design estimator.

Impacts of greater heterogeneity on the variance of each estimator are also not straightforward, as analysis of how the precision of each estimator changes with σ_β^2 is made much more complicated by the fact that β follows a truncated normal distribution. I simulate the effect of an increase in σ_β^2 on the MSE of each estimator in Figure 8, and then use the assumptions made about the distribution of β to gain some insight into the changes that are observed.

Figure 8: MSE and its Components as Functions of Spread of β

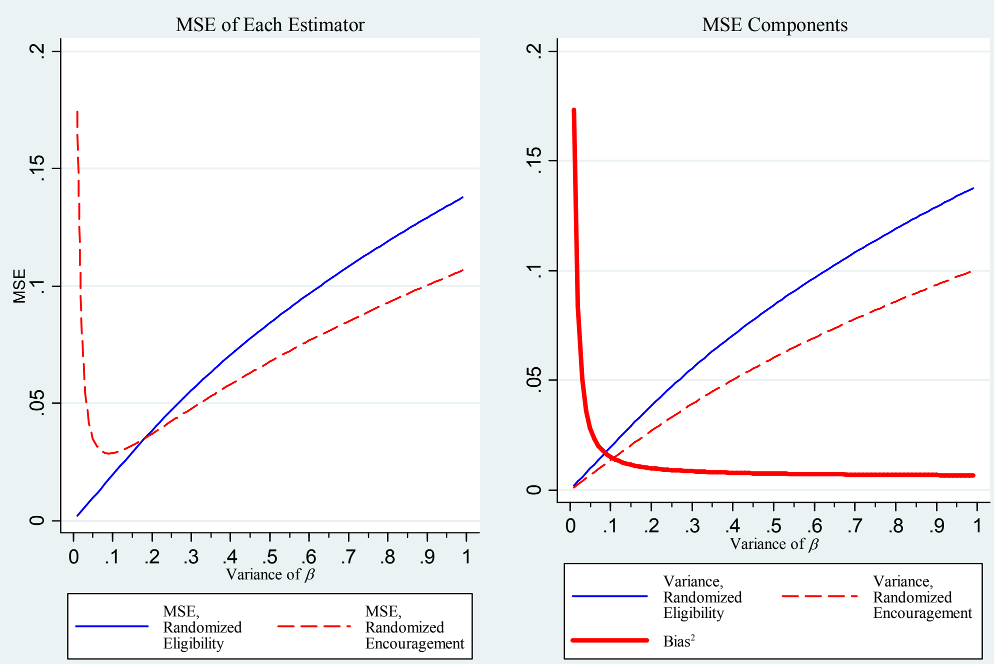

Figure 8 was generated by varying σ_β^2 while holding g fixed at 0.07, the value of the coupon at 9 (the value that drove the MSE of the encouragement design estimator below that of the randomization of eligibility estimator when $\sigma_\beta^2 = 0.25$), and all other parameters fixed at their levels given in Table 1. As shown in the left panel, the MSE of the estimator based on a randomization of eligibility increases steadily while that of the randomized encouragement design is minimized at $\sigma_{\beta}^{2} = 0.09$. While the variance of both estimators grows as the spread of the β distribution increases, the change in the MSE of the randomized encouragement design estimator is offset by a reduction in its bias.

Insight into what is going on in Figure 8 can be gleaned by examining the β distribution given different value of σ_{β}^{2}, along with the bounds on β_i for insurance purchasers under the two different research designs. This is shown below in Figure 9:

Figure 9: σ_{β}^{2} and the Distribution of β

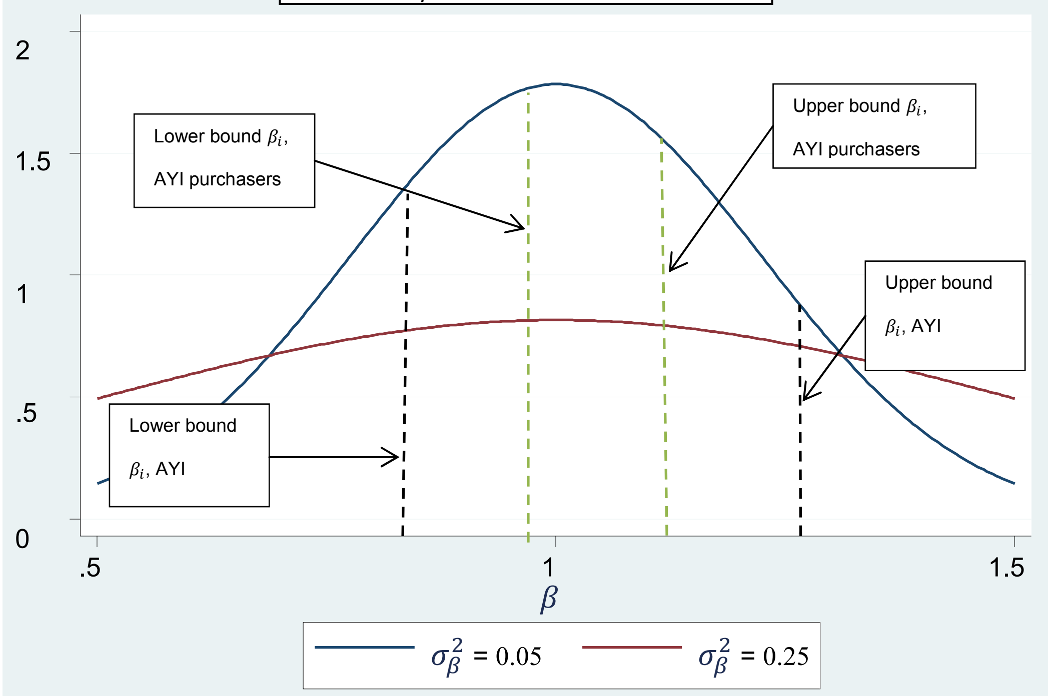

Farmers with β_i values between the two heavier innermost dashed vertical lines purchase AYI when made eligible to do so at the market price, and those with β_i values between the two outermost dashed vertical lines purchase it when given a coupon. The compliers, i.e., farmers who only purchase AYI when given a coupon, have values of β_i between the innermost and outermost dashed lines. The degree to which the average treatment effects identified by the randomized encouragement design and the randomization of eligibility coincide will depend on how similar the β_i values for compliers are to those of the group of farmers who purchase AYI at the market price when allowed to do so.

In the scenario depicted in Figure 9, shifting out the upper and lower bounds on β_i for insurance purchasers via the encouragement changes the mix of program participants in ways that offset. The movement of the upper bound brings in farmers who are also induced to switch from the subsistence crop to planting cotton, whereas the movement of the lower bound brings in farmers who plant cotton no matter what. A smaller value of σ_β^2 increases the mass of the β distribution contained between the two lower bounds on β_i for insurance purchasers relative to the mass contained between the two upper bounds; as shown in Figure 9, the lower bounds are closer to the mean than the upper bounds, and the gap between the former is nearly equal to that of the space between the latter. A further tightening of σ_β^2 will serve to enhance this disparity further, implying that a lower value of σ_β^2 results in greater differences in the pool of insurance purchasers brought in by the encouragement relative to the randomization of eligibility.

The implication of this is that due to the positions of the bounds on β_i for insurance purchasers in the β distribution, increasing the value of σ_β^2 results in a smaller impact of the encouragement design on the mix of insurance purchasers and a smaller bias of the encouragement design estimator relative to the PRTE. This result is interesting in that it is an example of greater unobserved heterogeneity with respect to program benefits not resulting in a larger discrepancy between different average treatment effects. But the result is context specific. If we were to alter the positions of the different bounds on β_i, greater variability in β could generate larger bias in the encouragement design estimator.

The impact of a larger value of σ_β^2 on the variance of each estimator is more straightforward. The identifying power of each design comes from the proportion of farmers induced to purchase AYI. Greater variability of the β distribution means that less mass is contained between the upper and lower bounds on β_i for insurance purchasers under the randomization of eligibility, resulting in a higher variance of that estimator. For the same reason, a large value of σ_β^2 will mean that the outward shift of the bounds on β_i for insurance purchasers generated by the encouragement will bring in a small amount of mass of the β distribution, resulting in low precision of the randomized encouragement design estimator. Lower variability of β has the opposite effect, enhancing the precision of both estimators.

5.4 Simulation implications

Several implications can be gleaned from the above simulation results. Firstly, larger coupons mean greater bias, and the bias-precision tradeoff involved with selecting the strength of the incentive offered by a randomized encouragement design cannot be avoided. However, bias does not necessarily translate into getting further away from the truth on average, as shown by the impact of larger coupons on the MSE in Figure 7. Randomized encouragement designs and other LATE estimators should not be dismissed out of hand because they are biased, particularly if the alternative is a research design that cannot produce much in the way of statistical precision.

Secondly, unobserved heterogeneity that is correlated with the decision to participate in a program and the outcome of interest will determine how strong the encouragement must be in order to recover statistically significant program effects, but the direction of this effect will depend on the context. For example, given the parameter values assigned to the model in Table 1, greater heterogeneity with respect to program benefits as represented by σ_β^2 would make it more costly to generate statistically significant estimates from a randomized encouragement design, as a given increase in incentives yields a smaller jump in participation. But this result depends on the particular situation being analyzed. If all of the bounds were located in the tails beyond where the two distributions pictured in Figure 9 cross one another, then an increase in σ_β^2 would result in greater precision of both estimators.[8]

This might be the case if the outcome of interest were something in which only farmers with very low sensitivity to covariate risk were involved. This would place the cutoff on the value of β_i for participants in the activity prior to the availability of insurance in the far left hand tail, and would likely also place the ranges of β_i values associated with purchasing AYI under the two competing research designs in this same region of the β distribution. While it probably would not be possible to estimate the distribution of β prior to program rollout in a developing country context, knowledge about the nature of the activities being affected by a program and the characteristics of the study region (e.g., geographic and climatic characteristics that might affect

[8] In fact, if β followed a non-truncated normal distribution, it could be shown that as long as the bounds are all within one standard deviation of the midpoint of the distribution prior to increasing its spread, raising the variance of β would decrease the precision of both estimators. The truncated normal is more realistic in this application but makes drawing analytical conclusions much less straightforward.

sensitivity to the common shock) could provide clues as to how different levels of encouragement are likely to influence statistical precision.

The implications of greater heterogeneity for the bias of the randomized encouragement design estimator will also vary depending on context. In the simulation above, greater heterogeneity actually led to lower bias of the encouragement design estimator. While is possible that this situation could be reversed, the same sorts of baseline data that could provide insight into how the level of encouragement can be expected to influence statistical precision can also by useful with respect to diagnosing possible implications for bias. If farmers tend to be similar, then participation in AYI or any other program could be very sensitive to the level of encouragement that is offered, and this sensitivity will be transmitted through the bias of the randomized encouragement design estimator.

This is all suggestive of the important role of learning about the selection and outcome processes underlying the program that is to be evaluated, through the collection of baseline data and the application of theory. Most obviously, baseline data can make it possible to determine whether or not we should expect low participation to be a problem, and if so, the factors depressing program uptake. This knowledge can be applied to the design of more effective encouragements. Secondly, learning about these underlying processes can shed light on how far the effects identified by a randomized encouragement design will stray from the average impacts on program participants who have not been offered an extra incentive, and just how strong the encouragement will have to be to make it possible to distinguish estimated impacts from zero.

6 Conclusion

This paper explored the choice between two different research designs in the context of a program affected by low participation rates. As development economists focus more on evaluating complex interventions, low participation by households in the programs being analyzed has become more common, the result of which will often be estimates so imprecise that it is impossible to distinguish what works from what does not. While low participation could be the result of optimizing households making the decision that yields the greatest level of welfare, it is not at all clear that this is the case in this context. In the example of index insurance, properly evaluating the benefits of participation requires accurate knowledge of the insured risk and how fluctuations in this risk affect household welfare. This may necessitate a high level of sophistication by potential participants. It is unrealistic to expect this from poor households with little education interacting in markets with minimal history in the area being studied.

Against this backdrop, I compared the potential of a randomization of eligibility and a randomized encouragement design to serve as the basis for an impact evaluation of an index insurance product. A randomization of eligibility yields can yield unbiased estimates of the impact of the insurance program on participants, which may be the program effect of greatest interest from the perspective of both research and policymaking. A randomized encouragement design featuring a suitably strong incentive will improve upon the precision of the randomization of eligibility. By changing the pool of participants in the insurance program, the average treatment effect identified by a randomized encouragement will differ from that of the randomization of eligibility. The presence of bias does not necessarily mean that randomized encouragement designs should be ruled out, however. When equal weight is given to bias and precision when selecting between these two research designs, as is done by a comparison of

Mean Square Error, using a randomized encouragement can indeed be preferable to a randomization of eligibility when low participation is an issue. Low participation in the absence of additional incentives is not sufficient reason to abandon the effort to measure program impacts.

References

Barrett, Christopher B., and Michael R. Carter. "The Power and Pitfalls of Experiments in Development Economics: Some Non-random Reflections." *Applied Economic Perspectives* 32, no. 4 (2010): 515-548.

Boucher, Steve, and Conner Mullally. "Case Study: Evaluating the Impact of Index Insurance on Cotton Farmers in Peru." Chap. 9 in *Designing Impact Evaluations for Agricultural Projects*, edited by Paul Winters, Lina Salazar and Alessandro Maffioli, 93-110. Inter-American Development Bank, 2010.

Bound, John, David A. Jaeger, and Regina M. Baker. "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak." *Journal of the American Statistical Association* 90, no. 430 (1995): 443-450.

Cai, Hongbin, Yuyu Chen, Hanming Fang, and Li-An Zhou. *Microinsurance, Trust and Economic Development: Evidence from a Randomized Natural Field Experiment.* Working Paper, National Bureau of Economic Research, 2009.

Carter, Michael, Peter Little, Tewodaj Mogues, and Workneh Negatu. "Poverty Traps and Natural Disasters in Ethiopia and Honduras." *World Development* 35, no. 5 (2007): 836-856.

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery. *Barriers to Household Risk Management in India.* Policy Research Working Paper, World Bank, 2010.

Deaton, Angus. *Instruments of Development: Randomization in the Tropics, and the Search for the Elusive Keys to Development.* Working Paper, National Bureau of Economic Research, 2009.

Dercon, Stefan. "Growth and Shocks: Evidence from Ethiopia." *Journal of Development Economics* 74, no. 2 (2004): 309-329.

Duflo, Esther, Rachel Glennerster, and Michael Kremer. *Using Randomization in Development Economics Research: A Toolkit.* Technical Working Paper, National Bureau of Economic Research, 2006.

Giné, Xavier, and Dean Yang. "Insurance, Credit, and Technology Adoption: Field Experimental Evidence from Malawi." *Journal of Development Economics* 89, no. 1 (2009): 1-11.

Giné, Xavier, Robert Townsend, and James Vickery. *Patterns of Rainfall Insurance Participation in Rural India.* Staff Report, Federal Reserve Bank of New York, 2007.

Guiso, Luigi, and Tullio Jappelli. *Financial Literacy and Portfolio Diversification.* Working Paper, Centre for Studies in Economics and Finance, 2009.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. "Trusting the Stock Market." *The Journal of Finance* 63, no. 6 (2008): 2557-2600.

Hazell, Peter. "The Appropriate Role of Agricultural Insurnace in Developing Countries." *Journal of International Development* 4, no. 6 (1992): 567-581.

Heckman, James J., and Edward J. Vytlacil. *Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Econometric Estimators to Evaluate Social Programs, and to Forecast Their Effects in New Environments.* Vol. 6B, chap. 71 in *Handbook of Econometrics*, edited by James J. Heckman and Edward Leamer, 4875-5143. Amsterdam: Elsevier B.V., 2007.

Imbens, Guido, and Josh Angrist. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62, no. 2 (1994): 467-475.

McKenzie, David. "Impact Assessments in Finance and Private Sector Development: What Have We Learned and What Should We Learn/." *World Bank Research Observer* 31, no. 4 (2009): 209-233.

Miranda, Mario. "Area-yield Crop Insurance Reconsidered." *American Journal of Agricultural Economics* 73, no. 2 (1991): 233-242.

Rosenzweig, Mark, and Hans Binswanger. "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments." *Economic Journal* 103, no. 416 (1993): 56-78.

A.1 Appendix

A.1.1 Demand for AYI and the parameters of the AYI contract

Recall that farmers who purchase AYI must have values of β_i lying within the following interval:

$$\max\left[\beta_1,\beta_2\right]<\beta_i<\beta_3 \tag{A.32}$$

or more explicitly:

$$\max\left[\frac{(2\mu-l)l+\gamma\sigma_I^2}{2\gamma\sigma_{I,\varepsilon_c}},\frac{\sigma_{I,\varepsilon_c}-x}{\sigma_{\varepsilon_c}^2}\right]<\beta_i<\frac{\sigma_{I,\varepsilon_c}+x}{\sigma_{\varepsilon_c}^2} \tag{A.33}$$

Where $x=\sqrt{\gamma\left(\sigma_{c,I}\right)^2+\sigma_c^2\left[\left(\mu-L\right)^2-w^2-\gamma\left(\sigma_c^2+\sigma_I^2\right)\right]}$. Assuming x is positive, which is quite reasonable given that the squared $\sigma_{c,I}$ term will be extremely large, the derivative of the upper bound of (A.33) with respect to $\sigma_{c,I}$ is:

$$\frac{1+\frac{\gamma\sigma_{c,I}}{x}}{\sigma_c^2}>0 \tag{A.34}$$

Depending on which lower bound is utilized, its derivative is either:

$$\frac{1-\frac{\gamma\sigma_{c,I}}{x}}{\sigma_c^2}<0 \tag{A.35}$$

which has an indeterminate sign, or assuming $\mu>L$:

$$\frac{-\left(\left(2\mu-L\right)L+\gamma\sigma_I^2\right)}{2\gamma\left(\sigma_{c,I}\right)^2}<0 \tag{A.36}$$

If the derivative is (A.35), the lower bound may be increasing, but at a rate slower than that of the upper bound; β_3 is always greater than β_1. Increasing the risk reduction potential of AYI via a higher value of $\sigma_{c,I}$ will weakly increase demand for insurance.

The derivative of the upper bound with respect to the variance of the indemnity, σ_I^2, is:

$$\frac{\gamma}{-2x} < 0 \tag{A.37}$$

The derivative of the lower bound with respect to σ_I^2 is then:

$$\frac{\gamma}{2x} > 0 \tag{A.38}$$

or

$$\frac{1}{2\sigma_{I,\varepsilon_c}} > 0 \tag{A.39}$$

depending on which lower bound is used. The size interval is therefore decreasing with respect to σ_I^2.

Lastly, we have the loading term L and its impact on the interval given in (A.33). The derivative of the upper bound with respect to L is:

$$\frac{-(\mu - L)}{x} < 0 \tag{A.40}$$

While the derivative of the lower bound with respect to L is:

$$\frac{(\mu - L)}{x} > 0 \tag{A.41}$$

or

$$\frac{\mu - L}{\gamma\sigma_{c,I}} > 0 \tag{A.42}$$

depending on which lower bound is used. The size of the interval is a decreasing function of the loading term, *L*.

A.1.2 Perceived cheating and the AYI contract

Suppose that farmers believe that the insurer exaggerates measured area-yields in every time period by an amount $g\mu$. Given g, the perceived covariate shock is $\tilde{\varepsilon}_t^c = \tilde{\mu} - \tilde{\tilde{q}}_t$. The AYI contract that replaces shortfalls in area-yields from μ is perceived as:

$$\tilde{I}_t = \max\left[0, \tilde{\varepsilon}^c - g\mu\right] \quad \text{(A.43)}$$

while the true contract is:

$$I_t = \max\left[0, \varepsilon^c\right] \quad \text{(A.44)}$$

The variance of area-yields given *g* is equall to that of true area-yields. Since the common shock is symmetric, this implies that the distributions of ε^c and $\tilde{\varepsilon}^c$ are identical. When we evaluate the moments of $\tilde{\varepsilon}^c$ as they pertain to the AYI contract, we can use the distribution of ε^c as long as we account for the fact that g will leads farmers to view the strike point of the contract as higher than it actually is.

While it is proved in the main text that incorrectly perceiving the mean of area-yields leads farmers to view AYI as more costly, we cannot unequivocally say that it leads AYI to be viewed as less risk-reducing. The true change in the variance of output due to purchasing AYI is:

$$\sigma_I^2 - 2\beta_i \sigma_{c,I} \quad \text{(A.45)}$$

Replacing *I* and ε^c with $\tilde{I}$ and $\tilde{\varepsilon}^c$ yields the perceived change in variance. Consider the covariance between the indemnity and the covariate shock. The true value of this parameter is:

$$
\begin{aligned}
&\sigma_{c,I} = E\left(\max\left[0,\varepsilon_t^c\right]\varepsilon^c\right) - E\left(\max\left[0,\varepsilon_t^c\right]\right)E\left(\varepsilon^c\right) = \\
&E\left[\left(\varepsilon^c\right)^2 \mid \varepsilon^c > 0\right] * P\left(\varepsilon^c > 0\right) = \\
&\left(\frac{\int_0^{\bar{\varepsilon}^c}\left(\varepsilon^c\right)^2 f\left(\varepsilon^c\right)d\varepsilon^c}{P\left(\varepsilon^c > 0\right)}\right)P\left(\varepsilon^c > 0\right) = \\
&\int_0^{\bar{\varepsilon}^c}\left(\varepsilon^c\right)^2 f\left(\varepsilon^c\right)d\varepsilon^c
\end{aligned}
\tag{A.46}
$$

Under the subjective distribution of area-yields, the covariance term is:

$$
\begin{aligned}
&\sigma_{c,\tilde{I}} = E\left(\max\left[0,\varepsilon_t^c - g\mu\right]\varepsilon^c\right) - E\left(\max\left[0,\varepsilon_{ct} - \mu g\right]\right)E\left(\varepsilon^c\right) = \\
&\left(E\left[\left(\varepsilon^c\right)^2 \mid \varepsilon^c > g\mu\right] - g\mu E\left[\varepsilon^c \mid \varepsilon^c > g\mu\right]\right)P\left(\varepsilon^c > 0\right) = \\
&\left(\frac{\int_{g\mu}^{\bar{\varepsilon}^c}\left(\varepsilon^c\right)^2 f\left(\varepsilon^c\right)d\varepsilon^c}{P\left(\varepsilon^c > g\mu\right)}\right)P\left(\varepsilon_c > g\mu\right) - g\left(\frac{\int_{g\mu}^{\bar{\varepsilon}^c}\varepsilon^c f\left(\varepsilon_c\right)d\varepsilon_c}{P\left(\varepsilon^c > g\mu\right)}\right)P\left(\varepsilon_c > g\mu\right) = \\
&\int_{g\mu}^{\bar{\varepsilon}^c}\left(\varepsilon^c\right)^2 f\left(\varepsilon^c\right)d\varepsilon^c - g\mu\int_{g\mu}^{\bar{\varepsilon}^c}\varepsilon^c f\left(\varepsilon^c\right)d\varepsilon^c
\end{aligned}
\tag{A.47}
$$

The last line of (A.46) is larger than the last line of (A.47); i.e., the true covariance is larger than the covariance given the incorrectly perceived mean.

Ambiguity with respect to the risk reduction potential of AYI comes from the impact of g on the variance of the indemnity. The true variance of the indemnity function is equal to:

$$
\begin{aligned}
&\sigma_I^2 = E\left(I^2\right) - E\left(I\right)^2 = E\left(I^2\right) - r^2 = \\
&E\left(\max\left[0, \varepsilon_t^c\right] \max\left[0, \varepsilon_t^c\right]\right) - r^2 = \\
&E\left[\left(\varepsilon^c\right)^2 \mid \varepsilon^c > 0\right] P\left(\varepsilon^c > 0\right) - E\left[\varepsilon^c \mid \varepsilon^c > 0\right]^2 P\left(\varepsilon^c > 0\right)^2 = \\
&\left[\left(\sigma_c^2 \mid \varepsilon^c > 0\right) + E\left(\varepsilon^c \mid \varepsilon^c > 0\right)^2\right] P\left(\varepsilon^c > 0\right) - E\left(\varepsilon^c \mid \varepsilon^c > 0\right)^2 P\left(\varepsilon^c > 0\right)^2 = \\
&\left(\sigma_c^2 \mid \varepsilon^c > 0\right) P\left(\varepsilon^c > 0\right) + \left[1 - P\left(\varepsilon^c > 0\right)\right] E\left(\varepsilon^c \mid \varepsilon^c > 0\right)^2 P\left(\varepsilon^c > 0\right) = \\
&\left(\sigma_c^2 \mid \varepsilon^c > 0\right) P\left(\varepsilon^c > 0\right) + \frac{P\left(\varepsilon^c < 0\right) E\left(\varepsilon^c \mid \varepsilon^c > 0\right)^2 P\left(\varepsilon^c > 0\right)^2}{P\left(\varepsilon^c > 0\right)} = \\
&\left(\sigma_c^2 \mid \varepsilon^c > 0\right) P\left(\varepsilon^c > 0\right) + E\left(\varepsilon^c \mid \varepsilon^c > 0\right)^2 P\left(\varepsilon^c > 0\right)^2 \\
&\left(\sigma_c^2 \mid \varepsilon^c > 0\right) P\left(\varepsilon^c > 0\right) + r^2
\end{aligned}
\tag{A.48}
$$

The second to last line follows from the symmetry of ε^c, i.e., $P\left(\varepsilon^c > 0\right) = P\left(\varepsilon^c < 0\right)$.

Using similar reasoning, it can be shown that given perceived cheating by the insurer of size $g\mu$, the variance of the indemnity is:

$$
\left(\sigma_c^2 \mid \varepsilon^c > g\mu\right) P\left(\varepsilon^c > g\mu\right) + \frac{\left(\tilde{r}\right)^2 P\left(\varepsilon^c < g\mu\right)}{P\left(\varepsilon^c > g\mu\right)} \tag{A.49}
$$

Since $P\left(\varepsilon^c < g\mu\right) > P\left(\varepsilon^c > g\mu\right)$, the second term in the right hand side of (A.49) may increase the variance of the indemnity, even though $\tilde{\tau}^2 < \tau^2$. However, raising the truncation point of a distribution decreases the truncated variance, and as a result, $\left(\sigma_c^2 \mid \varepsilon^c > 0\right) P\left(\varepsilon^c > 0\right) > \left(\sigma_c^2 \mid \varepsilon^c > g\mu\right) P\left(\varepsilon^c > g\mu\right)$. The perception of cheating by the insurer unequivocally lowers the covariance between the indemnity and the covariate shock, but it may also decrease the variance of the indemnity.

A.1.3 Perceived cheating by the insurer and demand for AYI

The impact of perceived cheating *g* on demand will depend on the derivatives of the bounds of the set of β_i values given in (A.33) for which purchasing AYI is optimal with respect to *g*. The derivative of the first possible lower bound of this set with respect to g is:

$$\frac{\sigma_{c,\tilde{I}}\left(2\frac{\partial h}{\partial g}\left(\mu-\left(L+h\right)\right)+\gamma\frac{\partial\sigma_{\tilde{I}}^{2}}{\partial g}\right)-\frac{\partial\sigma_{c,\tilde{I}}}{\partial g}\left(\left(2\mu-\left(L+h\right)\right)\left(L+h\right)+\gamma\sigma_{\tilde{I}}^{2}\right)}{2\gamma\left(\sigma_{c,\tilde{I}}\right)^{2}} \tag{A.50}$$

where *h* is the perceived increase the premium due to *g*. The sign of the left hand bracketed term in the numerator is ambiguous. If it is positive, then this possible lower bound is increasing with respect to *g*.

Now consider the other potential lower bound of the set, as well as the upper bound. Since $\sigma_c^2 > 0$ and $\left[\partial\sigma_{c,\tilde{I}}/\partial g\right] < 0$, the sign of derivatives of the other two bounds with respect to *g* will depend on $\partial\tilde{x}/\partial g$. This derivative is:

$$\frac{2\gamma\sigma_{c,\tilde{I}}\frac{\partial\sigma_{c,\tilde{I}}}{\partial g}-\sigma_c^2\left(2\left(\mu-\left(L+h\right)\right)\frac{\partial h}{\partial g}+\gamma\frac{\partial\sigma_{\tilde{I}}^{2}}{\partial g}\right)}{2\tilde{x}} \tag{A.51}$$

The sign of the bracketed term in the numerator is also ambiguous. Comparing the different derivatives upon which the signs of (A.50) and (A.51) depend (i.e., $\partial h/\partial g$, $\partial\sigma_{c,\tilde{I}}/\partial g$, and $\partial\sigma_{\tilde{I}}^2/\partial g$) yields no further information without explicitly assigning values to parameters.