

Public Disclosure Authorized

The Distributional Consequences of Group Procurement: Evidence from a Randomized Trial of a Food Security Program in Rural India*

Paul Christian[†]

July 1, 2015

Abstract

Public transfer programs that allow beneficiaries to choose the transferred good may be more efficient, but the poorest beneficiaries may not participate if the good chosen is too costly. A model shows that program targeting and consumption impacts are tied to selected quality of the provided good. Evidence from a randomized trial in rural India in which groups of beneficiaries choose the variety of rice to be offered as a subsidized loan confirms that choosing lower cost goods self-targets the program towards the poorest beneficiaries. Consumption impacts are biggest for wealthiest households and may be negative for moderately poor households.

*This paper is a product of the Social Observatory of the World Bank. Support and funding for data collection from the South Asia Food and Nutrition Security Initiative (SAFANSI) and from the Population Studies Training Center and Global Health Initiative at Brown University is gratefully acknowledged. I am grateful to Andrew Foster, Vijayendra Rao, Sriniketh Nagavarapu, Upamanyu Datta, Nethra Palaniswami, and Samrat Kumar Ghosh for their support and feedback. Valuable feedback was provided by Brian Knight, Anja Sautmann, and seminar participants at Brown, Tufts, Cornell, Occidental, Mathematica Policy Research, and NEUDC 2013 (Harvard). Special thanks are due to the Bihar Rural Livelihoods Project for their support, cooperation, and useful conversations, especially Ajit Ranjan, Archana Tiwari, and Nidhi Singh. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of The World Bank, its Board of Executive Directors, or the governments they represent. The World Bank does not guarantee the accuracy of the data included in this work. The author accepts all responsibility for errors.

[†]Dyson School of Applied Economics and Management, Cornell University and Social Observatory, Correspondence: 301H Warren Hall, Cornell University, Email: paul.christian@cornell.edu.

When designing distributive social programs, policymakers often attempt to balance two goals: 1. improving the welfare of people believed to have the highest return to the program while 2. minimizing total cost. One way that policies have attempted to address the first goal is by offering goods that only poor households want to consume, for example low quality staple foods that only poor households buy. Making transfers in the form of physical goods can lead to inefficiencies, however, because of corruption and wastage related to transportation, storage, and distribution, and because program implementers do not know the preferences of potential beneficiaries. To avoid these inefficiencies, programs may want to have communities choose the good themselves and procure the chosen good locally. This paper examines whether using public choice this way is consistent with self-targeting and shows that two program features which arise as natural attempts to deal with targeting problems (subsidizing staple grains and letting beneficiaries choose the grain to receive the subsidy) leads to outcomes that policy makers likely do not prefer. Specifically, I study two questions regarding a program that gives villages the choice of which variety of food grain to transfer to beneficiaries. First, will take-up be concentrated among the poor if villages jointly choose the good? Second, how does the quality of grain chosen determine heterogeneity of measured impacts on quantity of grain consumed by intended beneficiaries of varying wealth?

In general, tension often arises between the goal of improving welfare for the poor at the lowest cost, because while policymakers may deem poor households most in need of the program, wealthy households also would like to participate, and participation by relatively wealthy households increases the cost of the program. If program administrators can easily observe poverty or a proxy for poverty, they can allow or disallow participation based on that characteristic. But given the cost and uncertainty in measuring poverty, setting fixed eligibility rules often causes both exclusion errors (poor households are left out) and inclusion errors (wealthy households participate). (Alatas et al., 2012)

Given the challenge of determining exclusion directly on wealth, researchers and policymakers have attempted to design program features such that households with a particular characteristic favored by policymakers reveal themselves through their choice to participate or drop out. Influencing take up through program restrictions rather than explicit exclusion criteria is called self-targeting. Many of the major strategies for devising self-targeted programs were classified by Nichols and Zeckhauser. (Nichols and Zeckhauser, 1982) One class of mechanisms recently receiving attention are ordeal mechanisms, which hypothesize that making people complete a costly task can lead to better targeting, because only those with

the highest marginal return to participation will choose to pay the cost. (Alatas et al., 2012; Dupas et al., 2013) Implementing ordeal mechanisms has proven to be complicated however, because wealthy and poor members may have differential abilities to circumvent the ordeals, and because non-linear utility functions mean that fixed ordeals may impose bigger burdens on the people who have the highest returns to the program, undermining the potential for targeting. (Alatas et al., 2012)

An alternative example of a self-targeting mechanism was described by (Besley and Coate, 1991) and demonstrated empirically by (Jacoby, 1997). These models proposed that programs could focus on providing certain goods that would be taken up only by low income households. If policymakers subsidize only low-quality goods, poor beneficiaries for whom the marginal utility of consumption is high will choose to participate, but relatively wealthy households will willingly opt forgo the program in favor of consuming higher quality goods, even though the subsidy effectively raises the marginal cost of quality. Jacoby demonstrated the concept using a subsidy of a low-cost, bland staple food in Jamaica, finding that take-up of a school feeding program providing the bland food was higher among poorer households and those with more eligible children than among comparatively wealthy households. (Jacoby, 1997)

Directly observing poverty through expenditure surveys imposes a cost on programs who must fund expenditure surveys or proxy means tests, and ordeal mechanisms impose costs on beneficiaries that are by definition a burden to the potential beneficiaries. In contrast, in addition to the potential for self-targeting, in-kind transfers may have features that make them more desirable to both households and policymakers than equivalently valued cash transfers. First, in-kind transfers may be politically appealing if they are more easily portrayed and understood as using public funds only to help disadvantaged households buy the things they really need. (Barrett, 2002) Second, in-kind transfers can have differential consumption impacts than equivalently sized transfers of cash in ways that are useful for policy design. (Jacoby, 2002) demonstrates that in-kind transfers targeted to specific members have fly-paper effects which influence the consumption of targeted members more than the consumption of co-resident, but non-targeted members. Finally, policymakers might have paternalistic preferences for increasing consumption of food in particular. Currie and Gahvari's review of the literature cites numerous studies claiming that various food transfers or subsidies may increase food consumption more than equivalently valued cash transfers would. (Currie and Gahvari, 2008)

In-kind transfers can also have impacts which work against policy goals, and the ex-

pected effects on consumption of food subsidies can be theoretically ambiguous. Subsidy programs can even cause a reduction in estimated food consumption for certain sub-populations, either because of general equilibrium effects (Cunha, Giorgi and Jayachandran, 2011) or because of a Giffen effect, wherein income effects of subsidizing staple foods outweigh the cross-price effect of the subsidy, causing people to substitute to higher quality foods, which could ultimately reduce caloric intake for some households. (Jensen and Miller, 2011)

Unfortunately, implementing programs to directly transfer staple goods presents challenges for program design. Procurement and distribution of physical goods, especially food goods, is logistically difficult to implement by a centralized agency. Avoiding these logistical costs is especially important to policymakers in India, where the national government's primary food subsidy program, the Public Distribution System (PDS) has been criticized for leakage and corruption in the procurement and distribution process (Jha and Srinivasan, 2001). Local procurement obviates the need for these systems, since goods are purchased and stored in close proximity to beneficiaries, eliminating the need for transportation and storage. Comparisons of internationally transported food aid with locally or regionally sourced grains have found that buying grain locally is generally cheaper because of savings on transportation and storage, (Lentz, Passarelli and Barrett, 2013) and more preferred by beneficiaries because of differences in local varieties and cultivation methods. (Violette et al., 2013)

If policymakers choose to employ local procurement when implementing an in-kind transfer, additional challenges arise. First, in many contexts, even basic staple goods are heterogeneous, and their price and quality can vary widely. If a centralized policy maker chooses a good that is too costly, poor members may not be able to participate, but if the policy maker chooses a good that is too inferior, even relatively poor members targeted by the program may opt out. Evidence from migrants' food consumption patterns in India confirms that households forfeit caloric intake in favor of consuming goods which most closely match their preferences but may be more costly. (Atkin, 2013) Second, program implementers may not have good information about sources of food and prices, causing the program to overspend on grain. An attractive candidate to address both of these challenges is to incorporate a public choice mechanism into the policy design. Public choice can potentially incorporate information about the preferences of beneficiaries that is not easily available to a central planner, increasing satisfaction and take-up. Both theory and experience suggest that communities can leverage information to identify poor households and adapt programs to their preferences, but that doing so makes the program susceptible to elite capture. For a review, see Mansuri and Rao. (Mansuri and Rao, 2013) In addition, the exercise of public

choice has been shown to independently increase program participation and satisfaction. (Olken, 2010)

This paper asks whether the strategy of local procurement and public choice is compatible with the goal of self-targeting through the provision of a good generally preferred by the poorest members of the communities in a context where basic staple goods have some heterogeneity. A simple model shows the mechanisms through which price and quality of varieties of the staple good can generate separation of potential beneficiaries by wealth, and predicts the relative likelihood of participating over the wealth/income distribution. In the model, households have preferences for quality of the staple grains they consume, and quality is a normal good, meaning that optimal quality is increasing in household wealth/income. The model predicts that the surplus utility of participation in a program for a household of a given level of wealth is maximized when the program subsidizes a particular quality of food corresponding to household wealth, and that surplus utility is decreasing as quality is increased or decreased from this point. The implication of this model is that households who choose to participate in the program in a given village will come from particular sections of the wealth distribution, i.e. when the village subsidizes high quality grain, only wealthy people participate, and when the village subsidizes low quality grain, only poor people participate. Whether the program is targeted toward the poor then becomes an empirical question of what kind of grain the villages decide to provide. The model further shows that if some villages choose relatively high quality grain, the program may actually lead to decreases in the quantity of grain consumed relative to control villages, because some moderately poor households may decide to use the program to upgrade the quality of the grain they buy, even though the high quality rice from the program is still more expensive than the low quality grain they would have purchased otherwise.

The model used in this paper are applicable generally to public transfer programs where beneficiaries have some say in the quality of a good to be provided where the program only provides one or a limited, discrete set of options from a menu of choices where cost and quality are associated and purchasing other, non-selected varieties of the transferred good at market rates is an option. Other contexts where similar models have been applied include supplementary nutrition programs in developing countries (Jacoby, 1997), public school provision where private schools are an option (Besley and Coate, 1991).

Using data from a randomized evaluation of a food security program that provides in-kind subsidized loans, I show that communities participating in the program do not always choose the cheapest available staple good. The data suggests that the particular represen-

tative choice model for public choice employed by the program leads to village committees being more slanted to choosing higher price/quality goods and to lower participation by less wealthy members. Next, I show evidence that the variation in the price of the good chosen by each participating village predicts whether wealthier or poorer potential beneficiaries are more likely to participate. Finally, the paper turns to heterogeneous consumption impacts. The model predicts that the effects on quantity of rice consumed by the household are only unambiguously positive for the wealthiest households, since the program for them is subsidizing grain that is already at least as inexpensive as they would have purchased from the market. Moderately poor households who choose to purchase grain which is more expensive than what they would have chosen to buy from the market may actually decrease their consumption. The balance of treatment effects over the distribution of wealth in this particular case leads the program to show no impact on consumption across the distribution of wealth, even though the program: 1. Is at least weakly welfare improving for all households and 2. Does result in consumption increases for some groups, even if not the ones that the program designers might wish to target for consumption increases.

The next section describes the program and the particular context for the model. Section 3 provides data on the characteristics of rice in Bihar to set up the intuition of the model. Section 4 presents the model and demonstrates the results related to targeting. The implications of the model for choice of rice and targeting are taken to the data in Section 5 and Section 6 respectively. Section 7 demonstrates the application of the model for the analysis of consumption effects from the program. The consumption predictions are tested in Section 8 using the cluster randomization of the program. Section 9 summarizes the results, draws conclusions, and raises further questions and extensions.

1 The Program and Data

The data used was collected as part of an evaluation of a program operated by the Bihar Rural Livelihoods Project (BRLP), a joint project of the World Bank and the state government of Bihar in India. The BRLP's primary role is to organize groups of women called Self-Help Groups (SHGs). SHGs conduct microfinance activities and participate in the BRLP's programs aimed at increasing the earning potential of the women and their families. The SHGs are further grouped together into Village Organizations (VOs) which participate in programs designed to involve multiple SHGs. Each VO consists of 100-150 women, all of whom live within the same village.

The program is a food security intervention called the Food Security Fund (FSF). The FSF is designed to help credit-constrained VO members access large quantities of rice in a single purchase without exposing them to high interest rates. The BRLP grants 100,000 rupees (about \$2,000) to the VO, and assigns a staff member to use the funds to organize an in-kind lending program. First, members of the VO are interviewed about their food practices, and encouraged to think about how much food they may need to purchase or borrow over the following three months, based on the amount of rice they expect their family to produce and consume over that period. If members expect to purchase or borrow rice, they are asked if they would like to take a loan from the FSF fund. Once all members have expressed a desired amount to borrow, the BRLP staff adds up the aggregate loan demand. Members of the VO leadership take this amount to nearby grain vendors and solicit bids from each vendor. Samples of vendors' varieties of rice are taken, along with prices, and the VO leadership agrees on which type of grain to purchase. Once the rice is purchased, it is distributed to members, who are given three months to repay their loan with no interest rate charged.

Although this program does not directly subsidize the price of the rice, the project believes that it benefits participants in three ways: 1. Many members rely on loans to purchase food in absence of the program, meaning that the program is a direct subsidy on interest costs for these members; 2. Encouraging members to purchase three months' worth of food at one time, rather than smaller quantities at small intervals, allows individual members to leverage potential bulk discounts; and 3. Pooling the demands of the whole VO together when bargaining with shop keepers and wholesalers to solicit competitive bids could help them leverage further bulk purchasing discounts. Due to difficulties in comparing quality, the exact size of the bulk discount is not measured in this paper. However, program participation is voluntary, and high participation rates suggest that there is some benefit to the program relative to purchasing independently from the market. For the remainder of the paper, this benefit will be thought of as a small discount to the per-unit purchasing cost of rice.

The data was collected as part of a cluster-randomized control trial. Six blocks¹ were selected for the study, chosen because these locations had formed a sufficient number of VOs to be included in the study, but had not yet started the FSF program. The project divides each block into three units called clusters. Random assignment either to receive the FSF program (treatment) or not (control) was done at the cluster level, so that within a block,

¹A block is an Indian administrative unit similar to a county.

all of the VOs within one or two clusters were assigned to start the program, while the remaining cluster or clusters were assigned to wait until completion of the study.

After randomization, baseline data was collected from five VOs in each of the eighteen evaluation clusters (90 VOs: 45 treatment, 45 control) in July 2012 prior to initiation of the lending program in any VOs. A VO is comprised of ten to fifteen smaller units called Self Help Groups (SHGs), which each have ten to fifteen members, including three leaders (designated as President, Secretary, or Treasurer). To generate the sample of households, three SHGs were randomly chosen in each VO and seventeen members were selected from each of these SHGs. In each of the selected SHGs, two out of three leaders were selected, as well as either three or four non-leaders. If a respondent could not be located, a substitute was randomly selected from among the remaining members. Upon completion of the baseline, the lending program began in September 2012. A post-intervention survey was collected in February-March, 2013. The surveys covered detailed monthly food expenditure and weekly consumption, coarse categories of non-food expenditure, experiences of food insecurity, agricultural production, income, and child health measures. Data was collected from a total of 1,529 SHG members about their households in the baseline. In the post-intervention survey, the same members were resurveyed with 1,449 being relocated and agreeing to participate again. Attrition was evenly balanced across treatment and control ($p=.25$). A second post-intervention survey was conducted in September to October 2013, with 1,472 again successfully surveyed and attrition balanced by treatment status ($p=.68$). Out of the original 1,529 at baseline, 1518 households were re-surveyed in at least one one of the two follow-up rounds with balanced attrition ($p=.76$).

2 Rice as a Staple Good in Bihar

The literature on using in-kind provision of services as a second-best method of targeting (Nichols and Zeckhauser, 1982; Besley and Coate, 1991; Jacoby, 1997) suggests that choosing an inferior staple good can ensure that program participation is concentrated among disadvantaged members. This line of reasoning has encouraged programs to focus on providing staple grains, but neglects potential heterogeneity of quality even among staple goods.

Table 1 shows some summary statistics for the households of the members surveyed. The general pattern is one of disadvantage, with the average household earning little more than 1000 rupees per month per capita (about \$20) in the month before the baseline. Moreover, the average z-score for length-for-age and weight-for-age is near or below the WHO's cutoffs

for stunting and underweight. The incidence of poverty and related health indicators reflect both the circumstances of rural Bihar, as well as the project's intentional focus on recruiting disadvantaged households as beneficiaries. Since following results comparing consumption impacts rely on comparing treatment and control villages, the means are reported separately for these two groups, with a t-statistic for the significance of the difference. The t-statistics are calculated from a regression of each variable on a treatment dummy, clustering at the unit of randomization, and including block fixed effects to account for the stratification procedure used during randomization. The last column shows p-values calculated using Cameron, Gelbach, and Miller's wild cluster bootstrap procedure to account for the fact that only eighteen clusters were subject to randomization (Cameron, Gelbach and Miller, 2008). Given the small number of clusters, randomization did not quite achieve balance on all baseline characteristics, with six out of twenty variables not balanced across treatment and control. However, the primary outcome variable of interest relying on the randomization, consumption of rice, is well balanced. To be careful, the regressions identified through the randomization will use changes in consumption from the baseline, differencing out any discrepancy in levels between treatment and control following Jensen and Miller (2011).

(Table1)

In Bihar, the staple foods are mostly rice, wheat, or maize. These three grains account for 33% of food expenditures in the households surveyed, and 30% of these households cultivate at least one of these grains. Among the three available staple grains, rice accounts for the largest share of expenditure and consumption. Including home-produced goods valued at the median market rate in the village, rice alone accounts for 25% of expenditures on food. In addition to its importance in the diet of rural Bihar's population, rice is an inferior good. Table 2 shows the relationship between the share of food expenditure, including consumption from home production, and the log of monthly household expenditure. The correlation is negative, indicating that wealthier households spend a lower percentage of their total expenditure specifically on rice.

(Table2)

Given its importance in meeting the food needs of households and its relative importance for poor households in particular, rice appears to be the kind of candidate that policymakers would have in mind when trying to choose a food to subsidize that will be most valued by

the poor.

Other characteristics of rice, however, make it a less ideal candidate for a program seeking to raise the consumption of staple foods by the poor. The main issue is that rice is not a homogenous good. Even within a single market or shop, a diversity of varieties and qualities of rice are often available. Rice varies according to color, taste, texture, presence of impurities, whether grains are broken or whole, and other dimensions.² It is reasonable to expect that some of this variation in quality and variety is associated with prices paid for rice. Figure 1 shows the distribution of prices paid by households at the baseline for 1 kilogram of rice at the market prior to the baseline survey. The prices range from around 10 rupees per kilogram to over 20 rupees, suggesting that some members pay nearly twice as much for rice as others. Some of this variation is undoubtedly explained by geography, bulk discounting, and negotiating ability, but some is also likely explained by demand for quality.

(Figure1)

If there is a quality component underlying the price variation of market purchases and quality is positively associated with price, we would expect that wealth will be positively associated with quality choice and consequently with the price paid. Table 3 shows the linear relationship between the unit price paid at the market for rice (in the month prior to the baseline survey) with wealth measured in various ways. All four measures of wealth show a positive relationship between the price paid at the market and wealth. The two which are likely to be measured with the least error, household expenditures and expenditure on food (excluding rice and wheat), show the most significant relationship.

(Table3)

Another way to view the relationship between the price of rice and wealth is shown in Figure 2. The figure shows the kernel density plots of prices paid for rice (in the month prior to the baseline survey) for the lowest quartile of wealth and the highest quartile. The plots reveal that the sorting based on wealth is not perfect; nevertheless, those paying lower prices seem to be more likely to be from the poorest quartile, and those paying higher prices seem

²Variation in nutritional content is much more limited, however. Of the seven categories for rice listed in the publication Nutritive Value of Indian Foods, the difference in calories between the most and least calorically dense was less than 20% (?). Varieties fortified with micronutrients may be an exception, but interviews with households and shopkeepers suggested that people were not aware of these varieties. This suggests that either fortified varieties are not generally available or that their nutrient content is not a salient property affecting consumer choices.

more likely to be from the highest quartile.

(*Figure2*)

3 A Simple Model of Take-up Under Heterogeneity of the Staple Good

The descriptive evidence in the previous section suggests that rice is an important staple good in Bihar, and is particularly important as a component of the consumption basket for the poorest households. These characteristics suggest that, in this context, a program which attempts to target participation and benefits toward the poorest members should consider subsidizing rice. However, the evidence also suggests that rice is not a homogeneous good. When broadly defined, rice seems to be an inferior good, meaning that the relative share of rice consumed out of the total consumption basket is decreasing. However, the quality component of rice is likely a normal good, meaning that household demand for quality is increasing in household wealth. To explain why heterogeneity of the staple good may be problematic for targeting, the model below extends a model of take-up similar to those of Jacoby (1997) and Besley and Coate (1991). The model demonstrates that wealthier households' preference for higher quality foods causes programs that provide lower quality, less costly food items to attract beneficiaries who are relatively poor. Programs that offer higher quality goods attract wealthier participants. To determine which types of participants will choose to participate, I compare the utility that a household of a certain income would receive from participating in the program with the utility they would receive from choosing to forgo the program and purchase their food from the market.

The household consumes two goods, a staple food good c and a composite good n . The staple food has a quality component q . Both goods and the quality component can be chosen continuously, but q is bounded above and below, so that there is a minimum quality available

and a maximum.³ The utility function is:

$$U = \alpha \ln(c - a) + \beta \ln(q) + \delta \ln(n)$$

where α , β , and δ are all in the range (0,1). This utility function captures a few key intuitions about the choice of consumption and quality. First, the MRS between quality and the non-staple consumption good n does not depend on the amount of staple good consumed c (weak separability). In other words, the appeal of taste, appearance, ease of cooking, etc. for the different varieties of the staple good relative to consuming more non-staple goods does not depend on the amount of staple good consumed. Second, food consumption is subject to a subsistence constraint, which captures the intuition that as food consumption falls to levels needed for survival, the marginal utility of consuming food becomes very high.⁴

The budget constraint is simply:

$$y = pqc + nz$$

where y is income (or wealth in this single period model) and z is the unit price of n . The key feature of this model is that the price of c is assumed to increase in q , and q is endogenously chosen within the model along with c and n , with the choice of q directly determining the price paid. To keep the model simple, I assume that the unit price of c is a simple linear function of q :

$$price = pq$$

This function captures the idea that higher quality goods cost more per unit to buy. q is restricted to be in the range $q_{min} \leq q \leq q_{max}$.

Ruling out the case where every consumer chooses the corner solution $q = q_{min}$ requires one restriction on parameters, namely that $\beta > \alpha$, which I will assume for the remainder of the paper. In this simple setup, n and q are both increasing in y , but c is a constant in

³This assumption does not drive the results that follow, and only highlights the cases where all members with wealth above or below a certain cutoff are included or excluded. This restriction could be replaced by one in which wealth is bounded above and below so that the only minimum and maximum varieties available in a village are those preferred by the wealthiest and poorest members. Alternatively, quality could be unbounded, in which case any quality choice would create a threshold above and below which no one would participate. If quality is unbounded below, the restriction on parameters I will make to prevent all households from choosing the minimum quality instead ensures that the optimal choice of c and q at market rates is positive.

⁴This assumption does not drive targeting results, but allows for the possibility that the poor place the highest value on program participation when the program subsidizes their most preferred good. This intuition is part of the reason policymakers might view such a program as pro-poor.

preferences.⁵ Households consume a fixed quantity of the staple good but as their wealth increases, they both upgrade the quality of the staple good and consume more of the other good, which may include other non-staple foods.

Introducing a subsidy or transfer into the above model reveals the mechanism through which targeting can result from the choice of good to be subsidized. If the program offers a cash grant, the program would simply expand y in the budget constraint. Households would respond by increasing both q and n , and would consequently experience an increase in utility. A program that offers a voucher not restricted to one particular type of food could be modeled as a decrease in p , for at least up to a maximum quantity for all possible qualities. In the voucher case in this simple model, households will increase q , the quality of rice they desire.⁶ Again, the vouchers would be a welfare improvement for households of all wealth levels. Equivalently valued vouchers will create a higher welfare surplus to poor households who have a higher marginal utility of quality and quantity, but this alone would not lead to differential participation on the wealth dimension.⁷ In either the cash or voucher case, all households would be willing to take up the program, so if the policy maker wants to target the program to lower income households, exogenous rules would have to be put in place to restrict higher income households from participating or to encourage take up among poor households.

In contrast to flexible voucher or subsidy programs, a program that subsidizes the cost of one particular variety of food will have different welfare implications for different types of households. If the program only subsidizes one good, then for participants in the program, q is exogenously chosen and fixed at the level $q = q_p$.⁸ Households receive a discount for this variety, so that the price of a unit of c for the quality q_p is $\phi p q_p$, where $\phi < 1$ and the difference between ϕ and 1 represents the per unit value of the subsidy. Other than this

⁵Details of the optimal choices are in Appendix 1. Assuming that prices could be a more flexible function of quality could generate a solution where the optimal c is also increasing in wealth. This is not necessary for the intuition of the results, but would admit the possibility of the Giffen effects previously described by Jensen and Miller (2011)

⁶In a more sophisticated model where the price is a more general function of quality, households could also choose to change their consumption of c or n in response to a voucher program.

⁷It is possible that vouchers could lead to separation in take-up between rich and poor if vouchers are tied to specific shops that are more difficult to access than the ones households otherwise normally use. This variety of ordeal mechanism is explored by Dupas, Hoffman, Kremer, and Zwane (2013). But if vouchers are flexible enough to be used in the same places where households generally buy grain, they should be appealing to rich and poor alike.

⁸ q may be endogenously chosen at the village level in that the village choice is a function of the wealth of the residents of that village. For now, we assume that individual households cannot influence the village choice, so that from the perspective of a single household, the quality choice is exogenous.

price discount for the exogenously chosen q , the budget constraint and utility function for a program participant is the same as before, and households optimally choose n and c .⁹

To ensure that the optimal choices are non-negative, I assume throughout that $y > a\phi pq_{max}$ for all households, which means that every household is able to afford a quantity of the staple good at least as high as the subsistence constraint a at the prices the program would offer for the highest quality grain q_{max} . In contrast to the market case, where c was chosen endogenously, if a household participates in the program, optimal levels of both c and n will be increasing in income, y . To determine whether a given household would choose to participate or not, I define two indirect utilities. Define $u(c_p^*, n_p^*, q_p)$ as the indirect utility a household receives from taking the quality of their staple food as exogenously given by the program. $u(c_m^*, n_m^*, q_m)$ is the indirect utility the household would receive from purchasing all their goods from the market without the subsidy and endogenously choosing q . Assuming that the optimal q is in the range (q_{min}, q_{max}) ¹⁰, the surplus utility from participating in the program, relative to not participating, can then be given by the following quantity:

$$W = u(c_p^*, n_p^*, q_p) - u(c_m^*, n_m^*, q_m) = \alpha \ln \left(\frac{y - a\phi q_p}{\phi q_p} \right) + \beta \ln \left(\frac{q_p}{y} \right) + \delta \ln \left(\frac{y - a\phi q_p}{y} \right) + K$$

where K is a function of preference parameters and is constant in q_p and y .

Having defined the surplus utility from participating in the program, we can use this quantity to show that wealthier households will opt to participate when the program chooses a higher quality grain to provide. The relationship between household wealth and surplus utility from program participation can be stated in the following proposition:

Proposition 1: For a fixed quality chosen by the program q_p , there exist $y_L(q_p)$ and $y_H(q_p)$ with $y_L(q_p) < y_H(q_p)$, such that $W(q_p, y) > 0$ for all y in (y_L, y_H) , and $W < 0$ for all other y .

(Figure 3)

Proposition 1 holds for general parameters of the model as shown in Appendix 2, but Figure 3 provides a more intuitive depiction of the proposition by plotting $W(q_p, y)$ over y , choosing representative values for the preference and price parameters while keeping q_p fixed at various levels. The solid line shows the surplus utility from participating in the program

⁹The exact expressions for c and n are shown in Appendix 1.

¹⁰The optimal q may be a corner solution as shown in Appendix 1. The specific cases where corner solutions for quality apply are handled separately in the proofs.

relative to the market, $W(q_p, y)$, assuming that the program chooses the minimum available quality. For such a program, $W(q_{min}, y)$, the surplus value of participating in the program rather than the market, is highest for lowest income households. Intuitively, this is because poor households have the highest marginal utility of consumption (c), and they already optimally choose to purchase the lowest quality good available in the market without the subsidy. When the program is introduced, they opt into the program and benefit from moving away from the subsistence constraint. As consumers move up the income distribution, a gap opens up between q_m^* , the quality households would choose to purchase from the market, and q_{min} , the quality they would have to purchase through the program. For moderately low income households, the gain in utility from the extra consumption they can purchase at the lower prices makes it worthwhile for them to participate, and they are willing to switch from slightly above-minimum quality food to minimum quality food in order to participate. But as income increases, the marginal utility of consumption falls, while difference between quality they would choose at the market and the quality chosen by the program widens. For higher income households, the extra consumption they get by switching to q_{min} and participating in the program does not compensate them for having to consume a substantially lower quality good. When the program offers the minimum quality good, households with income in the range $[0, B]$ shown in the graph are willing to participate, and households with income greater than B opt out.

Comparing households with wealth lower than B in the graph with households with wealth above B demonstrates the result described by Jacoby (1997). By choosing the lowest quality good, low income households self-select into the program, and high income households self-select out without the policymaker otherwise needing to fix rules for who is allowed to participate.

The above shows how focusing a program on inferior goods can self-target the program toward relatively poor households. But given the heterogeneity in cost and quality of rice, a program which allows choice over the good is not guaranteed to produce this result. The dashed line in Figure 3 again shows $W(q_p, y)$, the surplus utility from participating in the program and buying the rice supplied by the program rather than their optimally chosen variety of market priced rice, over the income axis. This time W is shown assuming that the program supplies the highest quality rice available in the market, ie $W(q_{max}, y)$. For the figure, I assume that the subsidy is not so large as to make the highest quality rice under the program prices cheaper than the lowest quality rice. Choosing the highest quality rice available reverses the targeting result shown above. Now, households with wealth below the

point marked C in the graph would choose not to participate, since the utility they would get from consuming lower quality rice at a lower price would be more than the utility they would get from upgrading to the program's higher quality rice. For wealth levels above C , households experience a utility premium from buying the discounted program rice, rather than their optimal basket at market prices, and choose to participate.

Comparing the solid line and the dashed line demonstrates that, in the simplified case where the program can only buy either the highest or the lowest quality/price rice available, the choice of rice totally determines whether the program is targeted toward the poor or captured by relatively wealthy households. Since there is no reason to assume that markets will only make two kinds of rice available, the model assumes that grain varieties and prices available from the market are on a continuum between the highest and lowest available quality of rice. It is possible that the program could choose one of these intermediate varieties.

The dotted line in Figure 3 shows the surplus utility of program participation for the choice of an intermediate good. Not surprisingly, the take-up implications of this choice are halfway between what we see from the choice of the lowest or highest quality good. For values of wealth below the point A , households choose not to participate. Above the level of wealth D , wealthy households also self-select out. This curve demonstrates that there may be excluded populations at both the upper and lower ends of the income distribution, but such regions need not exist. If the income distribution in a village is bounded within the range $[A - D]$, the middle quality of grain shown by the dotted line would appeal to both low and high income types.

Proposition 1 says that when q is fixed at some q_p , the function $W(q_p, y)$ defines a range of $y = (y_L, y_H)$ such that $W(q_p, y) > 0$ for all q_p in this range where y_L, y_H are functions of q_p . In Figure 3, $(y_L(q_{min}), y_H(q_{min})) = (0, B)$, $(y_L(q_{max}), y_H(q_{max})) = (C, \infty)$, and $(y_L(q_p), y_H(q_p)) = (A, D)$ for another q_p for which $q_{min} < q_p < q_{max}$.

Figure 3 plots $W(q_p, y)$ for representative values of the preference parameters and prices, but Proposition 1 is true for general values of these parameters, which I prove in Appendix 2.

Having demonstrated that fixing the program quality q_p defines a range of income such that people within the range are willing to participate, I now turn to the relationship between this income region and the quality of the good provided. In Figure 3, I showed the surplus utility of program participation plotted for three different qualities selected by the program: q_{min}, q_{max} , and a quality q_p such that $q_{min} < q_p < q_{max}$. The range of income

such that people would be willing to participate in a program that offers grain of quality q_p is $[0, B]$. When the program offers a higher quality of grain q_p , the lower bound of the income range of participants goes up to A and the upper bound goes up to D , suggesting that the range participants' income is shifted to the right. As the program further increases the quality offered to q_{max} , the lower bound of participants income shifts up to C , and the upper bound extends to ∞ . The plot suggests that the as the quality of grain offered by the program increases, the range of participants' income will shift to the right. This feature is generally true for the assumed parameters of the utility function and can be stated more formally in the proposition below:

Proposition 2 : $\frac{\partial(y_L)}{\partial(q_p)} > 0$ and $\frac{\partial(y_H)}{\partial(q_p)} > 0$

Propositions 1 and 2 allow me to make predictions about the choice of q_p and the targeting implications for the program.¹¹ Typically, improving targeting refers to one or both of two goals: 1. Increasing the proportion of people below some cutoff who participate (desired inclusion) and 2. Decreasing the proportion of people above that cutoff who participate (desired exclusion). Focusing on this definition of targeting can obscure the relationship between quality choice and takeup, however; for a particular choice of program quality q_p , the income distribution between $y_L(q_p)$ and $y_H(q_p)$ determines the percentage of people above and below any particular cutoff who participate.

To see the relationship between q_p and the traditional measures of targeting, suppose that the program sets an income cutoff to assess targeting at the point marked C in Figure 3, so that the program intends to include households with income below C and exclude households with income above C . Then either a program that provides the lowest quality grain or one that offers the intermediate quality shown on the graph would be strictly better targeted than a program that provides the highest quality grain. Only the highest quality program out of the three would attract participants with income above C , and only the two lower quality programs attract any participants with income below C . But the comparison between the low and intermediate quality is ambiguous, because only the low quality program would attract those with income below A , while only the intermediate quality program would attract those with income between B and C . Which of the two programs attracts a higher percentage of those below the cutoff C will depend entirely on the number of households with income below A relative to the number of households whose income is between B and C .

More formally, the percentage of households below an income cutoff y^* can be written

¹¹The proof for Proposition 2 proceeds exactly as in the proof of Lemma 3 in Appendix 2.

as:

$$I = \int_0^{y^*} \pi(y)f(y)dy$$

where $\pi(y)$ is the probability that a household with wealth y will choose to participate in the program.¹² The distribution of income $f(y)$ is a feature of the data sampled or the population of interest, and y^* can be chosen either by the policy-maker or the econometrician. Assessing the simple comparative static $\frac{\partial I}{\partial q_p}$ depends on the shape of $f(y)$, the y^* chosen, and the function $\pi(y)$. Assessing targeting potential empirically by the proportion of members below an arbitrary cutoff who participate based on quality choice is problematic because such a measure would be very sensitive to both the specific income distribution and the cutoff chosen. Such a measure then is limited in external validity—because other contexts will have other income distributions—and prone to data mining—because the econometrician can cherry pick the cutoff that generates a desired result. To avoid confusion about whether results on targeting are driven by $f(y)$ or y^* , the remainder of the paper will focus on the quantity $\pi(y)$. This quantity can be thought of as expressing the proportion of people with income y who participate.

4 VO Choice of Quality as a Function of VO Characteristics

In the Bihar Rural Livelihoods Project, village organizations have a hierarchical structure. Each VO is composed of members of 10-15 Self Help Groups (SHGs), who participate together in VO-level programs like the food security loan. To choose the quality and price of grains provided through the program, members of the VO leadership collect bids from local merchants to supply the estimated amount of grain needed. At a meeting of VO leadership, samples of rice that represent each of the bids are assessed for quality and appearance. In some cases, the rice samples are cooked to allow members to taste the food before choosing a bid to accept. Although any member of an SHG is technically allowed to participate, the VO committees are comprised of leaders appointed by each SHG to represent them in VO-level decisions. Thus it is likely that these members have more influence on program outcomes, both through their attendance in the VO committee meetings, and through their influence as leaders. Such a representative model is a common approach for programs in developing

¹²Or the share of households with income y who participate.

countries to aggregate preferences and use local information in deciding policies and program allocations (Olken, 2010).

In this type of representative choice allocation, the preferences of the chosen VO committee members are more likely to be represented in the choice of goods, including the choice over the quality of grain selected. In the model, wealth is correlated with the price paid for rice is through the channel of preferences for quality, and empirical results in Section 2 showed a relationship between household expenditure and the price paid for rice at market rates prior to the baseline survey. Together, these results suggest that VO representation may underlie part of the choice of quality and explain some of the variation in the prices offered by the program.

(Table4)

Table 4 shows the relationship between VO committee status and household expenditure at baseline. An extra log point of household expenditure is associated with a 6 percentage point increase in the likelihood of being a VO committee member; this suggests that wealthier members are disproportionately represented on the VO committees that choose the quality of rice to provide through the program. Column 2 shows that VO committee members are also more likely to participate in the program than non-committee members. If VO committees are often made up of relatively wealthy members of the SHGs, we might expect these members' preferences to influence the program prices. Table 5 shows the relationship between the median price of rice paid by the VO committee members at the baseline of the survey, as well as the prices of the variety of rice ultimately chosen in the second column. The median market price paid at baseline is significantly correlated with the price the VO ultimately chooses, consistent with the idea that the median committee member is pivotal in choosing a variety of rice. The more this member values a higher quality variety of rice, the more likely it is that the committee will choose a higher cost variety.¹³

(Table5)

¹³Committee membership and prices offered at baseline are not randomly assigned. Any correlation between baseline prices and prices of available varieties when the committee makes this choice could explain the relationship.

5 Empirical Results on Targeting and Take-up

The model above outlines the mechanism through which quality and price of the provided good interact with wealth to determine which segments of the income distribution choose to participate. When communities choose their own grains, it is not clear which type of grain will be chosen, and which types of members will choose to participate as a result. Figure 4 shows the distribution of prices paid for grain by VO members purchasing through the community lending program against the distribution of prices paid by households purchasing from the market in the control villages during the midline survey. Comparing the prices paid by the VOs with prices paid by the households confounds the price differentials from two sources: those associated with quality differences and any discounts members receive by buying through the food security program rather than the market. This confounding means that the apparent leftward shift of the VO distribution relative to the household/market distribution cannot be necessarily attributed only either to a bulk discount or to VOs systematically choosing lower qualities.¹⁴ The main feature of the comparison is that there is significant variation in the prices that VOs pay. The VOs offering the most expensive grains offer the rice at a price about 50% higher than the VOs offering the grains at the lowest price. There seem to be VOs who offer grain at prices equivalent to those experienced by members purchasing the most expensive grains and other VOs offering rice at prices equal to the cheapest prices available in the market. Some of the variation may be spatial or temporal, so that all the VOs are in fact buying the same rice, which happens to be available at different prices in different times and places. However, given the correlation between wealth and market price seen earlier and the correspondence between market and VO prices, it seems reasonable to speculate that some of the difference in VO prices relates to the quality of grain chosen. It therefore seems unlikely that public choice of the staple good in this context will have led either to systematic targeting of the poor or elite capture by the wealthy, but that different VOs will have different participant profiles based on the price/quality choice made by that VO.

(Figure4)

To assess whether variation in take-up is correlated with the variation in prices of the grains offered by the VOs as predicted by the model, I estimate the quantity $\pi(y)$.

¹⁴There is also the issue of heaping at 20 rupees per kg in the market prices. It is not clear whether this heaping comes from stickiness of prices around a round number, measurement error, or both.

Starting from the simplest case, under the assumption that there are only two quality/price combinations of rice available in a particular village, the model predicts that take-up should vary based on whether the high or low quality rice is chosen. The model predicts that the surplus utility from participation should be positive for low levels of wealth and negative for high levels of wealth when the low quality of rice is chosen, and the opposite when the high quality rice is chosen. If the probability of participation is an increasing function of the surplus utility from participating in the program, relative to buying only from the market, the quantity $\pi(y)$ can be estimated by the following simple binary dependent variable model:

$$Y = f(\alpha + \beta_1 * \ln_expenditure + \beta_2 * vo_price + \beta_3 * \ln_expenditure * vo_price) + \epsilon$$

where the choice of the function f depends on the selection of an LPM, logit, or probit model. The model predicts that $\beta_1 < 0, \beta_2 < 0, \beta_3 > 0$, capturing the idea that when the VO chooses the low price good, wealth would be negatively correlated with the probability of participation.

In the above regressions, the price chosen and offered by the VO is not randomly assigned, and may be correlated with characteristics of the VO that are related to both wealth as measured by expenditure and take-up. To avoid this source of bias, I will instead implement a village fixed effects version of the above estimation equation. Since the program price only varies at the village level, we cannot separately identify β_2 , the effect of raising the VO price for the low-wealth households.

(Table6)

Table 6 shows the results of the above specification, using both a logit and LPM specification and including a village fixed effect. The first three columns show an LPM implementation of the model, while the following three columns show a conditional logit model. Including a village fixed effect means that the model can either separately identify within village variation in prices ¹⁵ over rounds or identify an effect of the round on participation, but can not separately identify both. For each model, the results are shown including: 1.No fixed effects for the round or village price, 2. Round dummies, and 3. Including the village price as a regressor. As predicted by the model, the coefficient on $\ln_expenditure$ is negative, suggesting that when the price of the grain chosen by the VO is low, the probability

¹⁵All prices are de-trended from a linear trend in prices where price is a linear function of the month in which the grain is purchased.

of take-up is highest for the lowest income members. This is a confirmation of the results predicted by Besley and Coate (1991) and demonstrated elsewhere by Jacoby (1997). But also consistent with the model allowing for the choice of a higher quality is the fact that the coefficient on $vo_price * ln_expenditure$ is negative. Solving the following equation for VO prices allows me to estimate the price at which household expenditure switches from being negatively correlated with take-up to being positively correlated:

$$\beta_1 * ln_expenditure + \beta_3 * ln_expenditure * vo_price = 0$$

The inflection points range from around the 65th percentile of the price distribution to the 98th percentile depending on the model choice. The interpretation of this fact is that for most of the range of the observed price distribution, participation is relatively associated with lower levels of wealth and expenditure. But when a price around the 90th percentile is chosen, wealth becomes positively correlated with participation. In the higher cost VOs, even when the program provides rice, a staple food which is an especially important component of food expenditure for the poorest members, wealthier members are more likely to participate. This reflects that the wealthier members perceive an additional benefit from participating in the program that the poorer members do not experience.

The above results restricted the probability of participation to be monotonically increasing or decreasing in wealth. In a version of the model where there are only two varieties of rice available in the local market, this restriction is consistent with the model. But when the model allows for a bounded continuum of quality choices, the model predicts that for a particular quality choice between the minimum and the maximum available quality, there can be regions at both the bottom and the top of the income distribution where members choose not to participate. If probability of participation is increasing in the surplus utility from participation, a more realistic representation would be to allow the probability of participation as a function of income to have both increasing and decreasing regions. To allow this loosening of restrictions, Figure 5 shows local polynomial plots of participation over the log expenditure distribution. The dependent axis is a dummy for participation removing VO and round fixed effects. Log expenditure is trimmed at the top and bottom 5 percent of expenditure to exclude outliers where the data is thin.

(Figure5)

The figure shows the relationship between wealth and participation for three different

wealth categories of prices of grains chosen by the VO, corresponding to terciles of the distribution of VO prices. The plots confirm the main features of the model and the logit results. Comparing the relationship for the quarter of VOs with the lowest cost grains (the light colored solid line) with the line for the third of VOs with the highest cost grains (dark solid line), the data show how the model connects to the relationship between participation and prices of rice chosen by VOs in the project. Choosing the least expensive grains is an effective way to self-target the program, because participation is highest among the poor, and the wealthy select out. When relatively expensive grains are chosen, however, the both wealthy and poor participate, suggesting that villages who choose these grains are spending more per kg than would be necessary to induce the poor to participate and encouraging relatively wealthy households to take advantage of benefits. Given that the subsidy on grain is small in this case, the fact that the poor do not select out when their village chooses expensive rice has important consumption implications, described in the following section.

6 Implications of the Model for Average Treatment Effects on Consumption of the Staple Good

To this point, the targeting of the program has been assessed by referring in the model to the wealth range of households who would choose to participate, and empirically referring to the probability of a household of a given wealth in taking up the program. But another way to view whether a program is well targeted is to assess the regions of the income distribution where the program changes the consumption patterns of households. A goal of policymakers might be not only that the program should have higher take-up among relatively poor households, but that the program should increase the consumption of relatively poor households. Additionally, the increase among poor households should be more than the increase among wealthy households.

Under heterogeneity of the good that can be selected for provision through the program, whether consumption is increased by the program and whether these increases are concentrated among the poor will depend on the quality of the good. Returning to the model, the optimal choice of consumption in the simple model when buying the staple good only from the market is:

$$c_m^* = \frac{\beta\alpha}{\beta - \alpha}$$

If they participate in the program, households optimally choose to consume:

$$c_p^* = \frac{\alpha \delta p \phi q_p + \alpha y}{(\alpha + \delta) p \phi q_p}$$

The role of income in determining optimal consumption is different depending on whether the household participates in the program or purchases their staple food from the market. This suggests that the difference between optimal consumption when purchasing from the market and consumption when purchasing from the program will be a function of income.

For an individual, the prograd effect on consumption is 0 if the household chooses not to participate in the program and $c_p^* - c_m^*$ if the household does choose to participate.

Proposition 3: For a fixed choice of program quality q_p among households who choose to participate, the size of the program effect on an individual household's consumption of rice is increasing in wealth.

Assuming that a household is willing to participate, the program effect is:

$$c_p^* - c_m^* = \frac{a \delta p \phi q_p + \alpha y}{(\alpha + \delta) p \phi q_p} - \frac{\beta a}{\beta - \alpha}$$

The derivative of this quantity with respect to income is:

$$\frac{\alpha}{(\alpha + \delta) p \phi q_p} > 0$$

which implies that if everyone were required to participate in the program, for a fixed program quality q_p , the size of the program effect on consumption is increasing in wealth.¹⁶

Proposition 3 says that the program effect in terms of increased consumption will be bigger for wealthier households than for poorer ones, which likely runs counter to the goals of policymakers interested in targeting programs to the poor. But even the policy goal of increasing consumption may not hold under the heterogeneity of the program good explored here. Others, in particular Jensen and Miller, have already demonstrated that consumption effects of a subsidy on inferior goods need not be positive over the income distribution, and can even be negative because of a Giffen effect (Jensen and Miller, 2011). However, the model

¹⁶It may appear that this result is only a product of the fact that at market rates consumption of the staple good is not a function of income, while in program participating households, income does influence optimal consumption. The fact that consumption at market rates results from the assumption that the price of the staple good is simply the quality times a constant (Price = pq). If the price is a slightly more complicated function of the quality (for example, Price = $P_0 + P_1q$), then c_m^* is an increasing function of y . Using these functional forms precludes closed form solutions, but numerical solutions suggest that the results are robust to such an extension.

of heterogeneity in staple goods presented above also introduces the possibility of negative consumption impacts because of a cross price effect within the choice of available goods. When the program affects the price of only one variety out of a continuum of varieties of rice available, some members may choose to upgrade to a higher quality rice (if the program offers a higher quality than they choose to consume at market rates) or to lower the quality of their grain in order to take advantage of the program discount and purchase more grain. This effect is stated in the proposition below:

Proposition 4: Suppose the program subsidizes a fixed quality q_p , where $\phi p q_p > p q_{min}$. For households whose income falls within the range of those willing to participate, there is a level of income such that all households with income below this level have lower consumption when participating in the program than they would if purchasing at market rates. All households with income above the range will consume more of the staple good when participating in the program than from market rates. ¹⁷

Proposition 4 states that for any choice of q_p such that the program price is higher than market price of the lowest available quality of grain,¹⁸ the lowest income households who choose to participate will respond to the program by lowering their consumption of the staple good, and only the highest income participants will show an increase in consumption.

The average treatment effect on consumption is the average of the individual program effects on consumption over the entire range of income. Proposition 4 then yields an interesting result with regard to consumption: the average treatment effects from a program that subsidizes only one quality of rice need not be positive. If the program does not subsidize the grain below the market price of the minimum available quality of rice, some households will respond by using the program to upgrade the quality of their consumption while actually reducing their consumption. Whether the average treatment effect is positive or negative for a given q_p will depend on the distribution of income in the village and the size of the treatment effects at each income level, since the average treatment effect for a fixed q_p can

¹⁷A proof of Proposition 4 is shown in Appendix 5.

¹⁸If the program offers grain of a quality higher than the minimum available at the market, but at a price lower than the market price available of the lowest quality grain, even the poorest households would respond by (weakly) increasing their consumption of the staple food. The reduction in consumption when participating in the program comes from households who respond to the program by choosing to upgrade the quality of the grain they buy, even though the program grain is more expensive than the market grain they would have bought (though not as much more expensive as it would have been without the program). If the program rice is cheaper than any market rice, no one faces higher prices by participating in the program.

be defined as:

$$\int_0^{\infty} (c_p(y) - c_m(y))dy$$

where $c_p(y) = c_m^*$ if a household chooses not to participate, $c_p(y) = c_p^*$ if the household does participate, and $c_m(y) = c_m^*$. This proposition does not imply that consumption impacts across villages making choices of different varieties of rice will be negative for the poorest, as some villages may select low cost varieties for which the only participation we observe comes from households who respond to the program by increasing their consumption.

These propositions demonstrate several important implications of the model: 1. Consumption impacts are heterogeneous on the income distribution, 2. The average treatment effect is only unambiguously positive in the case where the VO chooses to provide the lowest quality of rice available in the local market, 3. If a quality of rice higher than the lowest possible quality is chosen, treatment effects are only unambiguously positive for the wealthiest households, but may be only weakly positive for these households, 4. If higher than the minimum quality of rice is chosen, negative treatment effects are possible and are mostly likely to be observed for households who are not the wealthiest in the village, but only weakly negative for the very poorest.

7 Empirical Results on Consumption Impacts

Table 7 shows the average treatment effects (ATE) on consumption, estimated separately by quintiles of the expenditure distribution for changes from the consumption of rice at baseline to midline or endline. Following Jensen and Miller (2011), I compute the treatment effects as changes in weekly rice consumption in kg from the quantity at baseline in order to help account for any potential imbalances in baseline consumption levels. In particular I estimate the treatment effects as:

$$\Delta RiceConsumed_{ijt} = \alpha + \beta * Treatment_{jt} + \epsilon$$

The goal of policymakers is to increase the quantity of staple foods that households (especially the poor) are able to consume. The first column of Table 7 shows the treatment effects for the whole sample, which are positive, but small and not statistically significant. The model offers an explanation for how a subsidy on rice could fail to increase consumption of rice through the prediction that the only region of the wealth distribution for which the consumption impacts should unambiguously be positive is the right tail, while a region

of moderately poor households may reduce consumption. The far right column of Table 7 reveals that as predicted, the largest increases in consumption of rice occur among the wealthiest households, and for this group, the increase is statistically significant. The other prediction of the model is borne out by the third column, which reveals that for the second quintile of households, the consumption impact is negative. This is the surprising result of the model, that a subsidy on consumer staples reduces consumption of these goods for some members. This could result from Giffen behavior, but as shown in the model, can also arise from cross-price substitution between different varieties of rice.

(Table7)

The lowest quintile of households also shows a significant increase in consumption. However, this finding is not inconsistent with the model because we could expect to see increases in consumption among the poor if enough villages choose to provide low cost grains. As shown in Figure 5, the poorest households participate both when expensive and low cost rice is chosen. Whether the coefficient is positive or negative for this group depends on whether the increase in rice consumption when the village chooses low cost rice is greater than the potential decrease in villages that choose high cost rice. These results demonstrate the only predictions that can be made *ex ante*: consumption of rice will at least weakly increase for the wealthiest households, and consumption could increase or decrease for anyone else. Together these results demonstrate a surprising fact that a program designed to increase the consumption of staple foods by subsidizing rice may not achieve its goal.

8 Conclusion

Although rice seems to be a staple food and an inferior good, the fact that the program under study focused on providing rice did not seem to guarantee relatively high take-up by the poorest eligible beneficiaries. This paper suggests one reason that take-up is not uniformly concentrated among the poor could be that rice is actually a heterogeneous good, and some villages seem to have opted to buy high quality rice. Empirical results from the program evaluation show that the choice of rice seems to predict participation of relatively poor or wealthy households, consistent with the model. These results suggest that if programs want to self-target by offering inferior goods, and also allow for public choice, they need to have an additional mechanism to ensure that the preferences of the poorest members are given more weight than occurs in a representative candidate model like the one employed by the program

under study. Understanding the processes by which some VOs seem to successfully select goods that are preferred by the poorest members and designing mechanisms to encourage this choice could reveal mechanisms that make targeting and public choice more compatible.

There is an intuitive appeal in tying a subsidy to goods preferred by the poor in order to self-target the subsidy and maximize its impact in terms of increased consumption of staple foods by the poor. However, this paper shows how difficult it is to implement such a program when limited information about markets and preferences is available. Without a clear choice of a good to subsidize, decentralizing the choice to potential beneficiaries may seem like a way to incorporate local preferences. Unfortunately, allowing this choice unravels three goals of the program: 1. Wealthy households often participate because their village selects high quality rice that they prefer, 2. The program is able to purchase a smaller quantity of rice for a given subsidy grant than if all villages selected the lowest cost rice, and 3. The program does not increase the consumption of staples on average because many households end up purchasing higher cost, high quality rice. As in other investigations of self-targeting mechanisms, achieving self-targeting proves to be difficult.

References

- Alatas, Vivi, Abhijit Bannerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Po.** 2012. “Ordeal Mechanisms in Targeting: Theory and Evidence from a Field Experiment in Indonesia.” *CID Working Paper*, , (254): 175–191.
- Atkin, David.** 2013. “The Caloric Costs of Culture: Evidence Indian Migrants.” *NBER Working Paper*, , (19196).
- Barrett, Christopher B.** 2002. “Food Security and Food Assistance Programs.” In *Handbook of Agricultural Economics, Vol. II.* , ed. Bruce B. Gardner and Gordon C. Rausser, 2103–2190. Elsevier Science B.V.
- Besley, Timothy, and Stephen Coate.** 1991. “Public Provision of Private Goods and the Redistribution of Income.” *The American Economic Review*, 81(4): 979–984.
- Cameron, A. Colin, Jonah Gelbach, and Douglas Miller.** 2008. “Bootstrap-Based Improvements for Inference with Clustered Errors.” *Review of Economics and Statistics*, 90(4): 414–427.

- Cunha, Jesse M., Giacomo De Giorgi, and Seema Jayachandran.** 2011. “The Price Effects of Cash Versus In-Kind Transfers.” *NBER Working Paper*, , (17456).
- Currie, Janet, and Firouz Gahvari.** 2008. “Transfers in Cash and In-Kind: Theory Meets the Data.” *Journal of Economic Literature*, 46(2): 333–383.
- Dupas, Pascaline, Vivian Hoffman, Michael Kremer, and Alix Zwane.** 2013. “Micro-Ordeals, Targeting and Habit Formation.” *Unpublished Manuscript*.
- Jacoby, Hanan G.** 1997. “Self-Selection and the Redistributive Impact of In-Kind Transfers: An Econometric Analysis.” *The Journal of Human Resources*, 32(2): 233–249.
- Jacoby, Hanan G.** 2002. “Is there an Intrahousehold Flypaper Effect? Evidence from a School Feeding Programme.” *Economic Journal*, 112(476): 196–221.
- Jensen, Robert, and Nolan Miller.** 2011. “Do Consumer Price Subsidies Really Improve Nutrition?” *Review of Economics and Statistics*, 93(4): 1205–1223.
- Jha, Shikha, and P.V. Srinivasan.** 2001. “Taking the PDS to the Poor: Directions for Further Reform.” *Economic and Political Weekly*, 36(39): 3779–3786.
- Lentz, Erin C., Simone Passarelli, and Christopher B. Barrett.** 2013. “The Timeliness and Cost-Effectiveness of Local and Regional Procurement of Food Aid.” *World Development*, 49: 9–18.
- Mansuri, Ghazala, and Vijayendra Rao.** 2013. “Localizing Development: Does Participation Work?” *World Bank Policy Research Reports*.
- Nichols, Albert L., and Richard J. Zeckhauser.** 1982. “Targeting Transfers through Restrictions on Recipients.” *The American Economic Review*, 72(2): 372–377.
- Olken, Benjamin A.** 2010. “Direct Democracy and Local Public Goods: Evidence from a Field Experiment in Indonesia.” *American Political Science Review*, 104(2): 243–267.
- Violette, William J., Aur Elie P. Harou, Joanna B. Upton, Samuel D. Bell, Christopher B. Barrett, Miguel I. Gomez, and Erin C. Lentz.** 2013. “Recipients’ Satisfaction with Locally Procured Food Aid Rations: Comparative Evidence from a Three Country Matched Survey.” *World Development*, 49: 30–43.

TABLE 1: CHARACTERISTICS OF SAMPLE POPULATION AND BALANCE CHECK

Variable	(1) Mean in Treatment	(2) Mean in Control	(3) Difference T-C	(4) t-value for Difference	(5) CGM Wild Cluster BS p-value
Number of HH members	5.68	5.69	-0.01	-0.22	0.74
Dummy for Household is a General Cast Hindu	0.06	0.07	-0.01	-0.46	0.43
Dummy for Head Literate	0.32	0.38	-0.05	-0.79	0.18
Per Cap. Monthly Income	1258.7	1281.3	-22.6	-0.77	0.21
Monthly HH Expenditure on Non-Food Items	4965.1	4790.1	175.0	1.62	0.03
Monthly HH Expenditure on Food	3492.3	3385.4	106.9	0.03	1.00
Total Monthly HH Expenditure	8457.4	8175.5	281.9	1.56	0.05
Number of Assets Owned	6.42	6.70	-0.28	-2.60	0.01
Total Household Debt	13733	28175	-14441	-1.00	0.12
Acres of Land Cultivated in Past 1 Year	1.47	1.94	-0.47	-2.29	0.01
Dummy for Household Cultivates Grains	0.25	0.36	-0.11	-4.08	0.00
Kg. of Rice Purchased in Last 1 Month	54.22	51.88	2.34	0.82	0.27
Kg. of Rice Consumed in Last 1 Week	10.15	9.99	0.15	0.75	0.30
Price per Kg. of Rice purchased from Market	17.45	17.66	-0.21	-0.31	0.78
Kg. of Wheat Purchased in Last 1 Month	58.14	61.14	-3.00	-0.59	0.23
Kg. of Wheat Consumed in Last 1 Week	9.86	10.07	-0.22	0.52	0.40
Price per Kg. of Wheat purchased from Market	13.76	16.81	-3.05	-0.04	0.98
Z-score Length for Age for children < 5 years old	-3.27	-3.32	0.05	-0.88	0.25
Z-score Weight for Age for children < 5 years old	-1.96	-2.19	0.23	0.52	0.41
Z-score Weight for Length for children < 5 years old	-0.50	-0.96	0.46	2.86	0.00

Notes: Columns 1-3 are unadjusted means in the treatment and control groups. T-statistic (Col. 4) is computed clustering at the cluster level where clusters are groups of 5 VOs in a regression that includes block fixed effects to account for the fact that randomization was done by stratification at the block level. Col. 5 reports Cameron, Gelbach, and Miller's wild cluster bootstrap correction for the standard errors with 999 replications and the same regression as Col. 4.

TABLE 2: RELATIONSHIP BETWEEN WEALTH AND SHARE OF EXPENDITURE ON RICE

VARIABLES	(1) Percent of HH Food Expenditure on Rice	(2) Percent of HH Food Expenditure on Rice
Log of Total HH Expenditure	-0.0474*** (0.00339)	-0.0474*** (0.00321)
Block FE	NO	YES
Constant	0.417*** (0.0590)	0.403*** (0.0497)
Observations	1528	1528
R-squared	0.005	0.019

Note: Robust standard errors are clustered at the cluster level,
where clusters are groups of 5 VOs.

*** p<0.01, ** p<0.05, * p<0.1

TABLE 3: RELATIONSHIP BETWEEN UNIT COST OF RICE AND MEASURES OF HOUSEHOLD WEALTH

VARIABLES	(1) Log Price of Rice Rs/Kg	(2) Log Price of Rice Rs/Kg	(3) Log Price of Rice Rs/Kg	(4) Log Price of Rice Rs/Kg
Log HH Expenditure	0.0228*** (0.00711)			
Log Food Exp (Excluding Wheat and Rice)		0.0256** (0.0101)		
Asset Index			0.00738 (0.00821)	
Log of HH Income				0.00514 (0.00632)
Village FE	YES	YES	YES	YES
Constant	2.657*** (0.0627)	2.662*** (0.0768)	2.859*** (0.000858)	2.815*** (0.0534)
Observations	1,089	1,089	1,089	1,077
R-squared	0.006	0.006	0.001	0.001
Number of VOs	90	90	90	90

Note: Robust standard errors are clustered at the VO Level.

*** p<0.01, ** p<0.05, * p<0.1

TABLE 4: VO COMMITTEE MEMBERSHIP, WEALTH, AND PARTICIPATION

VARIABLES	(1) =1 if Member of VO Committee	(2) =1 if Member Purchased in a Given Round
Log Total HH Exp.	0.0602*** (0.0144)	
=1 if Member of VO Committee		0.0368* (0.0202)
Constant	-0.184 (0.129)	0.428*** (0.0276)
Observations	1,528	2,946
R-squared	0.009	0.001

Note: Standard errors are clustered at the VO level
 *** p<0.01, ** p<0.05, * p<0.1

TABLE 5: RELATIONSHIP BETWEEN PRICES PREFERRED BY
 COMMITTEE MEMBERS AND PRICE OF VARIETY CHOSEN

VARIABLES	(1) VO Price
Median Price per KG of Rice Among VO Committee Members at Baseline	0.267* (0.121)
Round FE	YES
Block FE	YES
Constant	15.72*** (2.069)
Observations	142
R-squared	0.203

Note: Standard errors are clustered at the cluster level,
 where clusters are groups of 5 VOs.
 *** p<0.01, ** p<0.05, * p<0.1

TABLE 6: MEMBERS' PARTICIPATION BY WEALTH AND PRICE OF VARIETY CHOSEN

VARIABLES	LPM Model			Conditional Logit		
	(1) =1 if Participated	(2) =1 if Participated	(3) =1 if Participated	(4) =1 if Participated	(5) =1 if Participated	(6) =1 if Participated
Tot. HH. Exp.	-0.0580*** (0.0207)	-0.0448** (0.0201)	-0.0645 (0.103)	-0.281*** (0.0968)	-0.220** (0.0926)	-0.285 (0.481)
Tot. HH. Exp. X VO Price	0.00241*** (0.000717)	0.00176** (0.000701)	0.00273 (0.00526)	0.0118*** (0.00322)	0.00874*** (0.00309)	0.0120 (0.0246)
Round 2		-0.149*** (0.0394)			-0.716*** (0.191)	
Round 3		-0.178*** (0.0410)			-0.823*** (0.200)	
Round 4		-0.154*** (0.0560)			-0.728*** (0.253)	
VO Price			-0.00285 (0.0438)			-0.00169 (0.208)
Observations	2,241	2,241	2,241	2,241	2,241	2,241
R-squared	0.009	0.033	0.009			
Number of VOs	45	45	45	45	45	45

Note: LPM Models include VO level fixed effects. Conditional logit models include VO as the group variable. All regressions are clustered at the VO level. Total Household Expenditure includes food and non-food expenditure estimated in the baseline survey. VO Price is the VO level median of the households' reported price offered by the VO for each round. Prices are de-trended from a linear time trend in months to account for inflation across rounds and across months within rounds.

*** p<0.01, ** p<0.05, * p<0.1

TABLE 7: CHANGE IN CONSUMPTION OF RICE BY WEALTH QUINTILE

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Δ Rice Cons. (kg)	Δ Rice Cons. (kg)	Δ Rice Cons. (kg)	Δ Rice Cons. (kg)	Δ Rice Cons. (kg)	Δ Rice Cons. (kg)
	ALL HHs	1 st Quintile	2 nd Quintile	3 rd Quintile	4 th Quintile	5 th Quintile
Treatment	0.131 (0.305)	0.946 (0.667)	-1.390* (0.724)	0.255 (0.879)	-0.00762 (0.467)	1.170* (0.615)
Constant	[0.536] -1.819*** (0.335)	[0.086] -1.628* (0.822)	[0.024] 0.538 (1.500)	[0.696] -2.130* (1.196)	[0.844] -1.594*** (0.529)	[0.082] -4.459*** (0.705)
Block FE	YES	YES	YES	YES	YES	YES
Round FE	YES	YES	YES	YES	YES	YES
Observations	2,818	559	569	562	561	567
R-squared	0.009	0.025	0.022	0.010	0.027	0.059

Note: Robust standard errors in parentheses, clustered at the cluster level, where clusters are groups of 5 VOs. P-values from Cameron, Gelbach, Miller's (2008) wild cluster bootstrap in brackets. The dependent variable is the change in kg of rice consumed in the previous week from the week prior to baseline data collection to the week prior to midline or endline data collection.

*** p<0.01, ** p<0.05, * p<0.1

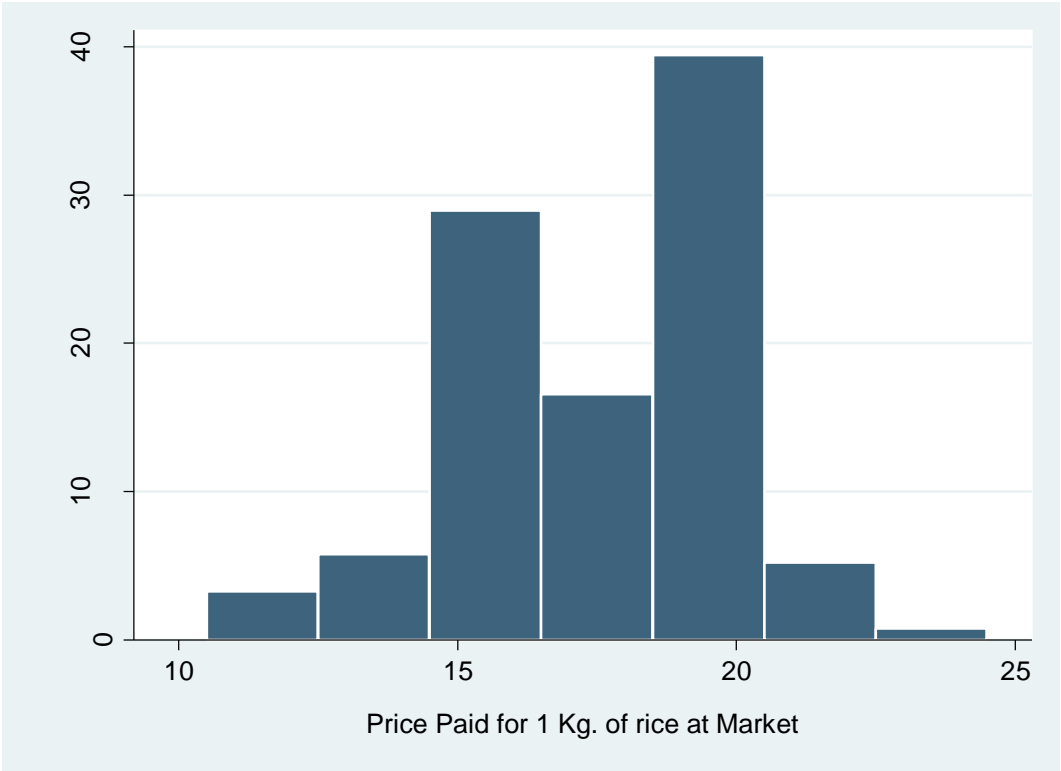


FIGURE 1: DENSITY OF UNIT COSTS FOR 1 KG OF RICE

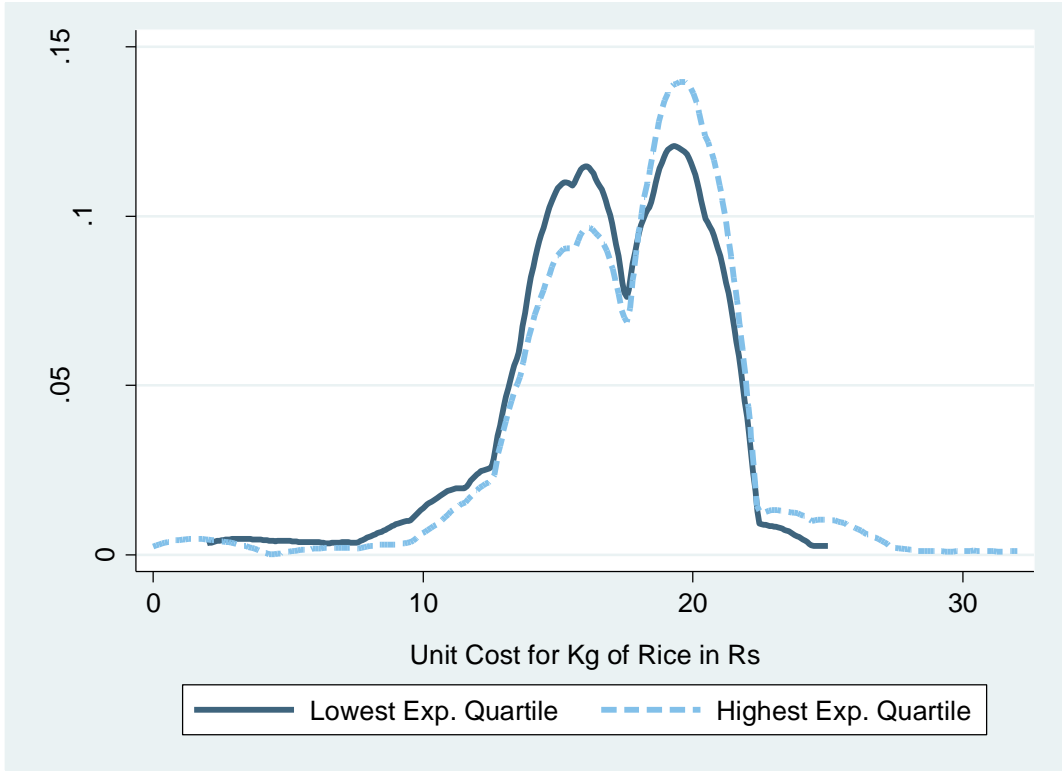


FIGURE 2: DENSITY OF MARKET UNIT COSTS FOR 1 KG OF RICE BY HH EXPENDITURE

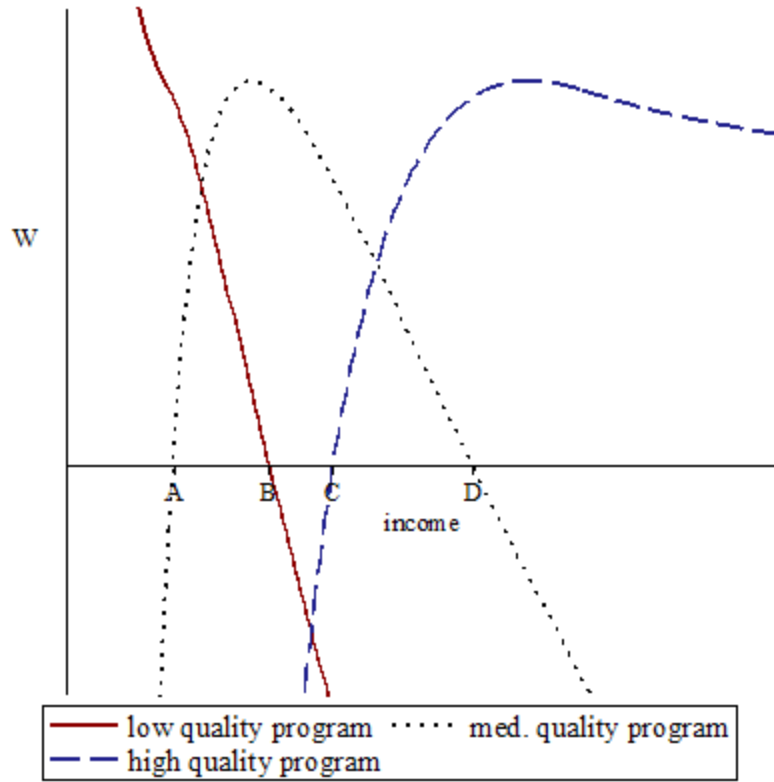


FIGURE 3: SURPLUS UTILITY FROM PROGRAM PARTICIPATION

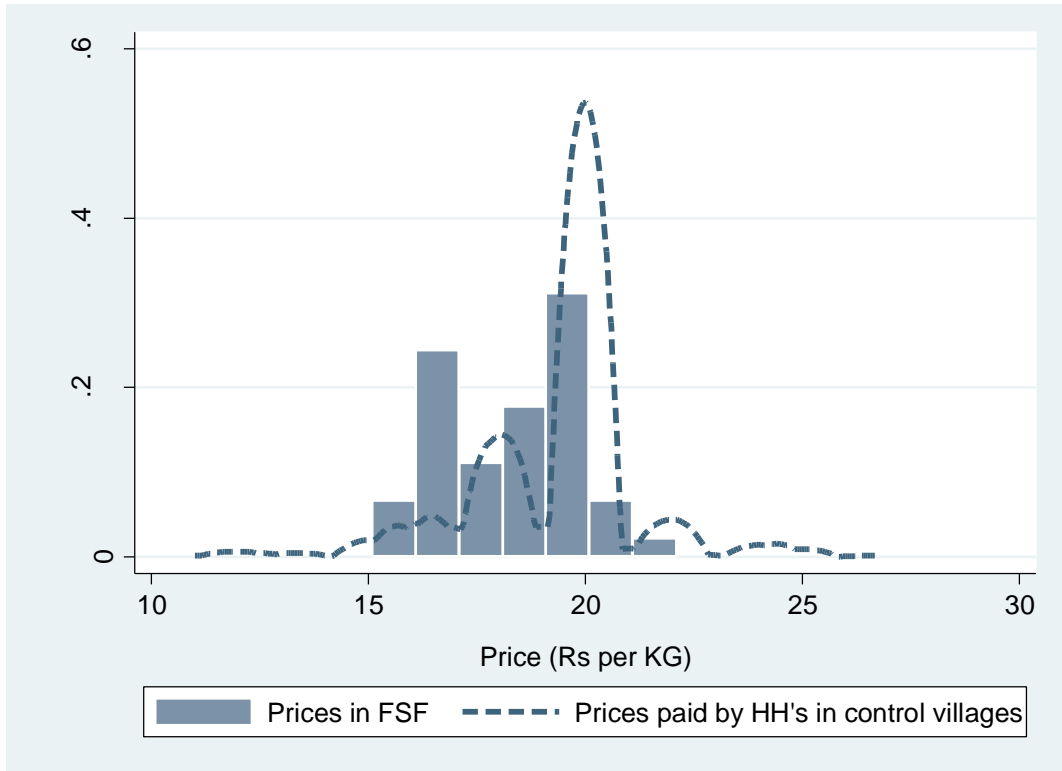


FIGURE 4: DENSITY OF PRICES FROM FSF PROGRAM (TREATMENT VILLAGES) AND MARKET (CONTROL VILLAGES)

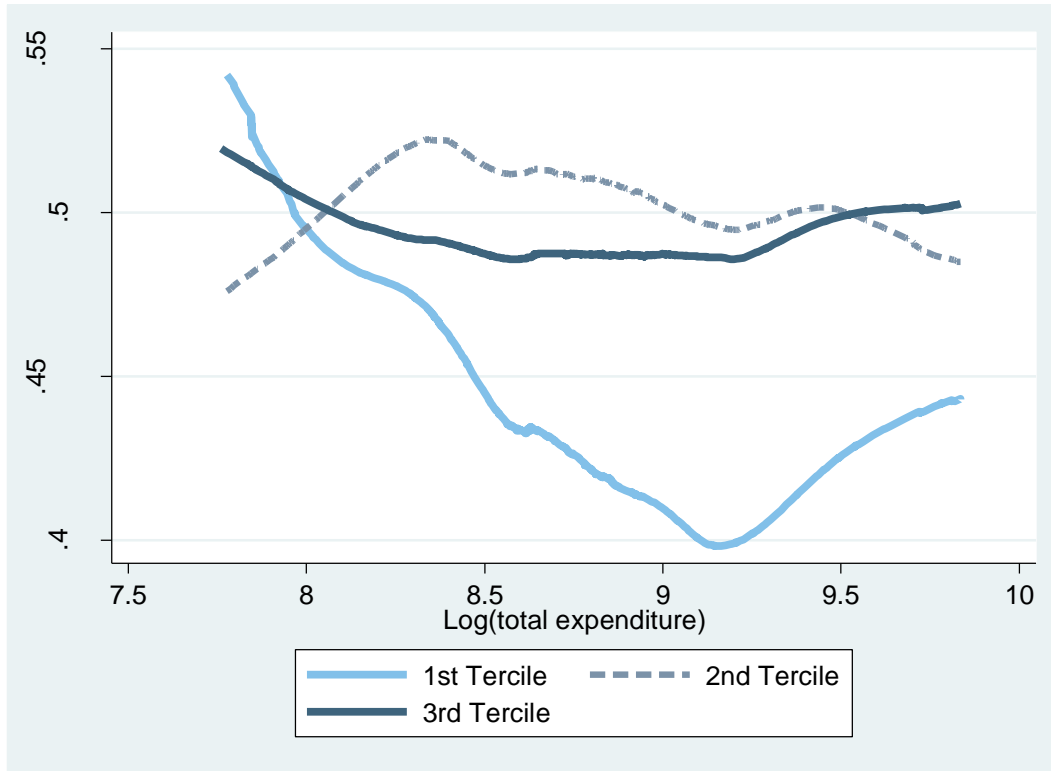


FIGURE 5: NON-LINEAR PARTICIPATION RATES BY LOG HH EXPENDITURE

Notes: Participation in the graph is demeaned at the village level and round level to control for village and round level factors that might be related to participation and wealth of village members. The overall mean is added back in to make the quartiles comparable. Log expenditure is trimmed at the top and bottom 5 percent of expenditure to exclude outliers where the data is thin.

Mathematical Appendices-For online publication only

Appendix 1: Optimal Choices of q , c , n in the Market and the Program

In the market, where q is optimally chosen by the household and the price of a unit of the staple good is pq , the optimal choices of q , c , and n are:

$$q_m^* = \frac{y(\beta - \alpha)}{ap(\beta + \delta)}, \quad c_m^* = \frac{\beta a}{\beta - \alpha}, \quad n_m^* = \frac{y\delta}{z(\beta + \delta)}$$

unless $q_m^* < q_{min}$ or $q_m^* > q_{max}$. If $q_m^* < q_{min}$, the optimal choices of q , c , and n are:

$$q_m^* = q_{min}, \quad c_m^* = \frac{a\delta pq_{min} + \alpha y}{pq_{min}(\alpha + \delta)}, \quad n_m^* = \frac{\delta(y - apq_{min})}{z(\alpha + \delta)}$$

Finally, if $q_m^* > q_{max}$, the optimal choices of q , c , and n are:

$$q_m^* = q_{max}, \quad c_m^* = \frac{a\delta pq_{max} + \alpha y}{pq_{max}(\alpha + \delta)}, \quad n_m^* = \frac{\delta(y - apq_{max})}{z(\alpha + \delta)}$$

If the household chooses to buy from the program, they take q as exogenously given as q_p . They get a discount on each unit of c purchased, so that the price of a unit of c through the program is ϕpq where $\phi < 1$. For households who purchase through the program, the optimal c and n are:

$$c_p^* = \frac{a\delta q\theta + \alpha y}{\theta q(\alpha + \delta)}, \quad n_p^* = \frac{\delta(y - aq\theta)}{z(\alpha + \delta)}$$

Appendix 2: Proof of Proposition 1

If $W(q_p, y)$ were a concave function with respect to y for fixed q_p , demonstrating Proposition 1 for the general case would only require that $W(q_p, y) > 0$ for some y and Proposition 1 would follow directly, but $W(q_p, y)$ is not concave over the entire range of y . So proving Proposition 1 proceeds in four steps. First, we show that for any fixed income $y = \bar{y}$, there exists at least one value of q_p such that $W(q_p, \bar{y}) > 0$. Second, we show that for a fixed value of income $y = \bar{y}$, there exist $q_L(\bar{y})$ and $q_H(\bar{y})$ with $q_L(\bar{y}) < q_H(\bar{y})$ such that $W(q_p, \bar{y}) > 0$ for all q_p in (q_L, q_H) , and $W(q_p, \bar{y}) < 0$ for all other q_p by showing that $W(q_p, \bar{y})$ is a concave function with respect to q_p . Third, we show $q_L(\bar{y}), q_H(\bar{y})$ are increasing functions of y . Finally, we use the second and third step to show the existence of the y_L, y_H described in Proposition 1.

The first step in proving Proposition 1 is to show the following Lemma:

Lemma 1: For a fixed $y = \bar{y}$, $W(q_p, \bar{y})$ is positive for at least one value of q_p .

Lemma 1 states that no matter which income is chosen, there is at least one quality of grain that the program could choose to offer and the household would be willing to participate in the program. To prove Lemma 1, imagine that the program chooses to subsidize the same quality of grain that the household would have chosen to consume at the market, so that the household consumes the same quality grain, but gets it a cheaper price if the grain is purchased from the program. In this case:

$$W(q_m^*, \bar{y}) = (\alpha + \delta) \ln\left(\frac{\phi\alpha + \beta(1 - \phi) + \delta}{\alpha + \delta}\right) + \alpha \ln\left(\frac{1}{\phi}\right) > 0$$

since $\phi < 1$ and we have assumed that $\beta > \alpha$, which proves the lemma.

The next step in proving Proposition 1 is to demonstrate the following:

Lemma 2: For a fixed value of income y , there exist $q_L(y)$ and $q_H(y)$ with $q_L(y) < q_H(y)$ such that $W(q_p, y) > 0$ for all q_p in $(q_L(y), q_H(y))$, and $W(q_p, y) < 0$ for all other q_p .

Lemma 2 can be proven by showing that $W(q_p, \bar{y})$ is concave. Concavity of $W(q_p, \bar{y})$, along with the existence of a global maximum with respect to q_p and Lemma 1, implies that for any choice of y there exist $q_1(y), q_2(y)$ such that $W(q_1, y) = W(q_2, y) = 0$.¹ Define $q_H =$

¹A detailed proof of this statement is in Appendix 3.

$\max(q_1, q_{min})$ and $q_L = \max(q_2, q_{max})$. The concavity of $W(q_p, \bar{y})$ implies that $W(q_p, \bar{y}) > 0$ for all q_p in (q_L, q_H) and $W(q_p) \leq 0$ for all q_p not in (q_L, q_H) , which proves Lemma 2.

Next, we show that $q_L(y), q_H(y)$ are increasing functions of y , by demonstrating the following:

Lemma 3:

Define y_H, y_L such that $y_H > y_L$.

3.a. The lowest quality offered by the program for which a household of income y_L is willing to participate is weakly lower than the lowest quality that would induce a household of income y_H to participate. The inequality is strict if the lowest quality for which a household of income y_H would be willing to participate is greater than q_{min} .

3.b. Similarly, the highest quality offered by the program for which a household of income y_L would be willing to participate is weakly lower than the highest quality which would induce a household of income y_H to participate, with strict inequality when the highest quality for which a household of income y_L would be willing to participate is less than q_{max} .²

Lemma 3 can be proven by focusing on three points on the range $q_p : q_1(y), q_2(y)$ and a point $q_p^* - \epsilon$ just to the left of the q_p where $W(q_p, y)$ attains its maximum and where $W(q_p - \epsilon, y) > 0$. Define y_L, y_H with $y_L < y_H$ and suppose that $q_1(y_L) > q_1(y_H)$. Then it would have to be the case that $\frac{\partial W(q_p, y_H)}{\partial q_p} < \frac{\partial W(q_p, y_L)}{\partial q_p}$, which contradicts the fact that $\frac{\partial W(q_p, y)}{\partial q_p \partial y} > 0$. A more detailed proof is in Appendix 4.

With Lemmas 1-3, we are ready to prove Proposition 1. Suppose $q_p = q_{min}$. By Lemma 2, $W(q_{min}, y) \geq 0$ for any y such that $q_1(y) \leq q_{min}$. By Lemma 3 and the continuity of $W(q_p, y)$ w.r.t. y , there exists y such that $W(q_{min}, y) = 0$. Define this y as $y_H(q_{min})$. For all $y < y_H(q_{min})$, $W(q_{min}, y) > 0$, and for all $y \geq y_H(q_{min})$, $W \leq 0$.

Suppose $q_p = q_{max}$. By Lemma 2, $W(q_{max}, y) \geq 0$ for any y where $q_2(y) \geq q_{max}$. By Lemma 3 and the continuity of $W(q_p, y)$ w.r.t. y , there exists y such that $W(q_{max}, y) = 0$. Define this y as $y_L(q_{max})$. For all $y > y_L(q_{max})$, $W(q_{max}, y) > 0$, and for all $y \leq y_L(q_{max})$,

²Lemma 3 can also be expressed in terms of the notation already introduced. For any y_H, y_L such that $y_H > y_L$:

3.a. $q_L(y_L) \leq q_L(y_H)$ and $q_L(y_L) = q_L(y_H)$ iff $q_L(y_L) = q_L(y_H) = q_{min}$

3.b. $q_H(y_L) \leq q_H(y_H)$ and $q_H(y_L) = q_H(y_H)$ iff $q_H(y_L) = q_H(y_H) = q_{max}$.

$W \leq 0$.

For q_p in (q_{min}, q_{max}) , fix $q_p = \bar{q}$. Following from the proof of Lemma 2, we can define functions $q_1(y)$ and $q_2(y)$ for which we must have $W(q_1(y), y) = W(q_2(y), y) = 0$. By Lemma 3, both functions $q_1(y)$ and $q_2(y)$ are continuous and monotonically increasing in y , which implies that the inverse of each function exists. Define $y_H(q_p)$ as the inverse of $q_1(y)$ and $y_L(q_p)$ as the inverse of $q_2(y)$. By Lemma 3, $y_L(q) < y_H(q)$. $W(y_L(\bar{q}), \bar{q}) = W(y_H(\bar{q}), \bar{q}) = 0$, which proves that if q_p is in (q_{min}, q_{max}) , there are at least two roots to the function $W(y, \bar{q})$.

Now suppose there are three points $y_1 < y_2 < y_3$ such that $W(\bar{q}, y_1) = W(\bar{q}, y_2) = W(\bar{q}, y_3) = 0$. By Lemma 2, we have that for fixed y , we can define $q_L(y)$ and $q_H(y)$ such that $W(q_L(y), y) = W(q_H(y), y) = 0$, and by Lemma 3, it must be the case that $q_L(y_1) < q_L(y_2) < q_L(y_3)$ and $q_H(y_1) < q_H(y_2) < q_H(y_3)$. Suppose that $\bar{q} = q_H(y_1)$. Then $q_H(y_2) > \bar{q}$ and $q_H(y_3) > \bar{q}$. Additionally, suppose that $\bar{q} = q_L(y_2)$. Then $q_L(y_3) > \bar{q}$. But by Lemma 2, $q_L(y_3)$ and $q_H(y_3)$ are the only values of q_p for which $W(q_p, y_3) = 0$, which contradicts that $W(\bar{q}, y_3) = 0$. Therefore, there cannot be more than two points in y such that $W(\bar{q}, y) = 0$.

All that remains to be shown is that for y in (y_L, y_H) , $W(\bar{q}, y) > 0$. The above argument shows that there cannot be more than two roots of the function $W(\bar{q}, y)$, which implies that either $W(\bar{q}, y) < 0$ or $W(\bar{q}, y) > 0$ in the region (y_L, y_H) . Therefore, if we can find one point y in the region where $W(\bar{q}, y) > 0$, the function must be positive over the entire region. In order to have two roots, the function must have a stationary point between the two roots. Solving $\frac{\partial W(q_p, y)}{\partial y}$ for 0 gives:

$$y^* = \frac{a\phi q_p(\beta + \delta)}{\beta - \alpha}$$

Since there is only one point where the derivative of the function is 0, this must be the stationary point, and therefore y^* must be in the region bounded by the two roots (y_L, y_H) . Substituting y^* into $W(q_p, y)$ gives:

$$W = \beta \ln\left(\frac{1}{\phi}\right)$$

Which is > 0 since $\phi < 1$.

We have shown: 1. That for a fixed $q_p = \bar{q}$ there exist $y_L(\bar{q}), y_H(\bar{q})$ with $y_L(\bar{q}) < y_H(\bar{q})$ such that $W(\bar{q}, y_L(\bar{q})) \geq 0$ and $W(\bar{q}, y_H(\bar{q})) \geq 0$ and 2. For all y in (y_L, y_H) , $W(\bar{q}, y) > 0$, which proves the Proposition.

Appendix 3: Proof that a continuous, concave function $y = f(x)$ with $y > 0$ at the function's global maximum has exactly two roots over the range $x = (-\infty, \infty)$.

By the fact that $f(x)$ is a concave function, we have that:

$$f(z) \leq f(x) + f'(x)(z - x)$$

Define x^* as the value of x such that $f(x)$ takes a global maximum. By assumption, $f(x^*) > 0$. Then there exists $\epsilon > 0$ such that $f(x^* + \epsilon) > 0$ and $f(x^* - \epsilon) > 0$.

$$\text{Define } x_1 = x^* - \epsilon \text{ and } z_1 = \frac{f'(x_1)x_1 - f(x_1)}{f'(x_1)}.$$

$$\text{Then } f(z_1) \leq f(x_1) + f'(x_1) \frac{f'(x_1)x_1 - f(x_1)}{f'(x_1)x_1}$$

The concavity of $f(x)$ implies that $f'(x_1) < 0$, since $f'(x^*) = 0$ and $x_1 < x^*$. Using this fact, we can simplify the above to:

$$f(z_1) \leq 0$$

Using again the fact that $f'(x_1) < 0$, we have that $z_1 < x_1 < x^*$, which proves that there is a point $z_1 < x^*$, such that $f(z_1) \leq 0$. If $f(z_1) = 0$, z_1 is a root. If $f(z_1) < 0$, then the fact that $f(x)$ is continuous and $f(x^*) > 0$ implies that there must exist a point x_{lbar} in the range (z_1, x^*) such that $f(x) = 0$, which proves the existence of one root where $x_{lbar} < x^*$.

To prove that there is a second root $x_{hbar} > x^*$, define $x_2 = x^* + \epsilon$ and the proof proceeds exactly as above.

Therefore there are at least 2 roots, x_{lbar} and x_{hbar} . To prove that these are the only roots, suppose that q_1, q_2, q_3 are all roots of W , so that $W(q_1) = W(q_2) = W(q_3) = 0$, and $q_1 < q_2 < q_3$. By concavity of W , $W(z) > W(q_1)$ and $W(z) > W(q_3)$ for all z in (q_1, q_3) , which contradicts that $W(q_2) = 0$. Therefore, q_2 cannot be a root of W , and there are at most 2 roots.

To apply this proof to Lemma 2, we have only to show that $W(q_p, y) > 0$ at the global

maximum. Setting the derivative of $W(q_p, y)$ w.r.t. q_p equal to 0 and solving for q_p gives q^* , the maximand of $W(q_p, y)$. Substituting q^* into $W(q_p, y)$ gives:

$$W(q^*, y) = \beta \ln \left(\frac{1}{\phi} \right) > 0$$

where the last inequality is true since $\phi < 1$.

Appendix 4: Detailed Proof of Lemma 3

We will do the proof only for 3.a. concerning $q_L(y)$. The proof for 3.b. concerning $q_H(y)$ is similar. The proof follows two cases: 1. $q_1(y_L) \leq q_{min}$ and 2. $q_1(y_L) > q_{min}$.

Case 1: $q_1(y_L) > q_{min}$

The proof of Lemma 3 uses relies on characterizing four points on the q_p axis:

1. q_p^* : The q_p which maximizes $W(q_p, y_L)$ for some fixed $y = y_L$
2. $q_p^* - \epsilon$: A point just to the left of q_p^*
3. $q_1(y_L)$: The point where $W(q_p, y_L) = 0$
4. $q_1(y_H)$: The point where $W(q_p, y_H) = 0$ with $y_H > y_L$

To clarify the interpretation of each point described in the proof, $W(q_p, y)$ is graphed below for two values of y . In the graph, the blue line is W plotted over q_p for $y = y_L$, and the red line is W plotted over q_p for $y = y_H$.

To avoid clutter on the graph, the 4 key points are labeled as:

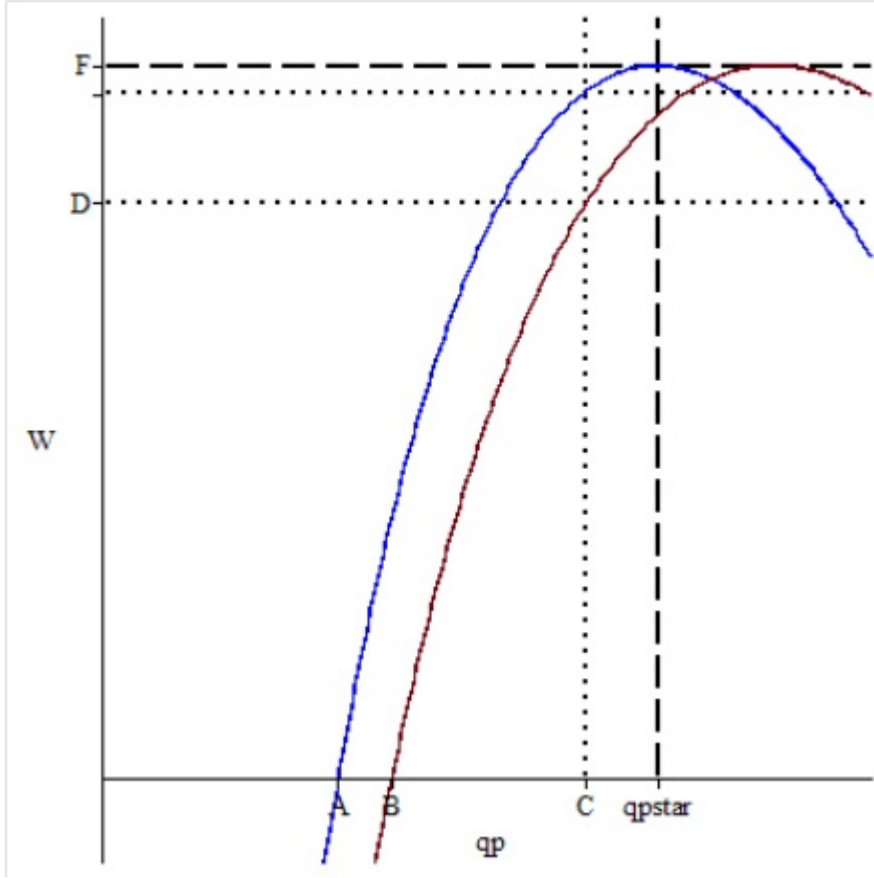
$$q_p^{**} = \text{qpstar}$$

$$q_p^{**} - \epsilon = C$$

$$q_1(y_L) = A$$

$$q_1(y_H) = B$$

Figure 1:



The proof follows three basic steps:

Step 1: Characterize the point $q_p^* - \epsilon$ which will be used in the proof. ϵ is chosen so that:

1. $q_1(y_L) < q_p^* - \epsilon < q_p^*$
2. $W(y_L, q_p - \epsilon) > 0$

Step 1 then shows that $W(q_p^* - \epsilon, y_L) > W(q_p^* - \epsilon, y_H)$.

Step 2: Assume that $q_1(y_L) > q_1(y_H)$. This assumption implies that $W(q_p, y_H)$ and $W(q_p, y_L)$ must cross at some q_p between the point $q_1(y_L)$ and the point $q_p^* - \epsilon$. We will use this crossing point in Step 3 to prove a contradiction that negates the assumption.

Step 3: At the crossing point described in step 2, it would have to be the case that $\frac{\partial W(q_p, y_H)}{\partial q_p} < \frac{\partial W(q_p, y_L)}{\partial q_p}$.

Show that this result contradicts the fact that the following cross partial derivative is positive:

$$\frac{\partial W(q_p, y)}{\partial q_p \partial y}$$

Detailed Steps of the proof:

Step 1: Characterize the point $q_p^* - \epsilon$ that we will use in the proof of Lemma 3.

For fixed y , there exists at least one point $z_1 > q_1(y)$ such that $0 < W(z_1, y) < W(q_p^*(y), y)$ and $\frac{\partial W(z_1, y)}{\partial y} < 0$.

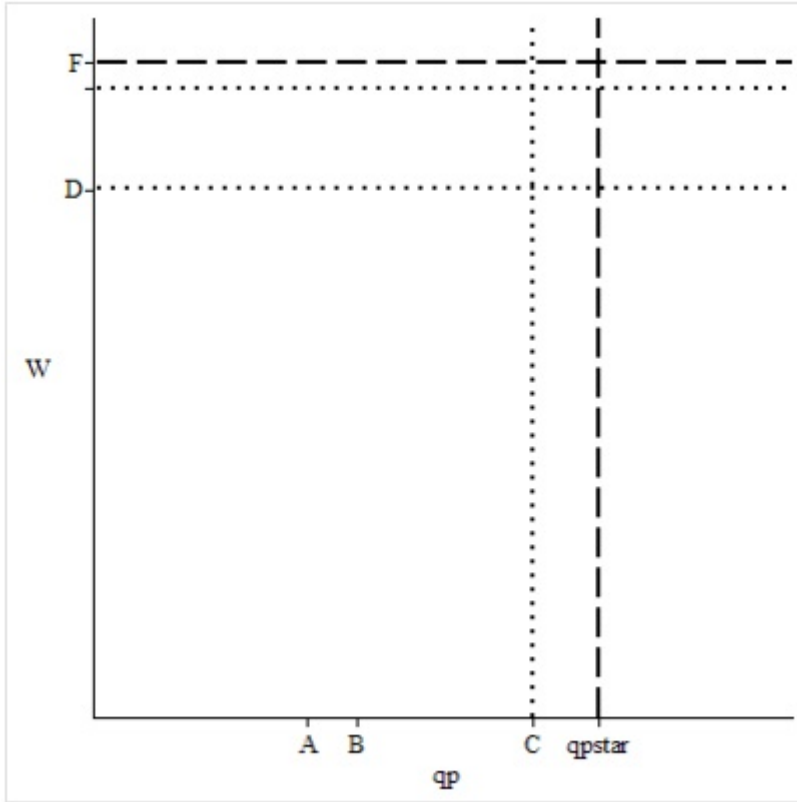
Proof: We have shown in the proof of Lemma 2 that there exists $q_1(y)$ such that $W(q_1(y), y) = 0$ and that $W(q_p^*(y), y) = \beta \ln\left(\frac{1}{\phi}\right) > 0$. Since q_p^* is the only stationary point of the function $W(q_p, y)$ and $q_1(y) < q_2(y)$, the fact that $W(q_p^*(y), y) > 0$ implies that $q_1 < q_p^*$. By the fact that $W(q_p, y)$ is continuous in q_p , there must exist $\epsilon > 0$ such that $q_p^* - \epsilon$ is in the range (q_1, q_p^*) and $W(q_p^* - \epsilon) > 0$. Evaluating $\frac{\delta W(q_p, y)}{\delta y}$ at the point $q_p = q_p^* - \epsilon$ gives:

$$-\frac{a\phi\epsilon(\beta^2 + 2\beta\delta + \delta^2)}{(a\beta\epsilon\phi + a\delta\epsilon\phi + \alpha y + \delta y)y} < 0$$

The negative sign on the above derivative implies that $W(q_p^* - \epsilon, y_L) > W(q_p^* - \epsilon, y_H)$ at the point $q_p^* - \epsilon$.

The graph below demonstrates the proof. q_p^* in the graph is the point where $W(q_p, y_L)$ attains its maximum of F . At $q_p = q_p^* - \epsilon$ (labeled C in the graph), $W(q_p^* - \epsilon, y_L) = E$, where $0 < E < F$. The fact that $E < F$ follows from the fact that $W(q_p, y)$ is a local maximum. The negative sign on the partial derivative at the end of the proof above implies that $W(q_p^* - \epsilon, y_H) < W(q_p^* - \epsilon, y_L)$ as shown by the fact that $D < E$ in the graph.

Figure 2:



Step 2: Assume that $q_1(y_L) > q_1(y_H)$. This assumption implies that $W(y_H, q_p)$ and $W(y_L, q_p)$ must cross at some q_p between the point $q_1(y_L)$ and the point $q_p^* - \epsilon$.

To obtain the contradiction that proves Lemma 3, suppose in the graph that $A = q_1(y_H)$ and $B = q_1(y_L)$ so that $q_1(y_H) < q_1(y_L)$. Step 1 demonstrated that $W(C, y_L) > W(C, y_H)$. Under the assumption, $W(q_p, y_H)$ must intersect both $(A, 0)$ and (C, D) and $W(y_L, q_p)$ must intersect $(B, 0)$ and (C, E) . Now imagine two lines that meet both conditions, there must be a point where both lines cross. This crossing must happen for any two continuous functions between these points, including the particular convex function $W(q_p, y)$. Therefore, the assumption that $q_1(y_L) > q_1(y_H)$ implies that there must be a point where $W(q_p, y_H) = W(q_p, y_L)$ in the range of q_p in (B, C) .

Step 3: Show that the result from step 2 contradicts the fact that the cross partial derivative of $W(y, q_p)$ w.r.t. y, q_p is positive.

Focus on the point where $W(q_p, y_H) = W(q_p, y_L)$ in the range of q_p in (B, C) . Since $W(y_H, q_p) > W(y_L, q_p)$ for any point to the left of this crossing point and $W(q_p, y_H) < W(q_p, y_L)$ to any point to the right, it must be the case that:

$$\frac{\partial W(q_p, y_H)}{\partial q_p} < \frac{\partial W(q_p, y_L)}{\partial q_p}$$

at the crossing point.

All that is left is to argue why the above implies a contradiction. This can be shown directly through taking the derivative of $W(q_p, y)$:

$$\frac{\partial W(q_p, y)}{\partial q_p \partial y} = \frac{\alpha \phi(\alpha + \text{delta})}{(\alpha \phi q_p - y)^2} > 0$$

The fact that the above derivative is positive implies that for all q_p , $\frac{\partial W(q_p, y_L)}{\partial q_p} < \frac{\partial W(q_p, y_H)}{\partial q_p}$ for all $y_H > y_L$, which contradicts

$$\frac{\partial W(q_p, y_H)}{\partial q_p} < \frac{\partial W(q_p, y_L)}{\partial q_p}$$

.

The assumption that $q_1(y_H) < q_1(y_L)$ leads to a contradiction of a known feature of W , namely that the cross partial derivative of W w.r.t. y, q_p is positive, so we must reject the assumption. It therefore follows that $q_1(y_H) > q_1(y_L)$, which proves Lemma 3.

Case 2: $q_1(y_L) \leq q_{min}$

By Lemma 2, if $q_1(y_L) \leq q_{min}$ then $q_L(y_L) = q_{min}$.

Suppose that $q_1(y_H) \leq q_{min}$, then again from Lemma 2, $q_L(y_H) = q_{min}$, and we have that $q_L(y_H) = q_L(y_L)$.

Now suppose that $q_1(y_H) > q_{min}$. Then by Lemma 2, $q_L(y_H) = q_1(y_H) > q_{min} \geq q_L(y_L)$ which implies that $q_L(y_H) > q_L(y_L)$.

Since these are the only two possible cases for $q_1(y_H)$, we have shown that $q_L(y_L) = q_L(y_H)$

iff $q_L(y_L) = q_L(y_H) = q_{min}$.

Appendix 5: Proof of Proposition 4

To show Proposition 4, first assume that all households are willing to participate. Solving $c_p^* - c_m^* = 0$ for y gives:

$$y_0 = \frac{a(\beta + \delta)p\phi q_p}{\beta - \alpha}$$

Given that the first derivative of the program effect with respect to income is positive (Proposition 3), for any income less than y_0 , the program effect on consumption would be negative and for any $y > y_0$, the program effect would be positive.

To prove the proposition, we have only to show that $W(y_0, q_p) > 0$, so that there are participants with income just above and just below y_0 . Substituting y_0 into W gives:

$$W = \beta * \ln\left(\frac{1}{\phi}\right)$$

which is > 0 since $\phi < 1$. At the point where the program effect is 0, households receive a level of surplus utility from program participation which is strictly positive, implying that they are willing to participate in the program. Since W is a continuous function, there exists $\epsilon > 0$ such that $W(q_p, y_0 - \epsilon) > 0$ and $W(q_p, y_0 + \epsilon) > 0$. By Proposition 3, at $y = y_0 - \epsilon$, $c_p^* - c_m^* < 0$ and at $y_0 + \epsilon$, $c_p^* - c_m^* > 0$, which proves Proposition 4.