On the Gains from Tradeable Benefits-in-Kind

Martin Ravallion[1]

Department of Economics, Georgetown University, and NBER

Abstract: Neither the purposive targeting of in-kind benefits nor randomized assignment is likely to yield a competitive equilibrium given private information on diverse personal gains. The paper characterizes equilibrium assignments and the implications for policy and evaluation. Using special-purpose survey data, the theory is applied to an antipoverty program providing jobs on rural public-works in a poor state of India. Large unexploited gains from trading assignments are evident. The potential gains exceed those from poverty targeting and stylized cash transfer options. Realizing the scope for poverty reduction will probably require complementary efforts to help poor people access the market for benefits-in-kind.

Keywords: Information; poverty; randomization; public-works; Bihar, India

JEL: H30, I38, O15

[1] The author thanks Nicholas Barr, Caitlin Brown, Ian Gale, Emanuela Galasso, Michael Lokshin, Milan Thomas and Dominique van de Walle for their comments on an earlier draft.

1. Introduction

In-kind benefits have long been popular in social policy making.[2] A common problem faced is how a limited number of such benefits-in-kind (BIKs) is to be assigned across a designated set of eligible individuals, not all of whom can be served. The decision maker presumably cares about the aggregate disbursement of the BIKs but also cares about the welfare gains, and those gains undoubtedly vary. Thus, the inter-household allocation of BIKs matters.

With information about the gains to all individuals, one could simply target the BIKs accordingly, with the first one going to those with highest gains and so on until the available stock is exhausted. What makes the problem difficult in practice is that how the gains differ is in large part unknown to the decision maker, although the personal gains may be known reasonably well at the individual level.

Two examples illustrate the problem. For the first, consider a training program with only so many slots available. The wage gains from training vary. Some plausible covariates of the gains may be observable, but some crucial ones are not, such as latent ability, although one can expect these variables to be reasonably well-known privately. The policy problem is how to allocate the limited number of slots, with little or no information about the individual gains.

The second example is a workfare scheme, providing extra work at a wage rate common to all. Here the BIK is the extra work, but not all who want that work can be accommodated. The gains to individuals joining the program vary, given differing forgone earnings from other available work (including self-employment). While each person probably has a fairly good idea of their best alternative at the time, this is not known to the decision maker in deciding how to ration the available jobs.

Administrative assignments have tended to rely on various proxies for the likely gains. The available evidence to date does not suggest that such targeting efforts are very good in practice.[3] The information problem looms large. In both the examples above (and other similar

[2] On the rationales for in-kind benefits see Moffitt (2006) and Currie and Gahvari (2008). Discussions of the policy choice between cash and in-kind benefits include Akerlof (1978), Coate (1989), Currie and Gahvari (2008), Cunha (2014) and Hirvonen and Hoddinott (2021).

[3] In the context of targeting based on the "poverty proxies" typically available in practice see Brown et al. (2018), using data for multiple countries in Africa. In the context of efforts to encourage in-kind consumption of food see Cunha (2014), using data for Mexico.

cases), the initial administrative assignment of BIKs is unlikely to be an equilibrium. There will be incentives for trade based on the private information about the individual gains, to the extent that some have larger gains than others. At the competitive market price, such trades will tend to re-allocate the BIKs toward those with larger gains, thus increasing the aggregate impact of the social policy.

We know very little about the potential gains from tradable benefits-in-kind. This is not surprising given that individual gains from trade are typically unobserved, including by researchers. This has made it hard to inform public discussion of the arguments for and against the frequent efforts by policy makers to restrict trade in assignments for in-kind benefits—indeed, re-sale is often explicitly forbidden, at least on paper. The information asymmetry also underlies the lack of evidence on the extent to which informal trades (outside administrative purview) can realize the aggregate gains in practice. When one person receives a BIK that has greater value to a non-recipient, the two parties may well be able to make a mutually beneficial trade, whether legal or not. Then we may find that there are only small differences in the remaining gains across the eligible population, such that the further trade brings little extra benefit; similarly, there may be little welfare cost to governmental rules aiming to restrict trade. This is no more than a conjecture, however; we may instead find that (governmental or other) restrictions on trade come with a high cost. We do not know. The knowledge gap exists for the same reason that it is a concern.

The lack of empirical knowledge has left much scope for debate. A common objection to market-like solutions in social policy design is that the gains may be captured disproportionately by the well-off. By this view, tradeable BIKs may bring less benefit to initially excluded poor people. Weitzman (1977) showed that the gains from a market-based allocation mechanism depend on how much individual gains differ <u>and</u> on the extent of income inequality. If one judges that incomes are too unequally distributed then one can also be concerned that a market mechanism for social programs will only make things worse. Yet the literature provides counterarguments. Sah (1987) demonstrates that, for poor people, allowing rationed BIKs to be tradeable (which he calls "convertible rations") can dominate the other allocation mechanisms he considers. Furthermore, Che et al. (2013) show that a competitive market allocation for an assignable good can attain higher utilitarian social welfare if it is introduced in the wake of an initially random assignment.

Recognizing the concern that a quasi-market assignment runs the risk of being captured by the non-poor, this paper studies the properties and performance of the Che et al. (2013) "randomization-with-resale" assignment mechanism when the welfare outcomes are to be judged by the pecuniary gains to poor people. The paper characterizes the competitive equilibrium of assignments to a program following initial randomization across a set of eligible people. This is the allocation we would observe if those eligible could freely trade. The model is key to the empirical analysis of the costs of restricting trade in BIKs. The model also carries some implications for the interpretation of randomized controlled trials (RCTs).

Based on this model, the paper simulates equilibrium allocations using a sample of surveyed workers in a large antipoverty program in India. The program provides low-wage, unskilled, jobs on rural public works projects. The task is to assess whether the gains from trade among participants have been realized in practice, to see if the program could be more effective in raising their incomes by promoting trade. The data studied here provide an unusual—indeed, unique (to my knowledge)—opportunity for addressing this issue, given that a plausible measure of the personal gains can be retrieved using survey data. The sample is treated as the universe from which artificial programs are simulated, consistently with the predictions of the theoretical model. The simulations are used to estimate the participant's mean monetary gains from trade, beyond what has been attained already, legally or otherwise. Various counterfactuals are considered, including a "needs-based" assignment, based on household consumption expenditure per person, as widely used for measuring poverty in India.

This is also a setting in which we can learn about how the potential gains from trade are distributed. The data come from a population of poor households; indeed, 75% come from families living below the World Bank's international poverty line, which identified about 15% of the world's population as poor at the same time. However, they are not all equally poor—indeed, the inequality in household consumption per person is similar to rural India as a whole. Gender inequality is also an issue. The paper looks at heterogeneity along both household consumption and gender dimensions.

It should be emphasized that the evidence presented here is for one specific type of in-kind benefit, in one specific setting. It may well be the case that providing jobs as a means of reducing poverty in poor areas of rural India is a situation in which the gains from trade are likely to vary a lot, pointing to large costs of impediments to trade. Would that also be true for other in-kind benefits such as when food rations are targeted to poor families? For some, the ration is no more food than wanted, given their food demand functions, but for others it is more than they want. Thus, the value of this BIK varies. It might be conjectured that this type of variation is not as great as that associated with a workfare scheme. That is a conjecture, however. Cunha (2014) finds substantial variation in the extent to which in-kind food rations are infra-marginal in data for a Mexican "food BIK" program.

The next section provides the model of equilibrium assignments, which carries the key insights needed for the empirical analysis. Section 3 applies the model to the survey data for workers on the public works scheme. Trade shifts the available jobs toward those with lower forgone earnings from other work opportunities. The paper's empirical results indicate large unexploited gains from trade. In equilibrium, tradeable assignments would increase the scheme's aggregate gains by a factor of 2 to 3. Similar gains are found when the comparison is with the needs-based assignment. The equilibrium assignment yields large gains to both workers from poor families and female workers. Tradeable BIKs can also have greater impact on poverty than feasible policy options for cash transfers. Section 4 identifies some potential impediments to realizing the gains in practice. The impediments relate to deeper features of the market and institutional/governmental environment that can be thought of as being among the reasons why poverty exists in this setting. Complementary policies are identified that may be necessary to realize the potential gains to poor people from tradeable BIKs. Section 5 concludes.

2. The equilibrium with tradeable assignments

A lumpy BIK is to be assigned across a pre-determined set of eligible recipients when there is not enough for everyone in that set. It is assumed that the BIKs are provided free of charge and that resale (if feasible) is costless. The nature of the BIK is such that nobody would want a second, and it cannot be stored for later use.

Let $D_i = 1$ if individual $i = 1,..,n$ (with n fixed) receives the BIK initially while D_i=0 if not, with mean $\bar{D} \equiv E(D)$, which we can call the coverage rate. There are both BIK recipients and non-recipients, so $0 < \bar{D} < 1$. There is a fixed number ($n\bar{D}$) of BIKs available (as determined by the budget), so $\bar{D}$ is exogenous. Using the (Neyman–Rubin) potential outcomes framework, whether or not an individual actually receives the BIK, one can define two numbers for all $i = 1, \dots n$, namely the outcome with the BIK, Y_{i1}, and that in its absence, Y_{i0}. The gain is $G_i \equiv Y_{i1} - Y_{i0}$, with cumulative distribution function $F(.)$ and mean $E(G)$. When it helps to simplify the analysis, G is treated as a continuous variable with a (strictly increasing) distribution function on the support $[G^{min}, G^{max}]$.

In the literature on the standard "evaluation problem" the main task is usually to estimate the mean gain $E(G)$ (or a conditional mean of interest such as $E(G|D = 1)$). No attempt is made to estimate the individual gains. The classic RCT randomly assigns the treatment and compares mean outcomes for those treated and those not. As is well known, under standard assumptions (including that there are no spillover effects, contaminating the controls), randomized assignment delivers an unbiased estimate of the mean gain (though, of course, any one trial will contain an experimental error).

Here we address a different problem: how should the program be assigned to maximize mean gain? Call this the "assignment problem." With perfect information, the solution is obvious: give the first BIK to the individual with G^{max} then to the next highest and continue until all the available BIKs have been allocated. Of course, information is far from perfect (as discussed in the Introduction). The individual gains are typically unknown to the decision maker, given the obvious difficulty in knowing outcomes in two different states of nature at the same time. However, each person clearly knows a lot more about his or her own likely gain. Indeed, in some settings (including the examples in the introduction) it can be expected that each individual is reasonably well-informed about the G_i, and acts accordingly.[4]

The competitive equilibrium: Trade will clearly change the assignments in an economically relevant way. Those who were assigned a BIK can sell it at a price P. Given that the personal gain from the program is known to each person, the sellers will be those who

[4] This specific information asymmetry is an example of what Heckman et al. (2006) call "essential heterogeneity."

receive the program initially but for whom $G_i < P$; they do better by selling it than keeping it. Buyers will be those who did not receive it initially, but with $G_i > P$.

When the decision maker knows nothing about the individual gains, a random (uniform in probability) assignment within the eligible set has obvious appeal. Initial randomization justifies assuming that the distribution of gains is the same for those who are initially assigned the program and those not. The share of the population that received the BIK and want to sell at the price P is $\overline{D} \cdot F(P)$. The corresponding share who did not receive the BIK but want to buy an assignment at price P is $(1-\overline{D})(1-F(P))$. Let us further assume that $F\left(G^{min}\right) < 1-\overline{D}$. (A sufficient condition for this to hold is that $F\left(G^{min}\right) = 0$ but a point mass at G^{min} is also allowed.) Then there is a positive excess demand for assignments at G^{min}. By definition $F(G^{max}) = 1$, so there must be a positive excess supply at G^{max}. Then, by continuity of $F(.)$, a unique equilibrium exists.[5] The market-clearing price solves $F(P) = 1-\overline{D}$, i.e., the equilibrium price is the quantile of gains corresponding to the share of the population not receiving a BIK ($P = F^{-1}(1-\overline{D})$).

There are four groups of people in this model:

1. The keepers: those assigned the BIK who do not want to sell it ($G_i > P$). The proportion of the population who are keepers is $\overline{D}\left(1-F(P)\right) = \overline{D}^2$ (in equilibrium) and their mean gain is $E(G_i|D_i = 1, G_i > P)$.
2. The sellers: those selected initially who would rather sell their assignment ($G_i < P$). Their population share is $\overline{D}(1-\overline{D})$ in equilibrium, with a mean gain of P.
3. The buyers: those initially excluded who expect a net benefit from buying access ($G_i > P$). Their population share is $\overline{D}(1-\overline{D})$ (in equilibrium) and their mean gain is $E(G_i|D_i = 0, G_i > P) - P$.
4. The rest, with population share $(1-\overline{D})F(P)$ and zero gain.

Notice that, in equilibrium, the share of the population participating in the market ($2\overline{D}(1-\overline{D})$) does not depend on the distribution of the gains; the price does all the adjustment

[5] Stability is assured under the usual condition that the price rises with excess demand and falls with excess supply.

given $\overline{D}$. Differences in that distribution do, of course, matter to the size of the aggregate gains from tradeable assignments.

Summing the gains across all four groups, weighted by population shares, the total gain per capita of the population is:[6]

$$E(G_i) = \overline{D}^2 E(G_i|D_i = 1, G_i > P) + \overline{D}(1 - \overline{D})E(G_i|D_i = 0, G_i > P) = \overline{D}E(G_i|G_i > P) \quad (1)$$

The first term on the LHS is the gain to keepers while the second term is the gain to the traders (the gain to sellers plus that to buyers).

Trade improves on the initial randomized assignment since the gains to those who buy an assignment (($E(G_i|D_i = 0, G_i > P)$) must exceed the gains to those who sell one ($E(G_i|D_i = 1, G_i < P)$).[7] Consider the maximum attainable aggregate gain with perfect information. For that allocation, there will be some threshold gain, Z, above which everyone receives the BIK, and below which no-one receives it. (Z is determined by the number of BIKs available.) The mean gain is $E(G_i|G_i > Z)$ and $1 - F(Z) = \overline{D}$, which implies that $Z = P$, giving the same mean gain as the market equilibrium attains after the initial randomized assignment. Thus, despite the decision maker knowing nothing about individual gains, the trade-based assignment mechanism attains the first-best optimum with perfect information.

A further comparison of interest is with the expected gain <u>without</u> trade, as given by $\overline{D}E(G_i|D_i = 1)$, which is the mean gain we would estimate using a RCT under standard assumptions (including no trade in assignments). The gain (per capita) from trade is then $\overline{D}(1 - \overline{D})[E(G_i|D_i = 0, G_i > P) - E(G_i|D_i = 1, G_i < P)] > 0$ (invoking randomization).

Three remarks: First, the above model can be readily adapted to allow stratification by categories of individuals defined by observed characteristics, taken as fixed. For example, this may be based on gender or a poverty map (showing poverty measures by area). Then the value of $\overline{D}$ varies by group, yielding different (group-specific) prices. The policy goal can then be thought

[6] The term in $\overline{D}(1 - \overline{D})P$ drops out as it is a pure transfer between groups 2 and 3.

[7] Note that randomized assignment implies that ($E(G_i|D_i = 0, G_i > P) = E(G_i|D_i = 1, G_i > P)$ (given that randomization assures that the assignment is uncorrelated with the potential individual gains). Also note that $E(G_i|D_i = 1, G_i > P) > E(G_i|D_i = 1, G_i < P)$. Thus, $E(G_i|D_i = 0, G_i > P) > E(G_i|D_i = 1, G_i < P)$.

of as maximizing aggregate gains for each category, which then assures a maximum of any fixed-weighted aggregate gain (given the partition into groups).

Second, the equilibrium assignment using tradeable BIKs does not have any obvious equivalence to a feasible cash-only transfer policy. The final assignment of the gains from tradeable BIKs will not (in general) be the same as for a cash-only policy since (of course) there will be no incentive to trade cash transfers; in contrast, trade in BIKs can be expected to modify the initial assignment. Nor would it be reasonable to argue that, with purposive targeting, a cash-only policy can implement the equilibrium assignment of BIKs since that would require a different information set, which should clearly be held constant for a valid comparison. The empirical section will comment further on the implications for comparisons of cash versus in-kind antipoverty policies in that specific context.

Third, there is an implication for the interpretation of RCTs aiming to evaluate the impact of BIKs using a pilot, to inform a government's decision about scaling up. Valid inferences from the RCT cannot of course ignore the scope for trade in assignments, since this is a source of spillover effects. Nor would it seem likely that the pilot will be able to prevent trade in assignments, as this would require laws and the power to enforce them. Yet, governments routinely prevent trade in assignments at scale, and are more likely to have the required power. So, we can imagine a scenario in which there is more trade in assignments at the pilot stage than at scale. Assuming that the mean gain is correctly calculated in the RCT (allowing for trade) the RCT will tend to over-estimate the impact of the scaled-up program, given that the gains from trade are lost on scaling up. The RCT will provide undue encouragement for scaling up.

This last argument assumes that the induced change in assignments is observable to the evaluator. That might not be the case. Suppose instead that the evaluator ignores the spillover effect to the control group implied by trade (as it is unobserved) and simply calculates the mean gain for those treated based on observable incomes. This calculation will also over-estimate the mean gain on scaling up (without trade) since the evaluator will over-estimate the gains to the sellers (attributing a gain of P per seller instead of $E(G_i|D_i = 1, G_i < P)$). Again, the RCT will deliver excessive encouragement for scaling up.

3. Simulating trade in public-works jobs in rural India

The rest of this paper implements the above model for a sample of workers participating in India's Mahatma Gandhi National Rural Employment Guarantee Scheme (MGNREGS). The scheme has a more-or-less explicit aim of reducing poverty by providing jobs on local small-scale public works projects at stipulated wage rates, typically set a little above prevailing wages rates for manual agricultural labor. As is often the case, requiring people to work for poverty relief is seen to have intrinsic merit. There is also a classic self-targeting argument, namely that non-poor people will not want to do such work, and nor will poor people with preferred options.[8]

Each household in rural India has a stipulated maximum number of days of work on the scheme, set at 100 days per year per household.[9] The survey data used here indicate that the 100-day limit was rarely reached. The mean number of days worked in the previous year was 17 and only reached or exceeded 100 days for 0.7% of households.[10]

While the scheme is intended to be demand driven, in practice, local village leaders are generally deciding what projects are done under the scheme and who is employed. There is evidence that the assignments are heavily rationed in practice, and more so in poorer states of India (Dutta et al. 2012, 2014; Desai et al. 2015). Using national survey data for 2010, Dutta et al (2012) report that, for India as a whole, 44% of those rural households who say that they wanted work on the scheme did not get it. In all but three of India's 20 larger states, the reported rationing rate was over 20%.[11] Furthermore, the rationing rate tended to be higher in states with a higher poverty rate. In one of India's poorest states, Bihar, the rationing rate was 79%; barely one-in-five of those workers who wanted work on the scheme got it. Similar rates of rationing were reported in 2020, during the pandemic (Kapoor 2020). Ravallion (2021) identifies reasons why rationing of the available jobs on MGNREGS can emerge as an equilibrium in the local political economy and argues that the conditions for this to occur are more likely in poorer states,

[8] On the incentive arguments for workfare versus cash transfers see Besley and Coate (1992). Alik-Lagrange and Ravallion (2018) generalize the theoretical model of Besley and Coate to allow for both unemployment and a welfare loss associated with the work requirement.

[9] Under the formal rules the state government can relax this, although that seems to be rare.

[10] The recorded days worked my exceed 100 in the data set due to differences in the definition of the household between the survey and the scheme, as well as measurement errors such as recall on when the work was done.

[11] Evidence of rationing is also reported by Ravallion et al. (1993) for the antecedent programs to MGNREGS, in the state of Maharashtra.

such as Bihar. The legislation has a provision for payment of an unemployment allowance if work cannot be provided, but such payments appear to be rare in practice (Dutta et al. 2014).

The scheme (implicitly) allows jobs to be transferred within households. A "job card" for each household identifies the work provided by its members, but the 100-day limit is only applied to the total, leaving the household free to choose how the available work is assigned among its members. It is likely that families make participation decisions with the intention of increasing the net income gain to the family as a whole by assuring that extra work opportunities go to family members with lower forgone earnings. (How the gains are distributed is another matter.) This is consistent with both observations from field work (interviewing participating families) and the econometric model of intra-household time allocation in Datt and Ravallion (1994), using data related to an antecedent program to MGNREGS, in the state of Maharashtra.

The scope for re-assignment between household is far less obvious. The legislation and operating rules of the scheme contain no provision to allow a household that does not want all of the 100 days of work to transfer the balance to another household that has reached 100 days. The 100 days is intended to be applied to every rural household and is supposed to be checked with reference to the job cards. As we have seen, the 100 days limit is generally not binding, and rationing of assignments is the norm. While there is no formal mechanism whereby households can trade their assignments, one cannot rule out the possibility that some form of inter-household trade in assignments is occurring, though it is hard to observe. What is being measured here should be interpreted as the remaining, unexploited, gains from competitive inter-household trade in assignments.

Policy discussions of this and similar schemes have often assumed (at least implicitly) that the worker has no other source of income while working on the scheme. Then the income gain is roughly uniform, and the aggregate gains from tradeable assignments would be small. However, it is implausible that there is no forgone income for participants. Poor people in poor places cannot afford to be idle. For a precursor to MGNREGS, in the state of Maharashtra, Datt and Ravallion (2004) used an econometric model of time allocation and found that forgone earnings accounted for about 25% of the gross wage rate. By observational means—comparing workers on a similar scheme in Argentina with matched comparators—Jalan and Ravallion

(2003) found that average forgone income was about 50% of the gross wage. In this context, however, we need the distribution of the net income gains.

Survey of workers on MGNREGS: The methods of estimating forgone income from the literature reviewed above do not give us the distribution of the gains, so they do not allow us to simulate the allocation in equilibrium implied by the model in Section 2. Dutta et al. (2014) designed a special purpose survey that gives an estimate of forgone earnings for each worker. The survey was done in two rounds over 2009/10.[12]

The simulations implementing the model in Section 2 are only possible for the surveyed sample of existing workers under the scheme since the gains are only observable for this sample. This is clearly a selected sample, rather than being representative of rural India, or even rural Bihar. 75% of this sample of workers live in households with consumption per person below the World Bank's international poverty line of $1.90 a day, at 2011 Purchasing Power Parity.[13] This is well above the corresponding poverty rate for rural India in the same year, which was 36%. That said, the workers in the sample are not all "equally poor." For example, the Gini index of household consumption per person among the surveyed workers is 0.27, which is only slightly lower than the corresponding Gini index for rural India at this time of 0.29 (Himanshu 2019).

The fact that this is a selected sample does not, of itself, reduce interest in these calculations. It is clearly of interest to know whether there is evidence of unexploited gains from trade within this sample of workers. However, in thinking about the policy implications, it should be noted that creating a market in BIK assignments can generate new incentives for participation. Section 4 returns to this issue in the context of the application studied here.

Surveyed MGNREGS workers were asked to report both their wages under the scheme and to estimate their forgone earnings, i.e., how many days work they think they would have found and at what daily wage rate. In this setting, the participants are likely to have a good idea of their options. Dutta et al. (2014) found that the answers given accorded well with prevailing earnings from the casual (mostly part-time) work available at the time. Response rates to the questions on forgone earnings were high (92% and 98% in the two survey rounds). These

---

[12] Workers' surveys in the two rounds are pooled, but standard errors are adjusted upwards by $\sqrt{2}$ to allow for re-surveying the same workers in different rounds. Since not all were re-surveyed, this adjustment is conservative.
[13] The consumption aggregate used here follows the same methods as for India's National Sample Survey.

questions were clearly no more difficult than the more familiar "objective" questions. The most common response to the question on what activity would have been forgone was "casual labor," which was the answer given by 42% of the respondents. This was casual manual work for a local landowner or some similar, relatively un-skilled, non-farm work (18% of respondents gave "casual agricultural labor" as their response, while 24% gave "casual non-agricultural labor"). "Work on own land" was the next most common (23%), followed by "remain unemployed" (19%) and "search for work" (14%). Very few (0.3%) of the respondents said that they "don't know" what activity they would have been doing.

The survey only allows us to measure the monetary gain from obtaining a job on the scheme. For such poor people, the income gains will undoubtedly be prominent in decision making about such a scheme. In principle, there may be non-pecuniary gains or losses. However, the work available on MGNREGS is manual labor that is very similar to the type of casual work normally available in this setting. So, one would not expect much difference in non-pecuniary aspects related to the work itself. Possibly the fact that the MGNREGS work is for the government makes it more attractive, though that is a conjecture at best.

Thus, in this setting, we have credible self-reported data on the individual gains, as given by the actual earnings less forgone earnings (both reported). This is unusual; we rarely have data on the individual gains from social programs, and only aim to estimate the mean gain. Furthermore, the gains are obviously known to participants, and it would seem reasonable to assume that they are the relevant gains from trade in assignments. (Since the survey asked actual participants, it is not likely that there would be an incentive to under-report forgone earnings to help gain access to the program.) Table 1 provides summary statistics on the gains, expressed as a proportion of the overall mean wage rate.

There are three groups of participants. One reported no available work at the time; the second group expected work to be available that would have covered some non-negligible share (around half on average) of the time they were working on the scheme; the third group reported on a small gain in extra work.[14] Figure 1 provides the kernel density functions for wage earnings

[14] Further details can be found in Murgai et al. (2016) which provides density functions for the ratio of forgone days and wages to the days and wages provided by the scheme. Three modes are indicated corresponding to these three groups of workers.

from the scheme and forgone incomes (both normalized by days worked under the scheme). The variance in forgone income per day worked is 2.5 times greater than the variance in wage rates.[15]

Simulations of the equilibrium and comparisons with policy options: Applying the model of Section 2, let $W_i(1)$ denote the wage received by worker i when participating in the scheme while $W_i(0)$ is her forgone earnings while on the program, so $G_i = W_i(1) - W_i(0)$ for all i. In the simulations, a worker who receives an initial (random) assignment sells it if $G_i < P$, or (equivalently) she sells if $W_i(0) + P > W_i(1)$. A worker who did not get assigned to the program initially buys one if $G_i > P$ (or, equivalently, $W_i(1) - P > W_i(0)$). The value of P is then derived that clears the market.

The gain is measured by total wages received under the program less forgone earnings, both normalized by the total days worked on the scheme.[16] Figure 2 plots the conditional means over the range of gains, i.e., $\varphi(X) \equiv \hat{E}(G_i|G_i > X)$ for $X \in [G^{min}, G^{max}]$. (Some high values are dropped as the sample sizes become too small to be considered reliable.) To aid interpretation, the gain is expressed as a proportion of the overall sample mean wage rate (84.28 INR per day in 2009/10 prices). A key number to focus on for now is the sample mean gain, $\bar{G} = 0.393$ (s.e.=0.017; N=2307), meaning that the average gain from the existing assignment to the program represents just under 40% of the mean wage rate. This is the status quo of the existing scheme, or, in expectation, the mean for a random subsample. We see that the conditional mean rises sharply once one includes the positives ($\hat{E}(G_i|G_i \geq 0) = 0.393$ but $\hat{E}(G_i|G_i > 0) = 0.644$). In the positive range, the conditional mean rises roughly linearly with X.

The implied values of the equilibrium price and implied mean gains are found in Table 2 for selected coverage rates ranging from $\overline{D} = 0.10$ to $\overline{D} = 0.50$. Recall that the equilibrium price is $P = F^{-1}(1 - \overline{D})$, i.e., the value of gains below which one finds $1 - \overline{D}$ of the workers.

We see that the equilibrium with tradeable assignments yields substantially higher mean gains relative to the status quo. The average gains in Table 2 are 2 to 3 times higher (depending

[15] Wages are fixed on a piece-rate basis for most jobs, but this still generates a variance in the wages received per unit time. The variance in wages is 44% of the variance in gains, while the variance in forgone incomes is 110% of the variance in gains. The positive covariance (x2) accounts for the rest.

[16] There are lags in actual wage receipts (Dutta et al. 2014, Chapter 4). I include wages owed. There were some cases where forgone earnings exceeded wages received (or owed). These were treated as measurement errors (probably reflecting some misunderstanding of the survey question); the net gain was then set to zero.

on the scale of the hypothetical program). Naturally, the mean gain per participant rises as one reduces the overall coverage rate since the BIKs tend to be picked up by those with higher gains.

Turning to the "needs-based" counterfactual, the obvious criterion is household consumption per person, as used for measuring poverty in India. In practice, this would probably be based on some form of proxy-means test, as one would not have a survey-based consumption aggregate; so we would expect larger errors of targeting in reality.[17] Table 3 gives mean gains for the same coverage rates, but now assigning from the lowest consumption per person upwards until the coverage rate is met. (Table 3 also gives the required thresholds.) We see that the "poverty-based" allocation achieves slightly higher mean gains than the actual mean ($\bar{G}$ =0.393), but that it falls far short of the gains attainable with trade in equilibrium.

A common policy option to workfare is cash transfers. Past research on this policy choice has treated the BIKs as non-tradeable. Relaxing this assumption can alter the evaluative comparison of workfare versus cash transfers. Murgai et al. (2016) provide revenue-neutral comparisons of MGNREGS in Bihar with hypothetical alternatives using cash transfers. They show that, on taking account of forgone earnings and other costs in implementation, MGNREGS in Bihar does no better than a (revenue-neutral) universal basic income. Nor does the workfare scheme do any better in reducing poverty than transfers paid to those holding one of the existing food ration cards, known as the "Below Poverty Line" cards. Murgai et al. (2016) find that the outcomes for poverty are essentially the same across these options. Although the self-targeting feature of the workfare scheme (whereby non-poor are unlikely to want this type of work) assures that more of the work tends to go to poor families, there is a trade-off against the extra costs incurred (by both participants and the government) under workfare. Since the comparison is a tie, the extent of the increase in mean gains found here with trade in assignment suggests that this would tilt the balance in favor of the workfare scheme. The self-targeting aspect would still be working, but aggregate forgone earnings would be lower.

Heterogeneity in the gains: A key aspect of heterogeneity in this context is the extent of poverty. We see from Figure 3 that the conditional mean gains tend to fall with higher household consumption, although the (non-parametric) regression line is being pulled up at the bottom by a

[17] See, for example, Brown et al. (2019).

few outliers. It is clear, however, that the mean gains from the equilibrium assignment are no lower for those workers coming from households living below the median (Table 2, Columns (5) and (7)). The potential gains are spread through the distribution of living standards.

Gender differences are also of interest. 29% of the sampled workers are women. The overall mean gain is slightly lower for women. The relationship with household consumption is similar by gender, though the tendency for the gains to fall as household consumption per person rises tends to be less evident for women (Figure 3(b)). When one calculates mean gains separately by gender (conditional on gains exceeding the market-clearing price), the differences are small at all levels of coverage (Table 2).

The gender split in Table 2 assumes a common market for both genders. Instead, one might split the market by gender (with trades only allowed among the same gender). This raises the market-clearing price for men, and lowers it for women, while the aggregate gains are similar to Table 2. For example, at $\overline{D} = 0.50$, the equilibrium price rises to 0.31 for men and falls to 0.18 for women, with mean gains of 0.760 and 0.724 respectively. The population weighted mean gain across genders is 0.750 (Table 2).

4. Impediments to realizing the gains in practice

The evidence in the previous section is at least suggestive of large unexploited gains to poor people from trade in their BIK assignments. This points to the potential for policy intervention. A minimalist form of that intervention would be that the government does nothing to restrict trade. However, there are reasons why the gains from trade will still not be fully realized. The reasons relate to market and institutional impediments that also play a role in creating poverty in the first place. Realizing the potential gains from facilitating tradeable assignments may require active public effort to facilitate and support the creation of a market, along with complementary policies to assure that poor people have access to that market.

Four main concerns can be identified. The first relates to credit-market imperfections, such that poor potential beneficiaries simply cannot afford to purchase a BIK assignment. Finding that the mean gains are no lower for poorer household does not, of course, imply that poor households who did not get an initial assignment could afford to buy one. Given liquidity constraints, the benefits may still be disproportionately captured by the relatively well off.

In support of this claim, Figure 4 gives the market price as a share of monthly consumption for a family of five. The cost rises to a (clearly) prohibitive share of consumption among poor families. Even the average shares for those living below the median are high; at a coverage rate of 0.5, the purchase cost is a little over 20% of consumption; it rises to over 80% at a coverage rate of 0.1. (Column 6 of Table 2 gives the share of monthly consumption for those living below the median and the standard errors.) Full payment of the market price up-front could represent a prohibitively large share of consumption for poor households.

If a government-regulated market was introduced, then this liquidity-constraint problem could be solved by introducing a "pay-as-you-go" option. One can anticipate resistance to such a step, as it might be interpreted as unequal (net) wages for the same work. That concern would need to be weighed against the potential benefits to poor workers in gaining access to the program. Good communications would clearly be needed about the reasons for the deduction from wages (to cover the initial purchase of an assignment).

Second, there will be scope for less liquidity-constrained, non-poor, people to participate in a new market for BIKs that strives to capture the unexploited gains from trade implied by this paper's results. This is a concern, though it could also be addressed by policy design features. In the context of explicitly targeted programs, one may want to maintain or even tighten existing eligibility criteria. The norm in antipoverty and other social programs appears to be under-coverage of those deemed eligible, with rationing among the set of eligible participants.[18] While the results of this paper indicate large potential gains from trade among current participants, there would be new risks in substantially expanding the set of those eligible since the potential gains from trading assignments would attract more affluent participants not facing the same liquidity constraints on participating in the market.

A means of addressing this second concern in the context of self-targeted programs is to introduce a minimum participation threshold and to only facilitate trade in assignments once that minimum is reached. While the self-targeting mechanism tends to discourage participation by the well-off, creating the market option could well impact selection into the program—attracting "speculators" and middlemen who do not intend to work, but only re-sell their assignment if they

[18] For an overview of evidence on this point see Ravallion (2017).

get one. However, one can avoid this by requiring a sufficiently long initial period of work before re-sale is available as an option.

Third, there may be social frictions in the flow of knowledge. Dutta et al. (2014) also surveyed participants' knowledge about MGNREGS. Most of the sample had heard of the scheme, but knowledge about the scheme's rules and provisions was poor, especially for women. The above calculations may overstate gains to women relative to men. Realizing the benefits from trade clearly requires that participants are reasonably well-informed about their personal gains and they know how to access the market. Creating a market for BIK assignments would presumably change the incentives for seeking and spreading information on the scheme. Even so, information dissemination efforts would probably be needed, to complement a switch to tradeable assignments in social programs.

There is evidence from the same setting that external intervention through "infotainment" can enhance individual knowledge about this program and (hence) the individual gains. Ravallion et al. (2015) report results from their RCT using an entertaining movie (produced for this purpose) to teach people their rights and the rules and administrative procedures under the scheme. The movie did enhance knowledge when assessed by a quiz given before and after seeing the movie, and especially so for women. However, in studying the impacts of this RCT, Alik-Lagrange and Ravallion (2019) find evidence of social frictions on information dispersal within villages—frictions that disadvantage lower caste and poorer individuals. Effort would be needed to widely advertise the scope for mutually beneficial trades in assignments.

The fourth concern is whether the administrative capacity will be present in poor places to implement an efficient, largely corruption free, secondary market in BIK assignments. This is no small matter. MGNREGA does have a quite sophisticated (public-access) web-based information system, and it would seem plausible that the software to support a market for assignments could also be developed, preferably integrated with the existing information system.[19] However, what actually happens on the ground could deviate appreciably from the program's formal operating rules. Dutta et al. (2014) identify a number of administrative

[19] The existing information system does not record the extent of rationing under the scheme. Under the scheme's rules, the state government is responsible for paying an unemployment allowance if work cannot be found. This gives the state government a strong incentive to report that all requests for work were honored. That is inconsistent with what has been observed based on surveys (Dutta et al. 2012, 2014; Desai et al. 2015).

performance problems in poor areas of India that make it hard for the scheme studied here to attain its potential impact on poverty even without tradeable assignments. These features can also be thought of institutional failures that help create poverty in the first place. Of course, enhancing public administrative capacity is an important element of development policy more broadly—a channel that is relevant to the efficacy of a wide range of policies.

5. Conclusions

Given the limitations of targeting efforts in practice—stemming from asymmetric information—letting people trade their assignments of benefits-in-kind in antipoverty programs can enhance the impact on poverty. Indeed, under ideal conditions, the competitive equilibrium will attain the maximum aggregate gain with perfect information even when the decision maker is ignorant of the individual gains. At the same time, the existence of private information is clearly one of the reasons why we know so little about how much scope there may be for poor people to gain from trade in assignments. In situations in which governments strive to prevent trade in assignments, and have the power to do so, this also means that we know rather little about the welfare costs of such restrictions.

This also has implications for how we interpret evaluative evidence on the impacts of antipoverty programs. Standard randomized controlled trials need not yield valid inferences about equilibrium assignments since the latter will not (in general) be random. Nor do these methods provide valid inferences when restrictions on the re-sale of benefits in kind are only enforceable in scaled up programs.

Based on a theoretical model of a competitive market for assignments, the paper has offered some evidence on the gains from trade in assignments to a large antipoverty program in a poor state in rural India. The program aims to reduce poverty by providing jobs on labor-intensive public works projects. A novel feature of the setting is that it is feasible to use surveys to measure individual pecuniary gains from the jobs provided, even though this information is unlikely to be available for implementing a scheme of transferable in-kind benefits. The surveyed workers were asked about both actual wages received and labor-market options at the time. Thus, we have an estimate of the distribution of the gains, not just the mean. And we have the joint distribution of the gains and other characteristics of the worker and household. Informed by a theoretical model of the equilibrium with tradeable assignments, the paper has used these novel data to simulate various stylized schemes, at different levels of coverage. Thus, the paper has been able to measure the extent of the unexploited gains from trade in this setting and to study their incidence by household levels of living and the gender of workers.

The calculations indicate that there are large unexploited gains to poor people from trade in assignments. A competitive market for tradeable assignments would generate aggregate gains that are orders-of-magnitude greater than the current mean gains to these (mostly very poor) families. Tradeable assignments would also have far greater impact (by a similar magnitude in terms of mean gains) than an allocation without the re-sale option that is targeted to workers from consumption-poor families. Facilitating tradeable assignments in this setting can also make workfare more effective against poverty than various options for (budget-neutral) cash transfers.

The results do not suggest that the potential gains from the market-based assignment would tend to favor more affluent families. The mean gains were very similar between the poorest half in terms of household consumption per person as for the rest. Indeed, the simulated allocations with tradeable assignments imply a tendency for somewhat larger gains among poorer households, not the opposite. Gains are similar between male and female workers.

The paper has pointed to some likely impediments to realizing the gains from trade, including credit market failures, capture by non-poor speculators, and imperfect information. There are complementary policies that may well be needed. A "pay-as-you-go" option could help address liquidity constraints stemming from credit market imperfections. Eligibility criteria and minimum participation requirements (to encourage self-targeting) could help dissuade non-poor speculators and middlemen from entering the market. And information campaigns are clearly desirable, and not so hard to do. Limitations on local-level administrative capacity and the (related) scope for corruption also warn for caution. The same features impede many other aspects of economic development. As is often the case, success in one policy effort—in this case assuring that the potential benefits to poor people of an antipoverty program are fully realized in practice—may require success on multiple fronts.

References

Akerlof, George, 1978, "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning." *American Economic Review* 68 (1): 8–19.

Alik Lagrange, Arthur, and Martin Ravallion, 2018, "Workfare versus Transfers in Rural India," *World Development* 112: 244-258.

Alik Lagrange, Arthur, and Martin Ravallion, 2019, "Estimating Within-Group Spillover Effects Using a Cluster Randomization: Knowledge Diffusion in Rural India," *Journal of Applied Econometrics* 34: 110-128.

Besley, Timothy, and Stephen Coate, 1992, "Workfare vs. Welfare: Incentive Arguments for Work Requirements in Poverty Alleviation Programs," *American Economic Review* 82(1): 249–61.

Brown, Caitlin, Martin Ravallion and Dominique van de Walle, 2018, "A Poor Means Test? Econometric Targeting in Africa," *Journal of Development Economics* 134: 109-124.

Che, Yeon-Koo, Ian Gale and Jinwoo Kim, 2013, "Assigning Resources to Budget-Constrained Agents," *Review of Economic Studies* 80(1): 73-107.

Coate, Stephen, 1989, "Cash versus Direct Food Relief," *Journal of Development Economics* 30(2): 199-224.

Cunha, Jesse, 2014, "Testing Paternalism: Cash versus In-Kind Transfers." *American Economic Journal: Applied Economics* 6 (2): 195–230.

Currie, Janet, and Firouz Gahvari, 2008, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature* 46(2): 333-383.

Datt, Gaurav, and Martin Ravallion, 1994, "Transfer Benefits from Public-Works Employment," *Economic Journal* 104: 1346-1369.

Desai, Sonalde, Prem Vashishtha and Omkar Joshi (2015), *Mahatma Gandhi National Rural Employment Guarantee Act: A Catalyst for Rural Transformation,* New Delhi: National Council of Applied Economic Research.

Dutta, Puja, Rinku Murgai, Martin Ravallion and Dominique van de Walle, 2012, "Does India's Employment Guarantee Scheme Guarantee Employment?" *Economic and Political Weekly* 48 (April 21): 55-64.

Dutta, Puja, Rinku Murgai, Martin Ravallion and Dominique van de Walle, 2014, *Right-to-Work? Assessing India's Employment Guarantee Scheme in Bihar.* World Bank.

Heckman, James, Serio Urzua and Edward Vytlacil, 2006, "Understanding Instrumental Variables in Models with Essential Heterogeneity," *Review of Economics and Statistics* 88(3): 389-432.

Himanshu, 2019, "Inequality in India: Review of Levels and Trends," WIDER WP 2019/42.

Hirvonen, Kalle and John Hoddinott, 2021, "Beneficiary Views on Cash and In-Kind Payments: Evidence from Ethiopia's Productive Safety Net Programme," *World Bank Economic Review* 35(2): 398–413.

Jalan, Jyotsna, and Martin Ravallion, 2003, "Estimating the Benefit Incidence of an Anti-Poverty Program by Propensity-Score Matching," *Journal of Business and Economic Statistics* 21(1): 19-30.

Kapoor, Mahima, 2020, "Only 20% Surveyed Found Jobs under NGREGA During Lockdown," Bloomberg.

Moffitt, Robert, 2006, "Welfare Work Requirements with Paternalistic Government Preferences," *Economic Journal* 116(515): F441-F458.

Murgai, Rinku, Martin Ravallion and Dominique van de Walle (2016), "Is Workfare Cost Effective against Poverty in a Poor Labor-Surplus Economy?," *World Bank Economic Review* 30(3):413-445.

Narita, Yusuke, 2018. "Toward an Ethical Experiment," Cowles Foundation Discussion Paper No. 2127, Yale University.

Ravallion, Martin, 2017, *Interventions against Poverty in Poor Places*, 20th Annual WIDER Lecture, World Institute of Development Economics, Helsinki.

Ravallion, Martin, 2021, "Is a Decentralized Right-to-Work Policy Feasible?" *Development, Distribution, and Markets,* Kaushik Basu, Kenneth Kletzer, Sudipto Mundle and Eric Verhoogen (eds), New Delhi: Oxford University Press.

Ravallion, Martin, Gaurav Datt, and Shubham Chaudhuri, 1993, "Does Maharashtra's 'Employment Guarantee Scheme' Guarantee Employment? Effects of the 1988 Wage Increase," *Economic Development and Cultural Change* 41: 251–75.

Ravallion, Martin, Dominique van de Walle, Rinku Murgai and Puja Dutta, 2015, "Empowering Poor People through Public Information? Lessons from a Movie in Rural India," *Journal of Public Economics* 132: 13-22.

Sah, Raaj Kumar, 1987, “Queues, Rations, and Market: Comparisons of Outcomes for the Poor and the Rich,” *American Economic Review* 77(1): 69-77.

Weitzman, Martin, 1977, “Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need it Most?” *Bell Journal of Economics* 8(2): 517-524.

Figure 1: Kernel density functions for wages and forgone incomes on rural public works projects in Bihar, India

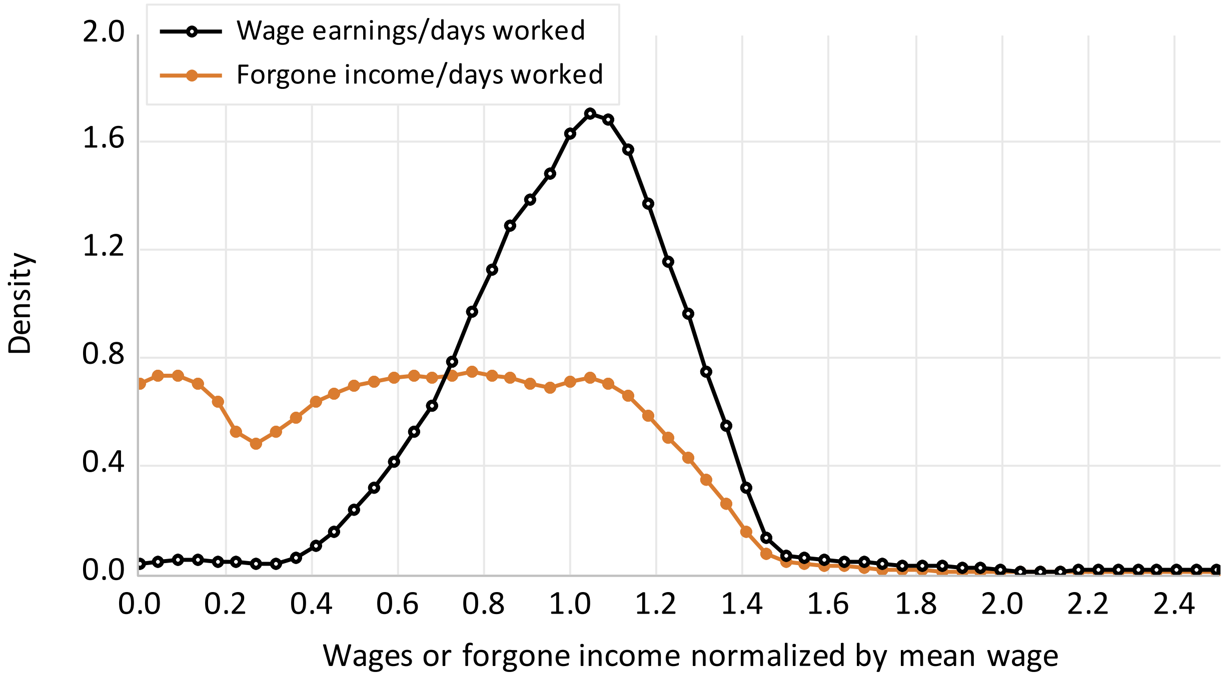

Note: Kernel densities with a bandwidth od 0.25. Top tail truncated at 2.5 (after estimation) to make the graph more readable.
Source: Author's calculations from the survey data collected in two rounds, 2009-10, by a team including the author and documented in Dutta et al. (2014). N=2307 after deleting 17 implausibly high outliers (exceeding 400 INR per day).

Figure 2: Conditional mean gains ($\hat{E}(G_i|G_i > X)$)

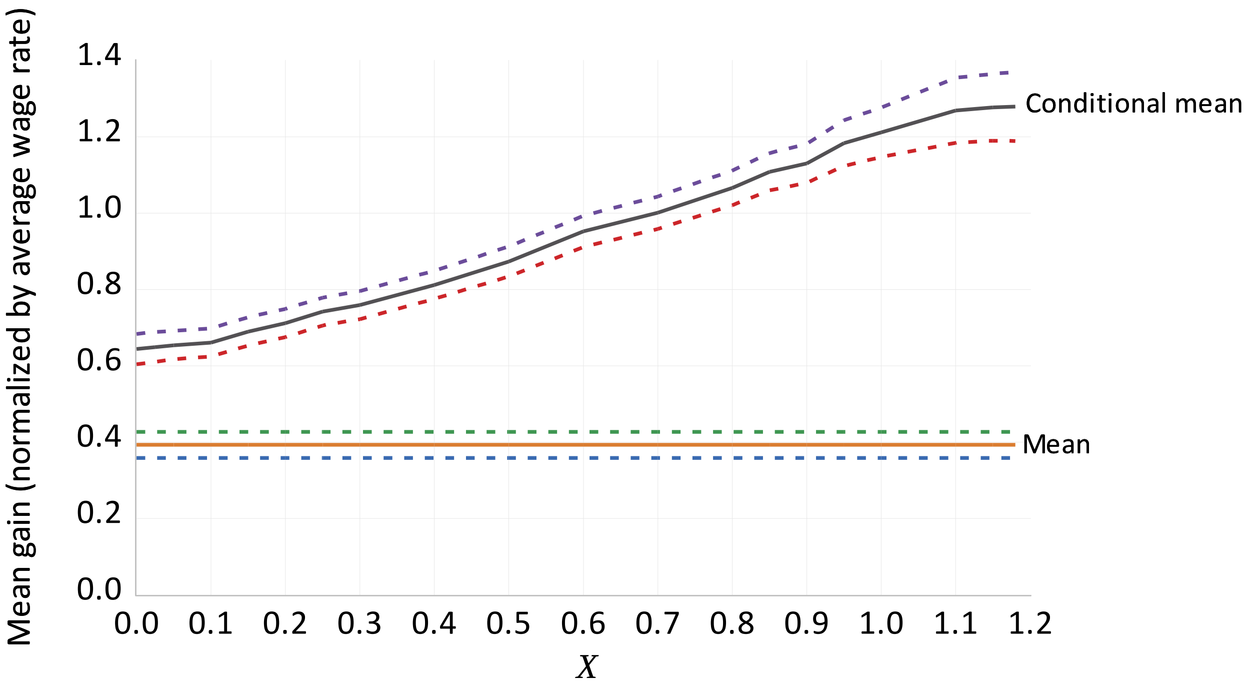

Note: Gain from employment in the National Rural Employment Guarantee Scheme in Bihar, India. The gain is normalized by the overall mean wage rate (84.28 INR per day). Lower and upper bounds of the 95% confidence intervals. Standard errors are clustered at the household level and scaled up by a factor of $\sqrt{2}$ to allow for the re-surveying of individuals across rounds.

Source: Author's calculations from the survey data collected in two rounds, 2009-10, by a team including the author and documented in Dutta et al. (2014). N=2307 after deleting 17 implausibly high outliers (exceeding 400 INR per day).

Figure 3: Gains plotted against household consumption per person

(a) All workers

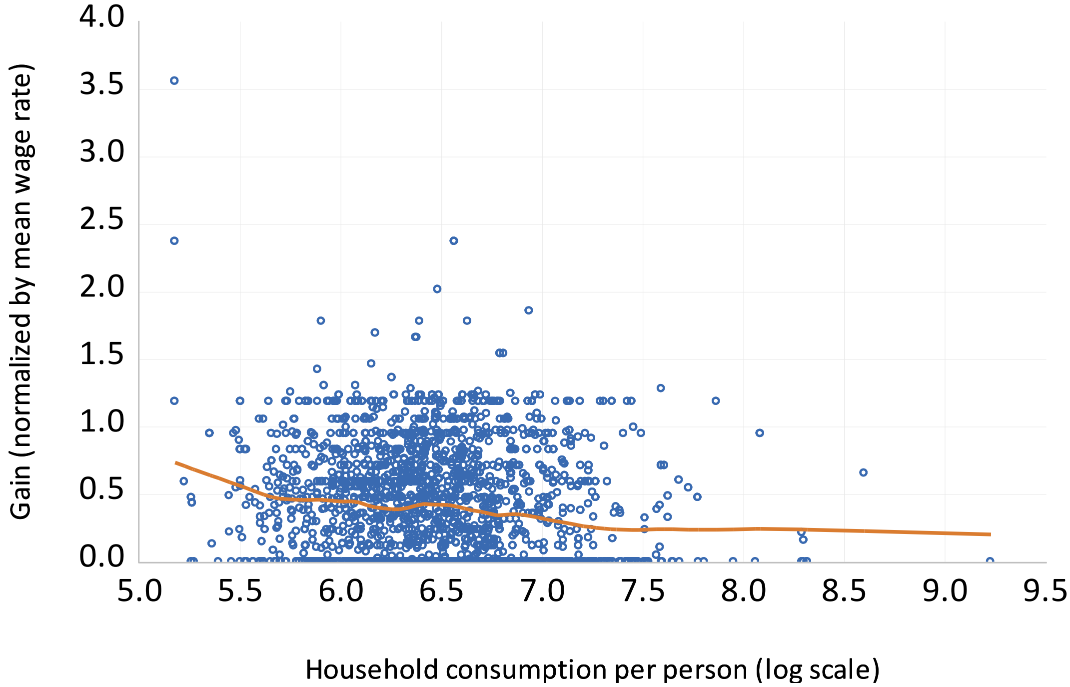

(b) Split by gender

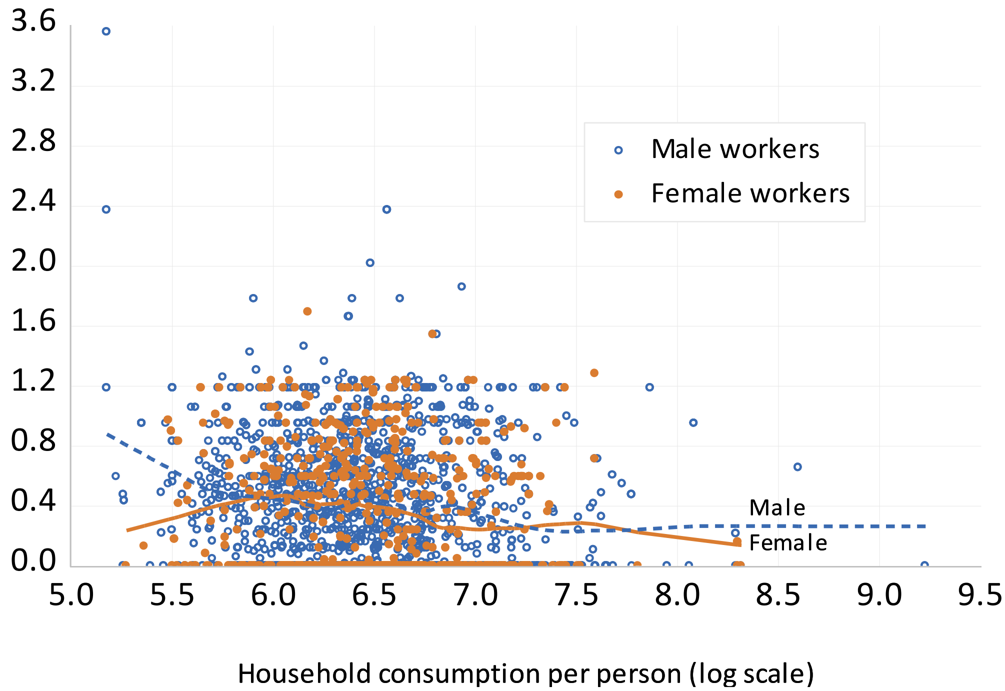

<u>Note</u>: See Figure 2. Nearest neighbor smoothed scatter plot.
<u>Source</u>: See Figure 2.

Figure 4: Market price of assignments as a share of household consumption

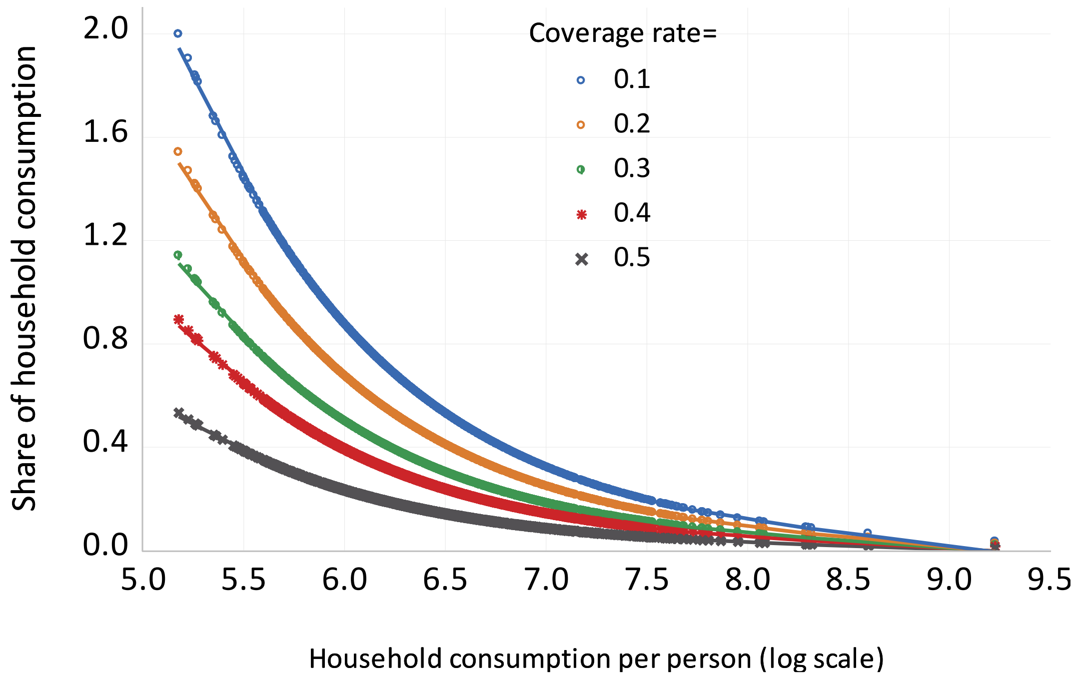

Note: Share of consumption calculated for a household of five people with one person selling an assignment for 20 days per month. Nearest neighbor smoothed scatter plots.
Source: See Figure 2.

Table 1: Summary statistics on monetary gains from public-works jobs in Bihar, India

	Sample mean	Standard error	N
Gain	0.393	0.017	2307
Gain for workers in households with below median consumption per person	0.436	0.024	1150
Gain for workers in households with above median consumption per person	0.350	0.031	1151
Gain for female workers	0.363	0.031	669
Gain for male workers	0.406	0.020	1638
Gain for female workers in household with below median consumption per person	0.435	0.047	315
Gain for male workers in household with below median consumption per person	0.437	0.027	835
Gain for female workers in household with above median consumption per person	0.299	0.037	352
Gain for male workers in household with above median consumption per person	0.373	0.028	799

Notes: Mean gains (wage rate less forgone earnings) normalized by overall mean wage rate. 17 extreme values for the gains (exceeding 400% of mean wage rate) are dropped. Standard errors are clustered at the household level and scaled up by a factor of $\sqrt{2}$ to allow for the re-surveying of individuals across rounds.

Table 2: Mean gains from competitively tradeable assignments

(1) Coverage rate ($\bar{D}$)	(2) Market-clearing price (P)	(3) Pop. share trading	(4) Mean gain for treated ($\hat{E}(G_i \mid G_i > P)$)	(5) Mean gain for those below median consumption	(6) Share of consumption for below median household	(7) Mean gain for those above median consumption	(8) Mean gain for female workers	(9) Mean gain for male workers
0.5	0.28	0.25	0.750	0.736	0.217	0.767	0.757	0.748
			(0.018; 1154)	(0.024; 654)	(0.004; 654)	(0.028; 497)	(0.030; 312)	(0.030; 842)
0.4	0.47	0.24	0.845	0.826	0.363	0.869	0.827	0.852
			(0.018; 920)	(0.024; 522)	(0.007; 522)	(0.030; 395)	(0.030; 262)	(0.023; 658)
0.3	0.60	0.21	0.952	0.944	0.463	0.963	0.931	0.961
			(0.021; 677)	(0.027; 367)	(0.011; 367)	(0.030; 307)	(0.028; 190)	(0.025 487)
0.2	0.81	0.16	1.068	1.059	0.642	1.078	1.043	1.078
			(0.023; 465)	(0.030; 249)	(0.020; 249)	(0.033; 214)	(0.025;129)	(0.030; 336)
0.1	1.05	0.09	1.215	1.217	0.834	1.214	1.188	1.225
			(0.033; 244)	(0.042; 126)	(0.038; 126)	(0.047; 117)	(0.018; 63)	(0.044; 181)

Notes: Mean gains normalized by overall mean wage rate. The status quo mean is 0.393 (s.e.=0.017). Equilibrium prices calculated numerically to nearest second decimal place. Share of consumption (column 6) calculated for household of five people with one person selling an assignment for 20 days per month. Median consumption per person=INR 647 per month. Standard errors in parentheses, followed by sub-sample size. Standard errors are clustered at the household level and scaled up by a factor of $\sqrt{2}$ to allow for the re-surveying of individuals across rounds.
Source: See Figure 1.

Table 3: Mean gains from needs-based versus market-based assignments

Coverage rate ($\bar{D}$)	Cut-off point for consumption (INR/person/month)	Mean gain for needs-based allocation	Mean gain using market (Table 1)
0.5	646	0.436	0.750
		(0.024; 1145)	(0.018; 1154)
0.4	580	0.438	0.845
		(0.027; 925)	(0.018; 920)
0.3	515	0.453	0.952
		(0.032; 682)	(0.021; 677)
0.2	447	0.468	1.068
		(0.042; 463)	(0.023; 465)
0.1	372	0.472	1.215
		(0.058; 230)	(0.033; 244)

Notes: Mean gains normalized by overall mean wage rate. The needs-based allocation goes to the poorest households based on consumption per person, with the cut-off determined by the coverage rate. Standard errors in parentheses, followed by sub-sample size. Standard errors are clustered at the household level and scaled up by a factor of $\sqrt{2}$ to allow for the re-surveying of individuals across rounds.

Source: See Figure 1.