Targeting with Agents*

Paul Niehaus
UC San Diego

Antonia Atanassova
The Nielsen Co.

Marianne Bertrand
University of Chicago

Sendhil Mullainathan
Harvard University

June 16, 2011

Abstract

Targeting economic assistance to the poor is a central problem in development. We study the problem of designing a proxy means test when the agent who implements it is potentially corruptible. Conditioning on more poverty indicators may strictly worsen targeting in this environment because of a novel tradeoff between the statistical accuracy of the means test and its enforceability. We assess the necessary conditions for this tradeoff in the context of Below Poverty Line card allocation in India. Using original data on rule-breaking and bribery we find that less eligible households are less likely to hold cards and pay larger bribes to get them, but that both effects are weak and consequently that rule violations are widespread. "Soft targeting" on unofficial measures of poverty is also weak at best, with the net result that the de facto allocation is much less progressive than the de jure one.

JEL codes: D02, D61, D73, H11, H40, H75, I38, K42, O1

Keywords: Targeting, Corruption

*Thanks to Nageeb Ali, Prashant Bharadwaj, Quy-Toan Do, Rocco Macchiavello, Mark Rosenzweig, Jeremy Shapiro, and seminar participants at NEUDC, UC San Diego, the World Bank, and MOVE Barcelona for helpful comments. Thanks also to the Institute for Financial Management and Research, Comat Technologies (P) Ltd., and the Government of Karnataka for their help, and to Jennifer Bryant, Kushal Tantry, and Qiuqing Tai for outstanding research assistance. The International Finance Corporation provided generous financial support for the data collection.

1 Introduction

Which households should be eligible for social assistance? Targeting is a central problem in public economics, particularly for developing countries. Because these countries do not have reliable data on the income or consumption of their citizens, they often rely instead on "proxy means tests," or categorizations of households into eligible and ineligible groups based on easier-to-observe characteristics. For example, households that own color televisions might be ruled ineligible. A large literature has developed showing how to design optimal PMTs by applying statistical decision theory to household survey data.[1] In this paradigm, "the optimal policy equates the marginal reduction in poverty from a further indicator being used with its marginal administrative cost" (Besley and Kanbur, 1990, 14).[2]

In practice, however, the rule implemented may be different from the rule designed. Research on corruption has provided many examples of ways in which the officials who implement social programs bend or break the rules. They divert transfers from the intended recipients (Reinikka and Svensson, 2004; Olken, 2006), inflate claims about program participation (Niehaus and Sukhtankar, 2009), demand bribes to issue permits to eligible recipients (Svensson, 2003), and take bribes to issue permits to ineligible recipients (Bertrand et al., 2007). The optimal response to such problems may involve not only tougher enforcement but also changing the very nature of the task assigned to officials (Banerjee, 1997; Banerjee et al., 2011).

Motivated by these observations, we ask a simple but important question: how should targeting rules be designed when they must be implemented by corruptible agents? We study a model in which a principal with progressive preferences defines a targeting rule to be implemented by a subordinate official. The official has distinct preferences: he is tempted to demand bribes, and may also want to give out slots to voters or friends. Because arbitrarily large punishment are not available the principal cannot perfectly discipline the official (Becker, 1968; Mirrlees, 1999). Instead the official sets a schedule of household-specific bribe-prices (possibly equal to 0) that optimally trade off his allocative preferences, bribe rents, and expected penalties. The price schedule then determines the allocation of slots and rents.

We use this framework to examine properties of optimal targeting rules. The most striking

[1]See Ravallion (1989), Ravallion and Chao (1989), Besley (1990), Besley and Kanbur (1990), Glewwe (1992), Ravallion and Sen (1994), Grosh and Baker (1995), Schady (2000), Park et al. (2002). See Grosh (1994) and Coady et al. (2004) for reviews of targeting in practice.

[2]Sophisticated applications may also take into account the distortion of household incentives that targeting generates (Mirrlees, 1971). An alternative approach to targeting is to impose requirements on beneficiaries that make the program "self-targeting" (Besley and Coate, 1992), generating a different set of agency problems (Niehaus and Sukhtankar, 2009).

conclusion is that conditioning eligibility on an additional poverty indicator may strictly worsen targeting. This is true even though from a purely statistical perspective the additional indicator can only help. Of course, if the indicator is not perfectly verifiable then one might expect these gains to be diluted because of the monitoring problem. What we show is that they may in fact be reversed. The reason is that the additional indicator affects not only who is eligible (the statistical effect) but also how verifiable the (in)eligibility of other inframarginal households is (the enforcement effect). If the enforcement effect is sufficiently negative it may trump the positive statistical effect.

A concrete example may help illustrate this. Suppose that households with paved floors are ineligible. Some of these households are in fact poor, so this rule is statistically imperfect. On the other hand, a third party can verify ineligibility simply by observing a paved floor. Anticipating this, the official will be reluctant to sell slots to ineligible households. Now consider refining the rule so that households with paved floors are eligible unless they also have a television set. Statistically speaking this will be an improvement provided the households newly made eligible are sufficiently poor. For a third party to verify ineligibility, however, he must now verify that a household has both assets. Verifying both facts is harder than verifying just the first, so the official will be less apprehensive about giving (or selling) slots to ineligible households. If this enforcement effect is strong enough it may more than offset the statistical gains.[3]

Certain conditions must hold for this tradeoff between statistical accuracy and enforceability to bind. The official's preferences must be misaligned with the principal's, so that mere delegation is not optimal. Enforcement must be imperfect, and it must imperfect in such a way that more obviously ineligible households are less likely to receive slots than marginally ineligible ones. In the second half of the paper we examine whether these conditions hold in a specific setting of interest, the allocation of Below Poverty Line (BPL) cards in India. BPL cards are India's most important targeting mechanism; participation in a wide range of public schemes, including the Targeted Public Distribution System (TPDS), is restricted to card-holders. Different states use different proxy means tests to allocate BPL cards. In Karnataka, where we work, the PMT consists of a series of exclusion restrictions. For example, a household that owns a water pump or an automobile is ineligible. Local officials are responsible for implementing this rule, subject to monitoring by back-checking teams.

One novel feature of our data is that we observe both households' statutory eligibility and their actual BPL status, letting us quantify rule violations. We find that violations are widespread: 70% of the ineligible households in our sample have BPL cards, while 13% of

[3]There is also a positive effect on enforcement for households that were previously eligible. Our formal result identifies cases where the net effect on enforcement is negative.

eligible households do not. Overall 48% of the households in our sample are misclassified, and eligible households are only 21% more likely than ineligible ones to hold cards. A household's probability of holding a card decreases monotonically with the number of eligibility rules it violates, consistent with the notion that degrees of ineligibility are important.

While inconsistent with strong enforcement, widespread rule-breaking need not imply poor targeting outcomes. It could be that officials bend the rules because they have better "soft" information about poverty than the rule-designer. There is little evidence of this in the data, however. While statutory eligibility is correlated -0.55 with log income in our sample, the correlation with the actual allocation of BPL cards is only -0.23. Strikingly, this is weaker than the correlation between income and a number of individual, readily observable asset criteria. For example, ownership of a water pump and having a gas connection are correlated -0.32 and -0.30 with income, respectively. The allocation of BPL cards would thus be more progressive if the government could consistently enforce an eligibility rule based exclusively on one of these criteria.

A second novel feature of our data is that 93% of BPL card recipients (and 73% of households overall) reported the prices they were charged for BPL cards, which lets us examine the role of bribery in the allocation process. Interestingly, bribes are both commonplace and small. 75% of households report paying a price above the statutory fee, but the mean (conditional) overpayment is only Rs. 14, smaller than the mean *monthly* subsidy of Rs. 133 that BPL card-holders report receiving at ration shops. When we estimate the pricing equations implied by our model we find that prices increase monotonically in the number of eligibility criteria a household violates, again consistent with the notion that degrees of ineligibility matter. This effect is small, however, with ineligible households paying only Rs. 3 more on average than ineligible ones and only 1% less likely to obtain a card. These results imply that enforcement is weak and suggest that non-monetary favors (e.g. votes) may play an important role in the allocation process.

Overall the data confirm that officials' preferences not well-aligned with those of the BPL rule designers, that enforcement is weak, and that degrees of (in)eligibility matter. In light of our theoretical results this suggests a role for more robust targeting rules. Interestingly, the policy debate within India has recently shifted in this direction. While early critiques of BPL policy focused on statistical accuracy (Sundaram, 2003) there is increasing recognition that mis-implementation is as important a constraint on performance (Hirway, 2003). Dreze and Khera (2010) have recently proposed using dramatically simpler targeting criteria, such that every household can attribute its inclusion or exclusion to a single criterion, on the grounds that this would reduce fraud. Our model provides a formal justification for their idea.

Our empirical results extend a line of work including Galasso and Ravallion (2005) and Bardhan and Mookherjee (2006c) that has documented village-level correlates of targeting performance (but has not measured rule performance and adherence directly). In a similar vein, Banerjee et al. (2009) contrast the targeting of government schemes with that achieved by a local NGO in West Bengal, and Alatas et al. (2010) provide an experimental comparison between rule-based and community-based targeting in Indonesia. The paper also relates indirectly to work on geographic targeting (Schady, 2000; Park et al., 2002) in that geographic targeting may be particularly attractive in weakly-enforced settings. Finally, our model has connections to the literature on decentralizing public programs, in which decentralization is seen as trading off the benefit of better local information against the risk of program capture by local elites (Bardhan, 2002; Bardhan and Mookherjee, 2000, 2005, 2006a,b). An analogous tradeoff arises in our model: more vigorous enforcement of a targeting rule may or may not lead to a more progressive allocation of benefits depending on how accurate the rule is and how progressive the official's own tastes are. In this sense our empirical comparison between the progressivity of de jure and de facto allocations contributes to the broader decentralization literature.

The rest of the paper proceeds as follows: Section 2 develops the theoretical apparatus necessary to think about targeting through agents; Section 3 describes the empirical context in which we work and the data we collected; Section 4 analyses targeting and rent extraction in this setting; and Section 5 concludes.

2 Targeting with Agents

The model describes a hierarchical relationship between a principal, an agent, and a set of households. The essence of the game is that the principal defines a targeting rule which the agent then implements, more or less faithfully. Section 2.1 characterizes the best response of the agent to a given targeting rule and enforcement environment; Sections 2.2 defines the principal's optimization problem taking this behavior into account. Sections 2.3 and 2.4 study delegation and rule design, respectively.

2.1 The Agent

A principal wishes to allocate slots among a set of households. Household i has income $y_i \in \{\underline{y}, \overline{y}\}$ and other characteristics $\mathbf{x}_i \in \mathbf{X}$ which are potentially correlated with income: for example, one component of $\mathbf{x}_i$ might indicate whether or not household i owns a color television.[4] Finally, households vary in their willingness to pay v_i for a slot; v_i is distributed exponentially with rate parameter $\frac{1}{\eta_i}$ where $0 < \underline{\eta} \leq \eta_i \leq \overline{\eta} < \infty$. Let $F(y, \mathbf{x}, \eta)$ be the joint distribution of these household attributes. We model variation in the elasticity of demand η_i in order to allow for unobservable heterogeneity when we turn to empirical applications of the model in Section 4, but we will abstract from it in presenting the main theoretical results in Sections 2.2-2.4.

The principal cannot observe income directly; instead he must define a proxy means test in terms of more readily observable characteristics. Formally, a targeting rule is a subset $R \subseteq \mathbf{X}$, with the interpretation that household i is eligible if and only if $\mathbf{x}_i \in R$. The rule is implemented by a relatively well-informed official who observes $(y_i, \mathbf{x}_i, \eta_i)$, though not the idiosyncratic valuation v_i. The official's payoff depends both on the allocation of slots and on his own net income Y. If $a_i \in \{0, 1\}$ indicates whether household i obtains a slot then the official's payoff is

$$U(Y, \{a_i\}) = Y + \underline{\alpha} \int_{y_i = \underline{y}} a_i di + \overline{\alpha} \int_{y_i = \overline{y}} a_i di \tag{1}$$

The parameters $(\underline{\alpha}, \overline{\alpha})$ measure the official's distributive preferences. An official with $\underline{\alpha} = \overline{\alpha} = 0$ simply maximizes his own income, while high $\underline{\alpha}$ ($\overline{\alpha}$) implies a strong preference for giving slots to the poor (rich). An official motivated primarily by electoral issues might place a high value on giving out slots to all voters and therefore have both high $\underline{\alpha}$ and high $\overline{\alpha}$. We can thus capture the uncertainty about the incentives and motives of local officials that has dominated the debate over decentralizing welfare. As Jean Dreze and Amartya Sen put it,

> "The leaders of a village community undoubtedly have a lot of information relevant for appropriate selection. But in addition to the informational issue, there is also the question as to whether the community leaders have strong enough motivation – or incentives – to give adequately preferential treatment to vulnerable groups. Much will undoubtedly depend on the nature and functioning of political institutions at the local level, and in particular on the power that the poor and the deprived have in the rural community." (Dreze and Sen (1989), quoted in Bardhan and Mookherjee (2006c))

In this spirit, the αs serve as combined measures of internal and external motives. Internal motives are simply the official's own intrinsic preferences over distribution, while external motives include social pressure and local political constraints. We will say that the official

[4] We use a binary indicator of poverty for simplicity, sidestepping issues of relative poverty measurement that are not central to the argument. This corresponds to the special case P_0 of the class of poverty measures defined by Foster et al. (1984).

has *progressive* preferences if $\underline{\alpha} > \overline{\alpha}$. Note that progressive preferences are still consistent with a strong preference for giving benefits to the rich ($\overline{\alpha} \gg 0$).

Because he has distinct allocative preferences and because he cares about his own profits, the official will be tempted to break the targeting rule R specified by the principal. If he does break the rule with respect to household i he is detected and punished by the principal with probability $\pi(a_i, \mathbf{x}_i, R)$, which reflects the structure of monitoring and the likelihood with which rule-breaking by the agent can be conclusively proved. We assume that $\pi(a, \mathbf{x}, R) = 0$ if $a = 1(\mathbf{x} \in R)$ so that the official is never punished for following the rule, while $\pi(a, \mathbf{x}, R) > 0$ if $a \neq 1(\mathbf{x} \in R)$ so that the official is always punished with some positive probability when he does break the rule.

When the official is detected breaking a rule he is fined a (monetized) amount $f > 0$ for rule-breaking. f should be interpreted broadly to include anything that motivates the official to adhere to the rules, including career concerns and psychic costs. It can be thought of as an (inverse) measure of discretion: as $f \to 0$ the official can choose the allocation of slots freely without any outside influence, while as $f \to \infty$ adherence to the rules becomes crucial. If the principal could make f arbitrarily large then (we shall see) he could perfectly enforce his targeting rule. In practice, however, there are limits on how harshly corrupt officials can be punished. There are typically norms of proportionate punishment (as opposed to penalties proportional to the inverse probability of detection, as advocated by many economists). Moreover in settings where corruption is an issue there are limits to the size of the penalty f that a supervisor or a court can be trusted to levy without themselves becoming vulnerable to subversion (Glaeser and Shleifer, 2003). It is common in India for higher-ranked officials to intervene and protect lower-ranked officials from punishment for corruption. Empirically, the evidence we present in Section 4 implies that limited enforcement is the relevant case for interpreting our data.

Taking these constraints into account, the official establishes a menu of type-specific prices $p(y_i, \mathbf{x}_i, \eta_i) \geq 0$ for slots.[5] Prices should be interpreted broadly as including non-monetary transfers: for example, an official in a repeated relationship with a friend might give him a slot in anticipation of having this favor returned in the future. Note also that "pricing" slots is entirely consistent with rule-abidance, as the official could set the price equal to 0 for eligible households and $+\infty$ for ineligible ones. His problem can be formulated as

$$\max_{\{p_i\}} \int (1 - G(p_i|\eta_i))[p_i - c(y_i, \mathbf{x}_i)]dF(y_i, \mathbf{x}_i, \eta_i) \text{ such that } p_i \geq 0 \,\forall\, i \tag{2}$$

[5] We do not model the creation of additional bureaucratic rules by the official as a screening device (see Banerjee (1997)).

where $G(\cdot|\eta)$ is the exponential CDF with rate parameter $\frac{1}{\eta}$ and the implicit marginal cost $c(y_i, \mathbf{x}_i)$ of providing a slot is

$$c(y_i, \mathbf{x}_i) \equiv f[\pi(1, \mathbf{x}_i, R) - \pi(0, \mathbf{x}_i, R)] - \underline{\alpha}1(y_i = \underline{y}) - \overline{\alpha}1(y_i = \overline{y}) \tag{3}$$

This cost consists of two components. First, allocating a slot to household i may either increase or decrease expected penalties, depending on whether or not $\mathbf{x}_i \in R$. The eligibility rule thus influences the allocation of slots but does so indirectly, by influencing the distribution of equilibrium bribe prices. Second, implicit costs are lower to the extent that the official derives a benefit from allocating a slot to a household with income level y_i.

Pointwise maximization of (2) yields a familiar monopolist's markup equation

$$p^*(y_i, \mathbf{x}_i, \eta_i) = \max\{0, c(y_i, \mathbf{x}_i) + \eta_i\} \tag{4}$$

From this it follows directly that the probability household i obtains a slot is

$$\mathbf{P}(a_i = 1|\mathbf{x}_i, y_i, \eta_i) = 1 - G(\max\{0, c(y_i, \mathbf{x}_i) + \eta_i\}|\eta_i) \tag{5}$$

A number of inferences follow directly from these equations.

Prices and Allocations. Evidently prices are increasing in income y_i (conditional on η_i, and $\mathbf{x}_i \in R$,) if and only if the official has progressive preferences ($\underline{\alpha} > \overline{\alpha}$). Similarly prices are weakly decreasing in eligibility $1(\mathbf{x}_i \in R)$ (conditional on η_i, and y_i,) and strictly so if and only if penalties are positive ($f > 0$). Since household-level demand is decreasing in price, conditional on η_i, corresponding opposite results on quantities follow directly. The targeting rule R thus continues to influence the final allocation of slots but in an indirect manner by influencing the official's willingness to accept payment from each household. For example, giving a slot to an ineligible household is potentially costly and the official must obtain a larger bribe to be willing to do so. Importantly, how much more the official needs to get will depend on the details of enforcement, which are summarized by π. This means that two ineligible households may face different prices if one is "riskier" from the officials point of view. We will return to this issue below when discussing rule design. Note also that prices may be low for all households if the official has strong political incentives to give out benefits to everyone ($\underline{\alpha}, \overline{\alpha} \gg 0$). In this case rule-violations will be widespread but reported bribes low.

Enforcement and Rule Adherence. For f sufficiently large all eligible households receive slots (at price 0), and as $f \to \infty$ the number of ineligible households that receive slots approaches 0.[6] This reflects the fact that there are no deep informational constraints in the model: since any particular rule violation is punished with some positive probability, the principal could obtain arbitrarily close compliance if arbitrarily harsh punishments were available (Mirrlees, 1999; Becker, 1968).

2.2 The Principal

The principal has progressive preferences and values transfers to the poor more than transfers to the rich. Formally, he values a unit of surplus transferred to a poor (rich) household at $\underline{\omega}$ ($\overline{\omega}$) with $\underline{\omega} > \overline{\omega}$. The cost of providing a slot to either type of household is normalized to 1. We also fix $\eta_i = \eta$ for the next three subsections; we will re-introduce heterogeneous demand elasticities in our empirical application. We assume $\underline{\omega} > 1/\eta > \overline{\omega}$ so that the principal's expected return from giving a slot to a poor (rich) household is positive (negative). The principal's payoff as a function of the price schedule $\{p_i\}$ charged to households is

$$V(\{p_i\}) = \int_{y_i=\underline{y}} 1(v_i > p_i)(\underline{\omega}(v_i - p_i) - 1)dF(y_i, \mathbf{x}_i, \eta_i) \\ + \int_{y_i=\overline{y}} 1(v_i > p_i)(\overline{\omega}(v_i - p_i) - 1)dF(y_i, \mathbf{x}_i, \eta_i) \quad (6)$$

By exploiting properties of the exponential distribution we can write this as

$$V(\{p_i\}) = (\underline{\omega}\eta - 1)\int_{y_i=\underline{y}} \exp\left\{-\frac{p_i}{\eta}\right\} dF(y_i, \mathbf{x}_i) + (\overline{\omega}\eta - 1)\int_{y_i=\overline{y}} \exp\left\{-\frac{p_i}{\eta}\right\} dF(y_i, \mathbf{x}_i) \quad (7)$$

This can be interpreted as a loss function parametrized by the cost $\underline{\omega}\eta - 1 > 0$ of a "Type I" error (excluding a poor household) and the cost $1 - \overline{\omega}\eta > 0$ of a "Type II" error (including a rich one). Notice that the loss function depends both on how well targeted benefits are (the proportion that go to the poor) and also on the overall scale of benefit provision. As Ravallion (2009) emphasizes it is important in evaluating targeting performance to balance both of these considerations.

The existing literature has studied the case where the agent is completely honest, which implies $p_i = 0$ for all eligible households and $p_i = +\infty$ for all ineligible ones. In that case

[6] Note that we have a limit result in the latter case because of the simplifying assumption that the demand shocks v_i are unbounded. If the v_i were bounded above then there would be some finite f that completely eliminates inclusion errors.

the principal's problem is

$$\max_{R \in \mathcal{P}(\mathbf{X})} (\underline{\omega}\eta - 1) \int_{y_i = \underline{y}} 1(\mathbf{x}_i \in R) dF(y_i, \mathbf{x}_i) + (\overline{\omega}\eta - 1) \int_{y_i = \overline{y}} 1(\mathbf{x}_i \in R) dF(y_i, \mathbf{x}_i) \tag{8}$$

Here the analogy to statistical decision theory is exact. When the principal *cannot* rely on the agent to behave honestly, however, he must take into account the more complex reactions of the agent's optimal price schedule $\{p_i^*\}$ to the choice of targeting rule. The next two subsections study the choice of penalties f and the targeting rule R taking these concerns into account.

2.3 The Costs and Benefits of Delegation

One can verify from Equation 4 that the principal's payoff increases with $\underline{\alpha}$ and decreases with $\overline{\alpha}$ and approaches its maximum as these approach $+\infty$ and $-\infty$, respectively. Intuitively the principal wants an agent with strong distributive preferences to offset the profit motive. With a sufficiently good agent the problem of rule-design becomes irrelevant, since the official shares the principal's preferences and has better "soft" information about household poverty.

An important corollary is that the allocation of slots may deviate substantially from that prescribed by the statutory targeting rule and yet be *more* progressive than the statutory allocation. This will be true when officials have strong incentives to use their own "soft" information about poverty to improve targeting. The empirical implication is that observing rule violations is not sufficient to say whether a rule needs reform, an issue our empirical analysis below must take seriously.

Understanding official's redistributive preferences is also important for understanding the optimal level of enforcement; enforcing an imperfect targeting rule could in fact be counter-productive if the implementing official has strongly progressive preferences.

Proposition 1. *Let the probability of detecting a violation be constant ($\pi(a_i, \mathbf{x}_i, R) = \pi > 0$ whenever $a_i \neq 1(\mathbf{x}_i \in R)$).*

- *If R perfectly targets the poor then $\frac{\partial V}{\partial f} \geq 0$.*
- *If R does not perfectly target the poor, so that there are some ineligible poor and some eligible rich, then there exist a scalar f^* and functions $\underline{\alpha}^*(f)$ and $\overline{\alpha}^*(f)$ such that if $f > f^*$, $\underline{\alpha} > \underline{\alpha}^*(f)$, and $\overline{\alpha} < \overline{\alpha}^*(f)$ then $\frac{\partial V}{\partial f} < 0$.*

Proof. All proofs are in Appendix A □

The mechanics of this result are that as $\underline{\alpha}$ and f increase and $\overline{\alpha}$ decreases nearly all the eligible poor will have slots and almost none of the ineligible rich will have them. The marginal effects of an increase in f are then concentrated on the ineligible poor and the eligible rich, groups for whom they have negative effects (from the point of view of the principal).

This result has analogs in the theoretical literature on decentralization. For example, Bardhan and Mookherjee (2000) identify various factors that determine whether national or local governments are likely to behave more progressively, and thus help determine whether decentralization is likely to improve targeting. Here the choice is not who should make targeting decisions but rather how much discretion a given agent should have.

2.4 Rule Design: Is More Information Better?

If the expected penalties for rule-breaking can be made sufficiently large then any rule can be enforced and rule design is a statistical problem. When expected penalties are bounded, however, the problem is more complicated because some rules may be more enforceable than others. To demonstrate this tension we examine here the effects of conditioning a targeting rule on additional poverty indicators. In the standard approach to targeting the costs and benefits of additional indicators are well understood: "more information is generally better than less, though there are diminishing returns" (Grosh and Baker, 1995, ix), while "the beauty of using just a few indicators is that administrative costs are kept low" (Besley and Kanbur, 1990, 13). In contrast we show that with agency constraints more indicators may yield strictly worse targeting outcomes.

To operationalize the idea of conditioning on "more indicators" let the space of household types be a product space $\mathbf{X} = \Pi_{n=1}^{N} X_n$ of N different household characteristics. $N = 2$ characteristics will be enough for our purposes. Recycling notation, let F be the joint distribution of x_i^1 and x_i^2, F_1 and F_2 the marginal distributions, and F_{12} the distribution of the sum $x_i^1 + x_i^2$. We assume that $\rho(x) \equiv \mathbb{P}(x^1 + x^2 \leq k | x^1 = x)$ is strictly decreasing in x for any k; this is true if x^1 and x^2 are independently distributed, for example, but rules out very strong negative correlations. The principal considers as poor agents whose total assets $x_i^1 + x_i^2$ fall below some threshold y^*. One interpretation is that these are income-generating assets measured in productivity units; they could alternatively be consumer durables. We will call them "land" and "jewelry."

As before the rule must be implemented by an agent who is tempted to charge bribes. For this section $\underline{\alpha} = \overline{\alpha} = 0$ so that the agent cares only about maximizing profits. To parameterize enforcement, suppose that the principal observes the value of characteristic $j \in \{1, 2\}$ for household i with independent probability ϕ_j. (The analysis that follows extends readily to the case where these events are not perfectly dependent at the cost of notational clutter.) If the principal observes enough to determine that the household has been incorrectly classified then he fines the agent f. Notice that we abstract from decisions about administrative costs by assuming that information is exogenously either free (with probability ϕ_j) or infinitely costly (with probability $1-\phi_j$).[7]

The statistically optimal rule in this setting is

$$R_{12} \equiv \{\mathbf{x} : x^1 + x^2 \leq y^*\} \tag{9}$$

which achieves perfect targeting in the absence of agency concerns. If we constrain ourselves to the class of rules that condition only on x^1 then the natural candidate is

$$R_1 \equiv \{\mathbf{x} : x^1 \leq x^{1*}\} \tag{10}$$

for some threshold value x^{1*}. Intuitively, this rule makes eligible all households that are sufficiently likely to be poor given their land-holdings x^1.[8] Note that both R_{12} and R_1 are examples of "scoring" rules, meaning that they can be written as $R = \{x : \sum_{n=1}^{N} h_n(x^n) < t\}$ for some collection of functions $\{h_n\}$ and some threshold t. Scoring rules are widely used in practice; the BPL rule we study below is an example.

One can show that the optimal threshold x^{1*} to use in rule R_1 is the same regardless of how effective enforcement is:

Lemma 1. *Fix any $\phi_1 > 0$ and let x^{1*} satisfy $[\rho(x^{1*})\underline{\omega} + (1-\rho(x^{1*}))\overline{\omega}]\eta = 1$ (or $x^{1*} = 0$ if that equation has no solution). Then the rule R_1 defined by threshold x^{1*} is uniquely optimal within the class of rules that condition only on x^1.*

This condition is obvious for the perfect-enforcement case, since it states that the expected welfare gain from giving a slot to a marginal household should just equal the cost of the slot. The important point is that it continues to be optimal with imperfect enforcement.

Figure 1 plots the typespace $X_1 \times X_2$ broken up into regions defined by the two candidate targeting rules. Solid lines separate the households that are eligible and ineligible under the two rules; dotted lines separate households whose eligibility is the same but who face different equilibrium prices. To summarize the differences in pricing in each region we have plotted the quantity $\pi(1, \mathbf{x}_i, R) - \pi(0, \mathbf{x}_i, R)$ for rule R_{12} and then for rule R_1 in each region. This quantity is positive if a household is ineligible in the given region and under the given rule;

[7] If households that were illegally denied benefits could complain about this then the probabilities of detecting inclusion and exclusion errors would be asymmetric. The present model suits the empirical setting we will study below, in which most households are unsure about their own eligibility and do not know what to do if they feel they have been miscategorized.

[8] Land-based targeting has been evaluated by Ravallion (1989) and Ravallion and Sen (1994), for example.

Figure 1: Targeting on Two Asset Measures

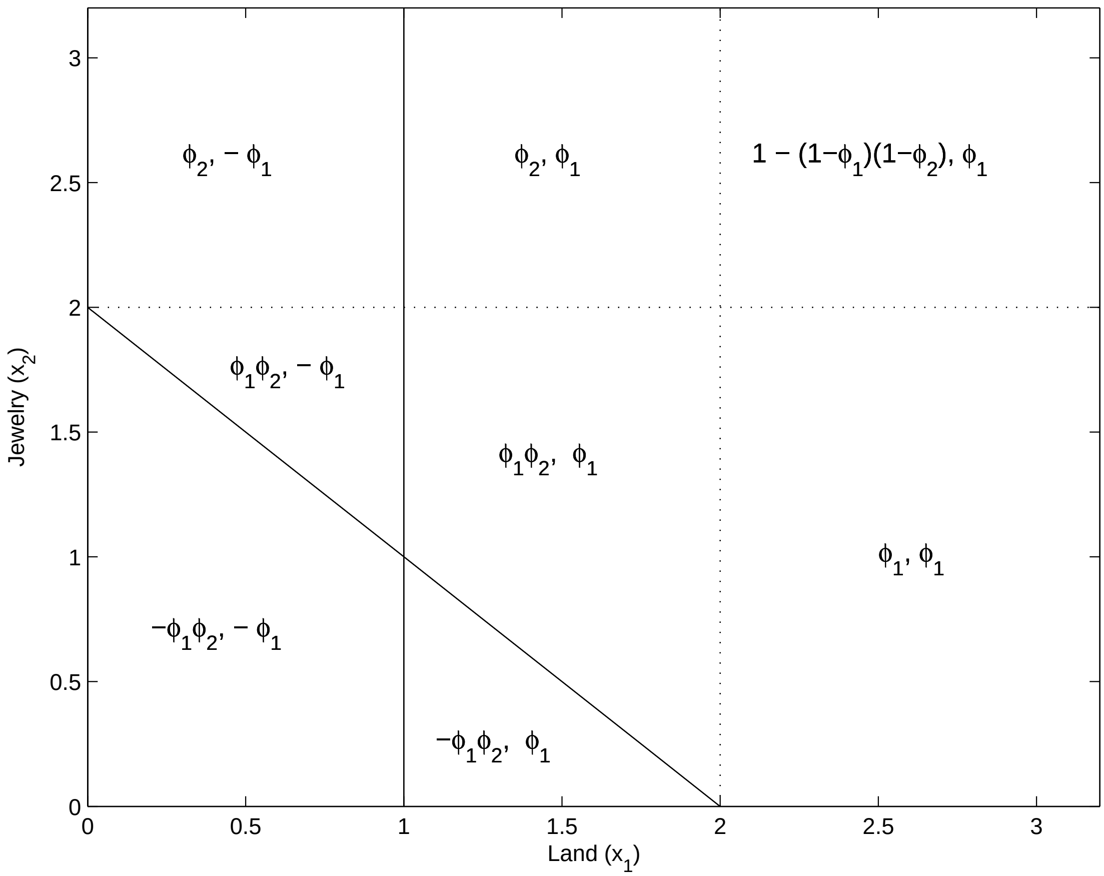

Displays $\pi(1, \mathbf{x}_i, R) - \pi(0, \mathbf{x}_i, R)$ first for rule R_{12} and then for rule R_1 by region of the type space.

in that case it is the probability of getting punished for allocating a slot to the household. It is negative if the household is eligible under the given rule; in this case it is minus the probability of getting punished for *not* allocating a slot to the household. The structure of π is determined by the audit technology defined above and thus these probabilities depend on ϕ_1 and ϕ_2.

The first point that emerges is that there are regions in which enforcement pushes prices in opposite directions, two in the upper-left and one in the lower-middle portion of the graph. These differences reflect the statistical inaccuracy of R_1: it defines as eligible some households that have low x^1 but such high values of x^2 that they are rich, and it makes ineligible some households with high x^1 but such low values of x^2 that they are in fact poor. Within these regions it is clear that R_{12} does better because it pushes prices in the right direction: lower for the poor, higher for the rich.

However, the two rules also generate different prices in the remaining regions (but one) in which they agree about eligibility. This is because of differences in the *enforceability* of the rules. For example, consider a household in lower-left corner of the type space which satisfies $x^1 < x^{1*}$ and $x^1 + x^2 < y^*$ and is thus eligible under both rules. Under R_1 the principal need only observe x^1 to verify that this household is eligible, and so the probability that the official will be punished for denying this household a slot is ϕ_1. Under R_{12}, however, this is insufficient since even if the principal observes $x^1 < x^{1*}$ the official can claim that x^2 was so large that the household is ineligible. A rule violation can only be conclusively proved in this case if the principal can verify holdings of both assets, and so the probability that the official is punished for not giving this household a slot is $\phi_1\phi_2 < \phi_1$. Intuitively, enforcement is more difficult under R_{12} because there are more criteria that need to be verified.

Similar logic applies to other regions of the typespace. When $x^{1*} < x^1 < y^*$ and $y^* - x^1 < x^2 < y^*$ rule R_{12} is harder to enforce; when $x^{1*} < x^1 < y^*$ and $y^* < x^2$ either rule may be harder depending on $\phi_1 \gtreqless \phi_2$; and when $y^* < x_1$ and $y^* < x^2$ rule R_{12} is easier to enforce since under that rule verifying either asset proves ineligibility. Across all of these regions, however, enforcement of rule R_{12} becomes relatively difficult as ϕ_2 shrinks. For ϕ_2 small enough the simpler rule R_1 will have an enforcement edge and there will be a non-trivial tradeoff between statistical accuracy and enforceability.

Figure 2 examines this issue numerically: it compares the payoffs generated by the two rules as the strength of enforcement f and the verifiability ϕ_2 of x^2 vary. When x^2 is relatively easy to verify the statistically superior rule R_{12} is also better overall. However, when x^2 is difficult to verify making the targeting rule contingent on x^2 not only does not improve targeting but actually worsens it. The enforcement difficulties created by using a more complicated rule more than offset the improvement in statistical targeting.

Figure 2 also illustrates how the tradeoff between the two rules shifts as enforcement f improves: the set of values of ϕ_2 for which R_{12} is preferable expands. Intuitively, high values of f compensate for low values of ϕ_2 as envisioned by Becker and Mirlees. We know that for any fixed rule R as $f \to \infty$ the principals payoff approaches his perfect-implementation payoff, and thus for any ϕ_2 there is an f^* large enough such that for $f > f^*$ R_{12} outperforms R_1. What makes R_1 attractive is precisely the fact that enforcement f is limited. More broadly, the need for statistically sub-optimal rules is a direct consequence of weak enforcement.

We find the comparison of R_{12} and R_1 compelling since R_{12} is statistically optimal. One can also show, however, that the conclusion extends to any rule that conditions non-trivially on x_2 in the sense that there exists a positive mass of households whose eligibility is sensitive to their holdings of asset 2, given their holdings of asset 1.

Proposition 2. *Fix any rule R that conditions non-trivially on x^2. If $\phi_1 > 0$ then there exists $\phi_2^*(R) > 0$ such that if $\phi_2 < \phi_2^*(R)$ then rule R_1 yields a strictly higher payoff for the*

Figure 2: Less Information can Yield Better Targeting

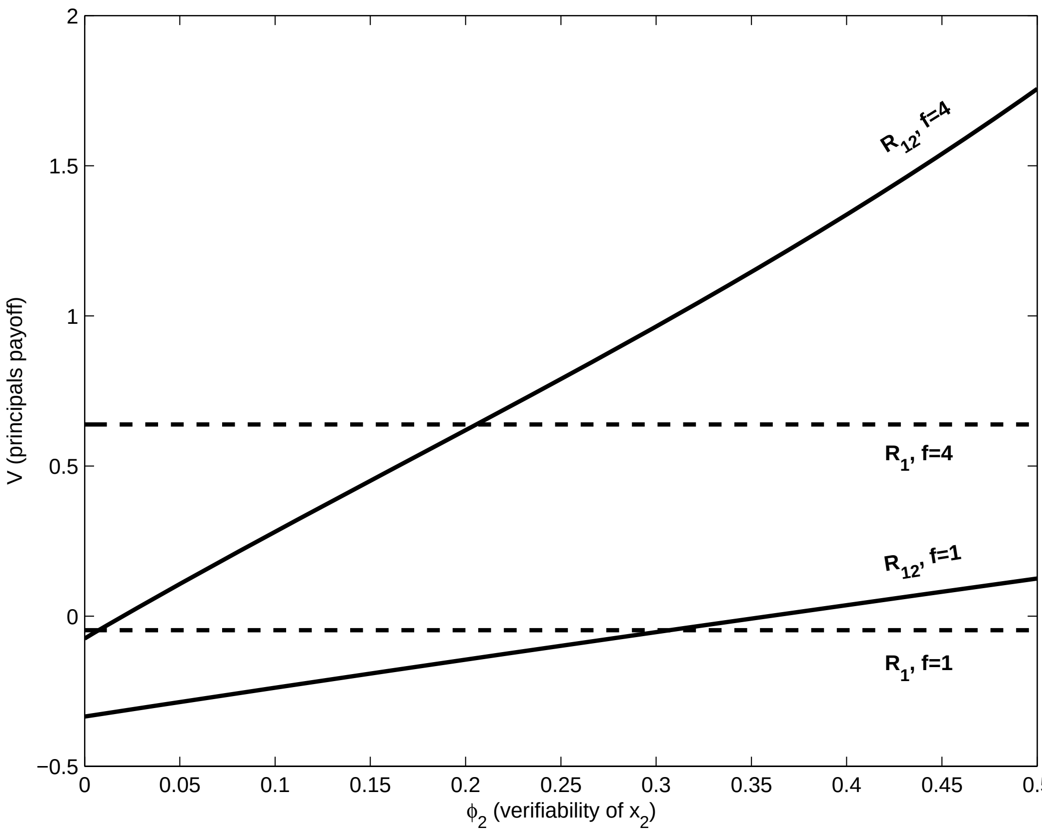

Plots the principals payoff for $\phi_2 \in [0, 0.5]$ and $f \in \{1, 4\}$ under rules R_{12} and R_1. Other parameters are fixed at $\phi_1 = 0.5$, $\eta = 1$, $\underline{\omega} = 2$, and $\overline{\omega} = 0.5$, and types are distributed uniformly on $[0, 3]^2$.

principal than R.

It is a corollary that within any finite set of alternative rules there exists a bound ϕ_2^* below which R_1 is optimal.[9]

The tradeoff captured here differs from those emphasized in the standard theory of incentive contracting. The issue is not the riskiness generated by attaching incentives to a noisy performance measure, as in the multi-tasking theory of Holmstrom and Milgrom (1991); here the agent is risk-neutral, and conditional on using rule R_{12} the principal would like to make incentives f as strong as possible. Nor is it an issue of misalignment between a contractible performance measure and the principal's objective per se (Baker, 1992); for any household type $\mathbf{x}_i$, the rule R_{12} generates incentives that point in the right direction. In these models the principal has an exogenously given performance measure and then optimally limits the strength of incentives placed on that measure. Here, in contrast, the key issue is which

[9] There does not exist a uniform bound ϕ_2^* below which R_1 performs better than *any* rule for the technical reason that one can construct infinite sequences of rules $\{R^t\}$ that approximate R_1 arbitrarily.

performance measure to choose given that the strength of incentives is constrained (f is bounded). The result is that the principal may choose a measure that is relatively sensitive to actions, thus increasing the effective strength of incentives, even if this measure is less well correlated with his own objective function.

Another difference between our analysis and standard principal-agent models is that we constrain the principal to offering two payoff levels, one if he catches the agent breaking a rule and one when he does not. We do so primarily because this seems most appropriate for the institutional contexts we aim to describe. We also note, however, that allowing penalty levels to vary with the nature of the rule violation would not qualitatively alter the results above; if anything, this makes R_1 more attractive relative to R_{12}.[10][11]

3 Empirical Context and Data Collection

We have seen that under certain conditions it may be better to choose a targeting rule that is less accurate statistically but more readily enforceable. Necessary conditions include that enforcement be weak and that the enforcement technology be such that officials perceive some rule violations as "safer" than others. Moreover, the implementing officials should not have strong progressive preferences themselves – if they did it would be better to simply delegate the targeting problem to them. In the following sections we test these conditions in the context of India's Below Poverty Line cards. This section describes the institutional setting and the data we will analyze in Section 4.

3.1 Targeting India's Poor

The origins of the current BPL system aptly illustrate the importance of targeting. Prior to 1997 India operated a universal Public Distribution System (PDS) intended to provide basic commodities to all Indian households at subsidized prices. To accomplish this the government created a vast system of procurement and distribution. The Food Council of India (FCI) purchased grain from farmers and stored it at government-owned warehouses; subsequently, these commodities were allocated to each state based on prior years' consumption levels and

[10]To see this, suppose the principal can define variable fines $f(a_i, \mathbf{x}_i, R)$ that depend on the nature of the rule violation. For a rule like R_{12} that perfectly targets the poor it will always be optimal to enforce as aggressively as possible, i.e. set fines at the upper bound for any violation. Reinterpreting the fixed f in this example as the upper bound, we can interpret the calculated performance of R_{12} as the best it can ever do, while the calculated performance of R_1 is a *lower* bound on how well it can do after possibly re-optimizing fines.

[11]Note also that nothing above contradicts the revelation principle, which states that the set of optimal mechanisms includes one in which the agent faithfully reports all his information but does not imply that transfers in an optimal mechanism are sensitive to every facet of this report.

distributed through a system of about 400,000 Fair Price Shops (FPS), each one servicing several villages. At the FPS, households purchased rice, wheat, sugar, and kerosene at uniform prices below those on the open market.

In 1997 the government judged the PDS too costly to support and introduced poverty targeting. Under the Targeted Public Distribution System, all households in India are classified as being below-poverty-line (BPL) or not. Each BPL household is entitled to defined quantities of basic commodities at subsidized prices set equal to about half of what it costs the government to purchase and distribute them. In contrast, above-poverty-line (APL) households pay prices approximately equal to the government cost, which are also very close to market prices. The Indian Planning Commission estimated that in 2001 the effective annual subsidy to BPL card holders in Karnataka from grain purchases alone was Rs. 294 (Programme Evaluation Organization, 2005). Many other social programs are also now targeted to BPL households – for example, the cards give access to advantageous loans for agricultural activities, education scholarships, medical benefits, housing schemes, and distributions of bicycles, books, clothes, soap, salt, oil, and tea. Dreze and Khera (2010) estimate that 33%-34% of Indian households held BPL cards as of 2005.

Identifying BPL households has thus become a critical task for welfare policy in India. The central government conducts surveys approximately every five years to identify the number of households it thinks are BPL in each state and then allocates funding for social programs in proportion to these numbers. The states can use their own criteria to actually allocate BPL cards, however (and generally each state estimates its own poverty count to be much higher than the central government's figure). Our empirical work is set in Karnataka, where the most recent round of BPL surveys was held in 2007. According to the state's rules a household was eligible for a BPL card if it did *not* have any of the following:

- Annual income more than Rs. 17,000 in urban areas or Rs. 12,000 in rural areas;
- A telephone (land line or mobile);
- A two-, three- or four-wheeler (e.g. motorcycle, auto-rickshaw, or car);
- A gas connection;
- A color TV;
- More than 5 acres of dry land;
- A water pump set;
- A household member who is a salaried government employee.

These criteria define the targeting rule (R). Note that this rule is a special case of a widely used class of *scoring* rules, which sum up functions of a range of indicators and define households as eligible if their score falls below a threshold. In this case households must receive a score of zero to be eligible. It is also interesting that while some of the eligibility criteria seem plausibly verifiable (e.g. land-holdings or status as a government employees) but others, particularly the income threshold, would be very hard to prove. The analysis in Section 2.4 suggests that this may adversely affect the enforceability of the rule and thus its overall performance.

The actual process of allocating BPL cards begins with a state-mandated survey to determine which households are eligible for BPL cards. Surveys are conducted by government officials at the level of the Gram Panchayat (GP), a collection of several villages. The official in charge was usually the village accountant, but may also in some cases have been the GP Secretary, Anganwadi (health) worker, or a local school teacher. Regardless of the exact identify of the official, he or she would typically have both "hard" and "soft" information about the poverty and other characteristics of households in the Panchayat.[12]

The legally mandated process for ascertaining BPL eligibility involved several additional steps. After the initial government survey was completed, the BPL eligibility list for each village was compiled at the taluk (sub-district) level. The compiled lists were then remitted to the corresponding GP's to verify that households disclosed asset ownership and wealth truthfully to the initial government survey team. GP officials were supposed to organize a meeting of all citizens registered on the GP's electoral rolls, called a Gram Sabha. At the Gram Sabha they should have read aloud the eligibility status of each household, giving the community a chance to disagree, and resolved all disputes on the spot. Finally, the revised list should have been posted at a well-known place in each village for several days before being finalized and remitted to the taluk, which then proceeded to issue BPL cards. In most of Karnataka temporary ration cards were issued in 2007 and households were in principal allowed an additional opportunity to appeal their eligibility status in the period before permanent ration cards were issued in 2008. Finally, in addition to these "grassroots" checks the state government sent out teams to re-survey a sample of households and check that the targeting rules were correctly implemented. Of course, how effective any of these mechanisms are in practice is an empirical question.

3.2 Cross-Checking BPL Allocations

Official policy in Karnataka specified both who should receive BPL cards and also the process by which this should be determined. Since BPL cards are valuable, however, officials had incentives to break both process and targeting rules. To understand how the BPL

[12]Note that de jure households do not need to apply for a BPL card. In programs with an application requirement targeting depends both on administrative decisions and on household's self-selection into the applicant pool (Coady and Parker, 2009; Baird et al., 2009).

targeting criteria work *in practice*, then, we need independently collected data on household characteristics and BPL status.

We collected such data as part of an independent quality of life survey in Karnataka in early 2008. We constructed our sample in two stages. First we selected villages; in most districts we drew a proportional random sample of villages, while in Raichur we sampled from among villages that had been part of an earlier experiment.[13] We then randomly selected 21 households from each village, sampling from the state governments list of all households that had been identified in the BPL survey. Our surveyors were not always able to complete interviews with all of the 21 assigned households, either because the household had migrated or because no one was at home during the day; in these cases we randomly selected replacement households for them to interview. In the event that both the originally sampled household and the backup could not be interviewed, fewer than 21 households were surveyed. In total we surveyed 14,074 households, or an average of 17 households per village.

The primary objective of the survey was to obtain independent measures of both BPL eligibility and BPL card ownership in order to measure the extent of misclassification. We structured the survey instrument carefully to encourage veracity. Questions about the BPL eligibility criteria were posed early in the survey along with other similar quality of life questions, and the surveyors did not refer to them as eligibility criteria. Questions about card ownership and other politically sensitive questions were located at the end of the survey to avoid influencing responses to the questions about eligibility criteria. In addition to regular BPL cards there are variants (Antyodaya Anna Yojane and Annapurna cards) for households that are not only below the poverty line but also disadvantaged in other ways, e.g. widowed or elderly. We collected data on each type of card but treat them symmetrically in the following analysis.

In addition to these core data on BPL eligibility and card ownership, we also collected information on the process through which cards were allocated and on respondent's understanding of the allocation rules. We were particularly interested in understanding the prices households paid for BPL cards. The state of Karnataka has fixed the fee for issuing a BPL card at Rs. 5, but given the discretionary scope local officials have we expected to see higher prices charged in practice, as in the model of Section 2. We therefore asked households about both the "official fee" necessary to obtain a card and also about any "extra fee" they were charged. Responses to the later question may need to be interpreted with some care, but

[13]The experiment involved providing a random sub-sample of villages with information about the BPL eligibility criteria. Sadly this treatment was found to have no effect on any measured outcome. We include village fixed effects in all our regression specifications, so any unnoticed effects of the experiment should not influence our results. Results are all qualitatively similar if we simply exclude the experimental villages (22% of sampled villages).

Table 1: Basic Household Characteristics

Variable	Percent	Observations
Religion		13717
Other	5%	
Hindu	95%	
Caste		13601
Scheduled Caste	26%	
Scheduled Tribe	14%	
General	60%	
HH Head Marital Status		13361
Married	81%	
Never Married	1%	
Widowed	18%	
Divorced	0%	
HH Head Education		13357
Illiterate	61%	
Less than Primary	5%	
Primary	10%	
Middle	13%	
Matriculate	7%	
Intermediate	2%	
BA/BSc	1%	
MA/MSc	0%	
Professional Degree	0%	
HH Head Gender		13381
Male	83%	
Female	17%	

as respondents' anonymity was assured they had no reason to be concerned about faithfully reporting the prices they faced.

Table 1 presents basic descriptive information about the households in our survey. The majority are illiterate (61%) and Hindu (95%), and a large minority come from a scheduled caste or tribe (40%).

4 BPL Targeting in Practice

This section presents evidence on the BPL allocation in Karnataka. We first document basic facts about the process used and the resulting allocation. We then turn to multivariate analysis to better understand the drivers of the BPL allocation, followed by our bottom-line estimates of the progressivity of the de facto allocation relative to the de jure one. In Section

Table 2: Households Are Unfamiliar With Eligibility Rules

Is a family eligible if it has...	Correct Answer	Percent Correct
Water pump	No	57%
Jewelry worth more than Rs. 8,000	Yes	38%
Motorized vehicle	No	41%
Bicycle	Yes	73%
Member with a monthly salary more than Rs. 1,000	No	41%
Electricity	Yes	77%
More than 3 hectares of dry land	No	35%
Black & White TV	Yes	65%
Telephone	No	19%
Are you...		Percent Agreeing
Familiar with the eligibility criteria?		35%
Aware how to disagree with your categorization?		17%

The first ten rows report the percentage of households that correctly identified whether or not the given condition makes a household legally ineligible for a BPL card. The latter rows report the percentage of households that agreed with the given statements.

4.3 we discuss whether the results are most consistent with corrupt behavior on the part of officials or fraudulent behavior by households.

4.1 How are BPL Cards Allocated?

Officials routinely circumvent the statutory procedures for BPL card allocation. Only 50% of respondents in our survey remembered being surveyed by someone to determine eligibility. Given the importance and salience of the BPL card it is implausible that this low figure is due to mere forgetfulness. Only 13% of respondents were aware of a Gram Sabha meeting held to discuss BPL eligibility; 25% remembered at least one Gram Sabha meeting held in the last two years but said it did not cover BPL eligibility, and the remaining 62% did not recall any Gram Sabha meeting having been held in the past two years. Moreover, conditional on being among the 13% of respondents who did recall a Gram Sabha held to discuss BPL eligibility, only 16% said that families had an opportunity to object to their eligibility status at this meeting. Finally, only 2% of respondents said that a list of eligibility assignments was posted somewhere in the village.

We also asked respondents about their familiarity with the eligibility criteria. The top panel of Table 2 gives the percentages of respondents that correctly answered the question "Is a family eligible for a BPL card if it has X" for various criteria. Not all of the criteria are

Table 3: Official Rules are Frequently Violated

	Ineligible	Eligible	Total
No BPL Card	2560	652	3212
	(30%)	(13%)	(24%)
Has BPL Card	5862	4419	10281
	(70%)	(87%)	(76%)
Total	8422	5071	13493

Column percentages in parenthesis, e.g. 70% of ineligible households have BPL cards.

actual eligibility criteria; several are placebos. Accuracy rates vary from 19% to 77% and the (unweighted) average accuracy rate across the criteria is 50%, exactly what respondents would have achieved by random guessing. Overall only 35% of respondents described themselves as familiar with the eligibility rules, and only 17% reported that they knew what to do if they disagreed with the way they were categorized. Overall, awareness of the eligibility criteria is low.

Procedural violations and lack of awareness suggest that the targeting rule is not being implemented as designed. To verify this, Table 3 cross-tabulates our measure of BPL eligibility and actual BPL card possession. Rule-breaking is rampant. 13% of households legally eligible for a BPL card do not have one, and fully 70% of households ineligible for a card have one nevertheless. In total we estimate that 48% of the household in our sample are misclassified.[14] Clearly, enforcement is weak.

4.2 What Drives the BPL Allocation?

Although the eligibility rule is often violated, the BPL allocation could in principle still reflect the preferences of the rule designers if local officials share those preferences and break the rules because they can better target the poor using their soft information. For example, an initially well-off household might have acquired a water pump but subsequently fallen on hard times and now be among the poorest in the village. If this were the case then we would expect to see that rule violations lead to a *more* progressive allocation of BPL cards.

[14] These figures are consistent with the findings of several smaller-scale studies of BPL allocations in other states. For six villages in Gujarat following the 1997 BPL Census, Hirway (2003) estimates that 10%-15% of eligible households did not receive cards, while 25%-35% of ineligible households received cards. For eight villages in Rajasthan following the 2002 BPL Census, Khera (2008) estimates that 44% of eligible households did not receive BPL cards while 23% of ineligible households did receive them. Ram et al. (2009) report high rates of violations of individual eligibility criteria in the India-wide National Family and Health Survey-3. In a distinct setting, Camacho and Conover (2009) document suspicious patterns in official records for Columbia's poverty-targeting scheme SISBEN which are consistent with manipulation, but do not directly measure rule violations.

Figure 3: Eligibility is More Progressive than BPL Status

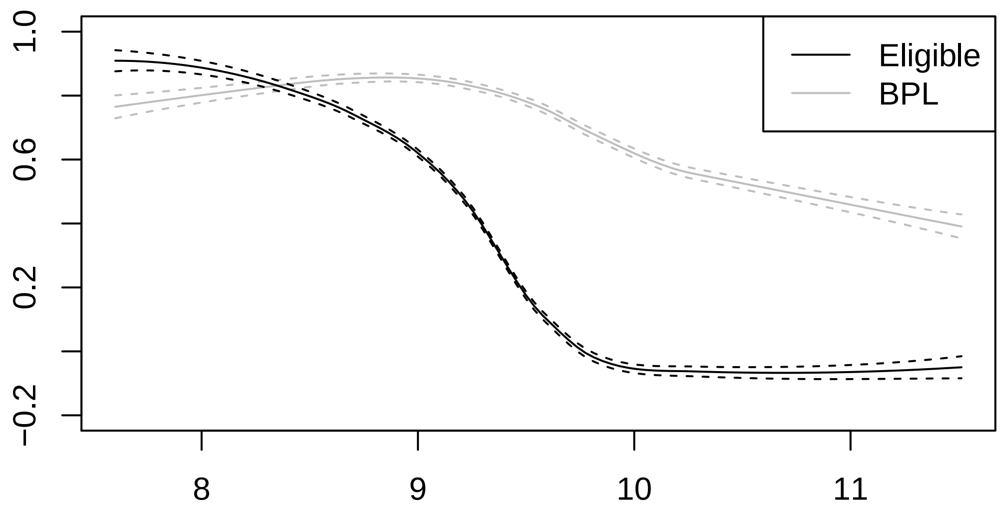

Plots nonparametric regressions of an indicator for statutory eligibility (Eligible) or actual BPL status (BPL) against log income. The domain of the plot is the 1%-trimmed sample income distribution. 95% confidence intervals are indicated by dashed lines.

The eligibility rule itself does a credible job of targeting the poor, in the traditional statistical sense. The raw correlation between eligibility (as constructed from our data on households' economic status) and log income is a healthy -0.55 and the within-village correlation is nearly as strong at -0.52. One might worry that even an honest official could not implement targeting on hard-to-observe criteria like the Rs. 12,000 income threshold, but after dropping this criterion the correlation is still a healthy -0.42. In contrast, actual BPL status is correlated -0.23 with log income. Figure 3 provides an analogous non-parametric comparison; it shows that the poorest households are slightly more likely to be eligible than to have BPL cards, while for richer households this relationship is reversed.

Strikingly, even drastically simplified rules that drop all but *one* of the eligibility criteria we still outperform the actual allocation: log income is correlated -0.37 with owning a phone, -0.32 with owning a water pump, -0.31 with owning more than 5 acres of land, and -0.30 with having a gas connection. As these are highly observable characteristics (in the case of land ownership and gas connections there exist independent records that could be used for cross-checking) it seems implausible that the BPL rule performs poorly soley because officials are unable to implement it. These patterns are inconsistent with the view that local officials have strong redistributive preferences and can be trusted to allocate welfare benefits unsupervised.

To further examine the determinants of the BPL card allocation we turn next to data on the prices households faced for cards. By law, officials were allowed to charge up to Rs. 5 to cover administrative costs. A large proportion of households in our survey – 73% of all households and 93% of BPL card recipients – reported the price they actually paid for their card. We define the total price as the sum of the reported "official fee" and any "extra" fee reported (7% of households reported an "extra" fee). Among those who reported a price, 75% reported one above the statutory maximum fee of Rs. 5 (0.2% reported one below Rs. 5) with a maximum bribe of Rs. 305. The mean bribe is small, however, at Rs. 9, or Rs. 14 conditional on being positive. These figures are small both in absolute terms and relative to our best estimates of the benefits of holding a BPL card. Using self-reported data on commodities purchased at the Fair Price Shop, prices paid, and corresponding market prices, we estimate that the mean BPL household in our sample receives an implied subsidy of Rs. 201 per month while the mean APL household receives Rs. 67, which puts the implied value of a BPL card at Rs. 133 per month.[15] This suggests non-monetary factors play an important role in the allocation of BPL cards – officials may trade benefits for votes, for example (which would appear as large $\underline{\alpha}$ and $\overline{\alpha}$ in our model). It also suggests that we should place more weight on inferences from the distribution of BPL cards than inferences from the distribution of prices.

To refine our understanding of mis-targeting we turn next to a multivariate analysis. Let h index households and v index villages. The key pricing equation of our model, Equation 4, can then be written

$$p_{hv} = f[\pi(1, \mathbf{x}_{hv}, R) - \pi(0, \mathbf{x}_{hv}, R)] + (\underline{\alpha} - \overline{\alpha})1(y_{hv} = \overline{y}) - \underline{\alpha} + \eta_{hv} \tag{11}$$

whenever prices are positive for household h in village v.[16] As we observe a continuous measure of household income y_{hv} we replace $1(y_{hv} = \overline{y})$ with $\log y_{hv}$. We know little about the enforcement term $[\pi(1, \mathbf{x}_{hv}, R) - \pi(0, \mathbf{x}_{hv}, R)]$ except that it should be negative for eligible households and positive for ineligible ones; we will therefore experiment with functions $h(\cdot)$ of the eligibility criteria including a simple eligibility indicator and the number of criteria violated.[17] We also augment the model with village fixed-effects λ_v, which absorb institutional variation across villages and isolate variation in decision-making by the same officials

[15] Inquality between the bribe price and the value of a good is a common feature of illicit markets and sometimes called the "Tullock paradox." See Bardhan (1997).

[16] We report OLS estimates that ignore the fact that self-reported prices are left-censored at Rs. 5, the true statutory fee. Tobit estimators that account for this suffer from an incidental parameters problem and are only consistent as the number of observations per village grows. Nevertheless we did estimate Tobit models and obtained estimated coefficients similar to and slightly larger than those reported below.

[17] We also estimated a variety of models in which specific violations and combinations of violations were allowed to have distinct effects. The conclusions we report below were robust to these variations (available on request).

within villages.[18] This yields

$$p_{hv} = fh(\mathbf{x}_{hv}) + (\underline{\alpha} - \overline{\alpha}) \log y_{hv} + \lambda_v + \eta_{hv} \tag{12}$$

Intuitively this equation says that bribe prices should be driven by eligibility if the official perceives rule-breaking as costly ($f > 0$) and by income if the official has redistributive preferences ($\underline{\alpha} > \overline{\alpha}$). Analogous opposite results follow for the probability that household hv holds a BPL card. This approach can be seen as a micro-founded analogue to the reduced-form specifications of Alderman (2002) who regresses welfare receipts on household expenditure while conditioning on a set of more readily observable attributes that could in principle have been included in a PMT.

The main outstanding concern with this approach is that there may be variation in households' willingness to pay η_{hv} that is observed by the official but not by us. With cross-sectional data we do not have plausible instruments for $\mathbf{x}_{hv}$ or y_{hv}. We can, however, implement a placebo test. If the eligibility criteria are predicting prices because they are correlated with an unobservable such as η_{hv} then we should find that ownership of other similar assets which are *not* eligibility criteria should predict prices in the same way. Our survey collected data on three such criteria: whether the household had electricity, a black and white television, and a bicycle. We will include these separately from the true eligibility criteria and see whether the estimated effects match.

Table 4 presents estimates of Equation 12 and analogous linear probability models for BPL card ownership. Panel A focuses on reported prices; Panels B and C focuses on BPL status, with Panel B restricting the estimation sample to households that reported prices for comparability with Panel A.

Columns 1 and 2 examine the role of ineligibility. Consistent with the model, ineligible households pay significantly higher prices for BPL cards. The point estimate is small, however: ineligible households pay Rs. 3 more on average and Rs. 1.3 per violated criterion, which suggests that while officials are cognizant of the costs of breaking the rules they perceive these as being small. Effects on quantities are similarly small, with ineligible households 1% less likely to hold BPL cards. Estimates for the full sample are larger, with ineligible households 21% less likely to hold cards, but still much smaller than the 100% difference that would obtain under perfect enforcement.

In Column 3 we test the idea that *degrees* of ineligibility matter. If officials perceive all

[18] An additional motivation for including village fixed effects is that, anecdotally, some villages in Karnataka were told that they should not allocate BPL cards to more than 60% of households. One can show that if the official faces a binding quantity constraint then optimal pricing is the same as in Equation 4 except that prices are augmented by the Lagrange multiplier λ on the quantity constraint, which would vary by village.

Table 4: Eligibility, Income, Prices and Allocations

Regressor	1	2	3	4	5	6	7
Panel A: Prices							
Ineligible	2.935		0.868	2.002			
	(0.509)***		(0.668)	(0.477)***			
# Violations		1.277	1.059		1.029	1.281	1.009
		(0.253)***	(0.343)***		(0.298)***	(0.255)***	(0.296)***
# Placebo Violations						-0.222	-0.35
						(0.343)	(0.366)
Log Annual Income				1.564	0.946		1.05
				(0.507)***	(0.577)		(0.611)*
N	9608	9608	9608	9608	9608	9608	9608
R^2	0.007	0.011	0.011	0.01	0.012	0.011	0.012
Panel B: Quantities							
Ineligible	-0.013		0.003	-0.006			
	(0.003)***		(0.005)	(0.003)*			
# Violations		-0.008	-0.009		-0.006	-0.008	-0.006
		(0.002)***	(0.002)***		(0.002)***	(0.002)***	(0.002)***
# Placebo Violations						-0.005	-0.004
						(0.003)*	(0.003)
Log Annual Income				-0.012	-0.006		-0.005
				(0.005)***	(0.005)		(0.005)
N	9608	9608	9608	9608	9608	9608	9608
R^2	0.002	0.006	0.006	0.004	0.006	0.006	0.006
Panel C: Quantities							
Ineligible	-0.215		0.008	-0.107			
	(0.01)***		(0.012)	(0.01)***			
# Violations		-0.097	-0.099		-0.079	-0.096	-0.081
		(0.003)***	(0.004)***		(0.004)***	(0.003)***	(0.004)***
# Placebo Violations						-0.038	-0.03
						(0.005)***	(0.006)***
Log Annual Income				-0.146	-0.062		-0.054
				(0.009)***	(0.009)***		(0.009)***
N	13183	13183	13183	13183	13183	13183	13183
R^2	0.065	0.145	0.145	0.109	0.151	0.15	0.154

Notes: (1) The unit of observation in all regression is a household. The outcome is the price the household faced for a BPL card in Panel A and an indicator equal to one if the household obtained a BPL card in Panels B and C. The estimation sample includes all households that reported BPL card prices in panels A and B, and all households in Panel C. (2) "Ineligible" is an indicator equal to one if the household violates any eligibility criteria; "violations" is the number of criteria it violates; "placebo violations" is the number of assets the household holds that do not disqualify it. All specifications include village fixed effects. (3) Robust standard errors clustered at the village level are presented in parenthesis. Statistical significance is denoted as: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 4: Prices and Allocations are Monotone in Violations

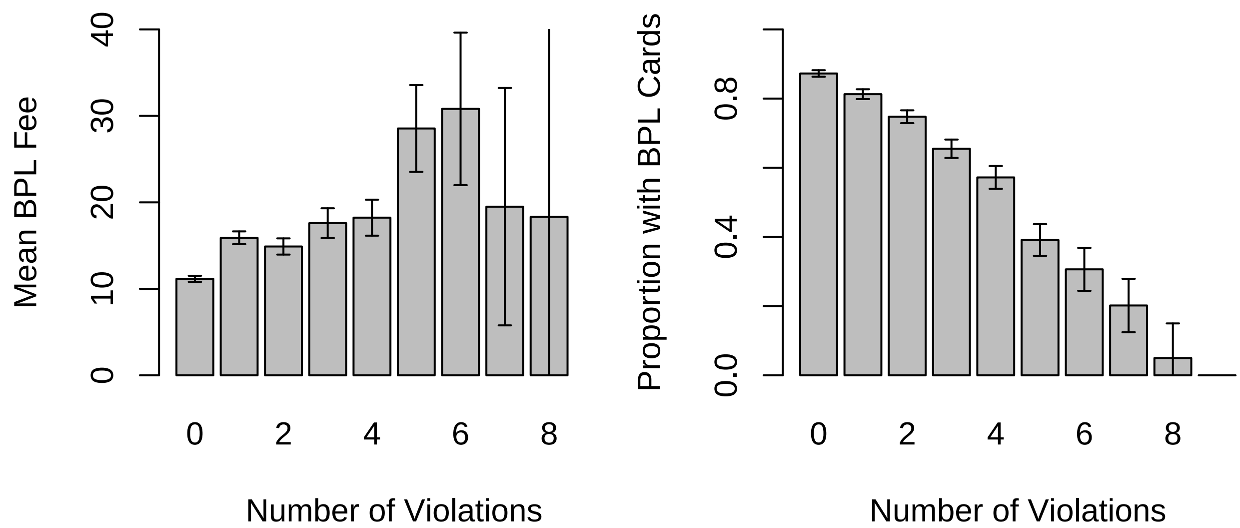

rule violations as being equally risky then we should find that the number of violations is unimportant once we control for whether or not this number exceeds one. In contrast, we find that moving from 0 to 1 violation has roughly the same effect as moving from 1 to 2, from 2 to 3, and so on. Figure 4 provides a graphical analogue to these regressions (unconditional on village effects). It shows that prices steadily increase and the probability of holding a BPL card steadily decrease as the number of eligibility criteria a household violates increases. This is crucial as in implies that it is not simply whether a household is ineligible but how ineligible it is that matters, which in turn implies that changes to the eligibility rule will affect inframarginal as well as marginal households.

Column 4 and 5 mimic 1 and 2 but also include the logarithm of annual household income as a predictor to test the joint hypothesis that officials have "soft" information about household poverty and use this to target BPL cards. The results are generally mixed. Among households who reported prices, higher income is associated with higher prices and a lower probability of holding a BPL card, but these results are insignificant once we control for the number of criteria violated (and not just an eligibility indicator). Only in the full sample does income consistently negatively predict BPL status. Even here the estimated effect is small: doubling log income has less of an effect than increasing by one the number of violated eligibility criteria. Thus while there is some evidence for soft targeting it appears insufficient to generate a progressive final allocation. We will return to this issue below.[19]

[19]Since income and ineligibility are positively correlated with each other and have similar effects on bribe prices, one interpretation concern is that an incorrect choice of functional form for one could generate

Table 5: Are Households Deceiving Officials?

Regressor	I	II
# Violations	-0.104	-0.098
	(0.004)***	(0.006)***
Visited by Official	0.052	
	(0.012)***	
Visited * Violations	0.012	
	(0.005)**	
# Violations Known		0.007
		(0.004)*
Violations Known * Violations		0.000
		(0.002)
N	13173	13145
R^2	0.152	0.147

Notes

1. The unit of observation in all regression is a household. The outcome is an indicator equal to one if the household obtained a BPL card.
2. Robust standard errors clustered at the village level are presented in parenthesis. Statistical significance is denoted as: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Columns 6 and 7 include the number of placebo criteria the household violates as a predictor. In contrast to real violations, placebo violations negatively predict prices and this effect is not statistically significant. This supports the view that rule violations matter because they raise the official's perceived cost of allocating a BPL card to the household. For the coefficient on true violations to be driven by some omitted variable, that variable would have to be correlated with prices and with true violations but uncorrelated with placebo violations. In the sample of households who reported prices placebo violations negatively predict BPL card ownership, but this is never more than marginally significant. Only in the full sample do placebo violations significantly predict allocations, and the magnitude of this relationship is about 1/3 that of the coefficient on true violations. To this extent that this does imply an upwards omitted variables bias in the estimated coefficient on true violations it further reinforces the point that enforcement is weak.

4.3 Do Rule Violations Reflect Corruption or Fraud?

To this point we have interpreted violations of Karnataka's BPL targeting rule as evidence of corruption. An alternative interpretation is that officials behave honestly but are either unable or unwilling to perfectly ascertain which households are legally eligible and which are not, allowing some ineligible households to fraudulently obtain cards.[20]

Some of the evidence we have already presented is inconsistent with this view. First, fraudulent misrepresentations by households should not lead to the exclusion of eligible households. Second, we have seen that most households must pay a bribe to obtain a BPL card and that bribe prices are systematically correlated *within villages* with household's eligibility status, consist with rent-seeking by officials.

To construct an additional test we use our data on which households were visited by a government official to ascertain their status. If the issue is that officials are not catching rule violations because they are not conducting proper inspections then we should see that (1) households that were visited are less likely to receive BPL cards, and (2) this effect is stronger for households that violate more rules. Column I of Table 5 shows that the opposite is true in both cases: households that were visited are *more* likely to obtain a BPL card, and this effect is stronger for less eligible households. This suggests that visits are less about inspection than about negotiation.

As a second test, we examine whether eligibility violations matter less for households that are better-informed about the eligibility criteria. The idea behind this test is that if households are fraudulently concealing characteristics that make then ineligible, then households that are better-informed should be able to do this more effectively. Column II of Table 5 shows that this is not the case. *Eligible* households that are better-informed about the rules, in the sense that they correctly identify more of the actual exclusion restrictions as such, are slightly more likely to obtain cards. This effect is no stronger, however, for ineligible households. Knowledge of the rules thus does not appear to be especially useful in allowing ineligible households to obtain cards.

---

a spurious result for the other. We experimented with a full set of non-parametric indicators for every possible combination of rule violations and obtained essentially identical results for income. Similarly, we experimented with higher-order polynomials in log income and obtained essentially identical results for violations. Functional form does not appear to be an issue.

[20]For example, Martinelli and Parker (2009) show that households' self-reported eligibility for Progressa/Opportunidades differs from eligibility as assessed by officials in follow-up visits. It is important to note that the BPL allocation process differs in that it does not include an initial self-report.

5 Conclusion

Accurately targeting resource transfers to the poor is one of the most pressing problems in international development. The predominant approach to targeting is to perform a statistical analysis of data from household surveys to define a proxy means test that is as tightly correlated with poverty as possible. This approach may fail to achieve the desired results, however, when the implementation of the targeting rule is delegated to corruptible agents. We study the problem of designing targeting rules subject to this agency constraint. Our main theoretical finding is that conditioning a targeting rule on an additional household characteristic, though it always improves statistical performance, may strictly reduce the principal's payoff because of novel effects on the *enforceability* of the rule.

Turning to data on the performance of a key proxy means test in India, we find overwhelming evidence of weak enforcement. Rule-breaking is widespread and the ultimate allocation of benefits is less progressive than it would have been had the rules been faithfully implemented. Targeting rules do appear to influence the bribe prices that officials charge to households, consistent with the existence of some enforcement, but the effects are small, consistent with enforcement being weak. We infer that this is an environment in which it may be important to design targeting rules that are relatively easy to enforce. Interestingly, Dreze and Khera (2010) have proposed simplifying targeting policy in India for exactly this reason.

References

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin Olken, and Julia Tobias, "Targeting the Poor: Evidence from a Field Experiment in Indonesia," Technical Report, MIT May 2010.

Alderman, Harold, "Do local officials know something we don't? Decentralization of targeted transfers in Albania," *Journal of Public Economics*, March 2002, *83* (3), 375–404.

Baird, Sarah, Craig McIntosh, and Berk Özler, "The Squeaky Wheels Get the Grease: Applications and Targeting in Tanzanias Social Action Fund," Technical Report, UC San Diego October 2009.

Baker, George P, "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, June 1992, *100* (3), 598–614.

Banerjee, Abhijit, "A Theory of Misgovernance," *The Quarterly Journal of Economics*, November 1997, *112* (4), 1289–1332.

_ , Esther Duflo, Raghabendra Chattopadhyay, and Jeremy Shapiro, "Targeting Efficiency: How well can we identify the poorest of the poor?," Technical Report 2009. MIT mimeo.

_ , Rema Hanna, and Sendhil Mullainathan, "Corruption," Technical Report January 2011.

Bardhan, Pranab, "Corruption and Development: A Review of Issues," *Journal of Economic Literature*, September 1997, *35* (3), 1320–1346.

_ , "Decentralization of Governance and Development," *The Journal of Economic Perspectives*, 2002, *16* (4), 185–205.

_ **and Dilip Mookherjee**, "Capture and Governance at Local and National Levels," *American Economic Review*, May 2000, *90* (2), 135–139.

_ **and** _ , "Decentralizing antipoverty program delivery in developing countries," *Journal of Public Economics*, April 2005, *89* (4), 675–704.

_ **and** _ , "Decentralisation and Accountability in Infrastructure Delivery in Developing Countries," *Economic Journal*, 01 2006, *116* (508), 101–127.

_ **and** _ , “Decentralization, Corruption and Government Accountability,” in Susan Rose-Ackerman, ed., *International Handbook On The Economics Of Corruption*, Edward Elgar, 2006.

_ **and** _ , “Pro-poor targeting and accountability of local governments in West Bengal,” *Journal of Development Economics*, April 2006, *79* (2), 303–327.

Becker, Gary S., “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, 1968, *76*, 169.

Bertrand, Marianne, Simeon Djankov, Rema Hanna, and Sendhil Mullainathan, “Obtaining a Driver’s License in India: An Experimental Approach to Studying Corruption,” *The Quarterly Journal of Economics*, November 2007, *122* (4), 1639–1676.

Besley, Timothy, “Means Testing versus Universal Provision in Poverty Alleviation Programmes,” *Economica*, February 1990, *57* (225), 119–29.

_ **and Ravi Kanbur**, “The principles of targeting,” Policy Research Working Paper Series 385, The World Bank March 1990.

_ **and Stephen Coate**, “Workfare versus Welfare Incentive Arguments for Work Requirements in Poverty-Alleviation Programs,” *American Economic Review*, March 1992, *82* (1), 249–61.

Camacho, Adriana and Emily Conover, “Manipulation of Social Program Eligibility: Detection, Explanations and Consequences for Empirical Research,” DOCUMENTOS CEDE 006211, UNIVERSIDAD DE LOS ANDES-CEDE August 2009.

Coady, David, Margaret Grosh, and John Hoddinott, *Targeting of transfers in developing countries: review of lessons and experience*, World Bank, 2004.

Coady, David P. and Susan W. Parker, “Targeting Performance under Self-selection and Administrative Targeting Methods,” *Economic Development and Cultural Change*, 04 2009, *57* (3), 559–587.

Dreze, Jean and Amartya Sen, *Hunger and Public Action*, Oxford University Press, 1989.

_ **and Reetikha Khera**, “The BPL Census and a Possible Alternative,” *Economic & Political Weekly*, February 2010, *45* (9), 54–63.

Foster, James, Joel Greer, and Erik Thorbecke, "A Class of Decomposable Poverty Measures," *Econometrica*, May 1984, *52* (3), 761–66.

Galasso, Emanuela and Martin Ravallion, "Decentralized targeting of an antipoverty program," *Journal of Public Economics*, 2005, *89* (4), 705 – 727.

Glaeser, Edward L. and Andrei Shleifer, "The Rise of the Regulatory State," *Journal of Economic Literature*, June 2003, *41* (2), 401–425.

Glewwe, Paul, "Targeting assistance to the poor : Efficient allocation of transfers when household income is not observed," *Journal of Development Economics*, April 1992, *38* (2), 297–321.

Grosh, Margaret, *Administering Targeted Social Programs in Latin America* Regional and Sectoral Studies, World Bank, 1994.

Grosh, Margaret E. and Judy L. Baker, "Proxy Means Tests for Targeting Social Programs," *Living Standards Measurement Study Working Paper*, 1995, (118), 1–49.

Hirway, Indira, "Identification of BPL Households for Poverty Alleviation Programmes," *Economic & Political Weekly*, November 2003, *38* (45), 4803–4808.

Holmstrom, Bengt and Paul Milgrom, "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics and Organization*, Special I 1991, *7* (0), 24–52.

Khera, Reetika, "Access to the Targeted Public Distribution System: A Case Study in Rajasthan," *Economic & Political Weekly*, November 2008, *43* (44), 51–56.

Martinelli, Csar and Susan Wendy Parker, "Deception and Misreporting in a Social Program," *Journal of the European Economic Association*, June 2009, *7* (4), 886–908.

Mirrlees, J A, "The Theory of Moral Hazard and Unobservable Behaviour: Part I," *Review of Economic Studies*, January 1999, *66* (1), 3–21.

Mirrlees, James A, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, April 1971, *38* (114), 175–208.

Niehaus, Paul and Sandip Sukhtankar, "Corruption Dynamics: the Golden Goose Effect," 2009. mimeo.

Olken, Benjamin A., "Corruption and the costs of redistribution: Micro evidence from Indonesia," *Journal of Public Economics*, May 2006, *90* (4-5), 853–870.

Park, Albert, Sangui Wang, and Guobao Wu, "Regional poverty targeting in China," *Journal of Public Economics*, October 2002, *86* (1), 123–153.

Ram, F, S K Mohanty, and Usha Ram, "Understanding the Distribution of BPL Cards: All-India and Selected States," *Economic & Political Weekly*, February 2009, *44* (7), 66–71.

Ravallion, Martin, "Land-contingent poverty alleviation schemes," *World Development*, August 1989, *17* (8), 1223–1233.

_ , "How Relevant Is Targeting to the Success of an Antipoverty Program?," *World Bank Research Observer*, August 2009, *24*, 205–231.

_ **and Binayak Sen**, "Impacts on rural poverty of land-based targeting: Further results for Bangladesh," *World Development*, June 1994, *22* (6), 823–838.

_ **and Kalvin Chao**, "Targeted policies for poverty alleviation under imperfect information: Algorithms and applications," *Journal of Policy Modeling*, 1989, *11* (2), 213–224.

Reinikka, Ritva and Jakob Svensson, "Local Capture: Evidence From a Central Government Transfer Program in Uganda," *The Quarterly Journal of Economics*, May 2004, *119* (2), 678–704.

Schady, Norbert R., "Picking the poor : indicators for geographic targeting in Peru," Policy Research Working Paper Series 2477, The World Bank November 2000.

Sundaram, K, "On Identification of Households below Poverty Line in BPL Census 2002: Some Comments on the Proposed Methodology," *Economic & Political Weekly*, March 2003, *38* (9), 896–901.

Svensson, Jakob, "Who Must Pay Bribes And How Much? Evidence From A Cross Section Of Firms," *The Quarterly Journal of Economics*, February 2003, *118* (1), 207–230.

A Proofs

A.1 Proof of Proposition 1

Proof. Given $\pi(a_i, \mathbf{x}_i, R) = \pi > 0$ whenever $a_i \neq 1(\mathbf{x}_i \in R)$ the official's problem amounts to choosing prices for the four categories defined by the product of rich/poor and eligible/ineligible. Call these categories $\{EP, IP, ER, IR\}$. His equilibrium choices are

$$p_{EP} = \max\{0, \eta - \pi f - \underline{\alpha}\} \tag{13}$$
$$p_{IP} = \max\{0, \eta + \pi f - \underline{\alpha}\} \tag{14}$$
$$p_{ER} = \max\{0, \eta - \pi f - \overline{\alpha}\} \tag{15}$$
$$p_{IR} = \max\{0, \eta + \pi f - \overline{\alpha}\} \tag{16}$$

Let m_x denote the mass of in category x and $\omega_x \in \{\underline{\omega}, \overline{\omega}\}$ the appropriate welfare weight for each group. The principal's payoff as a function of f satisfies

$$V(f) = \sum_{x \in \{EP,IP,ER,IR\}} \exp\left\{-\frac{p_x}{\eta}\right\}(\omega_x \eta - 1) \tag{17}$$
$$\frac{\partial V(f)}{\partial f} = -\frac{1}{\eta} \sum_{x \in \{EP,IP,ER,IR\}} \exp\left\{-\frac{p_x}{\eta}\right\}(\omega_x \eta - 1)\frac{\partial p_x}{\partial f} \tag{18}$$

If $\underline{\alpha} > \eta - \pi f$ then $p_{EP} = 0$ and so among the poor stronger enforcement can only hurt, by raising p_{IP}. If $\overline{\alpha} < -\eta + \pi f$ then all rich households face strictly positive prices and so the contribution of the rich to $\partial V / \partial f$ is

$$\frac{\pi}{\eta}\left[\exp\left\{-\frac{p_{ER}}{\eta}\right\} m_{ER} - \exp\left\{-\frac{p_{IR}}{\eta}\right\} m_{IR}\right](\overline{\omega}\eta - 1) \tag{19}$$

which is strictly negative provided that

$$\exp\left\{\frac{p_{IR} - p_{ER}}{\eta}\right\} > \frac{m_{IR}}{m_{ER}} \Leftrightarrow f > \frac{\eta}{2\pi}\log\left(\frac{m_{IR}}{m_{ER}}\right) \tag{20}$$

On the other hand if the rule perfectly targets the poor then $m_{IP} = m_{ER} = 0$ and it is easy to see that $\partial V / \partial f \geq 0$. □

A.2 Proof of Lemma 1

Proof. If the principal makes a household with $x_1 = x$ eligible then that household will receive a slot with probability $\exp\left\{-\frac{1}{\eta}(\eta - \phi_1 f)\right\}$ while if they are ineligible they will receive a slot with probability $\exp\left\{-\frac{1}{\eta}(\eta + \phi_1 f)\right\}$. The difference in the principal's payoff induced by making such a household eligible is proportional to

$$\left[\exp\left\{-\frac{1}{\eta}(\eta - \phi_1 f)\right\} - \exp\left\{-\frac{1}{\eta}(\eta + \phi_1 f)\right\}\right]\left((\underline{\omega}\eta - 1)\rho(x) + (\overline{\omega}\eta - 1)(1 - \rho(x))\right) \quad (21)$$

which is positive if and only if $\rho(x)\underline{\omega} + (1 - \rho(x))\overline{\omega} \geq \frac{1}{\eta}$. This along with the monotonicity of $\rho(x)$ implies that the strictly optimal rule among those that condition on x_1 only is $R_1 \equiv \{\mathbf{x} : x_1 \leq x_1^*\}$ where x_1^* is defined by $\rho(x_1^*)\underline{\omega} + (1 - \rho(x_1^*))\overline{\omega} = \frac{1}{\eta}$. □

A.3 Proof of Proposition 2

Proof. Consider any rule R which conditions non-trivially on x^2 in the sense that there is a positive-measure subset $S \subseteq X_1$ within which the eligibility status of households depends on x^2. Define E and I the (possibly empty) subsets of $X_1 \setminus S$ within which all households are eligible and ineligible, respectively. As $\phi_2 \to 0$ prices in regions E, S, and I approach $\exp\left\{-\frac{1}{\eta}(\eta - \phi_1 f)\right\}$, $\exp\{-1\}$, and $\exp\left\{-\frac{1}{\eta}(\eta + \phi_1 f)\right\}$ respectively.

The argument in Lemma 1 shows that there exists x^{1*} such that the principal obtains strictly positive expected utility from giving slots to households with $x_i^1 < x^{1*}$ and strictly negative expected utility from giving slots to households with $x_i^1 > x^{1*}$. Thus if there are any households in $(X_1 \setminus E) \cap [0, x^{1*})$ then he can do strictly better (asymptotically) by expanding E to $[0, x^{1*})$, raising these households' probability of obtaining slots from $\exp\{-1\}$ or $\exp\left\{-\frac{1}{\eta}(\eta + \phi_1 f)\right\}$ to $\exp\left\{-\frac{1}{\eta}(\eta - \phi_1 f)\right\}$. Similarly if there are any households in $(X_1 \setminus I) \cap (x^{1*}, \infty)$ he can do strictly better (asymptotically) by expanding I. Since S contains a positive mass of households at least one of these two modifications is possible; together they yield R_1 and a strictly higher payoff. □