

Motivating Knowledge Agents:
Incentive Pay vs Social Proximity*

JOB MARKET PAPER

For the latest version of this paper, go to:

www.erlend.biz

Erlend Berg (Oxford)[†] Maitreesh Ghatak (LSE)

R Manjula (ISEC) D Rajasekhar (ISEC)

Sanchari Roy (Warwick)

21 October, 2012

*We thank Oriana Bandiera, Arunish Chawla, Gerard Padro-i-Miquel, Torsten Persson, Vijayendra Rao, Yonah Rubinstein and seminar participants at Bristol, IGC Growth Week 2011, IGC Patna meeting 2011, ISI Delhi, LSE, Oxford and Warwick for helpful comments. We gratefully acknowledge funding from Improving Institutions for Pro-Poor Growth (iiG), a research programme consortium funded by DFID, as well as the International Growth Centre. Any views expressed are the authors' and not necessarily those of the funders.

[†]Corresponding author. Email: erlend.berg@economics.ox.ac.uk

Abstract

While economists have analysed problems relating to the supply of public services in developing countries, the demand side of public services is relatively under-studied. This paper presents a model which combines a motivated-agent framework with features of the multi-tasking model to analyse the impact of incentive payment for agents tasked with spreading awareness about a public service. The results are interpreted in the context of a randomized field experiment undertaken across 151 villages in South India. In a random subset of the villages, an agent was hired to spread awareness of a health insurance scheme and paid either a flat fee or a variable rate that depended on the level of knowledge about the scheme in the local population. It is found that the presence of an agent increases knowledge, and that this effect is driven by agents on incentive-pay contracts. Using the agents as an instrument for knowledge in the target population, it is shown that an increase in knowledge leads to an increase in the take-up rate; thereby establishing that information costs are an impediment to the take-up of the scheme. A measure of ‘social distance’ is constructed, and it is shown that information transmission is less effective when agent and household are more distant. Incentive pay does not increase effort with respect to the agent’s own group, but it does increase effort with respect to socially distant households. The typical agent seems to reach a maximum effort level with respect to her own group. Therefore, the extra effort related to the introduction of the bonus pay is purely in terms of effort with respect to the other group. It appears that modest incentives can overcome entrenched social barriers in specific contexts.

JEL Classification: I38, C93, M52, O15

Keywords: incentive pay, social proximity, knowledge

1 Introduction

In many developing countries, the quality of public services are poor. With the vast majority of the population in these countries being dependent on publicly provided services in areas such as health care and education, the human and economic costs are enormous. Although the real cause of this failure is often a combination of supply-side *and* demand-side constraints, the existing literature on the topic has mostly focused on supply-side problems. For example, a growing body of work looks at the role of incentives in inducing teachers and health workers to turn up for work and provide services at reasonable levels of quality (Duflo, Hanna and Ryan, 2012; Glewwe, Holla and Kremer, 2009; Muralidharan and Sundararaman, 2011; Banerjee, Glennester and Duflo, 2008).

The demand for public services, on the other hand, remains relatively understudied.¹ An important aspect of public services is to make intended beneficiaries aware of their entitlements. Even if there were no supply-side problems—if the quality of, say, schools and health centres were excellent and these facilities were widely available—the outcome would be disappointing if beneficiaries were unaware of them or did not value them. A recent World Bank report on public services in India shows programme awareness to be very low among target groups (World Bank, 2011, pp. 162–165). According to this report, the level of nationwide awareness regarding the National Rural Employment Guarantee, one of the flagship anti-poverty schemes of the Government of India, was around 57% in 2006, with some of the poorer states like Jharkhand and Madhya Pradesh, where one would expect demand for such schemes to be high, doing worse at 29% and 45%, respectively.²

This paper begins to fill the gap. It provides evidence from a randomized experiment designed to evaluate the impact of providing financial incentives to agents spreading information about a government welfare programme, on the level of programme knowledge and take-up among beneficiaries. The experiment was conducted across 151 villages in rural Karnataka, a state in South India. In a random subsample of villages (the treatment groups), local women were recruited to be our ‘agents’ in charge of information dissemination about the programme in her home village. The agents were randomly assigned to either a flat-pay contract or an incentive-pay contract. The latter agents’ pay was linked to a measurable outcome in the form of the average score on a knowledge test about the insurance programme, administered to eligible households in that village. Flat-pay agents were paid a fixed amount of 400 rupees (9 USD) every three months. Incentive-pay agents were paid a fixed amount of 200 rupees every three months, plus a variable component which depended on the outcome of knowledge tests in their village. The variable part of the pay scale was designed to equalise average payment across incentive-pay and flat-pay agents. This enables us to obtain causal estimates of the impact of incentive pay for knowledge agents on two key dimensions of public service participation: programme knowledge and enrolment.

The use of incentive-pay schemes for public sector workers is controversial, with several studies warning against possible perverse outcomes related to either multi-

¹Banerjee et al. (2010*a*) is a notable exception.

²In developed countries, too, information costs are often argued to be one of the main reasons for low take-up of various welfare programmes Hernanz, Malherbet and Pellizzari (2004). For example, in the US, Aizer (2007) finds that eligible children do not sign up for free public health insurance (Medicaid) because of high information costs, and Daponte, Sanders and Taylor (1999) find that randomly allocating information about the Food Stamp Program significantly increases participation amongst eligible households.

tasking mechanisms (Jacob and Levitt, 2003; Neal and Schanzenbach, 2010) or crowding out of intrinsic motivation (Gneezy and Rustichini, 2000; Mellstrom and Johannesson, 2008). This point assumes greater salience in the light of recent work identifying social identity matching (or social proximity) as an important determinant of programme outcomes. Existing empirical studies find evidence of strong own-group bias (Banerjee and Munshi, 2004; Kingdon and Rawal, 2010), which may have potentially adverse implications for resource allocation and the welfare of other groups. However, so far, little attention has been paid to the interplay between incentive pay schemes and social proximity in determining economic outcomes and associated welfare consequences. In particular, it is pertinent to ask whether incentive pay can overcome the negative consequences of social distance.

This paper develops a simple model of motivated agents embedded in a variant of the multi-tasking framework to generate predictions about the effect of incentive pay and its interaction with social-identity matching between agents and beneficiary groups.

Motivated by the theoretical predictions, the following questions are addressed empirically: (i) Does incentive pay for agents tasked with spreading information about public welfare programmes actually improve programme knowledge in the target population? (ii) Does improved programme knowledge, in turn, lead to higher programme take-up? (iii) Does matching agents with target households in terms of social characteristics have an effect on knowledge that is independent of the incentive-pay effect? (iv) Does incentive pay reinforce or weaken the social-matching effect?

The main findings are as follows: Hiring agents to spread knowledge about a specific welfare programme in South India has a positive impact on the level of programme knowledge, and the effect is entirely driven by agents on incentive-pay contracts. Households in villages assigned an incentive-pay agent score 0.25 standard deviations higher on the knowledge test than those in the control group.

Second, using the random assignment of incentivised agents as an instrument for knowledge, it is found that improved knowledge increases programme take-up. An increase of one standard deviation in knowledge score increases the likelihood of enrolment by 39 percentage points.

Third, social distance between agent and beneficiary has a negative impact on knowledge transmission. Putting agents on incentive-pay contracts appears to increase knowledge transmission by cancelling out (at our level of bonus pay) the negative effect of social distance. On the other hand, incentive pay has no impact on knowledge transmission for socially proximate agent-beneficiary pairs.

The findings suggest that, with respect to their own group, agents were already choosing the maximum effort \bar{e} and hence, introducing bonus pay has no impact. However, non-incentivised agents choose a lower level of effort with respect to the ‘other’ group. With bonus pay, effort goes up to the same level as for the agent’s own group. We do not observe crowding out empirically, but it could still happen outside of the observed parameter values.

To the best of our knowledge, this paper presents the first randomized evaluation of incentive pay for agents tasked with providing eligible households with information about a public service. The paper contributes to the growing literature on the importance of information in improving the efficiency of economic decisions. Banerjee et al. (2010*b*) study how information campaigns affect local participation and educational outcomes in India. Jensen (2010) finds that perceived returns to secondary schooling are low even though measured returns are high in the Dominican Republic, and that providing information on measured returns increases years of schooling. This is related to the findings of Dupas (2011) who provides information about HIV prevalence on the incidence of risky sexual behaviour among girls in Kenya, and of Duflo and Saez (2003), who find that providing incentives to attend information sessions affect retirement-plan decisions.

Our paper is also among the first to provide empirical evidence on the interplay between explicit financial incentives and social proximity. Our paper is therefore related to the rich literature on the role of monetary and non-monetary incentives on the performance of agents. This literature encompasses studies not only in the ‘standard setting’ of firms in developed countries where output or productivity is measurable but worker effort is not (Bandiera, Barankay and Rasul (2011) survey this literature), but also includes studies on incentives for teachers and health workers in developing countries (as surveyed in Kremer and Holla (2008) and Glewwe, Holla and Kremer (2009)). There is also a body of literature looking at the role of agents’ intrinsic motivation and identification with either the task at hand or the intended beneficiaries in reducing the need for explicit incentives (Akerlof and Kranton, 2005; Benabou and Tirole, 2003; Besley and Ghatak, 2005). However, as Bandiera, Barankay and Rasul (2011) point out, there is little field-experimental evidence in this area (Ashraf, Bandiera and Jack (2012) is a recent exception).

The rest of the paper is organised as follows: In Section 2, we provide a simple theoretical framework to analyse the impact of incentive pay on agents’ effort and its interaction with social identity matching. Section 3 describes the context, experimental design and data. Section 4 presents the empirical evidence and Section 5 interprets it. Section 6 concludes.

2 Theoretical Framework

In this section a simple model of motivated agents (Besley and Ghatak, 2005) is developed and extended to incorporate features of the multi-tasking model (Holmstrom and Milgrom, 1991). The goal is to provide a theoretical framework to generate predictions about the effects of incentive payment and how these might interact with the effects of social distance. The theoretical framework is also used to evaluate the welfare implications of the experiment as well as carry out a number of counter-factual exercises.

Suppose knowledge agents (henceforth, agents) exert unobservable effort in spreading awareness of a scheme to potential beneficiaries. The goal may either be the transmission of knowledge itself or it may be to increase programme enrolment. The principal can be thought of as a planner (say, the relevant government agency) who values either awareness of or enrolment in the programme among the eligible population. A given agent can interact with a fixed number of beneficiaries which we take to be exogenous.

2.1 Homogeneous Agents and Beneficiaries and the Role of Incentives

First, suppose agents and beneficiaries are homogeneous. Let e be the unobservable effort exerted by the agent. Let the outcome variable Y be binary and of value 0 or 1, with the former denoting ‘failure’ and the latter, ‘success’. For example, a group of beneficiaries doing well in the knowledge test (say, scoring above a certain threshold level), or enrolling in the programme, might be considered a success. The agent’s effort stochastically improves the likelihood of good outcomes. To keep things simple, assume that the probability of success is $p(e) = e$, so that attention is restricted to values of e that lie between 0 and 1. In particular, let us assume that the lowest value that e can take is $\underline{e} > 0$, and the highest value that e can take is $\bar{e} \in (\underline{e}, 1)$. This means that there is some minimum effort any agent supplies and that even with this minimum effort, there is some chance that the good outcome will happen. There is also a maximum level of effort but even at that level, the good outcome is not guaranteed to occur. Therefore, as is standard in agency models, there is common support, i.e., any value of the outcome is consistent with any value of effort in the feasible range. It is also assumed that both the principal and the agent are risk-neutral.

Let the agent’s disutility of effort be $c(e) = \frac{1}{2}ce^2$. If the project succeeds, the agent receives a non-pecuniary pay-off of θ —this is her intrinsic motivation for the

task—and the principal receives a pay-off of π , which may have a pecuniary as well as a non-pecuniary component. The planner’s pay-off incorporates both the direct benefit to the beneficiaries and how the rest of society values their welfare. Without any incentive problems, the problem is

$$\max_e (\theta + \pi) e - \frac{1}{2} c e^2.$$

It is assumed that $\theta + \pi$ is large enough to ensure that, in the absence of incentive problems, the efficient level of effort exceeds the minimum, $\frac{\theta + \pi}{c} > \underline{e}$. Then we have

$$e^{**} = \min \left\{ \frac{\theta + \pi}{c}, \bar{e} \right\}.$$

Without any incentive problems the principal can simply stipulate e^{**} . For the problem to be interesting, and for incentive pay to have an effect, assume there is moral hazard in the choice of effort. Also, agents have zero wealth and there is limited liability: the agent’s income in any state of the world must be above a certain minimum level, say, $\underline{\omega} > 0$. From the principal’s point of view, this creates a tension between minimising costs and providing incentives. In the absence of a limited liability constraint, the principal could have achieved the first-best outcome by imposing a stiff penalty or fine for failure. With limited liability, the only way the principal can motivate the agent, apart from simply relying on her intrinsic motivation, θ , is to pay her a bonus that is contingent on performance. In choosing the bonus for the agent, the principal has to respect the limited-liability constraint and the incentive-compatibility constraint (henceforth, ICC). There is also a participation constraint (henceforth, PC) which requires the agent’s expected pay-off to be at least as high as her outside option. To keep things simple, assume that the outside option is relatively unattractive so that the PC does not bind—the analysis is qualitatively unchanged if this assumption is relaxed.

Let \bar{w} be the pay the principal offers to the agent in the case of success, and let \underline{w} be the pay in the case of failure. Define $b \equiv \bar{w} - \underline{w}$, which can be interpreted as bonus pay with \underline{w} as the fixed wage component. Then the agent’s objective is

$$\max_e (\theta + \bar{w}) e + \underline{w} (1 - e) - \frac{1}{2} c e^2$$

which yields

$$e = \max \left\{ \min \left\{ \frac{b + \theta}{c}, \bar{e} \right\}, \underline{e} \right\}. \quad (1)$$

This is the incentive compatibility constraint (ICC). Since $b \leq \pi$, effort will, in

general, be lower than in the first-best scenario. This can be formally seen as follows. The principal's objective is³

$$\max_{\bar{w}, \underline{w}} (\pi - \bar{w})e - \underline{w}(1 - e),$$

subject to the ICC (1), the limited liability constraints (LLC) $\bar{w} \geq \underline{\omega}$ and $\underline{w} \geq \underline{\omega}$ and the participation constraint (PC)

$$(\theta + \bar{w})e + \underline{w}(1 - e) - \frac{1}{2}ce^2 \geq \underline{u}.$$

Since we ignore the PC (which is justified if \underline{u} is small enough), the optimal contract is easy to characterise (see Besley and Ghatak, 2005, for details). Since the agent is risk-neutral, \underline{w} will be at the lowest limit permitted by the LLC, namely $\underline{w} = \underline{\omega}$. The solution for optimal bonus then follows:

$$b = \max \left\{ \frac{\pi - \theta}{2}, 0 \right\}$$

Note that optimal bonus is strictly smaller than π .

The effort response is only observed experimentally for two values of b , so the focus here will be on the incentive constraint (1) rather than the optimal bonus. If there is no bonus pay and the agent is not intrinsically motivated, we may get a lower corner solution, namely $e = \underline{e}$. At the other extreme, if the agent is sufficiently motivated, $\frac{\theta}{c} \geq \bar{e}$, then even without any bonus pay the agent chooses the maximum level of effort \bar{e} . Otherwise, effort is increasing in bonus pay.

In Figures 1a, 1b and 1c, we summarize the key effect of bonus pay on three types of agents: those who are unmotivated ($\frac{\theta}{c} < \underline{e}$), those who are motivated ($\underline{e} < \frac{\theta}{c} < \bar{e}$), and those who are super-motivated ($\frac{\theta}{c} > \bar{e}$), respectively. In the first case, agents are not intrinsically motivated. Without bonus pay, effort is at the minimum level \underline{e} , and only when b exceeds some threshold does effort increase above \underline{e} . In the second case, the agents are motivated and so put in positive effort even without bonus pay. Effort is increasing in bonus pay, but for some critical level of bonus pay it hits the upper limit \bar{e} . For the third group, agents are super-motivated, i.e., they put in the maximum possible effort \bar{e} for all levels of bonus pay.

The main conclusions from this analysis are: (i) if the agent is unmotivated, bonuses will, in general, increase effort; however, they have to cross some minimum

³In the formulation presented here it is assumed that the principal does not put any direct weight on the agent's welfare but does take into account the welfare of the beneficiaries. An alternative formulation would be to put a weight λ on the welfare of the beneficiaries and a weight $1 - \lambda$ on the welfare of the agents.

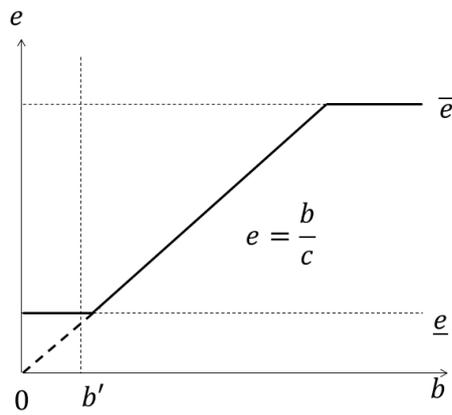


Figure 1a: Unmotivated agent

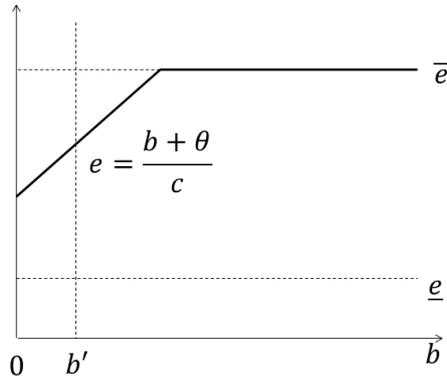


Figure 1b: Motivated agent

threshold for this to happen; (ii) if the agent is motivated, bonuses will increase effort but they have to be below some threshold for this effect to be observed, since otherwise the agent puts in the highest possible level of effort.

2.2 Heterogeneous Agents and Beneficiaries and the Role of Incentives vs Matching

Assume now the agent has two tasks, like in the multi-tasking model. The tasks may be thought of as the agent transferring knowledge to, or enrolling, two different types of beneficiary households. Unlike the multi-tasking model, the outcomes associated with the two tasks are assumed to be equally measurable. Instead, the differences between the two tasks lie in the agent's intrinsic pay-off for success and her cost of

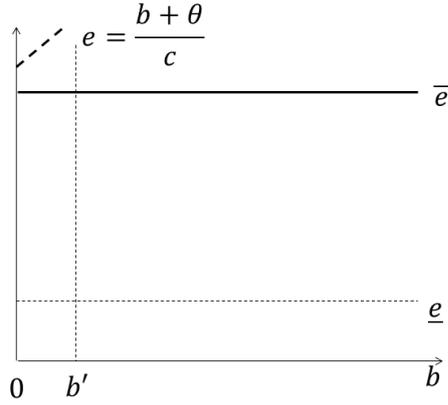


Figure 1c: Super-motivated agent

effort.

Extending the notation from the previous section, let Y_1 and Y_2 be the binary outcomes for the two tasks and e_1 and e_2 , the corresponding effort levels.

It is assumed that the principal is constrained to offer the agent the same conditional pay-outs for the two tasks. That is, the payment in the case of success must be the same for task 1 and 2, as must the payment in the case of failure. This is justified if the principal is politically, socially or legally constrained to offer the same pay rates for all tasks. The assumption is also justified if the relevant characteristics of the households are not observable to the principal. For example, in the context of teacher remuneration, a teacher may have favourite students, but if the principal cannot observe who they are, the remuneration scheme cannot vary along that dimension.

Let $0 < \underline{e} < \bar{e} < 1$ define lower and upper bounds for both e_1 and e_2 , and let θ_1 and θ_2 denote the non-pecuniary pay-offs to the agent from success in tasks 1 and 2, respectively. Let the agent's cost of effort be given by

$$c(e_1, e_2) = \frac{1}{2}c_1e_1^2 + \frac{1}{2}c_2e_2^2 + \gamma e_1e_2.$$

The parameter γ can be thought of as a measure of the substitutability of effort in tasks 1 and 2 in the cost function. To ensure that the marginal cost of effort in each task is always positive, it is assumed that $\gamma \geq 0$.

Note that if $c_1 = c_2 = \gamma = c$ and $\theta_1 = \theta_2 = \theta$, the set-up collapses to the single-task model. Abstracting from the special case $c_1 = c_2$ we can, without loss of generality, assume that $c_1 < c_2$ and refer to task 1 and 2 as the easier and the harder task, respectively.

The principal values the tasks equally and so receives the same pay-off π from success in both. Then the first-best is characterised by:

$$\max_{e_1, e_2} (\theta_1 + \pi) e_1 + (\theta_2 + \pi) e_2 - \left(\frac{1}{2} c_1 e_1^2 + \frac{1}{2} c_2 e_2^2 + \gamma e_1 e_2 \right)$$

The first-order conditions yield the following interior solutions:

$$\begin{aligned} \hat{e}_1(\pi) &= \frac{(c_2 - \gamma) \pi + c_2 \theta_1 - \gamma \theta_2}{c_1 c_2 - \gamma^2} \\ \hat{e}_2(\pi) &= \frac{(c_1 - \gamma) \pi + c_1 \theta_2 - \gamma \theta_1}{c_1 c_2 - \gamma^2} \end{aligned}$$

The second-order condition requires

$$c_1 c_2 > \gamma^2$$

for this to be a local maximum.

As before, corner solutions may be possible, and therefore we have the following characterisation of the first-best effort levels:

$$\begin{aligned} e_1^{**} &= \max \{ \min \{ \hat{e}_1(\pi), \bar{e} \}, \underline{e} \} \\ e_2^{**} &= \max \{ \min \{ \hat{e}_2(\pi), \bar{e} \}, \underline{e} \} \end{aligned}$$

There are two basic cases consistent with the second-order condition:

$$\begin{aligned} \text{Case 1} &: \gamma < c_1 < c_2 \\ \text{Case 2} &: \frac{\gamma^2}{c_2} < c_1 < \gamma < c_2 \end{aligned}$$

In Case 1, γ is smaller than both c_1 and c_2 , while in Case 2 the magnitude of γ is between those of c_1 and c_2 .

The second-best is characterised as follows. Let \bar{w} be the wage the principal offers to the agent conditional on success in a task, let \underline{w} be the wage conditional on failure, and define $b \equiv \bar{w} - \underline{w}$. The agent's objective is to maximise:

$$\max_{e_1, e_2} (\theta_1 + \bar{w}) e_1 + (\theta_2 + \bar{w}) e_2 + \underline{w}(1 - e_1) + \underline{w}(1 - e_2) - c(e_1, e_2)$$

The first-order conditions yield:

$$\begin{aligned}\hat{e}_1(b) &= \frac{(c_2 - \gamma)b + c_2\theta_1 - \gamma\theta_2}{c_1c_2 - \gamma^2} \\ \hat{e}_2(b) &= \frac{(c_1 - \gamma)b + c_1\theta_2 - \gamma\theta_1}{c_1c_2 - \gamma^2}\end{aligned}$$

Define the ‘intrinsically preferred task’ as the task in which the agent exerts the greatest effort when there is no bonus pay. Task 1 is the intrinsically preferred task iff $\hat{e}_1(0) > \hat{e}_2(0)$, or

$$\frac{\theta_1}{c_1 + \gamma} > \frac{\theta_2}{c_2 + \gamma}.$$

Intuitively, a higher θ_i and a lower c_i both contribute to the agent’s intrinsic task preference for task i . Note that it is possible that the harder task, task 2, is intrinsically preferred by the agent. This is the case if her intrinsic pay-off for the harder task (θ_2) is large enough to outweigh the cost disadvantage.

As in the single-task model, we expect effort levels to be lower than first-best because the participation constraint of the agent is assumed not to bind. Corner solutions are still possible, so we get the following characterisation of the second-best effort levels:

$$\begin{aligned}e_1^* &= \max\{\min\{\hat{e}_1(b), \bar{e}\}, \underline{e}\} \\ e_2^* &= \max\{\min\{\hat{e}_2(b), \bar{e}\}, \underline{e}\}\end{aligned}$$

It is straight forward to show that, for interior solutions, the marginal effect of bonus pay on total effort, $\hat{e}_1 + \hat{e}_2$, is always positive. It is also clear that the marginal effect of bonus pay on effort in task 1 is always positive and also greater than that on effort in task 2. However, the marginal effect on effort in task 2 can be positive or negative, depending on the case:

Case 1, $\gamma < c_1 < c_2$. In this case, effort in both tasks is increasing in bonus pay because the substitutability of the tasks in the cost function is relatively low.

Case 2, $c_1 < \gamma < c_2$. Now, effort in task 2 is decreasing while effort in task 1 is increasing in bonus pay. We can interpret this as a form of ‘crowding out’: the effect of increasing bonus pay on effort in the high-cost task is negative. The intuition is that when effort in the two tasks are relatively substitutable, providing monetary incentives leads the agent to substitute effort towards the easier task to a degree that causes effort in the harder task to fall.

The corner solutions introduce regions in which the marginal effect of bonus pay on the effort in one or both tasks is zero. Generally, there is a critical value $\bar{b}_1 \geq 0$

such that $e_1 = \bar{e}$ for all $b \geq \bar{b}_1$. In Case 1, there is also a critical value \bar{b}_2 such that $e_2 = \bar{e}$ for all $b \geq \bar{b}_1$. In Case 2, there is a critical value \underline{b}_2 such that $e_2 = \underline{e}$ for all $b \geq \underline{b}_2$.

Figures 2a-2d we present the four main combinations of intrinsically preferred task and whether crowding out is possible or not. For the sake of clarity, the upper and lower effort bounds are omitted from these figures.

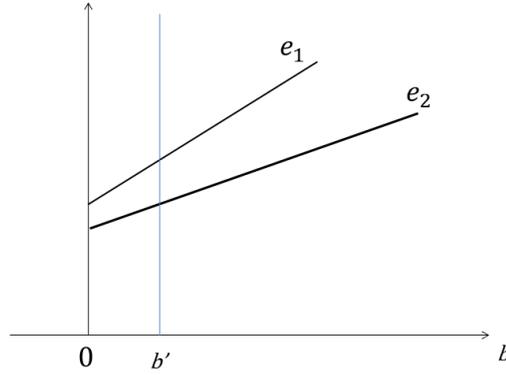


Fig 2a. Case 1: $\gamma < c_1 < c_2$
No crowding out; Group 1 own group

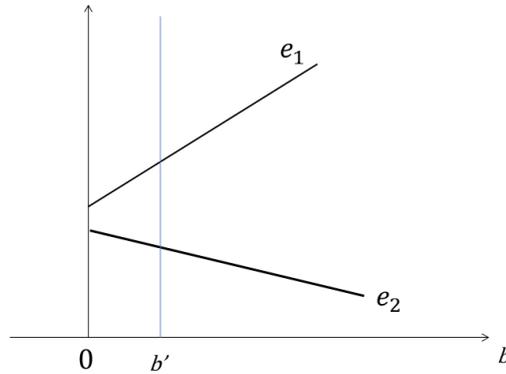


Fig 2b. Case 1: $c_1 < \gamma < c_2$
Crowding out; Group 1 own group

In Figure 2a, substitutability between the two tasks is relatively low ($\gamma < c_1 < c_2$); hence, both e_1 and e_2 are upward-sloping. Moreover, the intercept of e_1 being greater than that of e_2 implies that Task 1 (the easier task) is intrinsically preferred. There is no crowding out. In Figure 2b, however, the two tasks are more substitutable (i.e., $c_1 < \gamma < c_2$), so that e_1 is positively sloped but e_2 is negatively sloped.

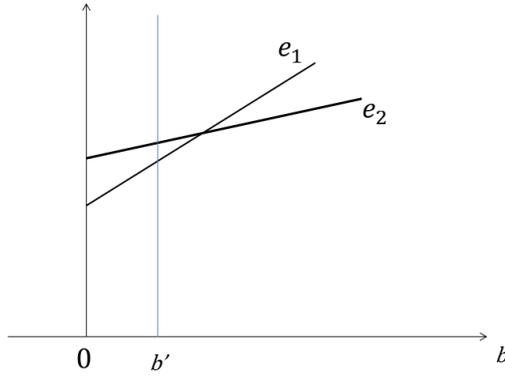


Fig 2c. Case 1: $\gamma < c_1 < c_2$
No crowding out; Group 2 own group

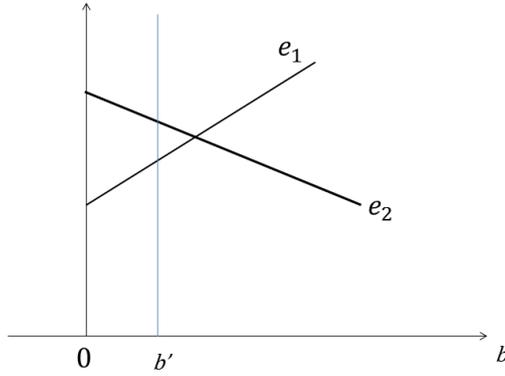


Fig 2d. Case 1: $c_1 < \gamma < c_2$
Crowding out; Group 2 own group

Task 1 is still the intrinsically preferred task. A higher bonus pay can increase effort in task 1 only by reducing effort in task 2, leading to crowding out.

Figures 2c and 2d are similar to 2a and 2b, except that here task 2 (the harder task) is intrinsically preferred. Figure 2c depicts the case with no crowding out while Figure 2d depicts the case with crowding out.

Mapping the theory to the experimental setting, each of the model's two tasks can be thought of as corresponding to a group of eligible households in the agent's village. It will be shown that, in the absence of bonus pay, agents tend to exert a greater effort with respect to households who are similar to themselves in terms of social characteristics. The model's 'intrinsically preferred task' therefore correspond to households who are socially proximate to the agent. These households will also be referred to as the agent's 'own group'. Households who are socially distant from

the agent (the ‘other’ group) correspond, in the model, to the task that is not intrinsically preferred.

Which task is intrinsically preferred depends on θ_i and c_i , which are unobservable. Therefore, while the ‘own’ group will be mapped to the intrinsically preferred task, it is not possible to know whether the ‘own’ group corresponds to task 1 (the easier task) or 2 (the harder task).

3 Context, Experimental Design and Data

3.1 The Programme

The experiment was conducted in the context of India’s National Health Insurance Scheme (Rashtriya Swasthya Bima Yojana—henceforth, RSBY). The scheme was launched by the government in 2007 with the aim of improving the ‘access of BPL [Below the Poverty Line] families to quality medical care for treatment of diseases involving hospitalisation and surgery through an identified network of health care providers’ (Government of India, 2009). Each state followed its own timetable for implementation, and a few districts from each state were selected for the first stage. In Karnataka, five districts were selected (Bangalore Rural, Belgaum, Dakshina Kanada, Mysore and Shimoga), and household enrolment in these districts commenced in February–March 2010 (Rajasekhar et al., 2011).

The health insurance policy covers hospitalisation expenses for around 700 medical and surgical conditions, with an annual expenditure cap of 30,000 rupees (670 USD) per eligible household. Each household can enrol up to five members. Pre-existing conditions are covered, as is maternity care, but outpatient treatment is excluded.

The policy is underwritten by insurance companies selected in state-wise tender processes. The insurer receives an annual premium per enrolled household,⁴ paid by the central (75%) and state (25%) governments. The beneficiary household pays only a 30 rupees (0.67 USD) annual registration fee.

Biometric information is to be collected from all members on the day of enrolment and stored in a smart card issued to the household on the same day.⁵ Beneficiaries are entitled to cashless treatment at any participating (‘empanelled’) hospital across India. Both public and private hospitals can be empanelled. Hospitals are issued

⁴The annual premium is determined at the state (and sometimes district) level, and is currently in the range of 400–600 rupees (9–13 USD). In Karnataka, the annual premium in the first year of operation was 475 rupees.

⁵According to RSBY guidelines, smart cards should be issued at the time of registration, but this is often not adhered to. For more detail, see Rajasekhar et al. (2011).

card readers and software. The cost of treating patients under RSBY are reimbursed to the hospital by the insurance company according to fixed rates.

3.2 Experimental Design

151 villages were randomly selected from two of the first-phase RSBY districts in Karnataka: Shimoga and Bangalore Rural. In the first stage of randomisation, some villages in our sample (112 out of 151) were randomly selected to be part of the treatment group, i.e. receive an agent, while the remaining form the control group. In each treatment village, our field staff arranged a meeting with the local Self-Help Groups (SHGs).⁶ All SHGs contacted were female-only. In the meeting, SHG members were given a brief introduction to RSBY and told that a local agent would be recruited to help spread awareness of the scheme in the village over a period of one year. They were told that the agents would be paid, but no further details on the payment were given. In each case, an agent was recruited on the same day, most often from the SHG itself; but in a small number of cases the selected agent was a non-member recommended by the SHG. In about a third of the cases, the president of the SHG became the agent. All agents were female.

Once the meeting was concluded and the agent selected, she was taken aside and given a more thorough introduction to the scheme, including details on eligibility criteria, enrolment, benefits and other relevant information. An agent background questionnaire was also fielded at this time.

At this point, the payment scheme was revealed to the agent. Each treatment village had been randomly allocated to a payment structure, but this information was kept secret. Even our field staff did not know about the contract type until after the agent had been selected. The day after recruitment, the agent would be called up and informed of her payment scheme. There were two payment schemes, defining the two treatment groups. Flat-pay agents were told that they would be paid 400 rupees every three months. Incentive-pay agents were told that knowledge of RSBY would be tested in the eligible village population every three months. The agent's pay would depend on the results of these knowledge tests. There would be a fixed payment of 200 rupees every three months, but the variable component would depend entirely on the outcome of the knowledge tests in the village.⁷

⁶Self-Help Groups are savings and credit groups of about twenty individuals, often all women, who meet regularly and receive subsidised loans.

⁷As part of the original experimental design, we also provided a second type of incentive pay to some agents based on programme utilisation by the beneficiaries in their village. But because the scheme was hardly operational during the period of our study, overall utilisation of RSBY across Karnataka was very low. See Rajasekhar et al. (2011) for details. These agents and the corresponding villages were excluded from the analysis presented here.

The bonus payments were determined as follows: A random sample of households eligible for RSBY were surveyed and orally presented with the knowledge test. A household was classified as having ‘passed’ the test if it answered at least four out of eight questions correctly. The proportion of passing households in a village was multiplied by the number of eligible households in that village in order to estimate the number of eligible village households that would have passed if everybody had taken the test. The bonus was calculated as a fixed amount per eligible household estimated to pass the test in a village, and set in such a manner that the average bonus payment across each of the two study districts would be 200 rupees per agent. The households taking the tests were not told how they scored, nor were they provided with the correct answers.

38 villages/agents were assigned to the flat-pay treatment group, and 74 to the incentive-pay treatment group. Agents were told that there would be other agents in other villages, but not that there was variation in the payment scheme.

The purpose of not revealing the payment scheme until after recruiting the agent was to isolate the incentive effect of the payment structure from its potential selection effect. None of the agents pulled out after learning of the payment scheme. However, four agents dropped out 6–12 months after recruitment. Three of these were in incentive-pay villages, while the fourth was in a flat-pay village. In each case, the reported reason was either childbirth or migration away from the village. The agents were replaced, but the villages in question are excluded from the analysis presented here. Hence, in the analysis presented here, there are 37 villages with flat-pay agents and 71 villages with incentive-pay agents, for a total of 108 agents in 108 treatment villages. The number of control villages remains 39, so the total number of villages in our final sample is 147.

The original plan was to set the variable part of the pay scale for incentive-pay agents in such a manner that average pay would equal 400 rupees in each of the two treatment groups. The aim of equalising average pay across the incentive-pay and flat-pay groups was to isolate the incentive effect of the contract structure (‘incentive effect’) from that of the expected payment amount (‘income effect’). The pay did in fact average 400 rupees for one district (Shimoga) in the first survey round and for both districts in the second round. But due to an administrative error, a majority of incentive-pay agents in Bangalore Rural were overpaid in the first round of payments. In spite of the error, the rank ordering of agents was preserved in the sense that better-performing agents were indeed paid more. Nevertheless, we also present results only for Shimoga district, where average pay in the knowledge group was equal to that of flat-pay agents (400 rupees) in both rounds.

3.3 Data

Following agent recruitment, three consecutive rounds of ‘mini-surveys’ were fielded. For each wave, randomly selected households in each sample village were interviewed to establish the state of their knowledge about the scheme, along with enrolment status. One purpose of these surveys was to provide information on agent performance so as to be able to pay the incentive-pay agents. The households were drawn at random for the first and second survey rounds, so that there was a partial overlap between the households in the first two rounds of surveys. The first and second rounds of mini-surveys were based on face-to-face interviews. For the third survey, the sample from the second survey was re-used, but this time the agents were contacted by telephone. Although not everyone could be reached by phone, the re-survey rate was significant. 2369 households were interviewed in the first mini-survey wave, 1933 in the second and 1348 in the third. In all, the mini-surveys cover 3296 households, of which 908 were interviewed twice and 723 were interviewed three times. However, using each household observation as an equally-weighted data point would give more weight to households that were observed more than once. Observation weights were introduced to take account of this, so that the total weight across observations is the same for all households. All empirical results presented here are based on weighted least squares regressions. In addition, standard errors are clustered at the village level. Since serial correlation is probably more severe within a household than across households within a village, clustering at the village level yields consistent, but not efficient, estimates.

After the completion of each mini-survey, the agents were revisited and paid. At the same time, the agents’ knowledge of the scheme was refreshed and added to.

Descriptive statistics on agents are presented in Table 1. Recall that all agents are female. The average agent is 35 years old. 88% are married. 59% of the agents’ household heads have completed primary school. 82% of agent households have a ration card, and 39% are from a forward or dominant caste.⁸ In 29% of the cases, the recruited agent was the president of a Self-Help Group.

The ‘female autonomy’ score was constructed on the basis of the following question fielded to all agents after recruitment: ‘Are you usually allowed to go to the following places? To the market; to the nearby health facility; to places outside the village.’ The answer options were ‘Alone’, ‘Only with someone else’ and ‘Not at all’. For each of the three destinations, agents were given a score of 0 if they were not

⁸In Karnataka, two castes formally classified as backward, Vokkaliga and Lingayath, tend to dominate public life. These two have therefore been classified together with the forward caste groups.

allowed to visit it at all, 1 if they were allowed to visit it only with someone else and 2 if they were allowed to visit it on their own. These three scores were added up to give an autonomy score ranging from 0 (least autonomous) to 6 (most autonomous). 82% of agents received the highest score, 6.

Table 2 presents summary statistics for households. The average household has 4.6 members. 19% are from a forward/dominant caste. In 26% of households, the household head has completed primary school. 93% have a ration card. It is interesting to note that agents are more likely than the average eligible household to belong to the forward/dominant caste category. Agent households are also more highly educated than the average eligible household.

The main outcome variable is the household ‘knowledge score’. A knowledge test was fielded to all households interviewed in each of the three mini-surveys. Each test consisted of eight questions about particulars of the RSBY scheme, including eligibility, cost, cover, exclusions and how to obtain care. The exact questions used in the knowledge tests are provided in Appendix A. Each answer was recorded and later coded as being correct or incorrect. The number of correct answers gives each interviewed household a score between 0 (least knowledgeable) and 8 (most knowledgeable).⁹

The test questions asked in the three surveys were different, so although the raw scores can be compared across households within a survey, they cannot easily be compared across surveys, even for individual households. The scores on each test were therefore standardised by subtracting the test-wise mean and dividing by the standard deviation.

4 Evidence

4.1 The Impact of Agents on Knowledge

Consider first the impact of knowledge agents on household knowledge score. The basic specification is

$$Y_{hv} = \alpha + \beta T_v + \epsilon_{hv}. \quad (2)$$

The outcome variable Y_{hv} is the test z-score for household h in village v . T_v is a binary variable equal to 1 if the household lives in a treatment village (a village with a knowledge agent of either type) and 0 otherwise. The coefficient β captures the

⁹Question 8 on the third test is difficult to mark as correct or incorrect, as there are several ways in which an RSBY member might plausibly check whether a particular condition will be covered ahead of visiting a hospital. For this reason the question is omitted when computing the overall score and the maximum score on the third test is taken to be 7.

average effect on test score of being in a treated village, and α is a constant. The standard errors ϵ_{hv} are clustered at the village level to take account of possible serial correlation and heteroscedasticity (Bertrand, Duflo and Mullainathan, 2004).

The results of regression (2) are presented in Table 3, column 1. Households living in a treatment village score 0.18 standard deviations higher on the knowledge test compared to households in the control villages. Column 2 indicates that this effect is robust to the inclusion of fixed effects for taluk (the administrative unit below district) and time (survey wave).

In column 3, the treatment effect is estimated separately for flat-pay and incentive-pay agents, while still including taluk and time fixed effects. Flat-pay agents have no significant impact on test scores. This is consistent with the argument that, since these agents are paid a constant amount irrespective of the outcome, they are not incentivised to exert any effort. In contrast, households in villages assigned an incentive-pay agent score 0.25 standard deviations higher on the knowledge test relative to those in the control group. Hence, providing agents with financial incentives leads to an improvement in knowledge about the scheme among beneficiaries. Moreover, the equality of these two coefficients is rejected with a p-value of 0.06. This suggests that the entire effect of knowledge-spreading agents in the village is due to the agents that are on incentive-pay contracts.

As already mentioned, an administrative error caused incentive-pay agents in one of the districts (Bangalore Rural) to be overpaid after the first survey. To allay concerns that our findings are driven by these higher rates of pay, Table 4 presents results using data only from Shimoga district, where no error was made. Overall, the qualitative findings concerning the main coefficients of interest are similar to those obtained in Table 3. Hence it appears that the main findings are not driven by the larger agent payments in Bangalore Rural in one of the rounds.

4.2 The Impact of Knowledge on Enrolment

Does increased knowledge about the programme cause higher enrolment rates? In order to estimate the causal impact of knowledge on enrolment we consider the following equation:

$$E_{hv} = \alpha_E + \gamma Y_{hv} + u_{hv}, \quad (3)$$

where E_{hv} captures the enrolment status of household h in village v . OLS estimation of equation (3) could lead to biased estimates of the causal impact of knowledge, because of either unobserved factors (for example, ability) that might influence both knowledge score and enrolment or reverse causality: being enrolled in the programme

may itself spur people to acquire more knowledge.

In order to address these concerns, we use the random assignment of our experimental treatment as an instrument for knowledge. Since we find that the households assigned a flat-pay agent did not exhibit a significantly different impact on knowledge scores compared to those in the control group (Table 3, column 3), we club these two groups together and use assignment to the incentive-pay group, compared to either flat-pay or pure control, as the instrument for knowledge.

The results are shown in Table 5. Column 1 presents the OLS estimate, which is positive and highly significant. Column 2 reports the reduced-form estimates obtained by regressing enrolment on the instrument, and the result indicates that households assigned an incentive-pay agent are 8 %-points more likely to enrol in the programme compared to those assigned a flat-pay agent or no agent. Column 3 reports the first stage of the IV regressions and indicates that the instrument has good explanatory power. This regression is similar to that reported in column 3 of Table 3, except that in this case the omitted group consists of not only the pure control villages but also those that were assigned a flat-pay agent.

Column 4 presents the second-stage regression, and the results indicate that an increase of one standard deviation in knowledge score increases the likelihood of enrolment by 39 %-points. Hence, we find robust evidence that improving knowledge about the welfare scheme had a positive impact on enrolment.

The fact that the co-efficient of interest nearly doubles between the OLS and IV regressions suggests that the effect of knowledge on enrolment is particularly strong for those whose knowledge levels are influenced by the presence of an incentivised agent. In other words, those households for whom the knowledge transfer is the most successful are also the households for which increased knowledge is most likely to lead to take-up.

A concern with this instrumental variable analysis might be that, in addition to imparting knowledge, the agents are associated with an ‘endorsement effect’. Having been selected by the Self-Help Group, the typical agent probably represents someone of considerable standing and respect among her peers. Therefore, the argument goes, it might be that in addition to getting knowledge of the scheme, the households perceive the agent’s involvement as a form of endorsement and might therefore be more likely to enrol irrespective of their level of knowledge.

To address this concern, the above analysis was repeated while excluding the pure control villages. Hence the comparison is now only between flat-pay and incentive-pay village, and the instrument is ‘putting the agent on an incentive contract’. Since the agents were selected before the contract type was revealed, and there was very

little attrition at any point, the flat-pay and incentive-pay agents are drawn from the same distribution and should not differ on average. In other words, if there is an endorsement effect associated with the agents, it will be the same for both types. Note that the difference between the two forms of payment contract changes the incentives to impart knowledge, but neither type is incentivised to enrol households into the scheme.

Dropping the control villages means that there is now only data on 71 incentive-pay agents/villages and 37 flat-pay agents/villages. In spite of the reduced sample size, the results are qualitatively similar to the above, although the reduced-form and first-stage estimates are significant only at the 10% level. The second-stage estimate of the effect of knowledge on enrolment is 0.54 and highly significant (not reported).

Another concern might be that the households learn of their agent's payment type and that it might somehow influence the enrolment decision. We believe it is unlikely that the households will learn of the agent's payment structure since this was revealed to the agent in private and she has little incentive to discuss it with the households. Even if she did, we believe that the most likely effect of such information would be to trust the incentive-agent relatively *less*. If so, the 'endorsement effect' of the incentive-pay should actually be negative so that the regression co-efficients are under- rather than over-estimates.

A final concern might be that the incentivised agent might try harder than the flat-pay agent to get the household enrolled *because* she thinks it will make it easier to boost their knowledge level. It is not possible to rule this out, although we find it unlikely that agents should play the game so strategically.

4.3 Agent Characteristics and Heterogeneous Treatment Effects

The questionnaire that was administered to the agents at the time of their appointment collected information on some individual and household characteristics including age, marital status, education of household head, caste, house ownership and personal autonomy.

Table 6 looks at how the impact of knowledge agents on knowledge test scores depends on agent characteristics. Column 1 replicates column 2 of Table 3 for ease of reference. In column 2, the main treatment variable (whether or not there is an agent in village) is interacted with variables on agent age, caste, education, ration-card status, home ownership, whether the agent is president of an SHG and her personal autonomy. None of these interacted effects are significant except for the

autonomy metric. (The autonomy variable is described in the data section.) It seems intuitive that an important factor determining the effectiveness of an agent is whether she is free to move around the village.

4.4 Pricing Out Prejudice: Financial Incentives and Social Distance

The results so far suggest that monetary incentives matter for agent performance, and that improving knowledge about the scheme increases enrolment. But previous work suggests that social identity is also an important determinant of insurance take-up. For example, Cole et al. (2010) find that demand for rainfall insurance is significantly affected by whether the picture on the associated leaflet (a farmer in front of either a Hindu temple or a mosque) matches the religion of the potential buyer.

This section asks whether matching agents with target households in terms of social characteristics has an effect on knowledge scores that is independent of the effect of incentive pay. Also, it is of interest to see whether the effects of social distance and incentive-pay are purely additive or whether they reinforce or weaken each other.

A simple metric of social distance is constructed as follows: First create four binary variables which capture basic social dimensions and for which we have data for both the agent and eligible households: forward/dominant caste status (0/1), whether the household head has completed primary school (0/1), ration-card status (0/1) and home ownership (0/1). In each of these four dimensions, define the social distance between an agent and a household as the absolute difference in the agent’s and the household’s characteristics. To take ration-card status as an example, ration-card distance is set to 0 if either both have a ration card or if neither of them does. Ration-card distance is 1 if one of them has a ration card and the other does not.

The composite social distance is the simple sum across the four individual distance measures. The composite social distance metric is normalised to lie between zero and one by dividing by four.

The empirical specification is of the following form:

$$Y_{hv} = \alpha + \beta D_{hv} + \gamma T_v + \delta D_{hv} * T_v + \pi X + u_{hv} \quad (4)$$

D_{hv} denotes social distance between household h in village v and the agent in village v . T_v is a binary variable indicating whether the agent in village v is on

an incentive-pay contract. (The control villages are necessarily dropped from this analysis.) X are control variables for each of the agent and household characteristics that are considered in the construction of the social distance metrics.

The coefficient β captures the effect¹⁰ of social distance on knowledge when the agent is not incentivised. The coefficient γ captures the effect of incentive pay for socially proximate (non-distant) agent-household pairs. Finally, δ captures the differential effect of incentive pay for socially distant agent-household pairs relative to socially proximate ones.

The results are presented in Table 7. Column 1 confirms that incentive-pay agents have a significant and positive impact on knowledge compared to flat-pay agents, even when controlling for agent and household caste, education, ration-card status and home-ownership as well as taluk and time fixed effects.

Column 2 presents results for the composite social distance metric. The un-interacted treatment effect is not significant, and the coefficients on social distance and the interaction of incentive pay with social distance are both highly significant and roughly opposite in magnitude. We interpret this in three steps: First, it confirms that social distance has a negative impact on knowledge transmission. Second, putting agents on an incentive-pay contract has a positive effect on knowledge transmission, but *only for socially distant agent/household pairs*. And third, the effect of providing financial incentives (at our level of bonus pay) is more or less exactly the level required to cancel out the negative effect due to social distance. In other words, the effect of incentive pay seems to be to cancel the negative effect of social distance, but no more.

In Indian villages, caste groups sometimes live in distinct sub-villages called hamlets. This means that the social distance between a pair of households in the village may be positively correlated with the physical distance between them. To the extent that this is the case, it is possible that the results so far confound the effect on knowledge transmission of social distance with that of physical distance. After all, it seems natural that the cost of knowledge transmission increases with the physical distance between the agent and a household.

While we do not have good measures of physical distance at the household level, a rough test can be constructed for a subset of villages for which we know whether caste groups tend to live apart or not. This information is available for 107 out of the 147 villages. Based on this information, a binary indicator is constructed which is equal to 1 if, in a given village, there is a walk of 5 minutes or more between the

¹⁰The words ‘effect’ and ‘impact’ are used for ease of exposition, but we cannot make the same claims of causality in this part of the analysis since social characteristics were not randomly allocated.

settlements of the major caste groups, and 0 otherwise. This indicator is 1 for 26 out of 107 villages. In column 3, this indicator and its interaction with the incentive-pay variable are included in the regression.

While the sample size drops, the results, in column 3, confirm that physical separation does have a negative effect on knowledge transmission and that this effect, like the one for social distance, is completely counter-acted by the introduction of incentive pay. But the results also show that the social distance indicator and its interaction with incentive-pay is still significant and qualitatively unchanged. (If anything the magnitudes have increased.) While the measure of physical distance is crude, and, therefore, it is conceivable that the the social-distance measure still captures some physical-distance effect, the fact that the co-efficients on the social-distance measure and its interaction with the treatment variable are unchanged suggests that this is not the case. Social distance matters, even after controlling for physical distance.

Columns 4–7 repeat the exercise of column 2 for each of the sub-component distance metrics. For distance in caste group, ration-card status and home ownership, the story appears to align with the findings for the overall distant metric presented above. However, for education, there appears to be no significant disadvantage due to social distance. In other words, agent-household communication appears to be hampered by differences in caste, poverty status and home ownership, but not by differences in education. Correspondingly, in this specification, un-interacted incentive pay has a large and positive co-efficient, though it is not statistically significant.

It is of interest whether the impact of social distance and its interaction with incentive pay is symmetric across the caste hierarchy. In other words, is the impact of social distance between agent and beneficiary household more severe when a lower-caste agent interacts with higher-caste households than vice versa? To test this, we compute differences-in-differences in mean effects by the agent caste group. Table 8 presents the mean comparisons and confirms that the qualitative findings are indeed symmetric: irrespective of whether the agent is of a forward/dominant caste or not, non-incentivised agents are significantly more effective at transmitting knowledge to her own group than to the cross-group, while for incentivised agents there is no significant difference in performance between the two groups. Furthermore, this is due to the effect on cross-group households responding positively to incentive pay while the effect on the own-group households does not respond significantly to incentive pay.

4.5 Relating the Empirical Results to the Theoretical Model

Let $e_s(b)$ denote the effort of a knowledge agent when dealing with her own social group, and let $e_o(b)$ denote the effort with respect to dealing with the other group. We observe four points empirically: $e_s(0), e_s(b'), e_o(0)$ and $e_o(b')$. The empirical findings can be summarised as follows:

1. $e_s(0) > e_o(0)$.
2. $e_s(0) = e_s(b') = e_o(b')$.

What this suggests is that with respect to their own group, agents were already choosing the maximum effort, and, therefore, bonus pay induces no additional effort. With respect to the other group, the agents were choosing a sub-maximal effort level without bonus, but with bonus pay the effort goes up to the maximum level.¹¹ We do not observe crowding out, but we cannot rule it out outside of the observed parameter values. Specifically, if we increased b enough, effort with respect to the matched group could decrease. From the four points we observe, we cannot tell whether or not we are in a crowding-out world.

Though the empirical reality seems to conform to the model's assumption of an upper limit to agent effort, the theory does not explain why such an upper limit should exist in the first place. One possibility is that households 'max out' on the knowledge tests, thereby creating an upper bound for agent performance. If households are attaining the maximum score, then any further effort would be unobservable and also, from the incentivised agents' point of view, futile. However, a quick look at the distribution of test scores reveals that the households are generally nowhere near the level of test scores where such saturation could become important.

Another, and in our view more likely, possibility is that the upper bound \bar{e} is not imposed by the test or the agent but by the household. The agent might be willing to sit with the households for long periods of time to teach them the intricacies of RSBY, especially if they are incentivised to do so, but households may have limited time or patience for this. Field anecdotes suggest that households think of the agent as a resource that can be used if the need arises: if a household member falls ill or otherwise needs health care, they will turn to the agent and ask her advice. If this perspective is widespread, it would not be surprising if the households' motivation for learning details of the scheme is limited. They only need basic knowledge about the scheme, and for this reason their patience with listening to details will probably

¹¹The two empirical 'facts' could also be explained by a horizontal (inelastic) own-group effort curve. However, in this case the fact that other-group effort increases to more or less the same level as own-group effort with incentive pay would be purely coincidental. This is a low-probability event, and since we observe this effort-level matching for three of our four simple distance metrics, we consider the maxing-out story to be a more likely explanation of our findings.

‘max out’ relatively quickly.

5 Discussion

In this section, we discuss our results in terms of the overall effects and their importance. However, undertaking a full cost-benefit analysis would require assumptions about the social value of knowledge about, or enrolment in, welfare programmes such as RSBY.

To some it may seem that the effect that has been identified here is implausibly large compared to the pay we offered the agents. After all, an average payment of 400 rupees (9 USD) for work over a period of several months is not that much, even for India’s poor. However, whether the job was well-paid or not is a function of the hours put in as well as of the period covered by each payment. While we do not have survey data to back this up, examining our field notes indicates that agents may spend in the region of 4-5 days of full-time work equivalents per payment period. This is a very rough ‘guesstimate’, and clearly there will be substantial variation around the mean, but if it is in the right order of magnitude then it would suggest that the average pay per day of work is around 100 rupees (2 USD), which is of the same order of magnitude as what agricultural labourers earn. 100 rupees per day is also the wage rate that was offered by government’s large-scale public-works programme, the National Rural Employment Guarantee, in Karnataka until recently.

In the experiment, the agents were paid a bonus of 8 rupees (0.18 USD) for each household that answered at least four out of eight knowledge test questions correctly. Since the average effect of incentive-pay is to increase knowledge levels by about 0.25 standard deviations or about 0.6 correctly answered questions on the knowledge test, crude extrapolation would suggest that a bonus of 13 rupees (0.29 USD) per household would increase the average number of correctly answered questions by one.

The reduced-form estimate in column 2 of Table 5 suggests that a 8 rupees (0.18 USD) bonus per household raises the enrolment rate by 8 %-points.

The findings concerning the relative importance of financial incentives and social distance have implications for contexts in which strong own-group bias can lead to adverse welfare effects. In India, caste and religious identities, in particular, have been found to create social divisions that impede the efficient functioning of markets (Anderson, 2011) and access to public goods (Banerjee, Iyer and Somanathan, 2005; Banerjee and Somanathan, 2007). This creates inequities and perpetuates disadvantage. In this context, our findings that incentive-pay may reduce own-group bias

may have implications on inequality and welfare.

It is interesting that even a modest amount of bonus pay (in our case, 8 rupees or 0.18 USD per household) to agents appears to completely wipe out the knowledge gap between beneficiary groups that are socially proximate to the agent and those that are not.

It would be hasty to extrapolate our findings from the current context of information transmission about welfare schemes to the wider societal effects of own-group bias, but our results do suggest that the interaction of social distance and financial distance may be a fruitful topic for future research.

The results may also indicate that agents are more sensitive to the fact that there is an incentive rather than the *size* of the incentive. This finding corresponds to results obtained in other recent work on conditional cash transfers, where the size of transfer was not found to matter beyond the fact that there is a positive transfer (Filmer and Schady, 2009) as well as in the context of preventive health behaviour, where uptake of such initiatives has been found to be sensitive to small incentives (Thornton, 2008; Banerjee et al., 2010*a*).

6 Conclusion

This paper sheds light on the role of financial incentives in motivating local agents to spread information regarding a public health insurance programme among the target beneficiaries. The results suggest that, first, hiring agents to spread knowledge about welfare programmes has a positive impact on the level of programme knowledge, but that the entire effect is driven by agents on incentive-pay contracts. Second, using the random assignment of our experimental treatment as an instrument for knowledge, we find that the improved knowledge in turn increases programme take-up. An increase of one standard deviation in knowledge score increases the likelihood of take-up by 39 %-points. Third, we find that social distance between agent and beneficiary has a negative impact on knowledge transmission, but putting agents on incentive-pay contracts increases knowledge transmission by cancelling out (at our level of bonus pay) the negative effect of social distance. On the other hand, incentive pay has no impact on knowledge transmission for socially proximate agent-beneficiary pairs.

Our results have broad implications for public service delivery in the context of developing countries where, aside from common supply-side problems like staff absenteeism, corruption and red tape; a lack of awareness and knowledge regarding available welfare schemes represents a demand-side problem that may short-circuit

the process of lifting the masses out of poverty. The experimental evidence from presented here points to a key mechanism that can address this problem.

In future work we hope to investigate the impact of spreading information about the health insurance scheme on potential health outcomes of the beneficiaries. Although utilisation of the scheme has been low so far, there are indications that it might pick up over time, and we hope to capture this in future follow-up surveys of our sample villages. In particular, we would like to know whether incentivising information dissemination, which was found to improve knowledge and enrolment of the programme, also leads to improved health outcomes for the beneficiaries.

References

- Aizer, Anna.** 2007. “Public Health Insurance, Program Take-up and Child Health.” *Review of Economics and Statistics*, 89(3): 400–415.
- Akerlof, George, and Rachel Kranton.** 2005. “Identity and the Economics of Organizations.” *Journal of Economic Perspectives*, 19(1): 9–32.
- Anderson, Siwan.** 2011. “Caste as an Impediment to Trade.” *American Economic Journal: Applied Economics*, 3(1): 239–263.
- Ashraf, Nava, Oriana Bandiera, and Kelsey Jack.** 2012. “No Margin, No Mission? A Field Experiment on Incentives for Pro-Social Tasks.” *Economic Organisation and Public Policy Discussion Papers*, 035.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul.** 2011. “Field Experiments with Firms.” *Journal of Economic Perspectives*, 25(3): 63–82.
- Banerjee, Abhijit, and Kaivan Munshi.** 2004. “How Efficiently is Capital Allocated: Evidence from the Knitted Garment Industry in Tirupur.” *Review of Economic Studies*, 71(1): 19–42.
- Banerjee, Abhijit, and Rohini Somanathan.** 2007. “The Political Economy of Public Goods: Some Evidence from India.” *Journal of Development Economics*, 82(2): 287–314.
- Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Dhruva Kothari.** 2010*a*. “Improving Immunisation Coverage in Rural India: Clustered Randomised Controlled Evaluation of Immunisation Campaigns with and without Incentives.” *British Medical Journal*, 340:c2220.

- Banerjee, Abhijit, Lakshmi Iyer, and Rohini Somanathan.** 2005. "History, Social Divisions, and Public Goods in Rural India." *Journal of the European Economic Association*, 3(2-3): 639–647.
- Banerjee, Abhijit, Rachel Glennester, and Esther Duflo.** 2008. "Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System." *Journal of European Economic Association*, 6(2-3): 487–500.
- Banerjee, Abhijit, Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani.** 2010*b*. "Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India." *American Economic Journal: Economic Policy*, 2(1): 1–30.
- Benabou, Roland, and Jean Tirole.** 2003. "Intrinsic and Extrinsic Motivation." *Review of Economic Studies*, 70(3): 489–520.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan.** 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119(1): 249–275.
- Besley, Timothy, and Maitreesh Ghatak.** 2005. "Competition and Incentives with Motivated Agents." *American Economic Review*, 95(3): 616–637.
- Cole, Shawn, Xavier Gine, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery.** 2010. "Barriers to Household Risk Management: Evidence from India." *American Economic Journal: Applied Economics*. forthcoming.
- Daponte, Beth, Seth Sanders, and Lowell Taylor.** 1999. "Why do Low-Income Households not Use Food Stamps? Evidence from an Experiment." *Journal of Human Resources*, 34(3): 612–628.
- Duflo, Esther, and Emmanuel Saez.** 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics*, 118(3): 815–842.
- Duflo, Esther, Rema Hanna, and Stephen P. Ryan.** 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review*, 102(4): 1241–1278.
- Dupas, Pascaline.** 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics*, 3(1): 1–36.

- Filmer, Deon, and Norbert Schady.** 2009. "Are There Diminishing Returns to Transfer Size in Conditional Cash Transfers?" *World Bank Policy Research Working Paper Series*, 4999.
- Glewwe, Paul, Alaka Holla, and Michael Kremer.** 2009. "Teacher Incentives in the Developing World." In *Performance Incentives: Their Growing Impact on K-12 Education*, ed. Matthew Springer. Washington, DC: Brookings Institution Press.
- Gneezy, Uri, and Aldo Rustichini.** 2000. "Pay Enough or Don't Pay At All." *Quarterly Journal of Economics*, 115(3): 791–810.
- Government of India.** 2009. *Tender document: Implementation of Rashtriya Swasthya Bima Yojana: Template for use by state governments*. Downloaded from www.rsby.gov.in.
- Hernanz, Virginia, Franck Malherbet, and Michele Pellizzari.** 2004. "Take-Up of Welfare Benefits in OECD Countries: A Review of the Evidence." *OECD Social, Employment and Migration Working Papers*, 17.
- Holmstrom, Bengt, and Paul Milgrom.** 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization*, 7(Special issue): 24–52.
- Jacob, Brian, and Steven Levitt.** 2003. "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating." *Quarterly Journal of Economics*, 118(3): 843–877.
- Jensen, Robert.** 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics*, 125(2): 515–548.
- Kingdon, Geeta, and Shenila Rawal.** 2010. "Akin to My Teacher: Does Caste, Religious and Gender Distance Between Student and Teacher Matter? Some Evidence from India." *Department of Quantitative Social Science Working Papers*, 10-18.
- Kremer, Michael, and Alaka Holla.** 2008. "Pricing and Access: Lessons from Randomized Evaluation in Education and Health." In *What Works in Development: Thinking Big and Thinking Small*, ed. William Easterly and Jessica Cohen. Washington, DC: Brookings Institution Press.

- Mellstrom, Carl, and Magnus Johannesson.** 2008. "Crowding out in Blood Donation: Was Titmus Right?" *Journal of the European Economic Association*, 6(4): 845–863.
- Muralidharan, Karthik, and Venkatesh Sundararaman.** 2011. "Teacher Performance Pay: Experimental Evidence from India." *Journal of Political Economy*, 119(1): 39–77.
- Neal, Derek, and Diane Schanzenbach.** 2010. "Left Behind by Design: Proficiency Counts and Test-based Accountability." *Review of Economics and Statistics*, 92(2): 263–283.
- Rajasekhar, D, Erlend Berg, Maitreesh Ghatak, R Manjula, and Sanchari Roy.** 2011. "Implementing Health Insurance: The Roll-out of Rashtriya Swasthya Bima Yojana in Karnataka." *Economic and Political Weekly*, XLVI(20): 56–63.
- Thornton, Rebecca.** 2008. "The Demand for, and Impact of, Learning HIV Status: Evidence from a Field Experiment." *American Economic Review*, 98(5): 1829–1863.
- World Bank.** 2011. *Social Protection for a Changing India: Volume 2.*

Appendix: The Knowledge Tests

In the **first survey**, the knowledge test consisted of the following eight questions (correct answers in italics):

1. Does the programme cover the cost of treatment received while admitted to a hospital (hospitalisation)?
Yes.
2. Does the programme cover the cost of treatment received while not admitted to a hospital (out-patient treatment)?
No.
3. Who can join this programme?
Households designated as being Below the Poverty Line. (Those who said 'the poor', 'low income' or similar were marked as correct.)
4. What is the maximal annual expenditure covered by the scheme?
30,000 rupees.
5. How much money do you have to pay to get enrolled in the scheme?
30 rupees per year.
6. How many members of a household can be a part of the scheme?
Up to five.
7. What is the allowance per visit towards transportation to the hospital that you are entitled to under the RSBY scheme?
100 rupees. (This was the expected answer, although strictly speaking the transportation allowance is subject to a maximum of 1000 rupees per year, i.e. ten visits.)
8. Is there an upper age limit for being covered by the scheme? If yes, what is it?
There is no upper age limit.

In the **second survey**, the knowledge test consisted of the following eight questions:

1. What is the maximum insurance cover provided by RSBY per annum?
30,000 rupees.

2. Does the beneficiary have to bear the cost of hospitalisation under the RSBY scheme up to the maximum limit?
No.
3. Are pre-existing diseases covered under RSBY?
Yes.
4. Are out-patient services covered under RSBY?
No.
5. Are day surgeries covered under RSBY?
Yes.
6. Does the scheme cover post-hospitalisation charges? If yes, up to how many days?
Yes, up to 5 days. (Anyone who answered 'yes' was marked as correct.)
7. Are maternity benefits covered?
Yes.
8. If a female RSBY member has given birth to a baby during the policy period, will the baby be covered under RSBY?
Yes.

In the **third survey**, the knowledge test consisted of the following eight questions:

1. Under RSBY, how many family members can be enrolled in the scheme?
Five.
2. What is the maximum insurance cover provided by RSBY per policy period?
30,000 rupees.
3. If hospitalised, does an RSBY cardholder have to pay separately for his/her medicines?
No.
4. If hospitalised, does an RSBY cardholder have to pay separately for his/her diagnostic tests?
No.
5. Is it compulsory for an RSBY cardholder to carry the smart card while visiting the hospital for treatment?
Yes.

6. If an RSBY cardholder is examined by a doctor for a health problem but not admitted to the hospital, will the treatment cost be covered under RSBY?

No.

7. What is the duration/tenure of the RSBY policy period?

1 year.

8. How can an RSBY cardholder check if a particular health condition is covered under RSBY prior to visiting the hospital for treatment?

Multiple correct answers, see text.

Table 1: Agent Summary Statistics

	Flat pay	Incentive pay
Agent age	34.76 (8.808)	34.77 (8.084)
Agent is married	0.811 (0.397)	0.915 (0.280)
Agent is of forward/dominant caste	0.432 (0.502)	0.352 (0.481)
Agent's household head has completed primary school	0.622 (0.492)	0.563 (0.499)
Agent household has ration card	0.892 (0.315)	0.789 (0.411)
Agent owns her home	0.865 (0.347)	0.873 (0.335)
Agent is Self-Help Group president	0.297 (0.463)	0.282 (0.453)
Agent autonomy score (the higher, the more autonomous)	5.568 (0.929)	5.676 (0.841)
Agent pay in round 1	400 (0)	507.7 (478.5)
Agent pay in round 2	400 (0)	403.0 (209.1)
Observations	37	71

Note: Standard deviations in parentheses.

Table 2: Household Summary Statistics

	Control villages	Flat-pay villages	Incentive-pay villages
Household is of forward/dominant caste	0.251 (0.434)	0.192 (0.395)	0.158 (0.365)
Household head has completed primary school	0.301 (0.459)	0.234 (0.424)	0.284 (0.451)
Household has ration card	0.936 (0.245)	0.931 (0.254)	0.918 (0.274)
Household owns its home	0.667 (0.472)	0.733 (0.443)	0.758 (0.428)
Observations	375	479	869

Note: Standard deviations in parentheses.

Table 3: The Effect of Knowledge Agents

	(1)	(2)	(3)
	Knowledge	Knowledge	Knowledge
Agent in village	0.175*** (0.0645)	0.187*** (0.0572)	
Flat-pay agent in village			0.0722 (0.0919)
Incentive-pay agent in village			0.246*** (0.0569)
Time fixed effects	No	Yes	Yes
Taluk fixed effects	No	Yes	Yes
Observations	5641	5641	5641
t-test: flat=incentivised (p-value)			0.0600

Notes: Weighted least squares regressions. Each household is given the same weight, divided equally between all observations of that household. Standard errors, in parentheses, are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.01

Table 4: The Effect of Knowledge Agents, Shimoga District Only

	(1)	(2)	(3)
	Knowledge	Knowledge	Knowledge
Agent in village	0.210** (0.0823)	0.191** (0.0743)	
Flat-pay agent in village			-0.0289 (0.121)
Incentive-pay agent in village			0.317*** (0.0683)
Time fixed effects	No	Yes	Yes
Taluk fixed effects	No	Yes	Yes
Observations	2885	2885	2885
t-test: flat=incentivised (p-value)			0.00700

Notes: Weighted least squares regressions. Each household is given the same weight, divided equally between all observations of that household. Standard errors, in parentheses, are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.01

Table 5: Knowledge Drives Enrolment: Two-Stage Least Squares

	(1)	(2)	(3)	(4)
	Enrolled (OLS)	Enrolled (Reduced form)	Knowledge (First stage)	Enrolled (IV)
Knowledge	0.206*** (0.00910)			0.390*** (0.128)
Incentive-pay agent in village		0.0816** (0.0362)	0.209*** (0.0618)	
Time fixed effects	Yes	Yes	Yes	Yes
Taluk fixed effects	Yes	Yes	Yes	Yes
Observations	5641	5641	5641	5641
t-test: flat=incentivised (p-value)				

Notes: Weighted least squares regressions. Each household is given the same weight, divided equally between all observations of that household. Standard errors, in parentheses, are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.01

Table 6: The Effect of Knowledge Agents, by Agent Characteristics

	(1) Knowledge	(2) Knowledge
Treatment (agent in village)	0.187*** (0.0572)	-0.499 (0.331)
Treatment x agent is 30+		0.0453 (0.0916)
Treatment x agent is 50+		-0.0826 (0.0938)
Treatment x agent of forward/dominant caste		-0.102 (0.0894)
Treatment x agent household head has completed primary school		-0.105 (0.0931)
Treatment x agent has ration card		-0.0642 (0.123)
Treatment x agent owns her home		0.148 (0.108)
Treatment x agent is Self-Help Group president		0.00918 (0.0865)
Treatment x agent autonomy		0.121** (0.0466)
Time fixed effects	Yes	Yes
Taluk fixed effects	Yes	Yes
Observations	5641	5641

Notes: Weighted least squares regressions. Each household is given the same weight, divided equally between all observations of that household. Standard errors, in parentheses, are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.01

Table 7: Incentives and Social Distance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge
Incentive pay	0.160* (0.0923)	-0.128 (0.139)	-0.254 (0.162)	0.00714 (0.114)	0.188 (0.118)	0.0851 (0.0943)	0.0758 (0.110)
Social distance		-0.655*** (0.214)	-0.766*** (0.249)	-0.384*** (0.100)	0.0887 (0.0943)	-0.274** (0.133)	-0.210* (0.120)
Incentive pay x social distance		0.794*** (0.217)	0.867*** (0.260)	0.366*** (0.129)	-0.0507 (0.109)	0.376** (0.144)	0.268** (0.120)
Castes live apart			-0.384** (0.182)				
Incentive pay x castes live apart			0.392* (0.200)				
Village size in thousands	0.196 (0.154)	0.205 (0.158)	0.288* (0.158)	0.212 (0.146)	0.193 (0.153)	0.171 (0.164)	0.211 (0.153)
Agent and household characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time and taluk fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Social distance metric	-	Composite	Composite	Caste only	Education only	Ration-card status only	Home ownership only
Observations	2900	2900	2327	2900	2900	2900	2900

Notes: Weighted least squares regressions. Each household is given the same weight, divided equally between all observations of that household. Standard errors, in parentheses, are clustered at the village level. In all columns, agent and household characteristics are binary indicators for whether the agent and household are of forward/dominant caste, whether the head has completed primary school, have a ration card and own their home.

The simple social distance metrics in columns 4–7 are binary variables equal to the absolute difference between the corresponding household and agent characteristic binaries. The composite social distance metric in columns 2 and 3 is the sum of the binary distance metrics for caste, education, ration-card status and home ownership.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Mean Comparisons by Caste Category

Dominant-caste agent			
	Household		
	Dominant	Non-dominant	Difference
Flat pay	0.22 (0.1)	-0.17 (0.05)	0.39*** (0.11)
Incentive pay	0.12 (0.09)	0.11 (0.04)	0.01 (0.09)
Difference	-0.1 (0.14)	0.28** (0.06)	-0.38*** (0.15)
Non-dominant-caste agent			
	Household		
	Dominant	Non-dominant	Difference
Flat pay	-0.2 (0.09)	0.09 (0.05)	-0.29** (0.11)
Incentive pay	0.08 (0.08)	0.17 (0.03)	-0.09 (0.08)
Difference	0.28** (0.12)	0.08 (0.06)	0.2 (0.14)