

Motivating Knowledge Agents: Can Incentive Pay Overcome Social Distance?*

Short title: Motivating Knowledge Agents

Erlend Berg, Maitreesh Ghatak, R Manjula, D Rajasekhar and
Sanchari Roy

11 July 2017

Abstract

This paper studies the interaction of incentive pay with intrinsic motivation and social distance. We analyse theoretically as well as empirically the effect of incentive pay when agents have pro-social objectives, but also preferences over dealing with one social group relative to another. In a randomised field experiment undertaken across 151 villages in South India, local agents were hired to spread information about a public health insurance programme. In the absence of incentive pay, social distance impedes the flow of information. Incentive pay increases overall agent effort and appears to cancel the negative effects of social distance.

JEL Codes: C93, D83, I38, J32, O15

Key words: public services, incentive pay, social distance, knowledge transmission, information flow, health insurance

*We thank the Editor, Frederic Vermulen, and three anonymous referees for detailed and helpful comments. We also thank Oriana Bandiera, Arunish Chawla, Clare Leaver, Gerard Padro i Miquel, Vijayendra Rao, E. Somanathan and seminar participants at Bristol, IGC Growth Week 2011, IGC Patna meeting 2011, ISI Delhi, JNU, LSE, Oxford and Warwick for helpful feedback. We gratefully acknowledge financial support from Improving Institutions for Pro-Poor Growth (iiG), a research programme consortium funded by DFID, as well as from the International Growth Centre. Any views expressed are the authors' own and not necessarily shared by the funders. Correspondence: Erlend Berg, University of Bristol, Department of Economics, Priory Road Complex, Bristol BS8 1TU, United Kingdom. Email: erlend.berg@bristol.ac.uk

Economists tend to believe in the power of incentives and prices to improve efficiency, whether the aim is to motivate workers or eliminate social ills such as discrimination.¹ Yet both theory and evidence suggest that there are circumstances in which there are grounds for caution: First, if there are multiple tasks or output is hard to measure, financial incentives may have undesirable consequences (Holmstrom and Milgrom, 1991; Gneezy *et al.*, 2011). Second, in jobs with an aspect of social service, as in public goods provision, or if reputation matters, workers may not be ‘in it just for the money’. It has been argued that financial incentives may interfere with or even ‘crowd out’ such intrinsic motivation (Bénabou and Tirole, 2006; Gneezy and Rustichini, 2000b). Third, theory suggests that aligning the identities of economic agents can increase efficiency (Besley and Ghatak, 2005; Francois, 2000), and there is evidence that ethnic fragmentation and ‘social distance’ can lead to worse economic outcomes (Easterly and Levine, 1997). Akerlof and Kranton (2005) argue that when group identity is salient, monetary incentives can be ‘both costly and ineffective’.

However, evidence on the effect of incentive pay on performance in pro-social tasks is still limited. Ashraf *et al.* (2014) find that both non-financial and financial rewards have stronger effects for socially motivated agents. Dal Bó *et al.* (2013) conclude that higher wages do not have adverse selection effects in terms of public-service motivation. Rasul and Rogger (2016) suggest that the use of incentives can negatively affect aspects of performance in the Nigerian Civil Service.²

Moreover, very little is known about the interaction of incentive pay and social distance, though the literature on discrimination has suggested that competitive markets can remove the effects of social distance where these cause inefficiencies (Lang and Lehmann, 2012).

In this paper we develop a theoretical model and provide empirical evidence on the role of incentive pay in spreading information about a public service in a socially heterogeneous population. We study whether incentive pay is effective in settings where output is noisy

¹Bandiera *et al.* (2011) and Prendergast (2015) review the evidence. Becker and Posner (2009) and Sandel (2012) provide different perspectives.

²Finan *et al.* (forthcoming) survey recent evidence on the role of incentives in the public sector.

and crowding out is a possibility, and whether incentive pay ameliorates or exacerbates the potentially detrimental effects of social distance.

A simple theoretical framework is developed which combines elements of a motivated-agent framework (Besley and Ghatak, 2005) with the multi-tasking model (Holmstrom and Milgrom, 1991). The framework predicts that when there is a single task and the agent is intrinsically motivated, effort is always weakly increasing in the part of the agent's compensation that is dependent on success (the 'bonus'). But when there are two tasks, which differ in terms of the agent's intrinsic motivation to succeed and in the marginal cost of effort, the effect of bonus pay will depend in part on the degree of substitutability in the cost of effort across the two tasks. If substitutability is low, increasing bonus pay will lead to an increase in the agent's effort with respect to both tasks. But if the two tasks are relatively substitutable in the cost function, an increase in bonus may cause effort in one task to decrease while effort in the other increases. This can be interpreted as incentive pay 'crowding out' intrinsic motivation for one of the tasks.

We then analyse data from a field experiment conducted across 151 villages in Karnataka, India, in the context of a government-subsidised health insurance scheme aimed at the rural poor. In a random sub-sample of the villages (the treatment groups), one local woman per village was recruited to spread information about the scheme. These 'knowledge agents' were randomly assigned to either a flat-pay or an incentive-pay contract. Under the latter contract, the agents' pay depended on how a random sample of eligible households in their village performed when surveyed and orally presented with a knowledge test about the scheme.

Our main empirical findings are as follows: First, hiring agents to spread information has a positive impact on the level of knowledge about the programme. The effect is driven by agents on incentive-pay contracts. Households in villages assigned an incentive-pay agent score on average 0.25 standard deviations higher on the knowledge test than those in the control group, and are also 8 percentage points more likely to enrol.

Second, social distance between agent and beneficiary has a negative impact on knowledge transmission. But putting agents on incentive-pay contracts appears to increase

knowledge transmission by cancelling (at our level of bonus pay) the negative effect of social distance. By contrast, incentive pay has no impact on knowledge transmission or enrolment for socially proximate agent–beneficiary pairs. This result appears to be symmetric across social boundaries, in the sense that it holds whether the agent is from a high- or low-status caste group. Our preferred interpretation is that, with respect to their ‘own’ group (socially proximate households), agents were already at a maximum effort level and hence, introducing bonus pay has no impact. However, non-incentivised agents choose a lower level of effort with respect to the ‘other’ group (socially distant households). With incentive pay, effort goes up to the same level as for the agent’s own group. One might say that incentives appear to ‘price out prejudice’, although social distance barriers can operate through channels other than prejudice. We do not observe crowding out empirically, but cannot rule it out for unobserved parameter values.

Third, incentivised agents appear to achieve higher knowledge scores by reallocating time away from socially proximate households (their ‘own group’) towards socially distant households (their ‘cross-group’), without increasing aggregate time spent. The findings are consistent with a story in which non-incentivised agents spend more time than needed with their ‘own group’ because it is enjoyable rather than productive (‘idle chatter’). Incentivised agents channel some of this time toward productive use with households in the ‘cross-group’.

The paper makes three main contributions. First, to the best of our knowledge, it presents the first randomised evaluation of incentive pay for agents tasked with providing information about a public service. An important aspect of service delivery is to make intended beneficiaries aware of their entitlements. Even if there were no supply-side problems—if the quality of schools and health centres were excellent and these facilities were widely available—the outcome would be disappointing if beneficiaries were unaware of the services or did not value them sufficiently (due to, say, a lack of information or present bias). While this is a recognised problem in rich countries,³ the issue has not

³Information costs are often argued to be one of the main reasons for low take-up of welfare programmes in developed countries (see Hernanz *et al.* (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 *et al.* (1999) find that randomly allocating information about the Food Stamp

received much attention in developing countries. There is, however, reason to believe that the problem is no less important there: a recent report on public services in India shows programme awareness to be low among target groups (World Bank, 2011).⁴ It is thus important to understand the role of incentives in raising awareness of social programmes, a context in which pro-social motivation is likely to feature.

Second, we contribute to the broader literature on financial incentives and performance by showing that incentives can matter, even in the context of a pro-social task, a soft objective and agents with possible intrinsic motivation. As Finan *et al.* (forthcoming) point out, we do not know enough about the effect of financial incentives in these settings; in particular, when incentives may cause agents to prioritise dimensions that are easy to measure to the detriment of those that are less so, or when incentives could crowd out intrinsic motivation.

Third, the paper extends our understanding of the interaction between incentive pay and social distance. While we are not the first to document the detrimental effects of social barriers, the question of whether incentive pay alleviates or exacerbates the negative consequences of social distance has not received much attention. This is particularly important in developing countries, many of which are highly stratified along socio-economic lines. The novelty of our findings is that what is ‘crowded out’ is an anti-social tendency to favour interactions with one’s own group, whereas most previous studies have focused on financial incentives crowding out pro-social tendencies, such as picking up one’s children on time from day care (Gneezy and Rustichini, 2000a).

Our paper is related to that of Bandiera *et al.* (2009), who study the interplay of social connections and financial incentives in the context of worker productivity in a private firm in the United Kingdom. They find that when managers are paid fixed wages, they favour workers with whom they are socially connected; but when incentive pay is introduced, managers’ efforts do not depend on social connections. But as Bandiera *et al.* (2011)

Program significantly increases participation amongst eligible households.

⁴According 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.

point out, provision of incentives for pro-social tasks raise different issues compared to private tasks for several reasons, including the possibility of crowding out.

Another paper looking at the effects of social distance in a for-profit setting is Fisman *et al.* (forthcoming). They analyse data from an Indian bank, and find that the volume of credit is larger, and repayment rates higher, when the borrower and the loan officer are matched on social identity. While they are able to exploit quasi-random variation in social distance, they do not study the interaction of social distance and pay.

There is a growing literature on the importance of information campaigns in economic decision-making and, in particular, in determining demand for public services. Previous work has explored how information campaigns affect local participation and educational outcomes in India (Banerjee *et al.*, 2010a), how providing information on measured returns increases years of schooling (Jensen, 2010) and how creating awareness about HIV prevalence reduces incidence of risky sexual behaviour among Kenyan girls (Dupas, 2011).⁵

There is substantial evidence that ethnic heterogeneity is linked to poor economic outcomes, including sub-optimal provision of public goods and poor governance (Easterly and Levine, 1997; La Porta *et al.*, 1999; Kimenyi, 2006). A possible explanation for this is that people prefer to interact with those who are similar to themselves, leading to fragmented markets, lower social mobility (Bertrand *et al.*, 2000) and reduced gains from trade (Anderson, 2011). Several studies find evidence of strong own-group bias (Banerjee and Munshi, 2004; Kingdon and Rawal, 2010), with potentially adverse implications for the flow of information. In the context of awareness campaigns, if people prefer to liaise with their own kind, information constraints on the demand for public services may be more severe in socially heterogeneous settings. However, micro-level evidence on the role of social distance in the spreading of awareness about public services is rare.

This paper is also related to the rich literature on the impact of monetary and non-monetary incentives on the performance of agents. This body of work encompasses studies in the ‘standard setting’ of firms in developed countries where output or productivity is

⁵In the context of the government health insurance scheme studied here, Das and Leino (2011) analyse an information campaign in North India and get mixed results.

measurable but worker effort is not (Lazear, 2000), as well as papers on incentives for teachers and health workers in developing countries as surveyed by Kremer and Holla (2008) and Glewwe *et al.* (2009). Muralidharan and Sundararaman (2011) and Duflo *et al.* (2012) study the effect of financial incentives for teachers on absenteeism and test scores, while BenYishay and Mobarak (2014) look at the role of incentives in leveraging peer learning to promote the adoption of agricultural technologies in Malawi.

There are also studies 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; Bénabou and Tirole, 2003; Besley and Ghatak, 2005). In a laboratory setting, Gneezy and Rustichini (2000b) find non-monotonicities in the effect of incentive pay on effort. As Bandiera *et al.* (2011) point out, there is little field-experimental evidence in this area, although Ashraf *et al.* (2014) is a recent exception.

The rest of the paper is organised as follows: In Section 1, a simple theoretical framework is presented with the aim of analysing the impact of incentive pay on agents' effort and its interaction with social-identity matching. Section 2 describes the context, experimental design and data. Section 3 presents the empirical evidence and Section 4 concludes.

1 Theoretical Framework

In this section we develop a simple model of motivated agents. It extends Besley and Ghatak (2005) by incorporating features of the multi-tasking model (Holmstrom and Milgrom, 1991). The aim is to provide a theoretical framework that can generate predictions about the effects of incentive pay and how these might interact with the effects of social distance.

Suppose agents exert unobservable effort in spreading awareness of a scheme to potential beneficiaries. The goal may be either the transmission of knowledge itself or 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 an exogenously fixed number of target households.

1.1 *A Single Task*

First, assume there is a single task. This could correspond to a situation in which the potential beneficiaries of the public service are relatively homogeneous. Let e be the unobservable effort exerted by the agent. Let the outcome variable Y be binary, with the value 0 denoting ‘bad performance’ or ‘failure’, and the value 1 denoting ‘good performance’ or ‘success’. For example, a household doing well in the knowledge test (say, scoring above a certain threshold level), or enrolling in the programme, might be considered a success.

Agent effort stochastically improves the likelihood of a good outcome. 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. Let us further assume that the lowest value e can take is $\underline{e} \in (0, 1)$, and the highest value e can take is $\bar{e} \in (\underline{e}, 1)$. This means that there is some minimum effort that 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, and even at that level, the good outcome is not guaranteed to occur. Therefore, as is standard in agency models, there is common support. That is, either outcome (0 or 1) is consistent with any level 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 π incorporates both the direct benefit to the beneficiaries and how the rest of society values their welfare. With perfect enforcement, the problem is

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

subject to $e \in [\underline{e}, \bar{e}]$. The solution is

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

It should be noted that the effect of θ and c on e are similar though opposite in sign: an agent puts in more effort when the disutility of effort decreases or the non-pecuniary payoff from success increases. This makes it hard to distinguish between the two empirically.

If effort is contractible, the principal can simply stipulate e^{**} . For the problem to be interesting, and for incentive pay to have an effect, assume that 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, beyond relying on her intrinsic motivation θ , is to pay her a bonus that is contingent on performance. When setting the bonus, the principal has to respect the limited-liability constraint (LLC) and the incentive-compatibility constraint (ICC). There is also a participation constraint (PC) which requires the agent's expected pay-off to be at least as high as her outside option. To keep things simple, it is assumed that the outside option is relatively unattractive so that the PC does not bind—the analysis would be qualitatively unchanged if this assumption were 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

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

subject to $e \in [\underline{e}, \bar{e}]$, which yields

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

This is the ICC. Since $b \leq \pi$, effort will, in general, be lower than in the first-best scenario.

This can be show formally as follows. The principal's objective is⁶

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

subject to the ICC (1), the LLCs $\bar{w} \geq \underline{w}$ and $\underline{w} \geq \underline{\omega}$ and the 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). As 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 .

Experimentally, we only observe outcomes for two given values of b , so the focus here will be on the ICC (1) rather than the optimal bonus. If there is no bonus pay and the agent is not sufficiently intrinsically motivated, we may get a lower corner solution, namely $e = \underline{e}$. This will be the case if $\underline{e} \geq \frac{\theta}{c}$. At the other extreme, if the agent is sufficiently motivated (namely, $\frac{\theta}{c} \geq \bar{e}$), then even without any bonus pay the agent chooses the maximum level of effort, \bar{e} . Otherwise, effort is increasing in bonus pay. The solution is illustrated in Figure 1. The slope of the interior-solution segment ($\frac{1}{c}$) is

⁶In 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. This would lead to higher incentive pay and higher effort.

positive and so is its intercept ($\frac{\theta}{c}$). However, depending on parameter values, the value of e for any given value of b could range from \underline{e} to \bar{e} . For example, the case of a relatively unmotivated agent is captured by the dashed vertical line marked by $c\underline{e} > \theta$. In this case, the vertical axis (at which $b = 0$) intersects the effort curve at a flat section where $e = \underline{e}$. Similarly, a case where the agent is relatively highly motivated is captured by the dashed vertical line marked by $\theta > c\bar{e}$, and an intermediate level of motivation is captured by the line marked $c\underline{e} < \theta < c\bar{e}$. In the former case, the agent is at the minimum effort level for $b = 0$ and initially the marginal effort with respect to bonus pay is zero. As bonus pay increases further, the marginal effort becomes positive, before returning to zero once the effort curve has hit the upper bound. If the vertical axis is at the right-most dashed vertical line, then the agent is already at the maximum effort level when $b = 0$ and effort will be unresponsive to incentive pay at any level. If the vertical axis is at the middle dashed line, effort level is at an interior value when $b = 0$ and the marginal effort with respect to bonus pay is positive.

1.2 Two Tasks

Assume now that the agent has two tasks, as in the multi-tasking model. The tasks may be thought of as the agent exerting effort to transfer knowledge to, or enrol, two different types of beneficiary households. However, unlike in the classic multi-tasking model, the outcomes associated with the two tasks are assumed to be equally measurable. Instead, the differences between the two tasks are in the agent's intrinsic pay-off from success and her cost of 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 payments 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, a knowledge agent may be biased in

favour of some social or economic group or may have purely idiosyncratic biases, but if the principal does not observe the relevant dimension, the remuneration scheme cannot correct for it.

Let \underline{e} and \bar{e} , where $0 < \underline{e} < \bar{e} < 1$, 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 task 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 cost-function substitutability of effort between tasks 1 and 2. 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

$$\max_{e_1, e_2} (\theta_1 + \pi) e_1 + (\theta_2 + \pi) e_2 - \left(\frac{1}{2}c_1e_1^2 + \frac{1}{2}c_2e_2^2 + \gamma e_1e_2 \right).$$

The first-order conditions yield the following interior solutions:

$$\begin{aligned} e_1(\pi) &= \frac{(c_2 - \gamma) \pi + c_2\theta_1 - \gamma\theta_2}{c_1c_2 - \gamma^2} \\ e_2(\pi) &= \frac{(c_1 - \gamma) \pi + c_1\theta_2 - \gamma\theta_1}{c_1c_2 - \gamma^2} \end{aligned}$$

For this to be a local maximum, the second-order condition requires

$$c_1c_2 > \gamma^2.$$

As before, corner solutions may be possible, and if e_i assumes a corner solution, then e_j

($j \neq i$) would take a different form.

Define the pair:

$$\hat{e}_1(\pi) = \begin{cases} \frac{\theta_1 + \pi - \gamma \underline{e}}{c_1} & \text{if } e_2(\pi) \leq \underline{e} \\ e_1(\pi) & \text{if } \underline{e} < e_2(\pi) < \bar{e} \\ \frac{\theta_1 + \pi - \gamma \bar{e}}{c_1} & \text{if } e_2(\pi) \geq \bar{e} \end{cases}$$

$$\hat{e}_2(\pi) = \begin{cases} \frac{\theta_2 + \pi - \gamma \underline{e}}{c_2} & \text{if } e_1(\pi) \leq \underline{e} \\ e_2(\pi) & \text{if } \underline{e} < e_1(\pi) < \bar{e} \\ \frac{\theta_2 + \pi - \gamma \bar{e}}{c_2} & \text{if } e_1(\pi) \geq \bar{e} \end{cases}$$

Now the complete first-best solution for the two-task model is given by:

$$e_1^*(\pi) = \max\{\min\{\hat{e}_1(\pi), \bar{e}\}, \underline{e}\}$$

$$e_2^*(\pi) = \max\{\min\{\hat{e}_2(\pi), \bar{e}\}, \underline{e}\}$$

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:

$$e_1(b) = \frac{(c_2 - \gamma)b + c_2\theta_1 - \gamma\theta_2}{c_1c_2 - \gamma^2}$$

$$e_2(b) = \frac{(c_1 - \gamma)b + c_1\theta_2 - \gamma\theta_1}{c_1c_2 - \gamma^2}$$

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. As in the the first-best case, corner solutions may be possible, and following the same steps as above, we can

derive $\hat{e}_1(b)$ and $\hat{e}_2(b)$:

$$\hat{e}_1(b) = \begin{cases} \frac{\theta_1 + b - \gamma \underline{e}}{c_1} & \text{if } e_2(b) \leq \underline{e} \\ e_1(b) & \text{if } \underline{e} < e_2(b) < \bar{e} \\ \frac{\theta_1 + b - \gamma \bar{e}}{c_1} & \text{if } e_2(b) \geq \bar{e} \end{cases}$$

$$\hat{e}_2(b) = \begin{cases} \frac{\theta_2 + b - \gamma \underline{e}}{c_2} & \text{if } e_1(b) \leq \underline{e} \\ e_2(b) & \text{if } \underline{e} < e_1(b) < \bar{e} \\ \frac{\theta_2 + b - \gamma \bar{e}}{c_2} & \text{if } e_1(b) \geq \bar{e} \end{cases}$$

The complete second-best solution for the two-task model is given by:

$$\tilde{e}_1(\pi) = \max\{\min\{\hat{e}_1(b), \bar{e}\}, \underline{e}\}$$

$$\tilde{e}_2(\pi) = \max\{\min\{\hat{e}_2(b), \bar{e}\}, \underline{e}\}$$

Several aspects of the solution are worth noting. First, effort in the easier task, e_1 , is always weakly increasing in b .

Second, e_2 is also non-decreasing in b , except when both tasks are at internal solutions and $c_1 < \gamma < c_2$, when it is decreasing in b . The intuition for the negative slope is that when effort in the two tasks are relatively substitutable and both effort levels are at internal solutions, providing a monetary incentive leads the agent to substitute effort towards the easier task to a degree that causes effort in the harder task to decrease. We view this as a form of ‘crowding out’ since increasing incentive pay leads the agent to work less in one of the tasks. However, it is not quite crowding out in the sense of Bénabou and Tirole (2006), where the term is taken to imply a decrease in effort overall. In our case, the sum of effort across the two tasks is always weakly increasing in b . This follows trivially from the above except when both efforts are internal. But then

$$e_1(b) + e_2(b) = \frac{(c_1 + c_2 - 2\gamma)b + (c_2 - \gamma)\theta_1 + (c_1 - \gamma)\theta_2}{c_1 c_2 - \gamma^2},$$

and $c_1 + c_2 - 2\gamma > c_1 + c_2 - 2\sqrt{c_1c_2} = (\sqrt{c_1} - \sqrt{c_2})^2 > 0$, where the first inequality follows from the second-order condition, $c_1c_2 > \gamma^2$.

Third, when both effort curves are internal, the slope of e_1 is always greater than the slope of e_2 .

Fourth, the slopes of all internal curves are completely determined by γ , c_1 and c_2 . The role of θ_1 and θ_2 is to shift the intercepts, and hence the lengths and meeting points, of the effort curves' constituent line segments.

Before classifying the types of possible solutions, it is helpful to define the 'intrinsically preferred task' as the task in which the agent exerts the greatest effort when there is no bonus pay, that is, at $b = 0$. Task 1 is the intrinsically preferred task iff $\hat{e}_1(0) > \hat{e}_2(0)$, or

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

Otherwise, task 2 is the intrinsically preferred task. (With equality in the above expression, effort in each task is equal at $b = 0$.) Intuitively, a higher θ_i and a lower c_i both contribute to the agent's intrinsic preference for task i . Note that it is possible that task 2, the harder task, 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.

The main types of solutions can be classified using the relative magnitudes of γ , c_1 and c_2 . Above, it was assumed without loss of generality that $c_1 < c_2$, and the second-order condition requires $c_1c_2 > \gamma^2$. The substitutability parameter γ must therefore be either less than both c_1 and c_2 , or equal to c_1 and less than c_2 , or lie between c_1 and c_2 .

Figures 2–4 illustrate representative cases⁷ where task 1 is intrinsically preferred (effort in task 1 is greater at $b = 0$), and moreover, e_1 is already at the highest possible level \bar{e} but e_2 has an interior solution. The latter corresponds to the condition $\frac{\theta_2 - \gamma\bar{e}}{c_2} < \bar{e} < \frac{c_2\theta_1 - \gamma\theta_2}{c_1c_2 - \gamma^2}$. As in the single-task model, other solutions can be generated by drawing the vertical axis just to the left of the crossing point of the two effort curves, in which case task 2 would be intrinsically preferred. Also illustrated are the 'kinks' in e_2 that arise as e_1 meets the upper or lower bounds.

⁷Appendix A discusses how these relate to the universe of possible cases.

Solutions with $\gamma < c_1 < c_2$ (relatively low task substitutability) are illustrated in Figure 2. In the centre of the figure, both effort curves are internal and positively sloped, while the slope of e_1 is greater than that of e_2 .

Figure 3 illustrates the case $\gamma = c_1 < c_2$. Here, effort in task 2 is temporarily satiated while both effort curves are internal. Again, which task is intrinsically preferred depends on the position of the vertical axis.

Figure 4 illustrates the case $c_1 < \gamma < c_2$ (relatively high task substitutability). This is the only case that permits ‘crowding out’, that is, a phase in which effort in one task (task 2) decreases with increasing bonus pay. As illustrated, crowding out can only happen when both effort curves are internal. Again, the intrinsically preferred task is determined by the position of the vertical axis.

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. In the empirical analysis we find 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 corresponds 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 , both of which are in principle unobservable. Therefore, while the agent’s ‘own’ group will be mapped to the intrinsically preferred task, it is not always possible to deduce whether this is task 1 (the easier task) or 2 (the harder task).

Note that we have modelled the agent’s effort but not her time use. Some of the results presented below suggest that these are not the same: agents appear to be able to hold effort constant while varying the time spent on a task. Our interpretation is that agents can control the intensity of effort (effort exerted per unit of time)—in particular, they may engage in enjoyable but unproductive ‘idle chatter’ with their friends.

2 Context, Experimental Design and Data

2.1 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 central 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 Kannada, Mysore and Shimoga), and household enrolment in these districts commenced in February–March 2010 (Rajasekhar *et al.*, 2011).

The policy covered hospitalisation expenses for around 700 medical and surgical conditions, with an annual expenditure cap of 30,000 rupees (652 USD) per eligible household.⁸ Each household could enrol up to five members. Pre-existing conditions were covered, as was maternity care, but outpatient treatment was excluded.

The policy was underwritten by insurance companies selected in state-wise tender processes. The insurer received an annual premium per enrolled household,⁹ paid by the central (75%) and state (25%) governments. The beneficiary household paid only a 30 rupees (0.65 USD) annual registration fee.

Biometric information was collected from all members on the day of enrolment and stored, along with photographs, in a smart card issued to the household.¹⁰ Beneficiaries were entitled to cashless treatment at any participating (‘empanelled’) hospital across India. Both public and private hospitals could be empanelled. Hospitals were issued

⁸Here and later, we use the currency exchange rate as per 1 July 2010 according to www.oanda.com (46 rupees/USD).

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

¹⁰According to RSBY guidelines, smart cards should be issued at the time of registration, but this was often not adhered to. For more detail, see Rajasekhar *et al.* (2011).

with card readers and software. The insurance companies reimbursed the hospitals for the cost of treating patients, according to fixed rates.

2.2 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).¹¹ All contacted SHGs were female-only. In the meeting, SHG members were given a brief introduction to RSBY and told that we were looking to recruit a local agent to help spread awareness of the scheme in the village over a period of one year. They were told that the agent would be paid, but no further details about payment were given at that time. In each case, a single candidate was nominated by the group and recruited on the same day. The nominated agent was a member of the SHG, except in two cases where 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.

The payment scheme was revealed to the agent only after recruitment. Each treatment village had been randomly allocated to a payment structure, which constituted the second stage of randomisation, 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 was called 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

¹¹Self-Help Groups are savings-and-credit groups of about 15–20 individuals, often all women, who meet regularly. All government-sponsored SHGs in the village were invited to the meeting.

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.¹²

The bonus payments were determined as follows: A random sample of households eligible for RSBY in each village was surveyed and orally presented with the knowledge test.¹³ 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 total 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.

Thirty-eight 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 (after at least two rounds of payments). Three of these were in incentive-pay villages, while the fourth was in a flat-pay village.¹⁴ In each case, the reported reason was either childbirth or migration away from the village.

¹²As 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 are excluded from the analysis presented here.

¹³For each survey wave, a fixed number of households per village were targeted, and on average 13 household were interviewed per village per wave. The average sample village had 50 eligible households.

¹⁴Assuming that attrition is Poisson-distributed, we are unable to reject the null hypothesis that the rate of attrition was the same across the two treatment groups.

The agents were replaced, but the villages in question are excluded from the analysis. 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.

One question of interest is whether eligible households knew the type of payment scheme the agent in their village was on. We do not have data on this, but on balance we believe that most did not. When asked to nominate an agent, the group was told the work would be remunerated but given no further details. We were careful to tell the agent about the type of contract in private, after selection, and away from the group. When we returned to make payments, these were also always made in private. The agents were of course free to tell others how they were paid, but people locally tend to be reticent in talking about money. Anecdotally, we know that at least some agents on incentive pay were careful not to reveal this information because, once spread, it might reduce their credibility as pro-social volunteers and thereby the villagers' willingness to listen. In any case, from a policy point of view one would probably want to capture the total effect of the contract types inclusive of any additional effect on eligible households who learn how their agent is paid.¹⁵

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,

¹⁵In this sense, side payments or doing favours to incentivised agents would be not regarded as possible threats to our story, but rather as mechanisms through which incentive pay might work. In practice, however, we believe that the scope for collusion was limited: Agents and households were never notified in advance of the knowledge tests. The households were not given the questions in writing, and they were not told whether they had answered the questions correctly or not, nor their overall score. The sample of households to whom the knowledge test was fielded was drawn independently each time. The agent did not know exactly how household knowledge test scores mapped into her bonus pay. Also, as pointed out by one referee, households similar to the incentivised agent would be the ones more likely to do her a favour by trying harder, yet we find the opposite: incentivised agents do relatively better with households who are socially different from her.

third and fourth rounds. 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 all rounds.

2.3 Data

Following agent recruitment, four consecutive rounds of ‘mini-surveys’ were fielded.¹⁶ In each wave, randomly selected eligible households in each sample village were interviewed to establish the state of their knowledge about the scheme and determine their test scores, as well as measure their enrolment status. An important 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 (with replacement) for the first, second and fourth survey rounds, so that there is a partial overlap between the households in these rounds. The first, second and fourth 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 households were contacted by telephone. Although not everyone could be reached by phone, the re-survey rate was significant. A sample of 2360 households were interviewed in the first mini-survey wave, 1931 in the second, 1346 in the third and 2093 in the fourth. In all, the mini-surveys cover 3998 households, of which 1068 were interviewed twice, 642 were interviewed three times and 460 were interviewed four times. As the tests were conducted in every sample village, there was no difference between incentive- and flat-pay agents in the intensity of monitoring.

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 equals 1 for all households. All regressions using data from more than one mini-survey are weighted

¹⁶The data and do-file used to generate the tables are available online.

least squares. In addition, standard errors are clustered at the village level (Bertrand *et al.*, 2004). 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 necessarily 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 around 35 years old. 88% are married. 58% of the agents' household heads have completed primary school. 82% of agent households have a ration card,¹⁷ and 38% are from a forward or dominant caste.¹⁸ In 29% of the cases, the recruited agent was the president of a Self-Help Group. We also constructed a 'female autonomy' score for the agents.¹⁹

Table 2 presents summary statistics for the villages. The average village has a little over 200 households, of which about 50 are eligible for the scheme. Only about a quarter of sample villages are GP headquarters, and the distance from the village to the nearest town is about 13 kilometres. None of the village-level variables differ significantly across treatment groups.

Table 3 presents summary statistics for households. The average household has 4.8 members. 17% are from a forward/dominant caste. In 27% of households, the household head has completed primary school. 92% 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.

¹⁷These cards entitle the holders to purchase certain foods at subsidised rates. The cards are intended for the poor, but because of mis-allocation issues they are an imperfect indicator of poverty.

¹⁸In Karnataka, two castes officially classified as 'backward', Vokkaliga and Lingayath, tend to dominate public life. These two have therefore been classified together with the forward caste groups in one category as 'dominant castes'.

¹⁹The 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 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.

The main outcome variable is the household ‘knowledge score’. A knowledge test was fielded, in each of the four mini-surveys, to all interviewed households across the three experimental arms. 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 B. 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).²⁰

The test questions asked in the four surveys were different, so although the raw scores can be compared across households within a survey, they are not necessarily directly comparable across surveys, even for a given household. The scores on each test were therefore standardised by subtracting the test-wise mean and dividing by the standard deviation.

3 Evidence

3.1 *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 average effect on test score of being in a treated village, and α is a constant reflecting the average test score in the control group.

The results of regression (2) are presented in Table 4, column 1. Households living in a

²⁰Question 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.

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. The estimated effect of flat-pay agents on test scores, while positive, is not statistically significant. This is consistent with the argument that, since these agents are paid a constant amount irrespective of outcome, they are not incentivised to exert any effort beyond the level determined by their intrinsic motivation. In contrast, households in villages assigned an incentive-pay agent score 0.25 standard deviations higher on the knowledge test than those in the control group. Hence, providing agents with financial incentives leads to an improvement in knowledge about the scheme among beneficiaries. Moreover, equality of these two coefficients is rejected. This suggests that, looking at the sample overall, the effect of knowledge-spreading agents is driven by the agents on incentive-pay contracts.²¹

As mentioned, an administrative error caused incentive-pay agents in one district (Bangalore Rural) to be overpaid after the first survey. To allay concerns that our findings are driven by these higher rates of pay, Table 5 presents results using only data from Shimoga district, where no error was made. Overall, the qualitative findings are similar to those obtained in Table 4, if not stronger. Hence, it appears that the main findings are not driven by the larger payments made to agents in one district for one of the four rounds.

Effects by survey round and by agent characteristics are presented in Appendix C.

3.2 The Impact of Agents on Enrolment

Next consider the impact of knowledge agents on programme enrolment. Although the agents were not incentivised to enrol household into the scheme, it is conceivable that

²¹In Appendix C we also check whether baseline agent characteristics play any role in explaining their efficacy in knowledge transmission. The only one we find to be significant at the 10% level is female autonomy.

increasing households' knowledge about the scheme might induce enrolment. Enrolment is also of primary policy relevance.

The results are presented in Table 6. Households living in a treatment village are on average 4.5 percentage points more likely to enrol into the scheme compared to households in the control villages, although this effect is not statistically significant (column 1). Controlling for taluk and wave fixed effects does not change this result (column 2). Once we disaggregate the treatment effect by agent contract type, we find that households living in a treatment village where the agent was on an incentive-pay contract are 7.5 percentage points more likely to be enrolled relative to control group, significant at 5%. No significant impact is found for those living in a treatment village where the agent was on a flat-pay contract. As with the knowledge score results reported above, we are able to reject the equality of the two coefficients at the 5% level. Incentivising agents to disseminate programme knowledge also boosted programme enrolment.

3.3 Incentives and Social Distance

The results so far suggest that monetary incentives matter for how effective agents are at disseminating information about the scheme and increasing enrolment. But previous work suggests that social identity is also an important determinant of insurance take-up. Cole *et al.* (2013) 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 investigates 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 was constructed as follows: First, we created 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 does. Ration-card distance is 1 if any one of them has a ration card and the other does not.²²

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

Before turning to the main specification, consider the basic difference-in-differences calculation in Table 7. We create a binary variable, 'socially proximate pair', indicating whether or not the composite social distance metric is equal to or less than 0.5. The mean knowledge scores for socially proximate and socially distant household-agent pairs are tabulated by agent contract type. Reading the table row by row, flat-pay agents are significantly more effective at transmitting knowledge to socially proximate households than to the socially distant households—average scores are higher by 0.14 standard deviations. For incentivised agents there is no significant difference in performance between close and distant pairs, and the difference in differences is significant. Alternatively, reading the table column by column, incentive pay seems not to affect knowledge transmission in proximate pairs, but it does increase knowledge transmission in distant agent-household pairs, up to about the same level as for proximate pairs. Note that the socially distant households who are assigned a flat-pay agent score significantly lower than any of the other three groups, which is indicative of the disadvantage created by social distance in the absence of incentive pay. In summary, the difference-in-differences analysis suggests that incentive pay has the effect of neutralising social distance as an impediment to the transmission of knowledge — a finding that will be corroborated in what follows.

²²The choice of variables for use in the calculation of social distance was severely constrained, as only a small number of variables were observed for both agents and eligible households. That said, our social distance metric does incorporate measures of caste, education, asset ownership and poverty status. A more sophisticated social distance metric for rural India would probably include measures of most, if not all, of these.

The main empirical specification for this section is

$$Y_{hv} = \alpha + \beta D_{hv} + \gamma T_v + \delta D_{hv} T_v + \pi X_{hv} + u_{hv}. \quad (3)$$

Here, 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 drop out from this analysis since the distance metrics are not defined when there is no agent.) X_{hv} are level 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 effect²³ 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 8. 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, while 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 the negative effect due to social distance. In other words, the effect of incentive pay seems

²³The 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.

to be to cancel the negative effect of social distance, but no more. Loosely speaking, we may say that incentive pay appears to ‘price out prejudice’, though the effects of social distance are not necessarily a function of prejudice alone. For example, socially proximate pairs may meet more often socially, reducing the cost of effort required to transmit information.

See Appendix E for a discussion of the magnitude of the main effect relative to the incentive.

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.

We control for the role of physical distance using two separate measures. The first is whether or not caste groups in a village tend to live apart. This information is based on enumerator recall and hence only available for 107 out of the 147 villages. Based on this information, we construct a binary indicator which is equal to 1 if, in a given village, the settlements of the major caste groups are physically separated, and 0 otherwise. The indicator is equal to 1 for 26 out of 107 villages. Returning to Table 8, this indicator and its interaction with the incentive-pay variable are included in the regression in column 3.

While the sample size drops, the results in column 3 confirm that physical separation does have a negative effect on knowledge transmission. Like for social distance, the coefficient on the interaction of incentive pay and physical distance is of opposite sign and roughly equal magnitude to the un-interacted physical distance term. While the interaction term is not statistically significant, this may indicate that incentive pay can overcome barriers to physical distance. But the results also show that the social distance indicator and its interaction with incentive-pay are still significant at the 10% and 5% levels, respectively, and do not drop much in magnitude. This indicates that social

distance matters, even after controlling for physical distance.

The measure of physical distance used above is crude, and it is conceivable that the the social distance measure still captures some physical distance effect. Hence, we use a second and more accurate measure of physical distance in the form of the actual straight-line distance in metres between the agent's and the household's dwellings, constructed from GPS coordinates collected in the field. Since these data were only collected during the fourth survey wave, the analysis can only be undertaken for a subset of the total sample. Using this second measure, we find results in column 4 that are qualitatively similar to those obtained with the crude measure in column 3. The coefficient on social distance indicator is no longer statistically significant at conventional levels, possibly due to the fact that the sample size drops by almost a quarter. But the coefficient continues to be of opposite sign to that of its interaction term with incentive-pay. In addition, both the coefficient on social distance and that on the interaction term are not significantly different from those obtained in column 2 without the inclusion of physical distance (p -values for the F -tests are 0.38 and 0.74 respectively), which increases our confidence that these results are less likely to be driven by confounding factors. Moreover, this measure of physical distance does not appear to exert any independent influence on knowledge scores.

Columns 5–8 repeat the exercise in column 2 for each of the component distance metrics. For distance in caste, ration-card status and home ownership, the story appears to align with the findings for the composite metric presented above, albeit not always with full significance. 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, ration-card status and possibly home ownership, but not by differences in education. Correspondingly, in this specification, un-interacted incentive pay has a large and positive co-efficient. That is, agents do not appear to be at a disadvantage when communicating with households with a different educational background (having completed primary school or not) from themselves, and the introduction of incentives correspondingly boosts results for socially proximate and distant households

alike.

It is of interest to examine whether the impact of social distance and its interaction with incentive-pay is symmetric across the caste hierarchy in the context of knowledge scores. In other words, is the impact of social distance between agent and beneficiary household more severe when a lower-caste agent interacts with a higher-caste household than vice versa? To test this, we compute differences in differences in mean effects by agent caste group. The results, presented in Table 9, suggest that the qualitative findings are symmetric: irrespective of the agent’s own caste group, the coefficient representing the effect of introducing incentive pay is greater with respect to the cross-group than to the own group. This is reassuring, though in this basic analysis the differences in differences are not statistically significant.

Table 10 presents results for our other main outcome variable, enrolment. The results are qualitatively similar to those obtained for knowledge scores, and with somewhat greater significance for the composite measure of social distance (columns 2–4). Once again the main effect of incentives appears to be to cancel the negative effects of social distance. This finding is robust to the inclusion of controls for either measure of physical distance. Taken together, the analysis in this section suggests that social distance can lower the efficacy of welfare programmes through reduced knowledge transmission and enrolment, and that the use of incentive pay can counter-act this effect.

3.4 Potential Mechanism: Time Spent with Households

One likely mechanism by which agents may boost the knowledge scores and enrolment rates of eligible households in their village is by spending time with them to talk about the scheme. It is, therefore, of interest to look at how agent time allocation varies with contract type and social distance. In the fourth survey round, households were asked how much time their agent had spent with them talking about the scheme over the past three months.²⁴

²⁴This question was introduced only in the fourth and final survey wave. At that time, the households were also asked to recall how much time the agent had spent with them earlier in the intervention. However, we focus on the most recent period as it is probably the most accurately recollected.

The results for time spent are presented in Table 11. The analysis is again restricted to treatment villages since the distance measures are not defined in villages without an agent. In column 1, the time spent by the agent with the household talking about the scheme is regressed on incentive pay, and the results suggest that total time spent with households by the agent does not depend on contract type. So while it was found above that incentivised agents are more successful at transmitting knowledge and inducing enrolment overall, this does not seem to be because she spends more time with households in aggregate.

Column 2 introduces social distance and the interaction of incentive pay and social distance. The coefficients on un-interacted incentive pay and social distance are negative and significant, while the coefficient on the interaction term is positive and significant. Moreover, the magnitudes are such that coefficient on the interaction is close to cancelling the sum of the un-interacted terms.

In column 3, physical distance measured in terms of separation of castes and its interaction with incentive pay are introduced. While the magnitudes suggest that physical distance may matter for time spent, the coefficients are not statistically significant. The coefficients on un-interacted incentive pay and social distance and their interaction lose significance but do not change much in magnitude. Using the GPS-based measure of physical distance, gives stronger results (column 4). While physical distance continue to matter little in terms of time spent, those on social distance and its interaction with incentive-pay are now larger in magnitude and highly significant. Thus, as in the case of knowledge scores, these results suggests that social distance matters in its own right as far as time spent by agent is concerned, and not just as a proxy for physical distance. Once again the magnitude of the interaction term comes close to cancelling the un-interacted terms.

We interpret these results as follows: The type of contract does not appear to make a great difference to the average time spent with each household across the sample. However, while agents on flat-pay contracts spend on average 8 minutes less with each socially distant households than they do with socially proximate households (based on

column 2 of Table 11), agents on incentive-pay contracts spend on average 3 minutes *more* on socially distant households compared to socially proximate households. In other words, incentives have little effect on the time spent with households overall, but they do appear to cause a large shift of agents' focus from socially proximate to socially distant households.

It was shown above that incentivised agents achieve superior results overall in terms of both knowledge scores and enrolment. The results presented in this section suggest that this was achieved without investing more time in talking to households overall. Instead, relative to flat-pay agents, incentivised agents spend less time with their 'own group' and more time with their 'cross-group'. One interpretation of these findings is that agent 'intensity of effort per minute' may vary. Spending time with one's friends may be pleasurable and hence, in the absence of incentives, agents chose to spend more time at a lower intensity with socially proximate households. The implication is that when incentives are introduced, agents are able to increase intensity and hence free up time to spend on socially distant households without sacrificing effort or output with respect to their friends.

Separate results (not reported) indicate that the time-use effect is exclusively on the 'intensive margin': incentive-pay agents did not talk to a greater number of households overall than did flat-pay agents.

3.5 Relating the Empirical Results to the Theoretical Model

The aim of this section is to tie the empirical findings back to the model. It should, however, be noted that what follows is subject to statistical inaccuracy. That is, while we cannot reject the equality of certain quantities, it is also possible that the true values of these quantities are different, but not different enough to be detectable by the econometric tests. While for simplicity we will proceed as if these equalities hold exactly, a full discussion would consider a broader range of cases in which the effort curve is *nearly* flat, effort across the two tasks *nearly* equal, and so on.

Let $e_f(b)$ denote the effort of a knowledge agent when dealing with her own social

group ('friends'), and let $e_o(b)$ denote the effort with respect to dealing with the other group. We observe four points empirically: $e_f(0)$, $e_f(b')$, $e_o(0)$ and $e_o(b')$; that is, the effort with respect to the agent's own and cross-group, with and without bonus. Using this notation, the empirical findings can be summarised as follows:

$$e_o(0) < e_f(0) = e_f(b') = e_o(b')$$

In words, the task of transmitting information to the agent's own group is intrinsically preferred. The introduction of bonus pay induces no change in effort in the intrinsically preferred task, but it does increase effort in the non-preferred task, up to the same level as for the intrinsically preferred task.

The most straightforward interpretation is that with respect to their own group, agents were already exerting 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. We do not observe crowding out, but we cannot rule it out outside of the observed parameter values. Specifically, given more variation in b , we might encounter a region in which effort with respect to one of the groups decreases with b . Unfortunately, from the four points we observe, we cannot tell whether or not we are in a 'crowding-out world'.

In Figures 2–4, the position of the vertical axes correspond to cases that are consistent with the empirical findings. At $b = 0$, e_1 has reached the maximum effort level while e_2 has not. A sufficiently high bonus b would bring e_2 up to \bar{e} where it would be equal to e_1 . If this reflects the empirical reality, then task 1, the easier task, corresponds to the agent's own group.

However, another possibility is generated by shifting \bar{e} in Figures 2–4 down until it meets, or crosses, the meeting point of the internal solutions. The vertical axis would now need to be placed to the left of the crossing point. This configuration would generate a solution in which e_2 , the harder task, corresponds to the agent's own group. For this to be the case, θ_2 , the intrinsic motivation for success in the own-group task would need to be not only greater than θ_1 but large enough to outweigh the cost disadvantage.

As an example of the latter, imagine that, irrespective of the agent’s own identity, it is easier to transmit knowledge to high-caste than low-caste households, perhaps because high-caste households tend to be better educated. Then, irrespective of the caste of the agent, task 1 (the easier task) corresponds to high-caste households and task 2 to low-caste households. If so, for a low-caste agent to intrinsically prefer the task of transmitting information to her own caste group, which is what we observe, her intrinsic motivation for the own-group task, θ_2 , needs to be large enough, relative to θ_1 , to outweigh the cost disadvantage. Whilst the cost-of-effort parameters c_i and the preference parameters θ_i are distinct concepts theoretically, empirically we are unable to distinguish the effect of low task preference (‘prejudice’) from high cost of effort.

It is also possible that the apparent convergence of the effort curves is not due to having reached the maximum effort level as assumed above, but rather that one of the effort curves is flat, as in Figure 3. If the vertical axis were to the left of the crossing point, and the positive bonus pay observation $b = b'$ were exactly *at* the crossing point, this could explain the empirical findings. However, we find this possibility less likely than the two described above, because it would require the arbitrarily chosen experimental value for bonus pay to have hit exactly the ‘sweet spot’ (the crossing point).

Though the empirical findings are supportive of 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 on agent performance. If households attain the maximum score, any further effort would be unobservable and hence, from the point of view of an incentive-pay agent, 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. In particular, only 5% of households answered seven or eight out of eight questions correctly.

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 households. 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 person who can be contacted if the need arises: if a household member falls ill or otherwise needs health care, they would turn to the agent and ask her advice on how to obtain treatment under the scheme. If this perspective is widespread, it would not be surprising if the households' motivation for learning details about the insurance policy is limited. They only need basic knowledge about the scheme, and for this reason their patience with listening to details will probably 'max out' relatively quickly.

Our model is mute on the relationship between effort and time. But the results suggest that, relative to flat-pay agents, agents on incentive pay are able to increase overall effort/output without spending more time in aggregate. Instead, we observe a reallocation of time away from socially proximate towards socially distant households. If *intensity of effort* can vary over time, non-incentivised agents may be spending more time with socially proximate households (their 'friends') than strictly necessary, presumably because they enjoy it. When incentives are introduced, agents are able to shift time away from socially proximate households towards time spent with socially distant households by increasing the intensity (i.e., reducing idle but enjoyable 'chatter') of the time spent with their 'friends'.

4 Conclusion

This paper sheds light on the role of financial incentives and social proximity in motivating local agents to transmit knowledge about a public service. The results suggest, first, that hiring agents to spread knowledge about welfare programmes has a positive impact on the level of knowledge, but that the entire effect is driven by agents on incentive-pay contracts. Second, agents on incentive-pay contracts also have a positive and significant impact on programme enrolment. Third, social distance between agent and beneficiary has a negative impact on knowledge transmission and enrolment, but putting agents on incentive-pay contracts increases knowledge transmission and enrolment by cancelling (at

our level of bonus pay) the negative effect of social distance. On the other hand, incentive pay has no impact on knowledge transmission or enrolment for socially proximate agent–beneficiary pairs. A likely mechanism is a reallocation of time spent by incentive-pay agents towards socially distant households at the expense of socially proximate ones.

Our results may have implications for public service delivery in developing countries, where, in addition to common supply-side problems like staff absenteeism, corruption and red tape, a lack of awareness and knowledge regarding available welfare schemes represents an important barrier to the take-up of government programmes. The experimental evidence presented here points to a key mechanism that may in some circumstances alleviate this problem.

Our 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 *et al.*, 2005; Banerjee and Somanathan, 2007). It would be hasty to extrapolate from our findings in the context of information transmission about welfare schemes to wider societal effects of own-group bias. Still, in the setting studied, a relatively small piece rate was sufficient to overcome the negative consequences of entrenched social barriers.

University of Bristol

London School of Economics and Political Science

Institute for Social and Economic Change

Institute for Social and Economic Change

University of Sussex

References

- Aizer, A. (2007). ‘Public health insurance, program take-up and child health’, *Review of Economics and Statistics*, vol. 89(3), pp. 400–415.
- Akerlof, G. and Kranton, R. (2005). ‘Identity and the economics of organizations’, *Journal of Economic Perspectives*, vol. 19(1), pp. 9–32.
- Anderson, S. (2011). ‘Caste as an impediment to trade’, *American Economic Journal: Applied Economics*, vol. 3(1), pp. 239–263.
- Ashraf, N., Bandiera, O. and Jack, K. (2014). ‘No margin, no mission? A field experiment on incentives for public service delivery’, *Journal of Public Economics*, vol. 120, pp. 1–17.
- Bandiera, O., Barankay, I. and Rasul, I. (2009). ‘Social connections and incentives in the workplace: Evidence from personnel data’, *Econometrica*, vol. 77(4), pp. 1047–94.
- Bandiera, O., Barankay, I. and Rasul, I. (2011). ‘Field experiments with firms’, *Journal of Economic Perspectives*, vol. 25(3), pp. 63–82.
- Banerjee, A., Banerji, R., Duflo, E., Glennerster, R. and Khemani, S. (2010a). ‘Pitfalls of participatory programs: Evidence from a randomized evaluation in education in India’, *American Economic Journal: Economic Policy*, vol. 2(1), pp. 1–30.
- Banerjee, A., Duflo, E., Glennerster, R. and Kothari, D. (2010b). ‘Improving immunisation coverage in rural India: Clustered randomised controlled evaluation of immunisation campaigns with and without incentives’, *British Medical Journal*, vol. 340:c2220.
- Banerjee, A., Iyer, L. and Somanathan, R. (2005). ‘History, social divisions, and public goods in rural India’, *Journal of the European Economic Association*, vol. 3(2-3), pp. 639–647.
- Banerjee, A. and Munshi, K. (2004). ‘How efficiently is capital allocated: Evidence from

- the knitted garment industry in Tirupur’, *Review of Economic Studies*, vol. 71(1), pp. 19–42.
- Banerjee, A. and Somanathan, R. (2007). ‘The political economy of public goods: Some evidence from India’, *Journal of Development Economics*, vol. 82(2), pp. 287–314.
- Becker, G.S. and Posner, R.A. (2009). *Uncommon sense: Economic Insights, from Marriage to Terrorism*, Chicago: University of Chicago Press.
- Bénabou, R. and Tirole, J. (2003). ‘Intrinsic and extrinsic motivation’, *Review of Economic Studies*, vol. 70(3), pp. 489–520.
- Bénabou, R. and Tirole, J. (2006). ‘Incentives and prosocial behavior’, *American Economic Review*, vol. 96(5), pp. 1652–1678.
- BenYishay, A. and Mobarak, A.M. (2014). ‘Social learning and communication’, National Bureau of Economic Research Working Paper 20139.
- Bertrand, M., Duflo, E. and Mullainathan, S. (2004). ‘How much should we trust differences-in-differences estimates?’, *Quarterly Journal of Economics*, vol. 119(1), pp. 249–275.
- Bertrand, M., Luttmer, E.F.P. and Mullainathan, S. (2000). ‘Network effects and welfare cultures’, *Quarterly Journal of Economics*, vol. 115(3), pp. 1019–1055.
- Besley, T. and Ghatak, M. (2005). ‘Competition and incentives with motivated agents’, *American Economic Review*, vol. 95(3), pp. 616–637.
- Cole, S., Gine, X., Tobacman, J., Topalova, P., Townsend, R. and Vickery, J. (2013). ‘Barriers to household risk management: Evidence from India’, *American Economic Journal: Applied Economics*, vol. 5(1), pp. 104–135.
- Dal Bó, E., Finan, F. and Rossi, M.A. (2013). ‘Strengthening state capabilities: The role of financial incentives in the call to public service’, *Quarterly Journal of Economics*, vol. 128(3), pp. 1169–1218.

- Daponte, B., Sanders, S. and Taylor, L. (1999). ‘Why do low-income households not use food stamps? Evidence from an experiment’, *Journal of Human Resources*, vol. 34(3), pp. 612–628.
- Das, J. and Leino, J. (2011). ‘Evaluating the RSBY: Lessons from an experimental information campaign’, *Economic and Political Weekly*, vol. XLVI(32), pp. 85–93.
- Duflo, E., Hanna, R. and Ryan, S.P. (2012). ‘Incentives work: Getting teachers to come to school’, *American Economic Review*, vol. 102(4), pp. 1241–1278.
- Dupas, P. (2011). ‘Do teenagers respond to hiv risk information? Evidence from a field experiment in Kenya’, *American Economic Journal: Applied Economics*, vol. 3(1), pp. 1–36.
- Easterly, W. and Levine, R. (1997). ‘Africa’s growth tragedy: policies and ethnic divisions’, *Quarterly Journal of Economics*, vol. 112(4), pp. 1203–1250.
- Filmer, D. and Schady, N. (2009). ‘Are there diminishing returns to transfer size in conditional cash transfers?’, World Bank Policy Research Working Paper 4999.
- Finan, F., Olken, B.A. and Pande, R. (forthcoming). ‘The personnel economics of the state’, *Handbook of Field Experiments*, forthcoming.
- Fisman, R., Paravisini, D. and Vig, V. (forthcoming). ‘Social proximity and loan outcomes: Evidence from an Indian bank’, *American Economic Review*.
- Francois, P. (2000). ‘Public service motivation as an argument for government provision’, *Journal of Public Economics*, vol. 78(3), pp. 275–299.
- Glewwe, P., Holla, A. and Kremer, M. (2009). ‘Teacher incentives in the developing world’, in (M. Springer, ed.), *Performance Incentives: Their Growing Impact on K-12 Education*, Washington, DC: Brookings Institution Press.
- Gneezy, U., Meier, S. and Rey-Biel, P. (2011). ‘When and why incentives (don’t) work to modify behavior’, *Journal of Economic Perspectives*, vol. 25(4), pp. 1–21.

- Gneezy, U. and Rustichini, A. (2000a). ‘A fine is a price’, *Journal of Legal Studies*, vol. 29(1), pp. 1–18.
- Gneezy, U. and Rustichini, A. (2000b). ‘Pay enough or don’t pay at all’, *Quarterly Journal of Economics*, vol. 115(3), pp. 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, V., Malherbet, F. and Pellizzari, M. (2004). ‘Take-up of welfare benefits in OECD countries: A review of the evidence’, OECD Social, Employment and Migration Working Paper 17.
- Holmstrom, B. and Milgrom, P. (1991). ‘Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design’, *Journal of Law, Economics, and Organization*, vol. 7(Special issue), pp. 24–52.
- Jensen, R. (2010). ‘The (perceived) returns to education and the demand for schooling’, *Quarterly Journal of Economics*, vol. 125(2), pp. 515–548.
- Kimenyi, M.S. (2006). ‘Ethnicity, governance and the provision of public goods’, *Journal of African Economies*, vol. 15(suppl 1), pp. 62–99.
- Kingdon, G. and Rawal, S. (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 Paper 10-18.
- Kremer, M. and Holla, A. (2008). ‘Pricing and access: Lessons from randomized evaluation in education and health’, in (W. Easterly and J. Cohen, eds.), *What Works in Development: Thinking Big and Thinking Small*, Washington, DC: Brookings Institution Press.
- La Porta, R., Lopez-de Silanes, F., Shleifer, A. and Vishny, R. (1999). ‘The quality of government’, *Journal of Law, Economics, and Organization*, vol. 15(1), pp. 222–279.

- Lang, K. and Lehmann, Jee-Yeon, K. (2012). ‘Racial discrimination in the labor market: Theory and empirics’, *Journal of Economic Literature*, vol. 50(4), pp. 1–48.
- Lazear, E.P. (2000). ‘Performance pay and productivity’, *American Economic Review*, vol. 90(5), pp. 1346–1361.
- Muralidharan, K. and Sundararaman, V. (2011). ‘Teacher performance pay: Experimental evidence from India’, *Journal of Political Economy*, vol. 119(1), pp. 39–77.
- Prendergast, C.J. (2015). ‘The empirical content of pay-for-performance’, *Journal of Law, Economics, and Organization*, vol. 31(2), pp. 242–261.
- Rajasekhar, D., Berg, E., Ghatak, M., Manjula, R. and Roy, S. (2011). ‘Implementing health insurance: The roll-out of Rashtriya Swasthya Bima Yojana in Karnataka’, *Economic and Political Weekly*, vol. XLVI(20), pp. 56–63.
- Rasul, I. and Rogger, D. (2016). ‘Management of bureaucrats and public service delivery: Evidence from the Nigerian civil service’, *Economic Journal*, forthcoming.
- Sandel, M. (2012). *What Money Can't Buy: The Moral Limits of Markets*, London: Allen Lane.
- Thornton, R. (2008). ‘The demand for, and impact of, learning HIV status: Evidence from a field experiment’, *American Economic Review*, vol. 98(5), pp. 1829–1863.
- World Bank (2011). *Social Protection for a Changing India: Volume 2*.




Figure 1: *The one-task solution*




Figure 2: *A solution without crowding out ($\gamma < c_1 < c_2$).*




Figure 3: *A solution with temporary satiation in e_2 ($\gamma = c_1 < c_2$)*




Figure 4: *A solution with crowding out ($c_1 < \gamma < c_2$)*

Table 1: *Agent Summary Statistics*

	Flat pay	Incentive pay	Difference
Agent age	34.8 (8.81)	34.8 (8.08)	0.018 [1.73]
Agent is married	0.81 (0.40)	0.92 (0.28)	0.10 [0.073]
Agent is of forward/dominant caste	0.43 (0.50)	0.35 (0.48)	-0.080 [0.10]
Agent's household head has completed primary school	0.62 (0.49)	0.56 (0.50)	-0.058 [0.10]
Agent household has ration card	0.89 (0.31)	0.79 (0.41)	-0.10 [0.071]
Agent owns her home	0.86 (0.35)	0.87 (0.34)	0.0084 [0.069]
Agent is Self-Help Group president	0.30 (0.46)	0.28 (0.45)	-0.016 [0.093]
Agent autonomy score (the higher, the more autonomous)	5.57 (0.93)	5.68 (0.84)	0.11 [0.18]
Observations	37	71	.

Notes: Standard deviations are in parentheses. Robust standard errors for the difference tests are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: *Village Summary Statistics*

	Control	Flat pay	Inc'tive pay	Flat –Control	Inc'tive –Control	Inc'tive –Flat
Village size (households)	205.1 (193.0)	222.6 (167.7)	237.2 (248.8)	17.5 [41.7]	32.2 [42.8]	14.6 [40.7]
Eligible population (households)	42.9 (45.1)	56 (51.5)	56.2 (53.9)	13.1 [11.1]	13.3 [9.6]	0.17 [10.6]
Village is GP headquarters	0.18 (0.39)	0.30 (0.46)	0.25 (0.44)	0.12 [0.098]	0.074 [0.081]	-0.044 [0.092]
Distance to GP headquarters in km (0 if headquarters)	2.92 (3.46)	2.26 (2.15)	2.35 (2.21)	-0.67 [0.66]	-0.58 [0.61]	0.088 [0.44]
Distance to nearest town in km	11.5 (6.20)	15.8 (12.9)	13.3 (10.4)	4.33* [2.33]	1.83 [1.59]	-2.49 [2.44]
Proportion of village land irrigated	0.69 (0.47)	0.59 (0.50)	0.63 (0.49)	-0.098 [0.11]	-0.059 [0.094]	0.039 [0.10]
Village has drainage sanitation	0.79 (0.41)	0.89 (0.31)	0.89 (0.32)	0.097 [0.083]	0.092 [0.075]	-0.0046 [0.064]
Average social distance between households in village	0.30 (0.081)	0.28 (0.083)	0.29 (0.082)	-0.017 [0.019]	-0.0089 [0.016]	0.0078 [0.017]
Observations	39	37	71	.	.	.

Notes: Standard deviations are in parentheses. Robust standard errors for the difference tests are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Household Summary Statistics

	Control	Flat pay	Inc'tive pay	Flat -Control	Inc'tive -Control	Inc'tive -Flat
Household is of forward/ dominant caste	0.22 (0.42)	0.18 (0.39)	0.15 (0.36)	-0.041 [0.046]	-0.068* [0.037]	-0.027 [0.036]
Household head has com- pleted primary school	0.29 (0.45)	0.23 (0.42)	0.30 (0.46)	-0.057* [0.032]	0.0042 [0.029]	0.061** [0.026]
Household has ration card	0.93 (0.26)	0.93 (0.26)	0.92 (0.27)	0.0040 [0.020]	-0.0037 [0.019]	-0.0077 [0.018]
Household owns its home	0.65 (0.48)	0.63 (0.48)	0.64 (0.48)	-0.024 [0.042]	-0.010 [0.034]	0.014 [0.038]
Household knowledge score, mean	-0.11 (0.74)	-0.028 (0.79)	0.11 (0.80)	0.077 [0.094]	0.21*** [0.077]	0.14 [0.087]
Household is enrolled, mean	0.67 (0.38)	0.68 (0.40)	0.73 (0.37)	0.0029 [0.051]	0.058 [0.039]	0.055 [0.044]
Social distance to agent	-	0.39 (0.24)	0.37 (0.26)	-	-	-0.023 [0.035]
Agent time spent with household, in minutes	-	15.9 (26.6)	15.6 (33.4)	-	-	-0.33 [3.23]
Observations	919	886	1634	.	.	.

Notes: Standard deviations are in parentheses. Standard errors for the difference tests, clustered at the village level, are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: *The Effect of Agents on Knowledge*

	(1) Knowledge	(2) Knowledge	(3) Knowledge
Agent in village	0.175*** (0.0662)	0.190*** (0.0611)	
Flat-pay agent in village			0.0677 (0.0842)
Incentive-pay agent in village			0.252*** (0.0623)
Survey wave fixed effects	No	Yes	Yes
Taluk fixed effects	No	Yes	Yes
Observations	7730	7730	7730
<i>t</i> -test: flat=incentivised (<i>p</i> -value)			0.0198

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: *The Effect Agents on Knowledge, Shimoga District Only*

	(1) Knowledge	(2) Knowledge	(3) Knowledge
Agent in village	0.222*** (0.0752)	0.202*** (0.0699)	
Flat-pay agent in village			0.0315 (0.106)
Incentive-pay agent in village			0.301*** (0.0694)
Survey wave fixed effects	No	Yes	Yes
Taluk fixed effects	No	Yes	Yes
Observations	3954	3954	3954
<i>t</i> -test: flat=incentivised (<i>p</i> -value)			0.0190

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 Agents on Enrolment*

	(1) Enrolled	(2) Enrolled	(3) Enrolled
Agent in village	0.0445 (0.0352)	0.0465 (0.0325)	
Flat-pay agent in village			-0.0102 (0.0434)
Incentive-pay agent in village			0.0754** (0.0332)
Survey wave fixed effects	No	Yes	Yes
Taluk fixed effects	No	Yes	Yes
Observations	7730	7730	7730
<i>t</i> -test: flat=incentivised (<i>p</i> -value)			0.0294

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: *Difference in Differences: Mean Knowledge Scores by Contract Type and Social Distance*

	Socially proximate pairs	Socially distant pairs	Difference
Flat pay	0.05 (0.03)	-0.09 (0.03)	-0.14** (0.05)
Incentive pay	0.12 (0.02)	0.13 (0.03)	0.01 (0.04)
Difference	0.07 (0.04)	0.22*** (0.04)	0.15*** (0.06)

Notes: Standard errors are in parentheses. ‘Socially proximate pairs’ is a binary variable equal to 1 if the composite social distance metric is equal to or less than 0.5 and 0 otherwise. These regressions use data only from the treatment villages, i.e. villages with agents, since social distance is not defined for villages without agents. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Incentives and Social Distance - Effect on Knowledge

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge
Incentive pay	0.156** (0.0781)	-0.0499 (0.114)	-0.0894 (0.133)	-0.100 (0.131)	0.0529 (0.0940)	0.220** (0.102)	0.0799 (0.0805)	0.0901 (0.0894)
Social distance		-0.429** (0.183)	-0.372* (0.210)	-0.328 (0.228)	-0.226** (0.0992)	0.0862 (0.0881)	-0.240** (0.118)	-0.123 (0.0798)
Incentive pay × social distance		0.559*** (0.194)	0.483** (0.212)	0.517** (0.234)	0.254** (0.122)	-0.123 (0.105)	0.398*** (0.136)	0.188** (0.0898)
Castes live apart			-0.278** (0.122)					
Incentive pay × castes live apart			0.260 (0.165)					
Physical distance				-0.0272 (0.0716)				
Incentive pay × physical distance				0.00673 (0.0834)				
Agent and household characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Survey wave and taluk fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Social distance metric	-	Composite	Composite	Composite	Caste only	Education only	Ration-card status only	Home ownership only
Observations	4854	4854	3877	3189	4854	4854	4854	4854

Notes: Weighted least-squares regressions. The regressions use data only from treatment villages, i.e. villages with agents, since the distance measures are not defined for villages without agents. 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, whether they have a ration card and whether they own their home. The simple social distance metrics in columns 5–8 are binary variables equal to the absolute difference between the corresponding household and agent characteristic binaries. The composite social distance metric in columns 2–4 is the sum of the binary distance metrics for caste, education, ration-card status and home ownership, divided by 4. ‘Castes live apart’ is a binary indicator for whether castes are separated into separate zones in the village, and ‘physical distance’ is the great-circle distance in metres between the agent and household dwellings. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: *Testing for Symmetry: Differences in Differences in Knowledge Score, by Agent Caste Category*

Dominant-caste agent			
	Dominant household	Non-dominant household	Difference (Dom-Non)
Flat pay	0.05 (0.09)	-0.11 (0.04)	0.16* (0.09)
Incentive pay	0.06 (0.07)	0.09 (0.03)	-0.03 (0.07)
Difference (Inc-Flat)	0.01 (0.11)	0.20*** (0.05)	-0.19 (0.12)
Non-dominant-caste agent			
	Dominant household	Non-dominant household	Difference (Dom-Non)
Flat pay	-0.13 (0.08)	0.06 (0.04)	-0.19** (0.09)
Incentive pay	0.07 (0.06)	0.15 (0.02)	-0.09 (0.07)
Difference (Inc-Flat)	0.20** (0.10)	0.10** (0.04)	0.10 (0.11)

Notes: Standard errors are in parentheses. This table uses only the sample of treatment villages, i.e. villages with agents, since distances are not defined in the control villages (where there are no agents). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Incentives and Social Distance - Effect on Enrolment

	(1) Enrolled	(2) Enrolled	(3) Enrolled	(4) Enrolled	(5) Enrolled	(6) Enrolled	(7) Enrolled	(8) Enrolled
Incentive pay	0.0648* (0.0388)	-0.0700 (0.0548)	-0.100 (0.0674)	-0.115 (0.0768)	0.0116 (0.0398)	0.0527 (0.0424)	0.00664 (0.0354)	0.0529 (0.0472)
Social distance		-0.238** (0.0983)	-0.212* (0.112)	-0.348** (0.138)	-0.0630 (0.0507)	-0.0286 (0.0355)	-0.237*** (0.0664)	-0.00297 (0.0420)
Incentive pay × social distance		0.368*** (0.132)	0.313** (0.134)	0.396** (0.174)	0.133** (0.0651)	0.0221 (0.0467)	0.306*** (0.0942)	0.0348 (0.0483)
Castes live apart			-0.0727 (0.0642)					
Incentive pay × castes live apart			0.173** (0.0820)					
Physical distance				-0.0544 (0.0446)				
Incentive pay × physical distance				-0.00202 (0.0592)				
Agent and household characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Survey wave and taluk fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Social distance metric	-	Composite	Composite	Composite	Caste only	Education only	Ration-card status only	Home ownership only
Observations	4854	4854	3877	3189	4854	4854	4854	4854

Notes: Weighted least-squares regressions. The regressions use data only from treatment villages, i.e. villages with agents, since the distance measures are not defined for villages without agents. 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, whether they have a ration card and whether they own their home. The simple social distance metrics in columns 5–8 are binary variables equal to the absolute difference between the corresponding household and agent characteristic binaries. The composite social distance metric in columns 2–4 is the sum of the binary distance metrics for caste, education, ration-card status and home ownership, divided by 4. ‘Castes live apart’ is a binary indicator for whether castes are separated into separate zones in the village, and ‘physical distance’ is the great-circle distance in metres between the agent and household dwellings. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: *Agent Time Spent with Household*

	(1)	(2)	(3)	(4)
	Minutes	Minutes	Minutes	Minutes
Incentive pay	-1.059 (1.683)	-5.178** (2.510)	-5.884* (3.178)	-6.369** (2.944)
Social distance		-8.298** (3.570)	-6.570 (4.123)	-9.965** (3.862)
Incentive pay \times social distance		10.96** (5.256)	9.112 (5.776)	13.04** (5.779)
Castes live apart			-4.297 (4.262)	
Incentive pay \times castes live apart			5.901 (5.258)	
Physical distance				-1.544 (1.831)
Incentive pay \times physical distance				-0.347 (1.983)
Agent and household characteristics	Yes	Yes	Yes	Yes
Taluk fixed effects	Yes	Yes	Yes	Yes
Observations	1539	1539	1221	1459

Notes: Ordinary least-squares regressions. Only data from treatment villages are used, i.e. villages with agents, since the distance measures are not defined for villages without agents. Standard errors, in parentheses, are clustered at the village level. The dependent variable is the number of minutes the agent spent with the household talking about the programme in minutes in the fourth round of the intervention, as reported by the household. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Cases Not Captured by Figures 2–4 for the Two-Task Model

Figures 2–4 illustrate the main classes of solutions for the two-task model. However, it is worth noting some features of the solution space that are not captured by the figures.

First, in all three illustrated solutions, when moving from left to right (that is, increasing b from minus infinity), task 2 is the first to move into the internal solution space. While it is possible for task 1 to ‘leave’ the lower bound first, the two effort curves cannot then cross. This is because task 1 would have a ‘head start’, and when both tasks are internal, the slope of task 1 is always greater so that effort in task 2 cannot ‘catch up’. These cases can, therefore, be generated graphically from Figures 2–4 by shifting the lower effort boundary, \underline{e} , up until it meets the crossing point of e_1 and e_2 , or beyond.

Second, there is a set of sub-cases of Figure 4 where e_2 ‘temporarily’ hits the upper and/or lower bound only to re-emerge into the internal solution space. Graphically, the first of these cases can be generated from Figure 4 by shifting the maximum effort level \bar{e} down until it crosses the first kink in the internal e_2 curve. The second is generated by shifting the minimum effort level \underline{e} up until it crosses the second kink in the internal e_2 curve. The third of these special cases is a combination of the first two; that is when e_2 is first temporarily saturated at \bar{e} and then at \underline{e} before it re-emerges and finally reaches \bar{e} .

Third, while Figures 2–4 show the kinks in e_2 when e_1 reaches the upper or lower bound, they do not illustrate that e_1 will be similarly kinked if e_2 reaches the upper or lower bound while e_1 is internal.

Fourth, it should be noted that, depending on the position of the vertical axis, not all regions of a solution type may be feasible when b is constrained to be non-negative, as is the case in our model.

B 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.

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

1. Can the beneficiary go to any hospital to get treatment under RSBY?

No, only RSBY-empanelled hospitals.

2. What document does the beneficiary need to take with him/her to the hospital to get treatment under RSBY?

RSBY card.

3. Mahesha is a construction worker who does not have a RSBY card. One day he fell from a roof while working and broke his leg. His cousin, who has a RSBY card, took him to a RSBY- empanelled hospital where Mahesh was treated well. The hospital bill came to Rs 12,000. Can the cousin's RSBY card be used to pay for Mahesh's treatment?

No.

4. A household has a RSBY card and has not used it since it was renewed. The wife falls ill and is treated in a RSBY-empanelled hospital. She has to stay in hospital for a long time, and the total bill comes to 36,000 rupees. Can this amount be charged to the card?

No, only Rs 30,000 can be charged to the card, the rest will have to be paid by the household.

5. If the beneficiary loses or damages the RSBY smartcard, can it be replaced?

Yes.

6. Chandrasekhar falls ill and is admitted to a RSBY-empanelled hospital. The cost of his treatment is covered by RSBY. Will the bed charges be covered, too?

Yes, as long as the total cost is less than Rs 30,000 for the year for the whole household.

7. Sangeeta is a housewife who signed up for RSBY with her husband. When giving birth to their second child, she was admitted to a RSBY-empanelled hospital for the delivery. Can the hospital costs be charged to the RSBY card?

Yes.

8. Ramappa signed up for RSBY but did not use the card the whole year. Therefore, he wants to get the Rs 30 registration fee back. Can he?

No, and he has to pay Rs 30 again if he wants to join for another year.

C Effects over Time and by Agent Characteristics

Given that the four survey waves were conducted at different times after the recruitment of agents, it is also possible to look at dynamics. For example, does the difference in knowledge score between households in the two experimental groups increase over time? This is examined in Table C1. In columns 1–4, we find that the gap between households in the incentive-pay group and those in flat-pay and control areas is established already in the first survey wave, and continues through the later waves. For wave 3, the two treatment coefficients are not individually significant, possibly attributable to the smaller sample, but we can still reject the equality of these two coefficient at 5% significance. In column 5, we regress knowledge score on a full set of interaction terms of the two experimental groups with the different waves. We are unable reject the null hypothesis that the marginal effects of flat-pay agents for the later waves are jointly zero, and similarly in the case of the incentive-pay agents. Therefore, we conclude that the impact of the two treatments is similar across all rounds.

The results suggests that the extra effort by the agents in the incentive-pay villages was exerted early on. This is consistent with rational behaviour on the part of the agents if knowledge is persistent: any extra effort should be exerted early on since knowledge imparted then would be rewarded in all subsequent payment rounds, whereas knowledge imparted in the last round would only be rewarded once. Alternatively, it could be that household knowledge does decay and that subsequent effort by the agent is required to maintain it. The finding is also consistent with households having a bounded ‘appetite’ for knowledge about the scheme; this possibility is also discussed in Section 3.5.

Table C2 looks at how the impact of knowledge agents on knowledge test scores depends on agent characteristics. Column 1 replicates column 2 of Table 4 for ease of reference. In column 2, the main treatment variable (whether or not there is an agent in the 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. While the un-interacted term is no longer significant, none of the interacted effects are

significant at the 10% level except for the autonomy metric. (The autonomy variable is described in the data section.) It seems plausible that an agent who is free to move around the village might be more effective.

Table C1: *The Effect of Agents on Knowledge, by Survey Round*

	(1)	(2)	(3)	(4)	(5)
	Knowledge	Knowledge	Knowledge	Knowledge	Knowledge
Flat-pay agent in village	0.00265 (0.115)	0.233* (0.119)	-0.0591 (0.109)	0.113 (0.128)	0.0273 (0.133)
Incentive-pay agent in village	0.206** (0.0794)	0.367*** (0.0886)	0.128 (0.0930)	0.268** (0.116)	0.214** (0.0845)
Wave 2 × flat-pay agent in village					0.182 (0.183)
Wave 3 × flat-pay agent in village					-0.0735 (0.150)
Wave 4 × flat-pay agent in village					0.0329 (0.198)
Wave 2 × incentive- pay agent in village					0.176 (0.146)
Wave 3 × incentive- pay agent in village					-0.0389 (0.120)
Wave 4 × incentive- pay agent in village					0.0164 (0.166)
Wave 2 fixed effect					-0.183 (0.118)
Wave 3 fixed effect					0.00690 (0.0895)
Wave 4 fixed effect					-0.0428 (0.142)
Taluk fixed effects	Yes	Yes	Yes	Yes	Yes
Survey wave	1	2	3	4	All
Observations	2360	1931	1346	2093	7730
<i>t</i> -test: flat=incentivised (<i>p</i> -val)	0.0696	0.244	0.0482	0.151	
Joint significance (<i>p</i> -value) of Wave 2 × flat – Wave 4 × flat					0.571
Joint significance (<i>p</i> -value) of Wave 2 × inc. – Wave 4 × inc.					0.504

Notes: Columns 1–4 are OLS regressions. Column 5 is a weighted least-squares regression in which 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 C2: *The Effect of Knowledge Agents, by Agent Characteristics*

	(1) Knowledge	(2) Knowledge
Treatment (agent in village)	0.190*** (0.0611)	-0.265 (0.310)
Treatment × agent is 30+		0.0359 (0.0905)
Treatment × agent is 50+		-0.0736 (0.132)
Treatment × agent of forward/dominant caste		-0.105 (0.0841)
Treatment × agent household head has completed primary school		-0.101 (0.0810)
Treatment × agent has ration card		-0.0529 (0.105)
Treatment × agent owns her home		0.0852 (0.128)
Treatment × agent is Self-Help Group president		-0.00685 (0.0682)
Treatment × agent autonomy		0.0899* (0.0462)
Survey wave fixed effects	Yes	Yes
Taluk fixed effects	Yes	Yes
Observations	7730	7730

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

D Village-Level Analysis

Table D1 checks that a few of our main results are robust the inclusion of village-level controls. The village-level variables included in these regressions are village size in households, eligible population in households, whether it is a GP headquarters village, distance to GP headquarters (0 if headquarters), distance to nearest town, proportion of village agricultural land that is irrigated, whether the village has drainage sanitation, and the average social distance between households in the village.

Columns 1 and 2 refer back to Tables 4 and 6, respectively. The main findings, that incentive-pay agents are associated with a higher level of knowledge and enrolment while flat-pay agents are not, are confirmed. Columns 3 and 4 do the same for the main social-distance results from Tables 8 and 10, and again the key findings are confirmed.

None of the village-level characteristics are systematically correlated with the outcome variables, conditional on the other variables included, with the exception of average social distance between households in the village. This is possibly an indication that the spread of knowledge may be impeded by social difference between households in a village, as well as between the agent and the households.

Table D2 presents results for key regressions at the village level. For this analysis, knowledge scores, enrollment and social distance were collapsed to a single observation per village by averaging across households and survey waves. The full set of village-level controls was included, as well as taluk (region) fixed effects. Robust standard errors were computed.

Columns 1 and 2 mirror the main results in Tables 4 and 6, and confirm that agents had an effect on knowledge as well as enrolment, and that the effect was driven by agents on incentive-pay contracts.

Columns 3 and 4 refer back to Tables 8 and 10 and look at the interaction of incentives and social distance. As before, control villages drop out in these specifications, since social distance is not defined in villages without agents. The main findings are confirmed: In the absence of incentives, social distance is detrimental to knowledge and enrolment. And

introducing incentives makes the agents ‘blind’ to social distance, in that sense the effects across social distance are statistically indistinguishable in the presence of incentives. (This is seen by confirming that the coefficients on incentive pay, social distance and their interaction sum to approximately zero for both knowledge and enrolment.)

Qualitatively, the main difference between the household-level results in Tables 8 and 10 and the ones in columns 3 and 4 of Table D2 is that in the latter, the un-interacted incentive variable has a negative and significant coefficient. That is, incentivised agents appear to spend less time on socially proximate households than do agents who are not incentivised.

But this is probably an artefact of collapsing the data to the village level. To see this, note that in the household-level regressions, the un-interacted incentive-pay coefficient captures the effect of an incentive-pay agent on a household with zero social distance to the agent, that is, a household that is identical to the agent in terms of each of our four characteristics. In our data, 17.5% of agent–household pairs satisfy this criterion. But since there is no village in which *every* household has a social distance of zero to the agent, social distance is never zero in the village-level data. That means that, in the village-level regressions, the un-interacted incentive-pay coefficient captures the effect of the agent in a village that does not exist in the data.

More realistically, we would want to test the effect of an incentive-pay agent for villages where the average social distance is low, but not zero. For a village with a social distance percentile of 17.5% from the bottom, corresponding to the percentile of the first non-zero distance observation in the household-level data, the difference in effects between flat- and incentive-pay agent is no longer significant at the 10% level (the p-values are 0.547 for knowledge and 0.131 for enrolment).

The difference in results may also relate to the lack of agent- and household-level controls in the village-level specification.

Table D1: Main Results, with Village-Level Controls

	(1)	(2)	(3)	(4)
	Knowledge	Enrolled	Knowledge	Enrolled
Flat-pay agent in village	0.0480 (0.0902)	-0.0361 (0.0455)		
Incentive-pay agent in village	0.248*** (0.0660)	0.0623* (0.0337)	-0.000710 (0.121)	-0.0555 (0.0588)
Social distance			-0.360* (0.195)	-0.224** (0.101)
Incentive pay \times social distance			0.481** (0.201)	0.374*** (0.134)
Village size (households)	0.000139 (0.000121)	-0.0000840 (0.0000598)	0.000200 (0.000147)	0.0000338 (0.0000661)
Eligible population (households)	-0.000551 (0.000542)	0.000221 (0.000275)	-0.00136** (0.000582)	0.0000430 (0.000323)
Village is GP headquarters	0.0148 (0.0866)	0.0204 (0.0405)	-0.0679 (0.111)	0.0122 (0.0525)
Distance to GP headquarters in km (0 if headquarters)	-0.0174 (0.0140)	-0.00406 (0.00720)	-0.0191 (0.0184)	-0.000397 (0.00922)
Distance to nearest town in km	-0.00378 (0.00265)	-0.000253 (0.00158)	-0.00364 (0.00337)	0.00102 (0.00179)
Proportion of village land irrigated	-0.106 (0.0659)	-0.0593 (0.0358)	-0.165** (0.0790)	-0.0557 (0.0433)
Village has drainage sanitation	-0.125 (0.0890)	0.00612 (0.0513)	-0.113 (0.125)	-0.0322 (0.0601)
Average social distance between households in village	-0.567 (0.386)	-0.445** (0.180)	-0.910* (0.496)	-0.378* (0.220)
Agent and household characteristics	No	No	Yes	Yes
Survey wave and taluk fixed effects	Yes	Yes	Yes	Yes
Observations	7582	7582	4732	4732

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. The regressions in columns 3 and 4 use data only from treatment villages, i.e. villages with agents, since social distance is not defined for villages without agents. 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, whether they have a ration card and whether they own their home. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D2: Village-Level Regressions

	(1) Knowledge	(2) Enrolled	(3) Knowledge	(4) Enrolled
Flat-pay agent in village	0.0868 (0.0927)	-0.0174 (0.0506)		
Incentive-pay agent in village	0.253*** (0.0633)	0.0664* (0.0374)	-0.349* (0.202)	-0.296*** (0.108)
Social distance			-0.885** (0.368)	-0.677*** (0.218)
Incentive pay \times social distance			1.331*** (0.425)	0.969*** (0.248)
Village size (households)	0.000191 (0.000126)	-0.0000346 (0.0000610)	0.000333** (0.000139)	0.0000772 (0.0000757)
Eligible population (households)	-0.000915 (0.000621)	0.000168 (0.000285)	-0.000943 (0.000677)	0.0000424 (0.000290)
Village is GP headquarters	0.00573 (0.0800)	0.0172 (0.0460)	-0.0475 (0.0999)	0.0141 (0.0567)
Distance to GP headquarters in km (0 if headquarters)	-0.00988 (0.0122)	0.00119 (0.00701)	-0.0114 (0.0198)	0.00161 (0.0109)
Distance to nearest town in km	-0.00416 (0.00283)	-0.000826 (0.00182)	-0.00460 (0.00339)	-0.00103 (0.00213)
Proportion of village land irrigated	-0.0995 (0.0670)	-0.0571 (0.0395)	-0.134 (0.0815)	-0.0609 (0.0471)
Village has drainage sanitation	-0.0508 (0.0789)	0.0259 (0.0579)	-0.0845 (0.121)	-0.0105 (0.0829)
Average social distance between households in village	-0.337 (0.389)	-0.348* (0.194)	-0.710 (0.491)	-0.473** (0.237)
Taluk fixed effects	Yes	Yes	Yes	Yes
Observations	144	144	105	105

Notes: Ordinary least-squares regressions. The data have been collapsed, across households and survey waves, to a single observation per village. The regressions in columns 3 and 4 use data only from treatment villages, i.e. villages with agents, since social distance is not defined for villages without agents. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E On Magnitude

In the experiment, agents in the incentive-pay group were paid a bonus of 8 rupees (0.17 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 (0.28 USD) per household would suffice to increase by one the average number of correctly answered questions. Column 3 of Table 6 suggests that a bonus of 8 rupees (0.17 USD) per household raises the enrolment rate by 8 percentage points. This modest rate of bonus pay also appears to completely wipe out the knowledge gap between beneficiary groups that are socially proximate to the agent and those that are not.

Given the policy relevance of social programme enrolment, it is of interest to estimate the marginal cost of enrolling an extra household into the programme. Column 3 of Table 6 suggests that incentive-pay agents increased enrolment rates by 7.5 percentage points above the control group. Given that the average village had 52.6 eligible households, our intervention increased enrolment by about 4 households per village. And since the mean incentive pay was 200 rupees per round by design, and there were four rounds in the intervention, the marginal cost of enrolling each household was approximately 200 rupees.

Given that most of the programme effect was achieved in the earlier rounds (Appendix C), one could argue that the programme should be stopped after, say, two rounds. Then the marginal cost of enrolling a household would be about 100 rupees.

Note that this does not take into account the fixed costs of running an incentivised-agents programme (which were substantial, but would matter less if scaled up as a policy), nor the marginal cost of enrolling another household beyond the incentive pay itself (relatively small).

It may appear that the effect of incentive pay is rather large compared to the rates of pay that were offered to the agents. After all, an average payment of 400 rupees

(9 USD) for work over a period of several months is not all that much, even for India's poor. However, whether the job was well paid or not is also a function of the hours put in. While we do not have survey data on total agent time use (only time spent with individual households), examination of field notes indicates that agents spend in the region of 4–5 days of full-time work equivalents per payment period. This is a rough estimate, and clearly there will be substantial variation around the mean. But if reasonably accurate, it would suggest that the average pay per day of work was around 100 rupees (2.17 USD), which is of the same order of magnitude as what agricultural labourers would have earned at the time. One hundred rupees per day is also the wage rate that was offered by the government's large-scale public-works programme, the National Rural Employment Guarantee, in Karnataka at the time of the surveys.

An alternative and more behavioural explanation for the finding that relatively small bonus rates can have large effects on outcomes is that agents may be more sensitive to the fact that there *is* an incentive, than to its size. This is in line with Bénabou and Tirole (2003), who suggest that the fact that incentive pay is offered can itself convey information to the agent about the task, the principal or the principal's view of the agent. Clearly there is no presumption from theory that these effects will always be positive, and indeed Gneezy and Rustichini (2000b) suggest they can be negative. Our finding corresponds with 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 demand for services were found to be sensitive to small incentives (Thornton, 2008; Banerjee *et al.*, 2010b).

F More on Agent Attrition (Online Only)

Four of the 112 recruited agents dropped out of the study.²⁵ Three of these four were from the incentive-pay group. However, we believe that non-random attrition is unlikely to be a major problem for our analysis, for the following reasons: First, recall that we started with 74 agents on incentive pay and 38 on flat pay. Since the number of agents dropping out from these groups was three and one respectively, the rates of attrition were 4.1% and 2.6%, respectively. Given the relatively small sample sizes, the difference in attrition rates appear to be well within what might be expected to occur spontaneously.

As a more formal exploration of whether we may have different attrition rates across the two groups, consider the following. Assume that attrition is driven by a Poisson process with potentially different parameters, a for incentive-pay agents and b for flat-pay agents. (While we cannot test whether attrition is Poisson-distributed, it is a natural starting point for count data.) The null hypothesis is $a = b$, and the question is whether, with 3 out of 74 dropping out for one group, and 1 out of 38 for the other, we can reject this hypothesis.

Under the null hypothesis, our best estimate for $a = b$ would be $(3 + 1)/(74 + 38) = 0.0357$. Given a single Poisson distribution with this parameter, how likely are we to observe what we have observed, or something more extreme? In a Poisson distribution with parameter 0.0357, the probability of observing 3 or more events out of 74 counts (an event rate of 0.041 or more extreme) is 0.49, and the probability of observing 1 or 0 events out of 38 (an event rate of 0.026 or more extreme) is 0.61. So we are far from being able to reject that attrition from the incentive-pay and flat-pay groups are distributed with the same parameter.

Second, note that no agent dropped out before the second survey wave, whereas our key findings hold even if we consider data from the first and/or second survey wave in isolation. This can be seen from Table C1 for the results on the overall effects of agents, and from Table F1 for the social distance results that includes data from the first and

²⁵This and the following online-only appendices contain additional analysis and discussion produced in response to comments from anonymous referees.

second survey waves only. The regressions are consistent with our main findings.

Table F1: *Social Distance Results Using Survey Waves 1 and 2 Only*

	(1) Enrolled	(2) Enrolled
Incentive-pay agent in village	-0.151 (0.152)	-0.113 (0.0722)
Social distance	-0.781*** (0.219)	-0.303** (0.145)
Incentive pay \times social distance	0.809*** (0.250)	0.508*** (0.174)
Agent and household characteristics	Yes	Yes
Survey wave and taluk fixed effects	Yes	Yes
Observations	2335	2335

Notes: Weighted least-squares regressions using data from the first two survey rounds only. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Finally, as noted in the main text, the reasons cited by all four agents for dropping out of the program were either childbirth or migration away from the village. While these reasons are self-reported, we would argue that they are plausible and also plausibly uncorrelated with the contract structure.

G For Whom is Information a Binding Constraint?

(Online Only)

The intervention described in the main text increased enrolment into the scheme by incentivising agents to spread knowledge about it. From a policy perspective, it is of interest to study which groups of households were constrained by a lack of information, in particular as 67% of the control group were enrolled.

While information constraints cannot be observed directly, it is possible to study which household characteristics predict enrolment in the absence of our intervention, i.e. in the control villages. (As explained in the main text, we only have data on a few household-level variables.)

In the regressions reported in Table G1, we are able to include village fixed effects. Standard errors are again clustered at the village level. As can be seen from the table, none of the four household characteristics we have data on are predictive of enrolment at the 10% level. It is perhaps surprising that enrolment does not seem to run along caste or educational lines (conditional on eligibility). It suggests that reducing programme cost by targeting specific demographics within the eligible population may have limited scope.

In column 2 we include the household's knowledge test z -score to confirm that, even conditional on household characteristics and inclusive of village fixed effects, knowledge and enrolment are positively correlated.

Table G1: *Predictors of Enrolment in Absence of Intervention*

	(1) Enrolled	(2) Enrolled
Household forward/dominant caste	0.00731 (0.0370)	0.0206 (0.0310)
Household head has completed primary school	-0.000685 (0.0321)	-0.0209 (0.0297)
Household has ration card	0.0273 (0.0565)	0.0317 (0.0564)
Household owns their home	0.0107 (0.0262)	0.000129 (0.0236)
Knowledge		0.166*** (0.0147)
Survey wave fixed effects	Yes	Yes
Village fixed effects	Yes	Yes
Observations	1756	1756

Notes: Weighted least-squares regressions using control-group data only. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

H Direct Comparison of the Flat-Pay and Incentive-Pay Groups with the Control Group (Online Only)

Regressions in Tables 8, 10 and 11 only use data from treatment villages, i.e. villages with agents, since the distance measures are not defined for villages without agents. A potential concern is that the difference in impact between socially proximate and distant households for agents who are on flat pay may be driven by pre-existing differences in knowledge and/or enrolment across groups of households. In particular, the concern is that the agent may be selected from a group that is more knowledgeable even in the absence of the programme.

However, the finding is unlikely to be driven by one sub-group in the village being better informed than others, since all regressions involving social distance controls for agent as well as household characteristics in levels. That is, flat-pay agents are more effective with respect to their “friends” than their “non-friends”, *holding fixed* caste, education, home ownership and ration card status for both agent and households. To the extent that the group with more pre-existing knowledge can be captured by these characteristics, the concern should be allayed.

As a further robustness check, we regressed our dependent variables, knowledge and enrolment, on the treatment variables (flat-pay agent and incentive-pay agent) as well as interactions between the treatment variables and an indicator for ‘low social distance’. This indicator is 1 if the composite distance measure is 0 or 0.25, and 0 otherwise. Crucially, the ‘low social distance’ indicator is set to 0 for the control villages, so that these do not drop out of the regressions. This permits a direct comparison of the two treatment groups with the control group. The analysis is also useful to rule out the possibility that a flat-pay agent actually makes outcomes for socially distant households worse, relative to the control group.

The results, presented in Table H1, corroborate the findings from our main analysis. Agents on flat pay are associated with a significant and positive effect on knowledge and enrolment for households that are similar to them. But for households who are socially

distant from the agent, a flat-pay agent is no more effective than having no agent at all. (This is now a direct comparison with the control group.) Agents on incentive pay are associated with significant and positive effects on knowledge and enrolment for both categories of households, compared to the control group, and the size of this effect is close in magnitude to the effect that flat-pay agents have on households that are similar to them. As elsewhere, the regressions include survey wave and taluk fixed effects, and the standard errors are clustered at the village level.

Note that, since ‘low social distance’ has been set to zero for all households in the control villages, the un-interacted ‘low social distance’ coefficient is not separately identifiable in these regressions.

Table H1: *Direct Comparisons of Flat-Pay and Incentive-Pay Groups with Control Group*

	(1) Knowledge	(2) Enrolled
Flat-pay agent in village	-0.0210 (0.0885)	-0.0546 (0.0489)
Flat-pay agent \times low social distance	0.231** (0.0889)	0.115*** (0.0440)
Incentive-pay agent in village	0.236*** (0.0644)	0.0946*** (0.0348)
Incentive-pay agent \times low social distance	0.0463 (0.0535)	-0.0472* (0.0262)
Survey wave fixed effects	Yes	Yes
Taluk fixed effects	Yes	Yes
Observations	7730	7730

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. The variable ‘low social distance’ has been set to zero for all control-group observations, so these are not dropped from the regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We would like to point out that while it is nice to benchmark the effects of the two treatment groups against the control group, this comes at the cost of not being able to control for agent characteristics in the regressions. Setting social distance to ‘high’ for the control villages, where there are no agents and hence no meaningful ‘distance’ between

agents and households, is also somewhat artificial.

I Triple Differences by Caste Segregation (Online Only)

It is of potential interest to examine whether our key findings on incentives and social distance vary by the level of caste segregation in the village. This question is analysed in Table I1. We capture caste segregation by whether or not ‘castes live apart’, a binary indicator for whether castes are separated into separate zones in the village. There is a drop in power here, associated with the fact that (as explained in the main text) the ‘castes live apart’ variable was constructed based on recall and hence not observed everywhere, as well as with the more demanding econometric specification. As elsewhere, the regression controls for agent and household characteristics in levels, as well as taluk and survey wave (time) fixed effects, and standard errors are robust and clustered at the village level.

The main results carry over: social distance appears to have a negative effect on knowledge transmission, and this effect is ‘cancelled out’ by the introduction of incentive pay. It does not appear that these effects are weaker or stronger depending on whether castes are segregated in the village i.e., the triple interaction is not statistically significant. In other words, our main finding on incentives and social distance holds in both types of villages.

Table I1: *Triple Interaction with ‘Castes Live Apart’*

	(1) Knowledge
Incentive-pay agent in village	-0.105 (0.138)
Social distance	-0.427* (0.225)
Castes live apart	-0.359 (0.259)
Incentive pay \times social distance	0.528** (0.241)
Incentive pay \times castes live apart	0.317 (0.272)
Social distance \times castes live apart	0.197 (0.453)
Incentive pay \times social distance \times castes live apart	-0.141 (0.481)
Agent and household characteristics	Yes
Survey wave and taluk fixed effects	Yes
Observations	3877

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