



Estimating within-cluster spillover effects using a cluster randomization with application to knowledge diffusion in rural India

Arthur Alik-Lagrange¹ | Martin Ravallion^{2,3}¹World Bank, Washington, DC²Department of Economics, Georgetown University, Washington, DC³NBER, Cambridge, MA**Correspondence**

Martin Ravallion, Department of Economics, Georgetown University, Washington, DC 20007.

Email: mr1185@georgetown.edu

Summary

Spillover effects within randomized clusters pose a challenge for identifying impacts of an individualized treatment. The paper proposes a solution. Longitudinal and intra-household observations are combined in estimating the direct knowledge gain from watching an info-movie in rural India, while randomized village assignment identifies knowledge sharing. Simulations on synthetic data and econometric tests provide support for the estimation method. We find evidence of information sharing, but far less so for disadvantaged groups, such as illiterate and lower-caste individuals; these groups rely more on actually seeing the movie. Our results are consistent with sizeable biases in ordinary least squares, matching or instrumental variable impact estimators that ignore within-cluster spillovers.

1 | INTRODUCTION

Many social programs take the form of an individually assigned intervention (such as a health program, a transfer payment, or access to credit). Some people take up the program and some do not. In evaluating such programs, a seemingly attractive option is to randomly assign access across clusters, such as geographic areas. Take-up within clusters is endogenous but the randomized assignment across them is used as the instrumental variable (IV) (Angrist, Imbens, & Rubin, 1996). This is often called cluster randomization with selective take up.

Spillover effects are an obvious threat to the validity of this approach and there are various examples in the literature.¹ In the context of the geographically assigned programs studied in this paper, the concern is not contamination effects between the selected areas but within them. To the extent that those within an assigned area who do not directly participate in the program are affected (positively or negatively), the exclusion restriction required by the IV estimator will clearly not hold and the estimates will be biased. If the investigator can do a double randomization then the problem can be avoided (Baird, Bohren, McIntosh, & Özler, 2018). Two control groups then allow identification of the spillover effect.

However, cluster randomization makes sense in situations in which randomization is not feasible within clusters. For example, in the application studied here, an info-movie provided information about a large anti-poverty program,

¹Spillover effects have been identified for: (i) health treatments within schools (Miguel & Kremer, 2004); (ii) schooling in the neighborhoods of transfer recipients (Angelucci & De Giorgi, 2009; Bobba & Gignoux, 2014; Bobonis & Finan, 2009); (iii) local government spending in response to geographically assigned programs (Chen, Mu, & Ravallion, 2009); (iv) crime displacement (Yang, 2008); and (v) a women's empowerment program to encourage child immunization (Janssens, 2011).

implemented under India's National Rural Employment Guarantee Act (NREGA).² The movie could be randomly assigned across villages but access to the movie could not be randomly assigned within villages. There would have been ethical concerns as to why some people got a ticket to see the movie and some did not. There was no plausible reason for rationing, such as due to budget or capacity constraints. And, no less importantly, the majority of villages lacked a venue that allowed exclusion based on not having a ticket. Identification then poses an econometric problem even with village-level randomization.

The paper proposes an original solution to the problem of distinguishing the direct effect from the spillover effect in cluster randomizations with selective take-up but in which double randomization is not feasible.^{3,4} We postulate that there is a latent individual effect on outcomes. In our application, this latent individual effect represents the individual's social "connectedness," which is taken to depend on long-standing networks of association between people, reflecting how each individual fits within the village social structure, including caste position and the ability of that individual to process new information. We also allow for household-specific time-varying knowledge, reflecting the information that is held collectively among individuals within a given household. Economic theories of the household have often assumed that information is shared fully *within* the household.⁵ We do not make that assumption here, although we do allow for some degree of knowledge sharing.

The essential idea is to combine a double difference method to estimate the direct impact (allowing for endogenous take-up) while using randomized assignment across villages to identify the spillover effect. For the purpose of identifying the direct effect, a survey design is required that combines a standard $T = 2$ panel structure with two separate adult interviews for each household. In our application, data were collected on personal knowledge about NREGA and whether the respondent has personally watched the movie. Having two observations within each household allows us to obtain estimates that are robust to latent heterogeneity in household factors, similar to the sibling-difference estimator that has been used in the literature to eliminate a confounding household effect in studying outcomes for children.⁶ By exploiting the differences over time, our nonexperimental estimator is also robust to time-invariant individual effects, such as latent social connectivity within the village.

Our key identifying assumption is that the intra-household differences in the probability of watching the movie are uncorrelated with any changes over time in the intra-household differences in knowledge about the program. While we recognize that such a correlation cannot be ruled out in the absence of a second (within-village) randomization, we provide a supportive test by estimating an augmented regression.

The paper finds evidence of spillover effects, which account for about one-third of the average impact of the information campaign.⁷ While knowledge sharing is evident, the paper finds that poorer people, by various criteria, benefit less from the spillover effect—relying more on direct exposure to the intervention. We demonstrate that there is a substantial bias in the IV estimator of the average treatment effect on the treated due to a failure of the exclusion restriction, stemming from knowledge spillovers. We also find that ordinary least squares (OLS) and matching estimators controlling for observable household and individual characteristics show indications of sizeable bias due to unobserved heterogeneity between those who did and those who did not watch the movie.

²Useful overviews of the arguments and evidence on factors relevant to the role of information and information campaigns in developing countries can be found in Keefer and Khemani (2005), Khemani (2007), and Mansuri and Rao (2013). The review by La Ferrara (2015) focuses specifically on entertainment media ("edutainment"). On social learning in Indian agriculture see (inter alia) Foster and Rosenzweig (1995) and Munshi (2004).

³For an overview of (experimental and nonexperimental) methods for estimating spillover effects see Angelucci and Di Maro (2015). The method used in the present paper is not one of those reviewed by Angelucci and Di Maro.

⁴The setting here assumes (i) a valid cluster-level randomization and (ii) no interference between treatment and nontreatment clusters. This differs from existing observational solutions to the existence of clustered interference between treated and untreated units (see, e.g., Barkley, Hudgens, Clemens, Ali, & Emch, 2017).

⁵This is the case for Chiappori's (1988) influential collective model of the household. Mazzocco (2007) develops an alternative model based on limited commitment that can yield inefficiencies in household decision making. Evidence consistent with incomplete knowledge sharing within households can be found in Bloch and Rao (2002), Ashraf (2009), and Ambler (2015). Udry (1996) found intra-household differences in farm productivity favoring plots controlled by men over those by women.

⁶Early examples of the sibling difference method of addressing household heterogeneity in estimating models for child outcomes include Rosenzweig and Wolpin (1988) and Duncan, Yeung, Brooks-Gunn, and Smith (1998).

⁷Using the same data, Ravallion, van de Walle, Murgai, and Dutta (2015) reports largely non-significant ITT effects for the same information campaign on other objectively-measured outcomes, such as participation in MNREGA. The present paper does not report on spillovers for such outcomes but focuses instead on knowledge about the scheme.

The following section outlines the proposed econometric method. Section 3 describes the setting and data. Section 4 presents the results. Section 5 concludes.

2 | THE MODEL AND ESTIMATION METHOD

The main task of the empirical analysis is to identify the knowledge spillover effect within villages and to see how this varies by socioeconomic group.

Whether a person decides to take up the intervention assigned at village level—in our application, to watch the movie or not—is obviously correlated with how much its members know already about the program and how much they expect to learn about it. Moreover, the decision on which member of the household watches the movie might well be determined by village, household, and individual characteristics potentially correlated with individual specific knowledge about the program or ability to learn. Recognizing this potential endogeneity, we develop an identification strategy allowing us to control for household–village time-varying fixed effect as well as time-invariant village–household–individual fixed effects. We provide a theoretical model supporting our econometric model.

2.1 | Theoretical model

For concreteness, we focus on our specific application, but the model and estimation method can be readily adapted to other relevant settings where a treatment can be administered to clusters—in our case, villages—but not to specific individuals within these clusters. Knowledge about NREGA is denoted by k_{ijvt} for person i in household j in village v at time t . The baseline survey at $t = 1$ includes the NREGA quiz for measuring k_{ijvt} . The treatment—a movie screening—is administered to a subsample of villages and a follow-up survey is done at $t = 2$. The movie is randomly assigned to villages, with the assignment denoted $M_{v2} = 1$ if village v got the movie and $M_{v2} = 0$ otherwise. The incidence of watching the movie is denoted by m_{ijvt} , with $m_{ijv2} = 1$ if person i in household j and village v saw the movie and zero otherwise. In $t = 1$ the movie has not yet been screened, so by definition $m_{ijv1} = 0$ and $M_{v1} = 0$ for all i, j, v .

To isolate knowledge sharing we must control for the direct (idiosyncratic) gains from watching the movie. Conditional on individual exposure to the movie (either directly or via the spillover effect), we also allow for a latent household and village effect in knowledge, denoted by η_{jvt} . This captures any collectively shared knowledge within the household.

To capture the idea of “social inclusion” in matters of knowledge sharing, we also postulate a time-invariant individual effect, denoted by δ_{ijv} .⁸ This can be taken to reflect how much each individual is able to tap into the general spread of knowledge given his or her social position. One can thus think of δ_{ijv} as reflecting the person’s “social connectivity” within the village. Idiosyncratic differences in baseline knowledge are also included in δ_{ijv} . Combining these observations, we postulate a knowledge production function of the form

$$k_{ijvt} = k(M_{vt}, m_{ijvt}, \eta_{jvt}, \delta_{ijv}) \quad (1)$$

Knowledge spillovers generate an effect on k_{ijvt} of going from $M_{vt} = 0$ to $M_{vt} = 1$ at given values of m_{ijvt} , η_{jvt} and δ_{ijv} .

Watching the movie is, of course, a matter of personal choice, so it is potentially endogenous to knowledge about NREGA. There can be a direct utility gain from watching the movie in addition to the benefit from the knowledge gained. (A person may watch the movie solely for its entertainment value.) In deciding whether or not to watch the movie, a person weighs the benefit against the cost. The benefit is the utility gain, where the utility function for period 2 is

$$u_{ijv2} = u(m_{ijv2}, k_{ijv2}, \delta_{ijv}, \eta_{jv2}). \quad (2)$$

When δ_{ijv} is low, the person will have a higher expected benefit from watching the movie but a lower level of knowledge about programs such as NREGA. The cost of watching the movie is mainly the opportunity cost of time, which will tend

⁸The literature on poverty has often ascribed a role for social exclusion, given that this can be expected to come with weak connectivity to relevant sources of knowledge. The existence of such constraints on access to public support has been seen as a defining characteristic of social exclusion since that term first appeared in debates on French social policy; see, for example, the discussion in Silver (1994).

to be lower for the poor, and especially the underemployed poor, who are likely to be especially interested in NREGA. There can also be idiosyncratic factors that influence the costs of watching the movie (such as the need to attend to a sick family member). For the present purpose, we can think of the cost of watching the movie as a random variable with distribution function F_{ijvt} . The probability of watching the movie is then a function of the utility gain from doing so, which depends on the gain in knowledge due to watching the movie, as well as δ_{ijv} and η_{jvt} :

$$P(m_{ijvt} = 1) = F_{ijvt} \left[u \left(1, k \left(1, 1, \delta_{ijv}, \eta_{jvt} \right), \delta_{ijv}, \eta_{jvt} \right) - u \left(0, k \left(0, 1, \delta_{ijv}, \eta_{jvt} \right), \delta_{ijv}, \eta_{jvt} \right) \right] \quad (3)$$

We see in Equation 3 that an identification strategy controlling for variations in δ_{ijv} and η_{jvt} would then get rid of the potential source of endogeneity. This is the objective of our main identification strategy described in the next section.

2.2 | Econometric model

We treat the knowledge gains from watching the movie as additive and we treat the values of δ_{ijv} and η_{jvt} as unobserved additive effects. To explain our econometric model for estimating the spillover effect (SE), consider first the levels of knowledge in the (pre-intervention) baseline ($t = 1$):

$$k_{ijv1} = \alpha + \eta_{jv1} + \delta_{ijv} + \varepsilon_{ijv1}. \quad (4)$$

Now introduce the information intervention. Let β_1 denote the knowledge gain for someone who watches the movie, and β_2 the indirect knowledge gain for someone who does not watch the movie but lives in a village where the movie was shown. Without knowledge sharing we have $\beta_2 = 0$ while (at the other extreme) with full sharing, $\beta_1 = \beta_2$. In the latter case, anything learned by someone who watches the movie is accurately and completely conveyed to all who did not do so within the same village.

Next, let k_{ijv2}^* denote the counterfactual knowledge in period 2—that is, knowledge in the absence of the intervention. Thus, by definition:

$$k_{ijv2} = k_{ijv2}^* + \beta_1 \text{ for } M_{v2} = 1 \text{ } m_{ijv2} = 1, \quad (5a)$$

$$k_{ijv2} = k_{ijv2}^* + \beta_2 \text{ for } M_{v2} = 1 \text{ } m_{ijv2} = 0, \quad (5b)$$

$$k_{ijv2} = k_{ijv2}^* \text{ for } M_{v2} = 0 \text{ } m_{ijv2} = 0 \quad (5c)$$

We do not, of course, observe k_{ijv2}^* . (While the movie is randomly assigned, the control group only reveals the mean value of k_{ijv2}^* .) Consistent with Equation 4, we assume that

$$k_{ijv2}^* = \alpha + \eta_{jv2} + \delta_{ijv} + \varepsilon_{ijv2} \quad (6)$$

Combining these assumptions, we have the following switching regression:

$$k_{ijv2} = \alpha + \beta_2 M_{v2} + (\beta_1 - \beta_2) m_{ijv2} + \eta_{jv2} + \delta_{ijv} + \varepsilon_{ijv2} \quad (7)$$

Two “benchmark” null hypotheses are of specific interest as restricted forms of this regression. The first is the null that knowledge is shared fully ($H_0 : \beta_1 = \beta_2$); the second null is that of no knowledge sharing ($H_0 : \beta_2 = 0$).

Following our discussion of the choice of whether to watch the movie or not, there are two confounding effects in the composite error term of Equation 7 ($\eta_{jv2} + \delta_{ijv} + \varepsilon_{ijv2}$), namely that $\text{Cov}(m_{ijv2}, \eta_{jv2}) \neq 0$ and $\text{Cov}(m_{ijv2}, \delta_{ijv}) \neq 0$. Clearly, the η_{jvt} effect cannot be eliminated using individual panel data alone.

Our identification strategy exploits the fact that we have two observations of individual knowledge within the sampled households, as well as panel data. Taking the difference over time eliminates the individual knowledge effect, δ_{ijv} ,

but still leaves the change over time in the household–village knowledge effect:

$$\Delta k_{ijv2} = \eta_{jv2} - \eta_{jv1} + \beta_2 M_{v2} + (\beta_1 - \beta_2) m_{ijv2} + \Delta \varepsilon_{ijv2} \quad (8)$$

(here Δ denotes the difference over time; also note that $M_{v1} = m_{ijv1} = 0$.) Taking the difference between the two adults interviewed in each household, we obtain

$$\nabla \Delta k_{.jv2} = (\beta_1 - \beta_2) \nabla m_{.jv2} + \nabla \Delta \varepsilon_{.jv2} \quad (9)$$

Here ∇ denotes the difference between two adults in the same household. By estimating this regression separately for different strata, we test whether intra-household knowledge sharing is stronger for some socioeconomic groups than others.

Given that the movie screenings are randomly assigned across villages, it can be assumed that $Cov(M_{v2}, \eta_{jv2}) = Cov(M_{v2}, \delta_{ijv}) = Cov(M_{v2}, \varepsilon_{ijv2}) = 0$ (for $t = 1, 2$). However, as already noted, the choice of whether to watch the movie (in an assigned village) is endogenous. Our key identifying assumption is that this endogeneity stems solely from the two effects: δ_{ijv} and η_{jvt} . While we allow $Cov(m_{ijv2}, \eta_{jv2}) \neq 0$ and $Cov(m_{ijv2}, \delta_{ijv}) \neq 0$ consistently with Equation 3, we assume that, once we have eliminated these two effects, there is no remaining endogeneity; that is,

$$Cov(\nabla m_{.jv2}, \nabla \Delta \varepsilon_{.jv2}) = 0 \quad (10)$$

We can recognize Equation 10 as a version of the double-difference assumption of “parallel trends”; we discuss this assumption further below.

Under our assumptions, OLS applied to Equation 9 gives a consistent estimate of $\beta_1 - \beta_2$, denoted by $\widehat{\beta}_1 - \widehat{\beta}_2$. On replacing $\beta_1 - \beta_2$ by this estimate, Equation 7 becomes

$$k_{ijv2} - (\widehat{\beta}_1 - \widehat{\beta}_2) m_{ijv2} = \alpha + \beta_2 M_{v2} + \eta_{jv2} + \delta_{ijv} + \varepsilon_{ijv2} \quad (11)$$

Given that the movie is randomly assigned, we can then estimate Equation 11 consistently by OLS.

Note that our balanced panel comprises one male and one female in almost all households. Thus we cannot identify the gender difference in the direct impact of watching the movie, although we can test for a gender difference in the spillover effect.

2.3 | Evaluation parameters

We define the following parameters based on the above model:

$$ATET = E[k_{ijv2} | M_{v2} = 1, m_{ijv2} = 1] - E[k_{ijv2} | M_{v2} = 0, m_{ijv2} = 1] = \beta_1, \quad (12a)$$

$$SE = E[k_{ijv2} | M_{v2} = 1, m_{ijv2} = 0] - E[k_{ijv2} | M_{v2} = 0, m_{ijv2} = 0] = \beta_2, \quad (12b)$$

$$ITT = E[k_{ijv2} | M_{v2} = 1] - E[k_{ijv2} | M_{v2} = 0] = \beta_2 + (\beta_1 - \beta_2) E[m_{ijv2} | M_{v2} = 1]. \quad (12c)$$

These are standard parameters adapted to the existence of spillover effects. The average treatment effect on the treated (ATET) is the impact on knowledge of watching the movie relative to the control villages where the movie was not shown. This combines the effect of seeing the movie and living in a village where the movie was shown. The spillover effect (SE) is the indirect effect alone—the knowledge spillover. The spillover has two components: between and within households, which we do not identify separately.⁹ ITT is the usual intent-to-treat parameter.¹⁰

⁹The between case is obvious. If one closed off all communication between households but people who did not watch the movie learn about it by talking to a family member who did, then we will still find that $SE > 0$. It can be shown that SE is the weighted sum of the spillovers within households and those between households within villages. A proof of this claim is available in Supporting Information Addendum 4.

¹⁰Note that we deviate here from the notation proposed by Baird et al. (2018), developed for the estimation of spillovers in a two-stage clustered randomized trial. Our method applies in any single-stage random clustered trial where the intent to treat units are clusters. Baird et al. describe our ITT parameter (Equation 12c) as the total causal effect (TCE).

2.4 | Test of the parallel trends assumption

The assumption in Equation 10 can be questioned in that different individuals within a given household may have different networks and information flows relevant to their program knowledge. The selection process determining which individual watches the movie may then reflect latent *time-varying* differences in knowledge within a given household.

Recognizing this concern, we provide a test. The idea is that if there are latent time-varying differences in knowledge relevant to whether one chooses to watch the movie then they should be contained in the differences in the probability of watching the movie unexplained by differences in observable characteristics. Therefore we test the parallel trends assumption using the within-household difference in individual residuals from a probit model of the probability to watch the movie as a function of covariates measured at $t = 1$. Specifically, we estimate the following augmented regression:

$$\nabla \Delta k_{.jv2} = (\beta_1 - \beta_2) \nabla m_{.jv2} + \gamma [\nabla m_{.jv2} - \nabla \hat{p}(m_{.jv2} | X_{ijv1})] + \nabla \Delta \varepsilon_{.jv2} \quad (13)$$

Here $\hat{p}(m_{.jv2} | X_{ijv1})$ is the predicted probability of watching the movie based on the baseline characteristics (i.e., the propensity score).¹¹ The residuals, $\nabla m_{.jv2} - \nabla \hat{p}(m_{.jv2} | X_{ijv1})$, should contain the confounding time-varying knowledge differences within households. We do this test for all our stratified versions of Equation 9, allowing for differences in the direct effect of watching the movie according to household characteristics.

2.5 | Comparison with other methods

We compare our estimates with a number of other approaches found in the literature. One of these is the OLS regression model motivated by the switching regression in Equation 7. We do this with and without controls. Another is a propensity score matching (PSM) estimator, which makes fewer assumptions than OLS—in particular, it does not assume a linear parametric model for knowledge.¹² We can write the PSM estimators for SE and ATET as follows:

$$SE = E[k_{ijv2} | M_{v2} = 1, m_{ijv2} = 0, p(X_{ijv1})] - E[k_{ijv2} | M_{v2} = 0, m_{ijv2} = 0, p(X_{ijv1})] \quad (14a)$$

$$ATET = E[k_{ijv2} | M_{v2} = 1, m_{ijv2} = 1, p(X_{ijv1})] - E[k_{ijv2} | M_{v2} = 0, m_{ijv2} = 1, p(X_{ijv1})], \quad (14b)$$

where $p(X_{ijv2})$ is the conditional probability for an individual in movie villages to watch the movie as a function of a vector of round 1 individual, household and village characteristics, X_{ijv1} . (This can be motivated by (3) but replacing δ_i, η_{jvt} by the vector of observables X_{ijv1} .) This estimator differs from the naïve OLS with controls but relies on the same assumption that the endogeneity issue can be fully addressed in controlling for observed characteristics. A bias will remain in the matching estimate if the decision to watch the movie reflects unobserved characteristics.

Another parameter of interest is

$$LATE = ITT / E[m_{ijv2} | M_{v2} = 1]. \quad (15)$$

This is the usual local average treatment effect; in this case, it is equivalent to estimating the average treatment effect of watching the movie using the randomized assignment at village level as the IV. This might help address the endogeneity issue, but it is expected to deliver biased estimates in the presence of knowledge spillovers within movie villages.

¹¹We assume here that $\hat{p}(m_{.jv2} | X_{ijv1})$ has the same functional form for both household members.

¹²PSM has been used often in estimating ATET and it is proposed by Angelucci and De Giorgi (2009) as a candidate for estimating spillover effects. One might also use PSM to create a comparison group (for those who watched the movie) among those who did not watch the movie within each village. However, the within-village samples are clearly too small for this purpose. Also note that such a PSM estimator makes a stronger conditional independence assumption than our estimator.

3 | BACKGROUND ON THE SETTING AND DATA

In rural India, caste creates special frictions in social interaction and (hence) information diffusion, especially in rural areas. Dalits (also called Scheduled Castes) have faced a long history of discrimination and exclusion.¹³ It is rare for lower- and upper-caste families to be close neighbors within the village, and the “Dalit only” area of the village is typically quite well defined and known to all. The degree of social and political connectivity in this context is greater for more advantaged castes (see, e.g., Desai et al., 2010).

In such a setting, knowledge about a new antipoverty program may diffuse rather poorly by word of mouth, especially if the information enters via the local elite. Unless a poor, lower-caste, individual comes into direct contact with the source of new information she may come to know little about a program intended to help people like her. Strategic behavior may act to worsen the information flow. For example, if the program has potentially adverse impacts for the village elite, who are better connected to reliable knowledge sources than the poor, then misinformation may be spread for strategic reasons. Similarly, eligible participants who anticipate rationing of the opportunities announced in an information campaign may rationally choose not to spread the word. Holding a public meeting as part of the campaign might well do a better job of knowledge diffusion than (say) an official letter to the village leader. However, it remains unclear just how effective a public meeting will be for knowledge sharing and coordination in unequal societies where socially nonneutral selection processes are at work in determining who participates and influences the agenda and deliberations; pre-existing inequalities are unlikely to vanish in such a public meeting (Heller & Rao, 2015).

These observations beg a number of questions: Does information about a public program spread reliably beyond those who learn about it directly? Does information flow more easily among some groups than others within a village? In particular, does poverty come with exclusion from information about programs designed to help poor people?

This paper focuses on the socioeconomic differences in the extent of knowledge sharing within households and villages. We use an information campaign to identify key aspects of how knowledge is shared, as one element in understanding program efficacy. An information campaign typically has both a direct channel of learning—direct exposure to the specific intervention—and an indirect channel through social interaction. In the absence of any social frictions to knowledge diffusion, the incidence of direct exposure should matter little. But if those frictions are strong then direct exposure will be the main channel of learning. The paper shows that an information campaign can be used to identify a parameter for knowledge sharing.

The campaign studied here used an entertaining fictional movie to teach people their rights under India's NREGA, which created a justiciable “right-to-work” for all rural households in India. The most direct and obvious way NREGA tries to reduce poverty is by providing extra unskilled manual employment in rural areas on demand. This requires an explicit effort to inform and empower poor people, who are encouraged to take deliberate unilateral actions to demand work on the scheme from local officials. The stipulated wage rate is often above the local market wage rate for similar work. This feature of the program may be seen as a threat to the landholding elites in traditional (primarily agricultural) villages; in particular, if the program worked well then it would put upward pressure on local wage rates, reducing farm profits. Against this, the program may discourage out-migration in the lean season, thus stabilizing labor supply across seasons to the benefit of labor-demanding farmers.

The setting is rural Bihar—a relatively poor state of about 100 million people in the northeast of India. In previous research, it was found that most men and three-quarters of women had heard about NREGA, but most were unaware of their precise rights and entitlements under the scheme (Dutta, Murgai, Ravallion, & van de Walle, 2014). With the aim of promoting better knowledge about NREGA in this setting, the movie was randomly assigned to sampled villages, with a control group not receiving the movie. Knowledge about NREGA was assessed in both treatment and control villages. Residents were encouraged to watch the movie; some did and some did not.

A previous paper studied the “intent-to-treat” (ITT) impacts of the movie on knowledge, perceptions, and outcomes for program participation (Ravallion et al., 2015). The present paper goes more deeply into the impacts on knowledge, and the channel of that impact—notably decomposing ITT impact in a direct effect of watching the movie and an indirect effect through knowledge sharing within villages. Importantly, the paper also explores whether this indirect effect varies between socioeconomic groups.

Bihar is one of the poorest two or three of India's states. Based on official Planning Commission poverty lines for 2009/10, 55% of its rural population of 90 million lived below the poverty line. Although one would hope that NREGA

¹³For example, in the majority of Indian villages today, Dalits are not allowed to share food with non-Dalits. A graphic account of the treatment of Dalits in much of rural India can be found in Human Rights Watch (2007).

worked well in India's poorest states (where it is presumably needed most), Bihar has one of the lowest participation rates of any state (Dutta et al., 2014). The scheme that implements NREGA in Bihar is the Bihar Rural Employment Guarantee Scheme (BREGS).¹⁴

3.1 | The survey data

The data were collected explicitly for the purpose of evaluating the scheme's performance, as documented in Dutta et al. (2014). Two survey rounds were done spanning 150 randomly chosen villages in rural Bihar.¹⁵ The first round (R1) was implemented between May and July of 2009 and the second (R2) during the same months 1 year later.¹⁶ The timings were chosen for being lean periods for agricultural work, and were thus expected to be peak periods for BREGS. The survey collected information at household and individual levels on a range of characteristics including caste, demographics, asset ownership, consumption, employment, and wages, as well as information on BREGS participation, knowledge of NREGA rights and rules, process-related issues, and questions related to perceptions about the scheme in the specific village context. In addition to the main household roster, two adults were interviewed, with preference given to public works participants when relevant.

The two-stage sampling design was based on the 2001 Census list of villages. In the first stage, 150 villages were randomly selected from two strata, classified by high and low BREGS coverage based on administrative data for 2008/9. In the second stage, 20 households per village were randomly selected, drawing from three strata based on an initial listing of all village members and a few selected attributes. All summary statistics reported in the study are weighted with appropriate sample weights to be representative at state level and the regressions allow for the survey design.

Our estimation method (discussed in detail in Section 2) requires that we focus on the subsample of panel households in which two members were interviewed in both rounds. This “2 × 2” subsample has 2,376 individuals. The fact that we exclude households for which only one adult was available for interview suggests the possibility of sample selection bias. The 2 × 2 subsample differs from the full sample in some respects. For example, the two-adult panel is more likely to be male headed, less likely to have a widowed respondent, and more likely to be a larger household.¹⁷ We test whether the sample selection affects our results for those evaluation parameters that do not require the balanced panel, which gives a sample of 4,792 individuals.

3.2 | Heterogeneity analysis

It is of interest to test for heterogeneity in the direct and spillover effects in the R1 data, though recognizing that stratification comes with a loss of power. In one case we stratify by three groups of households defined by actual or desired BREGS participation in R1, before the information campaign. First, there are the actual participants in BREGS. Second, we identify a group of “excess demanders,” defined as those who said they wanted BREGS work but did not get it; past research has indicated substantial unmet demand for work on the scheme (Dutta, Murgai, Ravallion, & van de Walle, 2012, 2014). The remainder forms the third group, identified as those who were not interested in participating. We expect the information intervention (discussed further below) to have more impact on the second group, since they express a desire to participate—which presumably motivates learning—but are not participating.

We also stratify by four socioeconomic indicators: caste, literacy, landlessness, and consumption poverty. Two caste groups are distinguished: the first is the less advantaged group—the Scheduled Castes comprising Dalits and Mahadalits (a local Bihar category, recognized by the state government); the second is mainly Other Backward Castes and the general caste group (not one of the commonly identified poorer caste/tribal groups). “Landlessness” is defined as owning no cultivatable land. “Poverty” is identified by a household consumption per person below the R1 median (which closely accords with the official poverty measure for Bihar). The sample is clearly not evenly spread across the combinations of these dimensions, as can be seen from Table 1, which gives the number of sample points in each cell combination. Of those who are both in the lower-caste grouping and consumption poor, 86% are landless, and in 57% of cases both

¹⁴The corresponding national program is called the Mahatma Gandhi National Rural Employment Guarantee Scheme.

¹⁵The survey instrument is available to readers online at <http://explore.georgetown.edu/people/mr1185/>. The data are available for replication purposes from the authors.

¹⁶The overall attrition rate between the two rounds was 8% and was not concentrated in any particular stratum (Ravallion et al., 2015).

¹⁷This was tested using a probit for whether the individual was in the balanced panel using covariates from R1. The probit is available from the authors.

TABLE 1 Distribution of sampled individuals across characteristics

		Literacy and land				Total
		Both illiterate + landless	Both illiterate + some land	1 or 2 literate + landless	1 or 2 literate + some land	
Caste and poverty	Lower caste + poor	212	22	110	30	374
	Lower caste + non-poor	172	30	98	32	332
	Upper caste + poor	190	136	156	236	718
	Upper caste + non-poor	204	168	186	394	952
Total		778	356	550	692	2,376

persons are also illiterate. Among those who are in our upper caste grouping and not consumption poor, we find that in 41% of cases at least one of the two persons is literate and the household has some land. But the associations are far from perfect; for example, the instances of both adults being illiterate are similar for the two caste groups. Also note that counts for some cells are quite low, suggesting caution about inferences for these cells.

3.3 | Measuring knowledge of NREGA rights and rules

In both R1 and R2, the respondents were asked whether they had heard of NREGA and, if so, they were asked 12 questions testing their knowledge of the scheme's functioning and their rights. We call this the NREGA quiz. This was administered separately to one male and one female member of each sampled household when feasible. The quiz covered the main provisions of NREGA, including wages, employment, and the facilities mandated for worksites. Knowledge about employment and wages is clearly more important to the lives of poor people in this setting than knowledge about facilities at worksites.

As an overall measure of knowledge about the scheme's employment aspects we calculate an "employment knowledge index" as the number of correct answers to the eight employment- and wage-related questions in the NREGA quiz. The average score on employment knowledge in R1 is 2.6 for men and 1.5 for women (out of a maximum of 8).

A second measure can be created for knowledge of the facilities and amenities that the scheme mandates must be provided at work sites (daycare, drinking water, shade, and first aid kits). Respondents were asked to identify what facilities were supposed to be provided. We call this the "facilities knowledge index." The mean number of correct answers on the facilities test in R1 was 1.4 and 1.0 for men and women, respectively, out of a maximum of 4. Clearly, knowledge about these aspects of the program is deficient.

The scores generally differed between the two respondents in each sampled household, which motivates our identification strategy. In R1, the scores on the employment quiz differed in 68% of cases, though falling to 44% for the facilities knowledge index; Table 2 gives the frequency distribution of the difference in scores. It is clear that the

TABLE 2 Difference in scores on the NREGA quiz

Absolute difference in scores between two adults in the same household	Employment/work knowledge (frequency, %)		Facilities knowledge (frequency, %)	
	R1	R2	R1	R2
0	31.67	23.22	55.65	49.93
1	25.93	31.85	18.20	25.58
2	16.20	21.41	12.65	16.22
3	10.65	12.34	7.69	5.84
4	9.01	6.49	5.82	2.44
5	4.57	3.12	n.a.	n.a.
6	1.74	1.04	n.a.	n.a.
7	0.21	0.52	n.a.	n.a.
8	0.02	n.a.	n.a.	n.a.

spread is quite large. Thus there is ample scope for explaining the differences within households in terms of whether the respondent watched the movie, though this is less so for knowledge about facilities than for employment and wages.

Table 3 gives the breakdown of mean scores on our knowledge tests between the three groups. Participants scored better than either excess demanders or others. Participants may well have learned by participating, so one cannot conclude from Table 3 that knowledge was the cause of participation. We also see from Table 3 that the participants and excess demanders are more likely to be illiterate and to have completed less schooling. There is an association with caste; participants are more likely to be Mahadalits, and the third group (neither participants nor excess demanders)

TABLE 3 Descriptive statistics by round 1 BREGS participation status

	BREGS participant			Excess demand			Rest			Total		
	N	Mean	SE (mean)	N	Mean	SE (mean)	N	Mean	SE (mean)	N	Mean	SE (mean)
Employment knowledge index	720	3.441	0.065	1097	2.644	0.054	596	2.432	0.072	2413	2.803	0.037
Facilities knowledge index	720	2.524	0.051	1097	2.032	0.041	596	1.749	0.049	2413	2.087	0.028
Gender	720	0.501	0.019	1097	0.504	0.015	596	0.510	0.020	2413	0.505	0.010
Education												
Illiterate	713	0.791	0.015	1094	0.716	0.014	595	0.472	0.020	2402	0.668	0.010
Literate (less than class 5)	713	0.116	0.012	1094	0.110	0.009	595	0.128	0.014	2402	0.117	0.007
Class 5 pass (primary)	713	0.039	0.007	1094	0.060	0.007	595	0.092	0.012	2402	0.063	0.005
Class 8 pass (middle)	713	0.024	0.006	1094	0.050	0.007	595	0.118	0.013	2402	0.062	0.005
Class 10th pass (secondary)	713	0.018	0.005	1094	0.047	0.006	595	0.105	0.013	2402	0.055	0.005
Class 12 pass (higher secondary)	713	0.010	0.004	1094	0.010	0.003	595	0.048	0.009	2402	0.020	0.003
More than higher secondary	713	0.002	0.002	1094	0.008	0.003	595	0.037	0.008	2402	0.015	0.002
Caste/tribe group												
Mahadalit	720	0.125	0.012	1097	0.037	0.006	596	0.016	0.005	2413	0.055	0.005
Scheduled tribe (ST)	720	0.031	0.006	1097	0.032	0.005	596	0.001	0.002	2413	0.023	0.003
Scheduled caste (SC) (other)	720	0.328	0.018	1097	0.263	0.013	596	0.066	0.010	2413	0.226	0.009
Other backward castes (OBC)	720	0.472	0.019	1097	0.555	0.015	596	0.648	0.020	2413	0.558	0.010
General caste	720	0.044	0.008	1097	0.113	0.010	596	0.268	0.018	2413	0.137	0.007
Religion												
Hindu	718	0.964	0.007	1097	0.876	0.010	596	0.853	0.015	2411	0.894	0.006
Muslim	718	0.036	0.007	1097	0.124	0.010	596	0.147	0.015	2411	0.106	0.006
Close to ...												
Ward member or panch of panchayat	720	0.618	0.018	1097	0.433	0.015	596	0.401	0.020	2413	0.475	0.010
Mukhiya or sarpanch of panchayat	720	0.478	0.019	1097	0.279	0.014	596	0.295	0.019	2413	0.338	0.010
Other elected member	720	0.153	0.013	1097	0.067	0.008	596	0.068	0.010	2413	0.091	0.006
BREGS worksite mate	720	0.257	0.016	1089	0.029	0.005	592	0.021	0.006	2401	0.090	0.006
Monitoring-vigilance committee member of BREGS	718	0.020	0.005	1089	0.006	0.002	591	0.005	0.003	2398	0.010	0.002
Gram Rozgar Sewak of panchayat	718	0.149	0.013	1087	0.042	0.006	592	0.069	0.010	2397	0.079	0.006
Gram Panchat Sewak/Panchayat Sachiv	718	0.076	0.010	1089	0.041	0.006	592	0.060	0.010	2399	0.056	0.005
Revenue official of Panchayat	718	0.023	0.006	1087	0.041	0.006	592	0.013	0.005	2397	0.029	0.003
PO/Block Dev. Officer of block	720	0.017	0.005	1097	0.003	0.002	596	0.020	0.006	2413	0.011	0.002
Any political worker	720	0.097	0.011	1097	0.013	0.003	596	0.006	0.003	2413	0.034	0.004
House asset index	716	-0.822	0.060	1097	-0.426	0.067	594	1.524	0.105	2407	0.008	0.048

Note. sample of households with two panel individuals.

are more likely to be general caste members. There are also signs that BREGS participants tend to have better political connections, as indicated by responses to survey questions on whether the respondent was “close to” each of a series of designated local officials.

3.4 | The information campaign

The majority of adults are illiterate, so a movie in the local language makes sense as an information tool. Dutta et al. (2014) report pilot tests of alternatives, such as reading a summary of NREGA provisions in gender-specific focal groups, which did not show much promise. But a pilot film (based on TV advertisements for NREGA) did show promise, judged by audience reaction and subsequent discussion in focal groups.¹⁸

A 20-minute movie was therefore produced to explicitly convey information about rights and entitlements under BREGS.¹⁹ The movie was tailored to Bihar's specific context and program guidelines. Professional actors performed in an entertaining and emotionally engaging story-based plot whose purpose was to provide information on how the scheme works, who can participate, and how to go about participating. The movie stresses the fact that all adults are eligible for the scheme and that potential workers need to demand work in order to get it, in addition to providing information on guidelines for time-bound responses from the government on providing work or an unemployment allowance, and for paying wages.

The storyline is centered on a temporary migrant returning to his village from the city to see his wife and baby daughter. He learns that there is BREGS work available in the village, even though it is the lean season, so he can stay there with his family and friends rather than return to the city to find work. It was intended that the audience would identify strongly with the central characters. While the lead actor is a man, and the main focus is on him throughout, the storyline includes a deliberate flow of supporting actors, including women, who indicated that the scheme is open to women who are supposed to receive the same wages as men for equal work.

The information campaign was conducted in February–March 2010 in 40 villages randomly selected from the baseline sample of 150 villages. Compliance at the village level was complete. The intervention was done 2–4 months prior to the follow-up survey (R2), so that information had a reasonable time to spread within the treated villages. Given the timing of the seasons and the nature of the program, the expectation was that if the intervention had impact it should be evident within this time period, given that this coincided with the lean season, when demand for BREGS should be high. Around the time that many people would be in need of extra work, the intervention tells them how to go about getting that work. It is also a season in which forgone income from watching the movie is likely to be relatively low.

The treatment and control samples are well balanced. The differences in sample means of the village variables used in the analysis (including village means of household and individual variables) are only statistically significant at the 5% level for three variables out of the 70 tested; Ravallion et al. (2015) provides details.²⁰

There was negligible contamination effect on the control villages; only 12 households (0.4%) in the R2 sample from control villages reported that they were aware that a film on BREGS had been shown elsewhere. However, spillover effects within the treatment villages are likely.

The movie was shown in two separate locations in each treatment village over one or two days. Typically, it was projected in common areas, such as open ground, a school building, or a community hall. The screenings were in an open space about half the time; school buildings were the venue for about half the remainder. The film was screened twice in each treatment village, followed by a question-and-answer session and distribution of one-page flyers that pictorially illustrated the main entitlements and processes under NREGA. On arriving in each village, the facilitators announced the upcoming screenings in advance in both the most central location in the village and in a central location within the lower-caste area of the village. Local officials such as the Mukhiya and Sarpanch, opposition leaders, and local BREGS officials were invited to attend.

The movie was clearly a big event in the treatment villages. On average, about 365 people attended either screening, roughly evenly split between the two screenings. About two-thirds of those attending were men. Only in 11% of the showings did people say that the information provided was not new. The average discussion time after the movie

¹⁸Videos have also been effective in agricultural extension, including in Bihar; see, for example, Kaushal (2015).

¹⁹The movie was commissioned by Dutta et al. and was produced by the local NGO, Praxis—the Institute for Participatory Practices. The movie can be seen there.

²⁰Of course, some significant differences using conventional “one-at-a-time” tests are to be expected by chance.

was 38 minutes and the movie was deemed by the facilitators to have generated “a lot of discussion” in 29% of the showings. Based on our survey, 86% of men and 77% of women in the treatment villages were aware that the movie had been shown. 55% of men in the sample had actually seen the movie, as compared to 43% of women.²¹ 27% of men and 33% of women had not seen the movie but reported that they had discussed it with others in the village. For 29% of the sample with two adults interviewed, one adult watched the movie and one stayed at home. In 76% of these cases it was the male who went to watch the movie.

Table 4 provides descriptive statistics for the balanced sample of individuals in treatment villages on knowledge of the movie and the incidence of watching the movie broken down by BREGS participation status, caste, literacy, landholding, and poverty, all measured in R1. 85% of the balanced sample knew about the movie, and this varies little across the groups. 42% decided to watch the movie, and this varies with individual, household, and village characteristics. BREGS participants in R1 were more likely to do so; illiterate individuals were less likely to do so. However, there are many people (half or more) in each category who did not watch the movie. We will see below whether the movie's messages reached them indirectly.

4 | RESULTS

Before implementing our method, we perform two tests. The first is to run Monte Carlo simulations to test our preferred identification strategy and the estimation procedure on synthetic data for which we know the true values of the parameters of interest. Full details are given in Supporting Information Addendum 1.a. The simulations suggest that all the estimation methods considered are reliable in the absence of spillover effects or endogeneity in the choice to watch the movie. However, with a spillover effect and endogeneity, our preferred identification strategy clearly dominates, with appreciably less bias than the other methods. The results give one confidence that the proposed method is to be preferred.

The second test relates to our identifying assumption that the selection process determining who watches the movie within a household is uncorrelated with idiosyncratic time-varying information about the program included in the error term. As discussed in Section 2, we test this assumption using the augmented regression (Equation 13), where the test is based on the parameter γ . Table 5 gives the results. We cannot reject the null that $\gamma = 0$ in any case.²²

Based on these test results, we proceed to implement our preferred estimation method. Column (1) of Table 6 gives our estimates for the employment knowledge scores, with $\beta_1 - \beta_2$ estimated using the regression in Equation 9 and β_2 from Equation 11; the table also gives the implied estimate of β_1 (ATET). We reject the null that the movie had no direct effect on knowledge about the scheme ($H_0 : \beta_1 = 0$). We also reject the null hypotheses of no knowledge sharing ($H_0 : \beta_2 = 0$) and full sharing ($H_0 : \beta_1 = \beta_2$).

In short, we find significant effects of the movie on knowledge and these effects are both direct and indirect. The spillover effect accounts for one-third of the direct effect (i.e., $\widehat{\beta}_2/\widehat{\beta}_1 = 0.32$). ATET is slightly more than one correct answer to the employment questions ($\widehat{\beta}_1 = 1.14$), representing a one-third increase in the number of correct answers in the control group ($\widehat{\alpha} = 3.46$).²³

Column (2) of Table 6 gives the naïve estimator, using OLS on the switching regression (Equation 7). This method incorrectly attributes the impact fully to the direct effect of watching the movie, and substantially overestimates that effect ($\widehat{\beta}_1 = 1.59$). Column (3) gives the naïve estimate with controls for individual and household characteristics; this is similar to the results for column (3). Regression controls appear to help little.

Column (4) gives the PSM estimates. We use a probit to predict the probability of watching the movie, $\widehat{p}(X)$, for all individuals in movie and non-movie villages.²⁴ In Figure 1 we provide the densities of the predicted propensity score in movie and non-movie villages. We ensure common support by dropping individuals j in non-movie villages such that $\widehat{p}(X_j) < \min(\widehat{p}(X_i))$ for i in movie villages, and individuals i in movie villages such that $\widehat{p}(X_i) > \max(\widehat{p}(X_j))$ for j in non-movie villages. We then match nearest neighbors in terms of $\widehat{p}(X_i)$. We do this by matching one by one each

²¹This information was self-reported during the second round of the household survey.

²²Supporting Information Addendum 2 reports the coefficients of the probit used in implementing this test.

²³In the case of our preferred estimate we report bootstrapped standard errors for 500 replications given that the coefficients come from different regressions on different samples. For OLS estimates we report the SE and p -values for a t -test of the coefficients' linear combination.

²⁴Supporting Information Addendum 2 reports the coefficients of the probit used for constructing the propensity scores; the estimates reveal some reasonably intuitive selection mechanisms.

TABLE 4 Descriptive statistics on compliance and knowledge about the movie, split by round 1 characteristics

BREGS status	N	Mean	SE (mean)
<i>Knows about the movie?</i>			
Participant	204	0.89	0.02
Excess demand	255	0.82	0.02
Rest	168	0.82	0.03
Caste			
Mahadalit/SC/ST	165	0.93	0.02
OBC or general	462	0.81	0.02
Literacy			
One or both literate	337	0.83	0.02
Both illiterate	290	0.85	0.02
Landholding			
Has some land	278	0.84	0.02
Landless	349	0.84	0.02
Poverty			
Poor	336	0.85	0.02
Non-poor	291	0.83	0.02
Total	627	0.85	0.01
<i>Watched the movie?</i>			
Participant	204	0.55	0.03
Excess demand	255	0.42	0.03
Rest	168	0.33	0.04
Caste			
Mushar/Mahadalit/SC/ST	165	0.54	0.04
OBC or general	462	0.38	0.02
Literacy			
One or both literate	337	0.44	0.03
Both illiterate	290	0.40	0.03
Landholding			
Has some land	278	0.40	0.03
Landless	349	0.44	0.03
Poverty			
Non-poor	336	0.42	0.03
Poor	291	0.42	0.03
Total	627	0.42	0.02

Note. sample of households with two panel individuals in villages where the movie was shown. Poor individuals are in households characterized by a monthly per capita expenditure below the whole population median in round 1 (585 rupees).

individual watching the movie with one counterpart in non-movie villages, and then matching each individual not watching the movie in a movie village with one counterpart in the non-movie villages.²⁵ We then compute \widehat{SE} and \widehat{ATET} reported in Tables 5 and 8. The PSM estimates are quite similar to OLS. Neither method accords well with our preferred method.

²⁵This is the simplest matching method, and more complicated methods may well give somewhat different results.

TABLE 5 Test of the parallel trends assumption

Estimates of γ in Equation 13	Employment knowledge index	Facilities knowledge index
<i>Whole sample</i>	0.06 (0.68)	-0.16 (0.47)
<i>NREGS status</i>		
Participant	-1.78 (1.56)	-0.42 (1.22)
Excess demander	-0.11 (0.78)	-0.54 (0.62)
Rest	1.00 (1.20)	0.26 (0.80)
<i>Caste</i>		
Mahadalit/SC/ST	-1.63 (1.15)	-0.06 (0.95)
OBC or general	0.50 (0.77)	-0.16 (0.55)
<i>Literacy</i>		
Both illiterate	-2.11 (1.58)	-1.98 (1.53)
One or both literate	0.64 (0.75)	0.19 (0.49)
<i>Landholding</i>		
Landless	0.27 (1.12)	-0.30 (0.83)
Has some land	-0.05 (0.76)	0.00 (0.50)
<i>Poverty</i>		
Poor	-0.91 (0.76)	-0.09 (0.63)
Non-poor	0.58 (0.94)	-0.27 (0.66)

Note. Robust standard errors in parentheses;

* $p < 0.10$;

** $p < 0.05$;

*** $p < 0.01$. The test for exogeneity consists in adding as a control the within-household difference in the residuals from the probit model reported in Supporting Information Addendum 2.

TABLE 6 Coefficients for employment knowledge index

	(1)	(2)	(3)	(4)	(5)	(6)
	Preferred	Naïve OLS	OLS with controls	PSM	IV	ITT
α	3.46*** (0.08)	3.46*** (0.08)	3.46*** (0.08)	3.12*** (.09)	3.46*** (0.08)	3.46*** (0.08)
β_2	0.36*** (0.13)	0.03 (0.15)	0.24 (0.15)	0.22 (0.14)		0.69*** (0.13)
$\beta_1 - \beta_2$	0.77** (0.36)	1.57*** (0.20)	1.32*** (0.19)			
β_1 (ATET/LATE)	1.14*** (0.24)	1.59*** (.17)	1.56*** (.17)	1.35*** (0.18)	1.63*** (0.29)	
N	1188/2376	2376	2219	1895/2016	2376	2376

Note. Robust standard errors in parentheses for OLS and ITT; for PSM Abadie–Imbens robust standard errors; for preferred estimate SE (ATET) is bootstrapped standard error (500 replications);

* $p < 0.10$;

** $p < 0.05$;

*** $p < 0.01$. Controls include individual, household and village characteristics observed in round 1. The estimate of α for OLS with controls is evaluated at the mean points of the controls.

Column (5) gives the IV estimate of the internal effect assuming no spillovers, and using the randomized assignment at village level as the IV. This is similarly biased to the OLS estimate. Column (6) gives the ITT estimate.

Recall that we are constrained to using the balanced panel of two individuals per household. To test for sample selection bias, we re-estimated ITT, the naïve OLS, the naïve OLS with controls, the PSM, and the IV estimator on

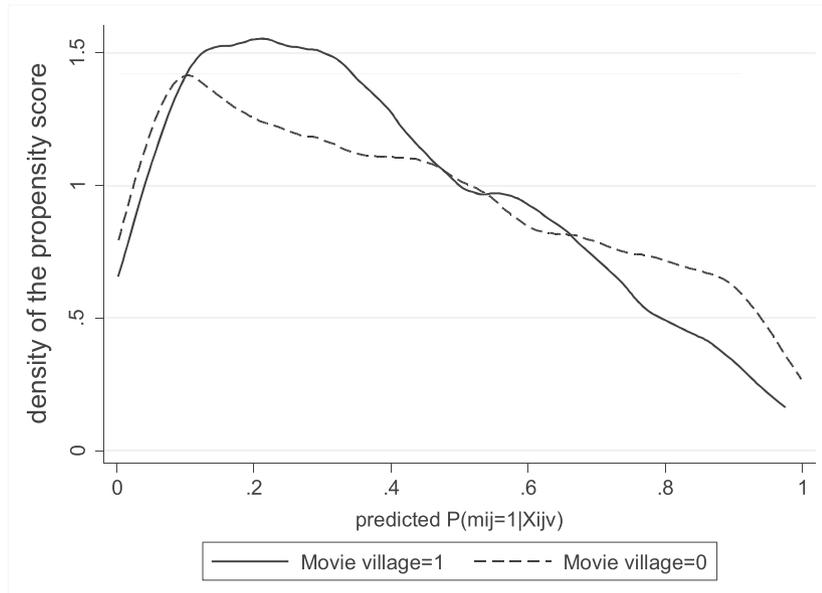


FIGURE 1 Density of predicted propensity scores in movie and non-movie villages. Univariate kernel density estimation

TABLE 7 Coefficients for employment knowledge index split by round 1 characteristics

	$\beta_1 - \beta_2$	β_2	β_1 (ATET)
<i>Gender</i>			
Male	0.77** (0.36)	0.36** (0.17)	1.13*** (0.27)
Female	0.77** (0.36)	0.36** (0.18)	1.14*** (0.26)
<i>NREGS status</i>			
Participant	-0.20 (0.63)	0.55** (0.22)	0.35 (0.40)
Excess demander	1.98*** (0.44)	-0.26 (0.19)	1.72*** (0.32)
Rest	-0.40 (0.45)	1.20*** (0.25)	0.80** (0.39)
<i>Caste</i>			
Mahadalit/SC/ST	1.47*** (0.44)	-0.10 (0.22)	1.37*** (0.32)
OBC or general	0.45 (0.48)	0.55*** (0.15)	1.00*** (0.34)
<i>Literacy</i>			
Both illiterate	1.47** (0.69)	-0.03 (0.18)	1.43*** (0.48)
One or both literate	0.31 (0.33)	0.67*** (0.18)	0.98*** (0.27)
<i>Landholding</i>			
Landless	1.06** (0.51)	0.00 (0.17)	1.07*** (0.35)
Has some land	0.33 (0.39)	0.82*** (0.19)	1.15*** (0.32)
<i>Poverty</i>			
Poor	0.83** (0.41)	0.29* (0.17)	1.12*** (0.30)
Non-poor	0.73 (0.55)	0.40** (0.18)	1.13*** (0.37)

Note. Standard errors in parentheses, SE ($\beta_1 - \beta_2$) and SE (β_2) are robust standard errors, SE (ATET) are bootstrapped standard errors (500 replications);

* $p < 0.10$;

** $p < 0.05$;

*** $p < 0.01$. Poor households are characterized by a monthly per capita expenditure below the whole population median in round 1 (585 rupees).

TABLE 8 Coefficients for facilities knowledge index

	(1)	(2)	(3)	(4)	(5)	(6)
	Preferred	Naïve	Naïve with controls	PSM	IV	ITT
α	2.32*** (0.06)	2.32*** (0.06)	2.32*** (0.06)	2.05*** (0.06)	2.32*** (0.06)	2.32*** (0.06)
β_2	-0.02 (0.10)	-0.22* (0.12)	-0.08 (0.12)	0.03 (0.10)		0.09 (0.10)
$\beta_1 - \beta_2$	0.26 (0.31)	0.73*** (0.17)	0.54*** (0.15)			
β_1 (ATET/LATE)	0.24 (0.22)	0.51 ** (0.14)	0.46*** (0.14)	0.54*** (0.13)	0.21 (0.24)	
N	1,188/2,376	2,376	2,219	1,895/2,016	2,376	2,376

Note. Robust standard errors in parentheses for OLS and ITT; for PSM Abadie–Imbens robust standard errors; for preferred estimate SE (ATET) is bootstrapped standard error (500 replications);

* $p < 0.10$;

** $p < 0.05$;

*** $p < 0.01$. Controls include individual, household, and village characteristics observed in round 1. The estimate of α for OLS with controls is evaluated at the mean points of the controls.

the full sample. The detailed results are reported in Supporting Information Addendum 3(a). The estimates are very similar to those we obtained for the selected subsample of two panel individuals.²⁶

4.1 | Heterogeneity

We turn to the results stratified by observable characteristics.²⁷ The first row of Table 7 shows estimates by gender. As noted, our identification strategy does not allow us to differentiate the direct effect by gender. We can, however, differentiate the spillover effect. We find that this is almost identical between men and women.

Next, Table 7 gives the results by R1 participation status—namely BREGS participants in R1, excess demanders in R1 (who wanted work on BREGS but did not get it), and the rest. We cannot reject the null of full knowledge sharing ($H_0 : \beta_1 = \beta_2$) for either BREGS participants or the rest. We find that the significant direct effect of watching the movie is confined to the excess demanders, and for them it entails an extra two correct answers to the questions on employment provisions of the scheme—almost doubling the number of correct answers. The spillover effect is not, however, found in this group but in the other two. Indeed, for the excess demanders we cannot reject the null hypothesis of no knowledge sharing ($H_0 : \beta_2 = 0$).²⁸ While excess demanders may well be more socially excluded, there is also the possibility that they do not want to share the information, as they want to get work that they know is being rationed.²⁹

Table 7 gives stratifications by the various indicators of disadvantage. Considering caste first, we find that the direct effect is among the “lower caste” (Mushar/Mahadalit/SC/ST) group, not among the other castes, while the spillover effect is found for the latter. The weaker spillover effect for the lower-caste group is consistent with the hypothesis that they are less well connected to the information flows within the village.³⁰ Another possibility is that lower-caste groups have lower ability to convey information. For the higher-caste group we cannot reject the null of full knowledge sharing. Interestingly, ATETs are not very different between the two caste groups, but this hides a marked difference in how much the impact is direct versus indirect.

Similarly, illiterate, landless, and consumption-poor individuals had strong direct effects of watching the movie, but saw far weaker spillover effects. Indeed, the parameter for the spillover effect (β_2) is not significantly different from zero

²⁶We also report in Supporting Information Addendum 3(b) the estimates obtained for a consistent sample across methods. The estimates are again very similar.

²⁷Supporting Information Addendum 5 provides balance checks for each stratum of the heterogeneity analysis.

²⁸In further calculations (not reported but available on request) we also found that the naïve estimator performs quite well for the excess demanders, capturing 90% of the direct effect of watching the movie on employment knowledge. The IV estimator does less well. However, we found that the naïve estimator performs poorly for the other two groups, substantially overestimating the direct effect and underestimating the spillover effect.

²⁹Evidence on the existence of extensive rationing in NREGA jobs in Bihar can be found in Dutta et al. (2012).

³⁰This should not be interpreted as saying that the lower-caste group is less inclined to share information among its members. Here we are testing the diffusion of information that is provided village wide.

TABLE 9 Coefficients for facilities knowledge index, split by round 1 characteristics

	$\beta_1 - \beta_2$	β_2	β_1 (ATET)
<i>Gender</i>			
Male	0.26 (0.31)	0.02 (0.15)	0.28 (0.25)
Female	0.26 (0.31)	-0.07 (0.14)	0.19 (0.23)
<i>BREGS status</i>			
Participant	-0.99 (0.66)	0.08 (0.18)	-0.92** (0.39)
Excess demander	1.16*** (0.38)	-0.18 (0.15)	0.98*** (0.29)
Rest	-0.07 (0.20)	0.31 (0.19)	0.24 (0.26)
<i>Caste</i>			
Mahadalit/SC/ST	0.32 (0.38)	0.11 (0.17)	0.43 (0.27)
OBC or general	0.24 (0.42)	-0.07 (0.13)	0.16 (0.32)
<i>Literacy</i>			
Both illiterate	0.10 (0.65)	-0.13 (0.14)	-0.03 (0.44)
One or both literate	0.37 (0.28)	0.06 (0.14)	0.43 (0.24)
<i>Landholding</i>			
Has some land	0.24 (0.48)	-0.06 (0.13)	0.18 (0.34)
Landless	0.30 (0.25)	0.02 (0.16)	0.32 (0.25)
<i>Poverty</i>			
Poor	-0.10 (0.32)	0.14 (0.13)	0.04 (0.22)
Non-poor	0.54 (0.47)	-0.16 (0.15)	0.38 (0.35)

Note. Standard errors in parentheses, SE ($\beta_1 - \beta_2$) and SE (β_2) are robust standard errors, SE (ATET) are bootstrapped standard errors (500 replications);

* $p < 0.1$;

** $p < 0.05$;

*** $p < 0.01$. Poor households are characterized by a monthly per capita expenditure below the whole population median in round 1 (585 rupees).

for illiterate, landless, or lower-caste individuals. By contrast, there is a strong indication of knowledge sharing for the more advantaged groups. In each case, we cannot reject the null of full knowledge sharing within households for the more advantaged group, but we can for the other.

Table 8 gives the results for the facilities knowledge score (corresponding to Table 6 for employment). Our preferred estimates in column (1) do not indicate any significant direct or indirect effect of the movie on knowledge about facilities. By contrast, the OLS and PSM estimators suggest a significant positive effect of watching the movie, although this vanishes once one addresses the endogeneity using the randomized village assignment as the IV.

When we stratify by the three aforementioned participation groups we again find a significant direct effect for the excess demanders but not others (Table 9). We find no difference according to caste.

5 | CONCLUSIONS

The paper has proposed a method of identifying knowledge sharing for an individualized treatment in a cluster randomization when double randomization is not feasible. Differences in learning between individuals within the same household are used to identify the direct impact of an information campaign, while randomized assignment of the campaign at village level is used to identify the spillover effect between people, within or between households. The absence of a second randomization (in addition to that at village level) means that we require that any time-varying intra-household differences in program knowledge are uncorrelated with intra-household differences in the probability of watching the movie. A Monte Carlo simulation exercise validates the method and various robustness checks fail to reject our identifying assumptions.

We have applied the method in studying the impacts of an entertaining “information movie” that aimed to teach poor people in rural Bihar their rights under India’s National Rural Employment Guarantee Act. We find evidence of significant knowledge spillover effects, though there is far from full knowledge sharing in the data as a whole. On average, the knowledge gains to people who did not actually watch the movie but lived in a village where it was shown represent one-third of the gain to those who actually watched the movie. Substantial biases are evident in impact estimation methods that ignore spillover effects.

A key finding is that the knowledge diffusion process is far weaker for disadvantaged groups, defined in terms of caste, landholding, literacy, or consumption poverty. There are clear signs of knowledge sharing for the more advantaged group (in terms of literacy, landholding, caste, or household consumption). For these better-off strata it matters little which adult in the household watched the movie, and the village assignment is then the key factor. By contrast, for poor people, it appears that the direct effect of watching the movie is all that really matters to individual learning about NREGA, with much less sign of knowledge sharing for the poorer groups.

These results are consistent with the view that poverty persists, in part at least through weak social connectivity, leading to limited gains to poor people from knowledge diffusion—including knowledge about public programs intended to help make them less poor. Public information campaigns in this setting need to be targeted to poor groups, rather than relying on existing knowledge diffusion processes within villages.

ACKNOWLEDGMENTS

For comments the authors are grateful to Pascal Lavergne, Michael Denly, and seminar participants at Georgetown University, the Tinbergen Institute, Amsterdam, the Toulouse School of Economics, the Paris School of Economics, and the University of Southern California. The authors are also grateful to the journal’s three anonymous referees for their helpful comments.

OPEN RESEARCH BADGES



This article has earned an Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available at [<http://qed.econ.queensu.ca/jae/2019-v34.1/alik-lagrange-ravallion/>].

REFERENCES

- Ambler, K. (2015). Don’t tell on me: Experimental evidence of asymmetric information in transnational households. *Journal of Development Economics*, 113, 52–69.
- Angelucci, M., & De Giorgi, G. (2009). Indirect effects of an aid program: How do cash injections affect ineligibles’ consumption? *American Economic Review*, 99(1), 486–508.
- Angelucci, M., & Di Maro, V. (2015). Program evaluation and spillover effects. *Journal of Development Effectiveness*, 8(1), 22–43.
- Angrist, J., Imbens, G., & Rubin, D. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91, 444–455.
- Ashraf, N. (2009). Spousal control and intra-household decision making: An experimental study in the Philippines. *American Economic Review*, 99(4), 1245–1277.
- Baird, S., Bohren, A., McIntosh, C., & Özler, B. (2018). Optimal design of experiments in the presence of interference. *Review of Economics and Statistics*. https://www.mitpressjournals.org/doi/abs/10.1162/REST_a_00716
- Barkley, B. G., Hudgens, M. G., Clemens, J. D., Ali, M., & Emch, M. E. (2017). *Causal inference from observational studies with clustered interference*. arXiv:1711.04834 [stat.ME].
- Bloch, F., & Rao, V. (2002). Terror as a bargaining instrument: A case study of dowry violence in rural India. *American Economic Review*, 92(4), 1029–1043.
- Bobba, M., & Gignoux, J., (2014). *Policy evaluation in the presence of spatial externalities: Reassessing the Progresa program* (PSE Working Papers n2011–37). Paris, France: Paris Jourdan Sciences Economiques.
- Bobonis, G., & Finan, F. (2009). Neighborhood peer effects in secondary school enrollment decisions. *Review of Economics and Statistics*, 91, 695–716.

- Chen, S., Mu, R., & Ravallion, M. (2009). Are there lasting impacts of aid to poor areas? Evidence from rural China. *Journal of Public Economics*, 93, 512–528.
- Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica*, 56(1), 63–89.
- Desai, S., Dubey, A., Joshi, B., Sen, M., Shariff, A., & Vanneman, R. (2010). *Human development in India: Challenges for a society in transition*. New Delhi, India: Oxford University Press.
- Duncan, G. J., Yeung, W. J., Brooks-Gunn, J., & Smith, J. R. (1998). How much does childhood poverty affect the life chances of children? *American Sociological Review*, 63, 406–423.
- Dutta, P., Murgai, R., Ravallion, M., & van de Walle, D. (2012). Does India's employment guarantee scheme guarantee employment? *Economic and Political Weekly*, 48(April 21), 55–64.
- Dutta, P., Murgai, R., Ravallion, M., & van de Walle, D. (2014). *Right-to-work? Assessing India's employment guarantee scheme in Bihar*. *Equity and Development Series*. Washington, DC: World Bank.
- Foster, A., & Rosenzweig, M. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103, 1176–1209.
- Heller, P., & Rao, V. (2015). Deliberation and development. In P. Heller & V. Rao (Eds.), *Deliberation and development: Rethinking the role of voice and collective action in unequal societies* (pp. 1–26). Washington, DC: World Bank.
- Human Rights Watch (2007). *Hidden apartheid: Caste discrimination against India's "Untouchables"*. New York, NY: Human Rights Watch and Center for Human Rights and Global Justice, New York University School of Law.
- Janssens, W. (2011). Externalities in program evaluation: The impact of a women's empowerment program on immunization. *Journal of the European Economic Association*, 9(6), 1082–1113.
- Kaushal, K. K. (2015). A good farmer and neighbor. Digital Green Blog, Patna, Bihar, India.
- Keefer, P., & Khemani, S. (2005). Democracy, public expenditures, and the poor: Understanding political incentives for providing public services. *World Bank Research Observer*, 20(1), 1–28.
- Khemani, S. (2007). Can information campaigns overcome political obstacles to serving the poor? In S. Devarajan, & I. Widlung (Eds.), *The politics of service delivery in democracies: Better access for the poor* (pp. 56–69). Stockholm, Sweden: Ministry of Foreign Affairs, Government of Sweden.
- La Ferrara, E. (2015). *Mass media and social change: Can we use television to fight poverty?* (Working Paper 564). Milan, Italy: IGIER, Università Bocconi.
- Mansuri, G., & Rao, V. (2013). *Localizing development: Does participation work?* Washington, DC: World Bank.
- Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *Review of Economic Studies*, 73(3), 857–895.
- Miguel, E., & Kremer, M. (2004). Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1), 159–217.
- Munshi, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the Indian green revolution. *Journal of Development Economics*, 75(1), 185–213.
- Ravallion, M., van de Walle, D., Murgai, R., & Dutta, P. (2015). Empowering poor people through public information? Lessons from a movie in rural India. *Journal of Public Economics*, 132, 13–22.
- Rosenzweig, M., & Wolpin, K. (1988). Heterogeneity, intrafamily distribution, and child health. *Journal of Human Resources*, 23(4), 437–461.
- Silver, H. (1994). Social exclusion and social solidarity. *International Labour Review*, 133(5–6), 531–578.
- Udry, C. (1996). Gender, agricultural production, and the theory of the household. *Journal of Political Economy*, 104(5), 1010–1046.
- Yang, D. (2008). Can enforcement backfire? Crime displacement in the context of customs reform in the Philippines. *Review of Economics and Statistics*, 90(1), 1–14.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Alik-Lagrange A, Ravallion M. Estimating within-cluster spillover effects using a cluster randomization with application to knowledge diffusion in rural India. *J Appl Econ*. 2019;34:110–128. <https://doi.org/10.1002/jae.2658>