Review of Economic Studies (2024) **91**, 1884–1922 https://doi.org/10.1093/restud/rdad068

Advance access publication 7 July 2023

When Less Is More: Experimental Evidence on Information Delivery During India's Demonetisation

Abhijit Banerjee
Department of Economics, MIT; NBER; JPAL, USA
Emily Breza
Department of Economics, Harvard University; NBER; JPAL, USA
Arun G. Chandrasekhar
Department of Economics, Stanford University; NBER; JPAL, USA
and
Benjamin Golub
Department of Economics, Northwestern University, USA

First version received May 2019*; Editorial decision December* 2022*; Accepted July* 2023 *(Eds.)*

In disseminating information, policymakers face a choice between broadcasting to everyone and informing a small number of "seeds" who then spread the message. While broadcasting maximises the initial reach of messages, we offer theoretical and experimental evidence that it need not be the best strategy. In a field experiment during the 2016 Indian demonetisation, we delivered policy information, varying three dimensions of the delivery method at the village level: initial reach (broadcasting versus seeding); whether or not we induced common knowledge of who was initially informed; and number of facts delivered. We measured three outcomes: the volume of conversations about demonetisation, knowledge of demonetisation rules, and choice quality in a strongly incentivised policy-dependent decision. On all three outcomes, under common knowledge, seeding dominates broadcasting; moreover, adding common knowledge makes seeding more effective but broadcasting *less so*. We interpret our results via a model of image concerns deterring engagement in social learning, and we support this interpretation with evidence on differential behaviour across ability categories.

Key words: Information diffusion, Social learning, Common knowledge, Demonetisation

JEL codes: D8, G41, O33

1. INTRODUCTION

How should new information that is potentially valuable to a large population be delivered? For example, during an epidemic such as zika or COVID-19, there is a useful list of do's and don'ts; how does a government or an NGO get that information to the relevant population? In practice, there are two commonly used strategies: (1) broadcasting information widely to all (*e.g.* radio, television, newspaper, or a Twitter feed) and (2) delivering information to a select few "seed"

The editor in charge of this paper was Adam Szeidl.

individuals and relying on subsequent diffusion (which we see in viral marketing, agricultural extension services, or the introduction of microcredit).[1]

It might seem evident that the dissemination of the information is maximised by delivering it to the maximum number of people from the start. However, the success of dissemination strategies often depends on community members' own engagement in social learning. If the mode of dissemination affects engagement, then it is not clear which policy option ultimately generates the most knowledge. This is the question we tackle in this paper.

To fix ideas, suppose people need to talk to others to understand some of the information they receive but worry about revealing their lack of comprehension. Prior surveys and experiments show that efforts to maintain a desirable image can inhibit engagement in otherwise useful activities.[2] In our setting of information delivery, people may hesitate to ask for clarification of public broadcasts that others know they received; asking is a signal of not understanding or not paying attention. Other strategies, such as giving the information to a few people, may allow engagement in social learning with less reputational concern.

These considerations highlight a potentially vital role for meta-knowledge—what people know about the information that has been shared with others—in shaping the effectiveness of information campaigns. In this paper, we examine how authorities' meta-knowledge choices affect the value of providing more people with information. We focus in particular on the possibility that too much meta-knowledge crowds out the benefits of greater initial reach. This has direct implications for the design of information policies and also raises new theoretical and empirical questions concerning situational incentives for engagement in learning.

To investigate whether potentially perverse effects of meta-knowledge play an important role in a policy-relevant setting, we conducted a randomised experiment in 200 villages in Odisha, India, during the 2016 Indian demonetisation, approximately six weeks after Prime Minister Narendra Modi announced the demonetisation of all Rs. 500 and Rs. 1000 notes. The policy was unexpected and far-reaching, affecting 86% of India's currency. While there was near-universal awareness of the broad outlines of the policy, its chaotic implementation led to widespread confusion and misinformation. For example, in our baseline sample, 15% of respondents thought, falsely, that the Rs. 10 coin was also being demonetised, and 25% did not understand that demonetised currency could only be deposited into a bank account (as opposed to being exchangeable for new bills over the counter). In this setting, individuals needed help interpreting policy information, as we document in Section 2. Indeed, our survey evidence indicates that image concerns deterred people from asking questions: respondents worried that they would appear ignorant or unintelligent if they asked for clarification on aspects of the policy they did not understand. They also negatively assessed others who revealed ignorance about the widely publicised policy. These facts motivate our experiment.

In our experiment, described in Section 3, we explore how to deliver official policy information. We vary how many people are informed, implementing seeding (S) and broadcasting (BC) policies; we also vary whether meta-information is provided, implementing common knowledge (CK) and no common knowledge (NCK) policies. Combining these arms, we thus compare the following four dissemination strategies.

1. See, *e.g.* Leskovec *et al*. (2007), Ryan and Gross (1943), Conley and Udry (2010), Miller and Mobarak (2014), Banerjee *et al*. (2013), Beaman *et al*. (2021), and Cai *et al*. (2015).

2. Chandrasekhar *et al*. (2018) showed the existence of such a friction in a lab-in-the-field setting in a similar population. This ties into the image concerns literature, which explores behaviours ranging from retail transactions to educational choices to tax compliance to vaccinations. See, for example, Goldfarb *et al*. (2015), Bursztyn *et al*. (2019), Butera *et al*. (2019), Perez-Truglia and Troiano (2018), Gerber *et al*. (2008), and Karing (2018). Bursztyn and Jensen (2017) provide a comprehensive survey.

(1) (BC, CK): information is broadcast widely to all households in a village, and this fact about the information policy is itself made evident to all (as in many standard broadcasting methods).
(2) (S, CK): information is delivered to five "seed" individuals, and this is again made evident to the community (similarly to how extension services publicise the identities of model farmers).
(3) (S, NCK): information is again delivered to five individuals, but this is not publicised (as in viral marketing).
(4) (BC, NCK): information is dispersed widely, but in a way that does not generate public awareness of the delivery strategy (for example, through private mailers).

The experiment also had a third dimension of variation: the volume of information delivered in the printed pamphlets. In some villages, we delivered two facts (a short pamphlet), while in others we delivered twenty-four (a long pamphlet). This dimension is of obvious practical relevance, since policymakers need to decide how much information to deliver. Moreover, as we will discuss, it enables important auxiliary tests of our theory.

We present a simple theoretical framework in Section 4 to analyse our results, beginning with a version that excludes situational frictions such as image concerns. An individual decides whether to engage in social learning. The value of not engaging depends on information that the individual has at baseline, as well as information received from an intervention. Engaging yields information, which is more valuable on average if other people in the community are more informed. To decide whether to seek information, the individual assesses the probability that others are informed based on the explicit and implicit content of announcements of the policymaker. Engaging also has opportunity costs: for instance, the time and energy needed to find people and have conversations. What makes the model frictionless is that these costs are taken to be independent of a potential advisor's beliefs about the seeker, which excludes image costs. The key testable implication we derive is that (BC, NCK) should have less engagement than (BC, CK) because the latter treatment makes it clear that information is available. For similar reasons, in the frictionless model, (S, NCK) should do worse than (S, CK). Other comparisons in the frictionless model depend on the parameters. For instance, (S, CK) may involve more seeking than the (BC, CK) treatment,[3] but only if the non-image costs of seeking exceed the incremental value of clarification for someone informed in (BC, CK). We term this the *high-cost hypothesis*, and we will discuss it more below when we interpret our main results.

Next, we add to the model the possibility of a friction that distorts the seeking decision. The friction comes from a reputation-management concern: people caring about how others assess their ability to understand the information they were given. We analyse how this friction affects social learning, and how the effect depends on the information-delivery strategy. Public announcements about who has information increase the perceived value of seeking information. However, they also create reasons for an observer to make negative inferences about people's understanding based on their seeking behaviour, and thus may decrease the incentive to seek information. Our model examines how a particular strategic reputation-management motive is differentially activated by both the reach and the meta-information aspects of information deliveries. In this our approach builds on and goes beyond prior work studying the costs of being

3. This also implies that (S, CK) has more seeking than (BC, NCK) by what we have said above.

viewed negatively.[4] The model establishes that in the presence of image concerns, (BC, NCK) can generate more conversations and more learning than (BC, CK), overturning the main prediction of the frictionless model. We also use the model to show when (S, CK) outperforms (BC, CK), despite the fact that the latter delivers more information and offers people more opportunities to seek clarification. The force here is that under (S, CK), seeking information is normal since those who are not seeds have little information, whereas under (BC, CK), seeking is a sign that they did not understand what they were told. The model also provides more detailed predictions that allow us to assess whether the frictionless model under the high-cost hypothesis can also account for (S, CK) outperforming (BC, CK). By examining how changes in seeking rates depend on an individual's ability, we will be able to assess the image concerns model against the main alternatives.

Our experiment was conducted in the ten days starting on 21 December 2016, when banks stopped accepting demonetised notes. Each household receiving information was delivered a printed pamphlet of facts from the Reserve Bank of India's (RBI's) 19th December circular, containing up-to-date information that the policymakers chose to communicate to the public. In each village, the same pamphlet was provided to all households receiving information. Two points about this process are especially relevant to our study. First, the information was unlikely to cover everything the villagers needed to know. Even our long lists of twenty-four facts fell short of a full policy description and contained only national information rather than local implementation details. Second, people were inundated with information, not all of it credible. As a result, consulting others was likely beneficial; indeed, we hoped the pamphlets would initiate conversations and encourage the sharing of information, including on topics that were not covered in the pamphlets.

We returned to the study villages three days after the intervention and measured three primary outcomes: engagement in social learning, policy knowledge, and an incentivised choice. To measure engagement, we asked about the number of people with whom villagers had discussed demonetisation over the prior three days (the "volume of conversations"). To measure knowledge, we asked questions about the demonetisation rules and calculated an overall knowledge score. For an incentivised measure of choice quality, we asked the subjects to select one of three lottery prizes: (a) same-day receipt of a Rs. 500 note (worth 2.5 days' wage) in the old currency; (b) an IOU ("I owe you", *i.e.* a promissory note) for Rs. 200 in notes unaffected by demonetisation, redeemable 3–5 days later; and (c) an IOU for dal (pigeon peas) worth Rs. 200, again redeemable 3–5 days later. At the time of this choice, subjects still had time to deposit the Rs. 500 note at the bank, no questions asked, and we show that it was very easy to do so.

Our results, presented in Section 5, show that contrary to the frictionless model, (BC, NCK) outperforms (BC, CK) in the extent of social learning. Furthermore, (S, CK) dominates both (S, NCK) and (BC, CK). A striking observation is that (S, CK) does as well as (BC, NCK): social learning occurring in the former is sufficient to match the value of informing everyone without activating signalling concerns.

4. Work in psychology on shame experimentally establishes that feeling negatively judged can lead to general withdrawal (Gruenewald *et al.*, 2007; de Hooge *et al.*, 2010), typically through single-subject lab manipulations with self-reported measures of psychological outcomes (*e.g.* depressive symptoms and happiness). Chandrasekhar *et al.* (2018) formalise a distinction between these bad feelings *per se* and the strategic motive to manage others' beliefs; it is the latter that is important in the present paper. More distantly related is stereotype threat (Steele and Aronson, 1995; Pennington *et al.*, 2016; Spencer *et al.*, 2016): some groups (differentially) becoming discouraged when stereotypes are salient. There, the focus is typically on cognitive performance rather than equilibrium effects of shifting the information content of an action.

First, we look at participation in social learning. Adding common knowledge to a seeding strategy makes for more conversations; going from (S, NCK) to (S, CK) increases the number of conversations by 103% ($p = 0.04$). Among broacasting strategies, however, we find the reverse: (BC, CK) generates 63% fewer conversations ($p = 0.02$) than (BC, NCK). This reversal, as we noted above in discussing theoretical predictions, should not happen in a frictionless model. In addition, going from (S, NCK) to (BC, NCK) increases the number of conversations by 113% ($p = 0.048$), but (BC, CK) leads to 61% *fewer* conversations ($p = 0.029$) than (S, CK). While potentially unintuitive, this alone does not refute the frictionless model; per our earlier discussion, this model can produce such an effect if broadcasting delivers enough information and seeking is sufficiently costly (this is the "high-cost" hypothesis). However, the frictionless model does attribute this difference to non-image costs of seeking being sufficiently high, and we use other evidence to probe this prediction below.

Second, we turn to whether the changes in participation correspond to changes in knowledge. Going from (S, NCK) to (S, CK) increases the knowledge index by 5.6% ($p = 0.0142$). Within broadcasting strategies, (BC, CK) has a 3.8% lower knowledge index than (BC, NCK), though the effect is not statistically significant ($p = 0.17$). Finally, whereas going from (S, NCK) to (BC, NCK) corresponds to a 4.9% increase in the knowledge index ($p = 0.053$), going from (S, CK) to (BC, CK) leads to a 3.1% *reduction* in knowledge ($p = 0.062$). Though this (as noted before) is not on its own inconsistent with a frictionless model, the fact that seeding five people generates more knowledge overall than seeding everyone is nevertheless striking. In particular, it implies that social learning is crucial for generating knowledge of the policy.

Third, we look at the incentivised decision—whether subjects choose the Rs. 500 note over an IOU worth Rs. 200. We again see a similar pattern. Going from (S, NCK) to (S, CK) leads to an 81% increase in the probability of choosing Rs. 500 ($p = 0.037$) but going from (BC, NCK) to (BC, CK) leads to a 48% decline in the Rs. 500 choice probability ($p = 0.041$), paralleling the effects on other outcomes. There is a 114% increase in the probability of choosing the note when going from (S, NCK) to (BC, NCK) ($p = 0.014$). In contrast, going from (S, CK) to (BC, CK) leads to a 38.5% decline in the probability of choosing Rs. 500 ($p = 0.104$), showing that the knowledge outcomes documented above translate into corresponding choice-quality differences.

We then present additional evidence supporting specific predictions of our signalling model. First, the distinction between high and low types is essential to the theory: the model results are driven by low-ability types seeking more in general and the high-ability types cutting back on seeking more when reputational concerns are activated. To empirically assess this mechanism, we classify each respondent in our endline sample as having either high or low predicted ability, in the language of the model.[5] We find that when common knowledge is added to broadcasting strategies, high types reduce seeking more than the low types. This is exactly the compositional effect that underlies a signalling explanation of differential seeking across treatments. The additional evidence also helps us assess the "high-cost" hypothesis—that, within the frictionless model, people seek less in (BC, CK) relative to (S, CK) simply because they are already endowed with information. This hypothesis cannot explain why the composition of those seeking information is so different across (BC, CK) and (BC, NCK).

A second examination of mechanisms is afforded by the arm that varies the length of information delivered. We find that the perverse effects of meta-information are more strongly observed in the short treatments. Each person has 1.396 fewer conversations in (BC, CK) compared to (BC, NCK) when pamphlets are short ($p = 0.00428$). When pamphlets are long, the same

5. Since we did not collect a household panel, we use information from baseline surveys in a separate sample and a random forest approach to construct a mapping from demographics to predicted baseline policy knowledge.

change results in a smaller reduction in conversations ($p = 0.07$ in a test of whether the size of the reduction is the same). This provides support for the signalling model. It is more compromising not to understand short, simple messages than long, complicated ones, and asking questions is correspondingly more revealing about one's ability when the message is short.

In Section 6, we discuss several alternative models in light of the empirical evidence. An important class of alternatives makes the main endogenous decision *sharing* rather than seeking information.[6] For example, we consider a theory where seeds exert effort to share information widely when it is commonly known they have it (*e.g.* because they feel responsible for distributing it); however, we find little difference in seed engagement across treatments. We then consider a different theory of ability-signalling in sharing. This theory posits that high-ability people tend to share information that is scarce/interesting, and this deters sharing in (BC, CK); we argue, however, that this has a hard time explaining low sharing in (BC, NCK). We also argue that well-known social learning frictions cannot explain our results. Finally, we consider several other more elaborate alternative behavioural models that could be devised to explain our findings. While we cannot rule out all combinations of alternative explanations, we argue that a mechanism based on image concerns has substantial advantages in explaining the data parsimoniously.

Taken together, the evidence shows that in a policy-relevant context, perhaps counterintuitively, (S, CK) is the best of the typically available strategies, and as good as (BC, NCK) when that is feasible.[7] Our results highlight meta-information's central role in mediating the success of information campaigns and its perverse interaction with the scope of dissemination activities. Consistent with a model of image concerns, but not a frictionless model, removing common knowledge under broadcasting leads to increased learning. In other words, even if contacting all households is feasible, the policymaker might do just as well by publicly informing a few seeds. Moreover, if broader outreach is carried out, then it may be important not to publicise its breadth. Moreover, by emphasising the decision to engage in conversation and the image concerns involved in doing so, this paper highlights the importance of a force relevant for the large and growing literature on social learning that has not typically been studied in social learning models.

2. CONTEXT AND SETTING

2.1. *Demonetisation*

On 8 November 2016, Indian Prime Minister Narendra Modi announced a large-scale demonetisation. At midnight after the announcement, all outstanding Rs. 500 and Rs. 1000 notes (the "specified bank notes" or SBNs) ceased to be legal tender. Demonetisation affected 86% of circulating currency (in terms of value), and individuals holding SBNs had until 30 December 2016 to deposit them in a bank or post office account. Modi intended for the surprise policy to curb "black money" and, more broadly, to accelerate the digitisation of the Indian economy. The policy affected almost every household in the country, either directly because they held the SBNs or indirectly through the cash shortages that resulted from problems in printing and distributing enough new bills fast enough.[8]

6. See Niehaus (2011) for a seminal model of this.

7. For example, when a policymaker can regularly deliver messages to everyone without it being widely known. We discuss this type of strategy in our discussion of results and the conclusion.

8. Chodorow-Reich *et al.* (2020) estimate an aggregate decline in employment and night lights from the policy.

TABLE 1
Summary statistics

	Mean	SD	Obs
Female	0.32	(0.47)	1082
SC/ST	0.50	(0.50)	1082
Age	39.18	(11.88)	1079
Casual labourer	0.21	(0.41)	1082
Farmer: landed	0.16	(0.37)	1082
Domestic work	0.16	(0.37)	1082
Farmer: sharecropper	0.09	(0.29)	1082
Unemployed	0.02	(0.14)	1082
Bank account holder	0.89	(0.31)	1078
Literate	0.80	(0.40)	1047

Notes: This table gives summary statistics on the endline sample used for analysis.

The implementation was chaotic. The initial roll-out revealed a number of ambiguities, loopholes, and unintended outcomes. As a result, the government changed the rules over fifty times in the seven weeks following the announcement. The changes concerned issues such as the time frame for over-the-counter exchange of SBNs, the cash withdrawal limit, the SBN deposit limit, and various exemptions—*e.g.* for weddings, which tend to be paid for in cash. See Supplementary Material, Appendix A for a timeline of these rule changes.

2.2. *Setting*

Our study took place in 225 villages across nine sub-districts in the state of Odisha, India. The baseline was conducted starting 21 December 2016, the intervention on 23 December 2016, and the endline ran from 26–30 December 2016. All survey activities were completed before the 30 December bank deposit deadline.

Our study villages have two or more hamlets, each dominated by a different caste group. Typically one hamlet consists of scheduled caste and/or scheduled tribe (SCST), commonly referred to as lower-caste. The other hamlet consists of general or otherwise-backwards caste individuals, commonly referred to as upper caste. The hamlets are typically 1/2 to 1 km apart. Given the hamlet structure of the study area, all of our treatments and outcomes were focused on only one randomly chosen hamlet in each village.

Basic sample statistics are provided in Table 1. Of individual respondents in the sample, 89% had some kind of formal bank account, 80% of respondents were literate, and major occupations included casual labourer (21%), domestic worker (16%), landed farmer (16%), and sharecropper (9%).

2.3. *Baseline knowledge of demonetisation rules*

Using responses from our baseline survey, we first explore the beliefs of villagers about the rules prior to our intervention. While villagers almost universally understood that the Rs. 500 and Rs. 1000 notes were being taken out of circulation, panel A of Table 2 documents that many households had inaccurate beliefs about other aspects of the policy. For example, approximately 15% of the population thought (inaccurately) that the Rs. 10 coin was also being taken out of circulation; [9] 25% of villagers believed (falsely) that, at the time of our baseline survey,

9. This specific rumour spread across much of the country and was reported in the Indian press (*e.g.* http://www.thehindu.com/news/national/tamil-nadu/Rs.10-coins-pile-up-as-rumours-take-toll/article16966261.ece).

TABLE 2
Baseline error statistics

	Mean	SD	Obs
Panel A: error rates			
10 rupees coin	0.15	(0.36)	965
Notes being demonetised	0.17	(0.38)	965
Over-the-counter exchange	0.25	(0.44)	965
Exchange locations other than banks	0.50	(0.50)	966
Weekly withdrawal limits from bank accounts	0.78	(0.41)	965
Withdrawal limits on Jan Dhan accounts	0.87	(0.33)	965
Daily withdrawal limits on ATMs	0.90	(0.30)	965
Panel B: incidence of "don't know" responses			
Notes being demonetised	0.01	(0.11)	966
Exchange locations other than banks	0.30	(0.46)	966
Weekly withdrawal limits from bank accounts	0.33	(0.47)	966
Withdrawal limits on Jan Dhan accounts	0.78	(0.41)	966
Daily withdrawal limits on ATMs	0.32	(0.47)	966

Notes: Panel A gives error rates on knowledge about demonetisation in the baseline sample. Panel B gives the incidence of "don't know" responses for the relevant questions. All respondents giving a "don't know" response were asked to make their best guess.

they could still exchange notes at the bank without first depositing them into an account. Moreover, only a small fraction of respondents could accurately tell us the deadline for being able to exchange the demonetised notes; only 50% of respondents could tell us that the notes could be deposited at post offices, RBI offices, or village government offices. Our subjects were particularly uninformed about some of the economically important details, such as the weekly withdrawal limits from banks. Thirty-three percent of respondents reported that they did not know the limit, and, in total, only 22% of respondents could tell us the correct answer (Rs. 24,000). Respondents also had very poor knowledge about limits on ATM withdrawals (10% accuracy) and withdrawal limits on the low-documentation *Jan Dhan* accounts used by the poor (13% accuracy). Given the widespread penetration of bank accounts noted above, the low levels of knowledge are not due to limits to financial inclusion in the study setting.

Panel B of Table 2 shows the incidence of the respondent reporting to us that they "don't know" the answer to the question.[10] While almost all respondents believed they knew which notes were being demonetised, more than 30% of respondents reported that they did not know about the withdrawal limits or how to deposit the demonetised notes anywhere besides a bank branch. So a large fraction of individuals were willing to acknowledge to us (and thus, to themselves) that they were uninformed about important aspects of the policy.

One might ask whether it was important for relatively poor households with limited formal savings to understand various details of the policy. One major implementation problem associated with demonetisation was that there simply were not enough notes to meet demand, which ended up affecting the lives of most people. For example, microfinance borrowers were not able to service their loans, and demand for cash purchases at small shops fell. Even for individuals without bank accounts, properly understanding the rules would have been useful for a variety of decisions: *e.g.* whether to accept a promissory note from an employer or customer, or how much inventory to order for a small business.

10. If respondents answered "don't know" to any of the questions, they were then asked to make their best guess. These guesses are included in our measures of errors in panel A.

More importantly, the policy change took place during Odisha's primary agricultural harvest, when labour demand is high and when rural households receive a large share of their annual labour and agricultural self-employment income. Many employers reported not having enough cash to pay workers. This would have affected the majority of households in our sample.

2.4. *Motivating evidence for image concerns*

Our motivating hypothesis—that people's desire to seek out clarification, even when it is needed, may conflict with their desire to signal desirable attributes—came out of conversations about demonetisation during the field-scoping phase of the project and was also motivated by prior work (Chandrasekhar *et al*., 2018). That paper develops a theory of image concerns and the decision to seek information that we build on here. It also provides supporting evidence from both an experiment and a field survey conducted in Indian villages. The survey asked villagers how they seek information on several topics: farming, health, and household finance. Among our respondents, 88% reported feeling constrained in seeking advice from others; of these, 64% felt the reason they were constrained was that they did not want to appear "weak" or uninformed. In the field experiment, when image concerns are activated, there is a 55% decline in the probability of a low-ability subject seeking out information that has a high monetary return.

Our fieldwork suggests that the types of seeking frictions documented in Chandrasekhar *et al*. (2018) were also relevant during the demonetisation. In 2018, we conducted interviews and surveys with 102 randomly selected subjects from four villages in rural Karnataka, India. We first include some representative quotes and then summarise the survey results.

Consistent with the chaotic policy implementation and the low levels of baseline knowledge we document above, respondents recall feeling confused or knowing others were confused during the period of the demonetisation. They also report that, because information was abundant, asking for clarification was potentially compromising.

> There was confusion about where to deposit money, how much to deposit, where to withdraw from, where all money could be deposited and last date. People hesitate to ask because they may think, "even after showing so much on TV, if I ask, what will they think of me. They will think I don't understand". (Respondent 1)

> People with more money hesitate to ask because they will worry what others will think about them [...] Others will think, "Don't they know anything? People with money should know more. But if they are still asking, they must be less intelligent". (Respondent 2)

Relatedly, individuals who understood the key points of the policy report judging others for not understanding them.

> If someone didn't exchange money till December, they must definitely be the biggest *bewakoof* (fool) in the world. (Respondent 3)

> Not everyone knew the deadline and application process. In December if someone comes and asks even after showing on TV, I will think they are dumb. They can't understand so they must be unintelligent. Fearing that others will think like [me], some people who were confused didn't ask. (Respondent 4)

And this of course reinforced the hesitancy to ask in the first place. That is, people are indeed cognizant of such judgments.

> I came to know a little later that I had 2 old notes with me. I didn't exchange because I didn't know when the last date was. If I ask someone, I was worried what they will say about me.

Survey Questions

1. Were you confused about any of of these changes?
2. Do you think anyone in your village was confused about these changes?
3. Do you think you completely understood the details such as the amount to exchange, the exchange location, the exchange method, the deadline for exchange, the application process, etc.?
4. Do you think everyone in your village completely understood the details such as the amount to exchange, the exchange location, the exchange method, the deadline for exchange, the application process, etc.?
5. After the policy was introduced, did you ever hesitate to ask someone from your village about the note-ban policy because you were concerned about what they might think about you?
6. After the policy was introduced, did you ever hesitate to ask **an acquaintance** from your village about the noteban policy because you were concerned about what they might think about you?
7. If yes, did you hesitate because you were concerned they would think you are:
   (a) dumb?
   (b) irresponsible?
   (c) lazy?
   (d) dealing in black money?
8. If someone from your vil1age were to ask about the note-ban policy in December, after it was heavily broadcast on TV, how do you think others might perceive them?
   (a) dumb for not understanding even after the broadcast?
   (b) irresponsible for not checking earlier?
   (c) lazy?
   (d) dealing in black money?
9. In December, since the news about the note-ban policy was being extensively broadcast on TV, do you think it was the responsibility of people in your village to know everything/to be completely informed about the note-ban policy?
10. In December, after the changes were extensively broadcast on TV, do you think some people in your village asked less about the note-ban policy even though they were still confused because they were worried that they would be judged as dumb, lazy, irresponsible, or dealing in black money?

Survey Results

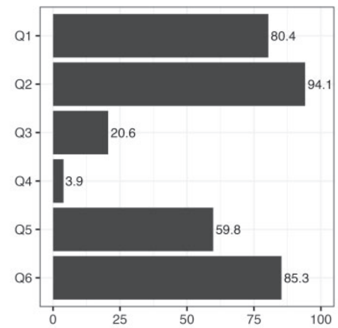

Subfigure A

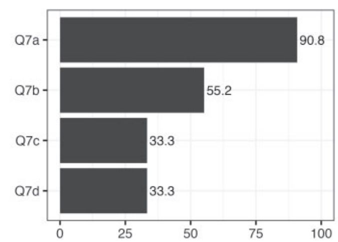

Subfigure B

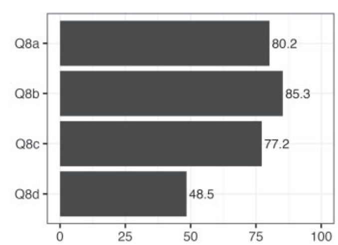

Subfigure C

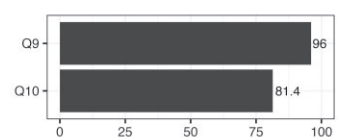

Subfigure D

FIGURE 1

Survey results: percent of respondents answering “Yes” to each question/subquestion

Note: The sample consists of 102 randomly sampled respondents across four villages in Karnataka.

> What will people think? They will say, "Were you lazy? Were you sleeping till now? Everything was shown on TV". (Respondent 5)

Figure 1 displays the survey results. We find that 80% of respondents said they felt confused, and 79% felt that even at the end of the demonetisation period they did not understand the note-ban's policy-relevant implications completely. Nearly everyone—94% of the sample—reported that others in the village were confused as well. At the same time, 96% of the individuals felt that people were responsible for understanding the policy. If someone in the village asked about the policy in December (after extensive public information campaigns), 80% of respondents said that the individual would seem unintelligent, while 85% said the individual would appear irresponsible. Finally, 85% said that even if they were confused, they held back from asking questions of acquaintances for fear of being judged.

In short, this is a setting in which the policy implementation made it hard to learn; individuals felt confused; they felt that confusion was associated with being unintelligent or irresponsible and thus they worried that seeking out information would look bad; and they therefore reduced their information-seeking. Though a large fraction of people were somewhat confused themselves, they readily admitted they were willing to pass judgment on others who did not understand how to behave. This qualitative evidence motivates analysing treatment differences through the lens of image concerns.

3. EXPERIMENT: DESIGN AND IMPLEMENTATION

We designed our experiment to explore how to convey policy information to communities in the context of a real-world, high-stakes setting where image concerns may be relevant. Central to the design is the observation that meta-information might have a perverse impact on social learning if it activates image concerns and changes the willingness of individuals to participate. The ultimate goal was to examine the tradeoffs in the choice of how to deliver information—to understand, for example, whether the common policy of delivering information to everyone by loudspeaker might actually do worse than simple alternatives.

3.1. *Treatments*

All of our experimental treatment arms involved distributing pamphlets with information about demonetisation to the study villages. Our goal was to spread the official policy rules, and thus all information came from the RBI circulars released up until 19 December 2016. We took this official information, published by the central bank, and subdivided it into 30 distinct policy rules. Because we implemented our experiment over the last week before the 30 December deadline, the rules that our pamphlet touched on did not change over the course of our experiment. Through informal conversations in pilot villages, we identified the ten most useful rules for a typical villager in the study area.[11] Our experimental protocol involved giving a randomly selected set of facts to each village—below we describe how the selection was done. All individuals receiving lists of facts in a village received the same list.

Our main focus is on a 2×2 design that varies how many people got information and the extent of common knowledge. Because another important dimension for information policies concerns how much information to disclose, we added an arm varying whether villages received

11. For example, one rule that we do not classify as "useful" explained how foreigners could exchange SBNs.

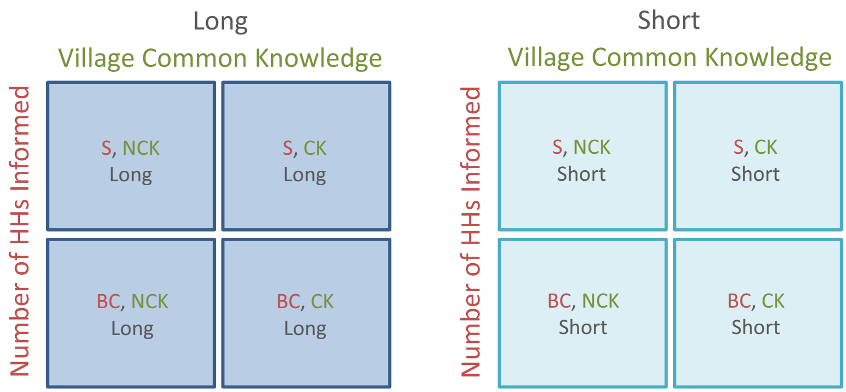

FIGURE 2
Experimental design

long or short lists of facts. Prior work has shown that more information can overwhelm individuals and harm learning (Abaluck and Gruber, 2011; Beshears *et al.*, 2013; Carvalho and Silverman, 2019), so we wanted to examine whether similar effects would be present in our social learning setting. The variation in the length of information also permits a useful test of our signalling model, discussed when we interpret our results below. Figure 2 summarises the design. The treatments are as follows:[12]

(1) Information dissemination:

- *Broadcast:* information pamphlets were provided to all households in the hamlet.
- *Seed:* information pamphlets were provided to five seed households in the hamlet, chosen as the individuals best situated to spread information in the village according to the "gossip survey" methodology of Banerjee *et al.* (2016).[13]

(2) Common knowledge:

- *No CK:* we did not tell any subject that we were providing information to anyone else in the community.
- *CK:* we provided common knowledge of the information dissemination protocol. In "BC" treatments, every pamphlet contained a note that all other households received the same pamphlet. (Thus, if subjects understood and believed us, then they had common knowledge of the pamphlet's distribution.) In the "seed" treatments, every household received a notification that five individuals in their community (who were identified) were provided

12. We also attempted to get data from thirty villages where we did not intervene whatsoever and instead only collected endline data. We call these "status quo" villages. Unfortunately, these villages are not entirely comparable to our core set due to implementation failures that led to violations of randomisation. We detail this in Supplementary Material, Appendix M.

13. Seed households were not told that they were chosen for any particular reason. To select these households, we asked each respondent at baseline, "If we want to spread information about the money change policy put in place by the government recently, whom do you suggest we talk to? This person should be quick to understand and follow, spread the information widely, and explain it well to other people in the village. Who do you think are the best people to do this for your hamlet?" and we allowed them to nominate anywhere from zero to four individuals. The results reported in Banerjee *et al.* (2016) show that this methodology identifies people in the village good at spreading information—informing gossips led to three times as many people being reached as informing random people or informing prominent people.

information about demonetisation by us and that the seeds were informed that we would identify them to everyone.

(3) Information volume:

- *Long:* twenty-four facts were provided. The "long" lists of facts were drawn uniformly from the overall list.
- *Short:* two facts were provided.The "short" lists of facts contained one of the ten most "useful" facts, drawn uniformly at random, and a second fact drawn uniformly at random from the remaining twenty.[14]

Supplementary Material, Appendix B provides the total list of facts from which we selected the list for each pamphlet, and Supplementary Material, Appendix C provides examples of the pamphlets we handed out.[15] We simplified official facts from the RBI circular into ordinary language.

3.2. *Sample*

Our data collection was constrained by the fact that the fieldwork went from conception to completion in less than a month. Demonetisation was announced on 8 November 2016 and banks stopped accepting the demonetised notes after 30 December 2016. We saw that there was a need to provide information that also offered a research opportunity and came up with an implementable research design as quickly as possible. However, by the time we were ready to launch the intervention, it was already early December and the study had to be completed by the end of the year. This imposed constraints on what we could do and led to some implementation errors.

We started with a list of 276 villages in an area where one of us had previously worked.[16] We required that all villages in the study have multiple hamlets (the predominant village organisation in the study area) and that each hamlet have at least twenty households. One hamlet in each village was supposed to be in our study; in half of the villages, chosen at random, this was the SCST hamlet while in the other half, it was the non-SCST hamlet. To facilitate planning, we randomised villages into treatments before we verified that each village met our criteria. As a result, only a set of 221 villages were eventually treated. Sixteen villages in a new subdistrict were then added to increase the sample to 237.[17]

We collected a repeated cross section (rather than a panel, due to the time cost of tracking each respondent multiple times) in each village. A baseline survey was administered for five randomly chosen individuals in each study hamlet. We also implemented an endline survey, after treatment, with a total sample of 1248 households.[18] Given the rush of implementing 200+

14. Thus, on average, in the long treatment, eight facts were useful. In the short treatment, at least one fact was always useful, and the additional fact was useful with probability $1/3$.

15. Supplementary Material, Appendix H contains a version of our main analysis, looking separately at the endline knowledge of useful facts, facts that were reported in that particular village, and facts that were omitted from that village's pamphlets.

16. Our list included some places where the research team had been before over the course of work on Breza *et al*. (2017), Breza *et al*. (2021), Breza *et al*. (2020), and Kaur *et al*. (2019) though the presence of researchers in these villages had ended many months before the baseline survey was conducted for this study.

17. Supplementary Material, Appendix L repeats our main analysis dropping these new villages and shows that our conclusions remain the same.

18. We use clustered standard errors as described below. Simonsohn (2021) notes that the appropriate heteroskedasticity and autocorrelation consistent estimators perform as well as randomisation inference. Given the complex nature of implementation and attrition, coding clustered standard errors is considerably simpler.

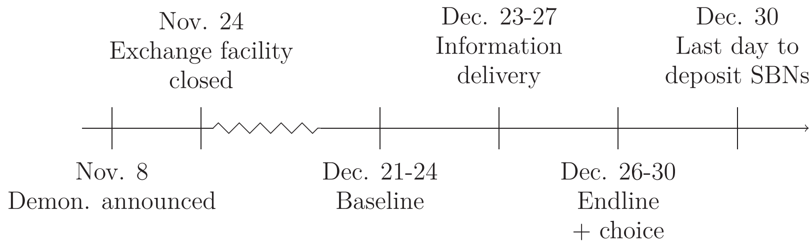

FIGURE 3

Intervention timeline

interventions in a matter of days, some additional field errors were made. Endline data were not collected in six villages and the intervention did not happen in five villages (we also did not collect endline data there). In two villages, the elders refused entry to our surveyors. We examine village-level attrition, with a particular focus on differential attrition by treatment status, in Supplementary Material, Appendix N. Ultimately, we have a sample of 225 villages that were treated and received endline surveys.[19] Figure 3 presents a timeline of the roll-out.

In each survey round, the enumerators selected households using standard circular random sampling. We asked to speak with any adult permanent resident of the household. Almost all of the survey refusals were from households in which no adult permanent resident was home at the time of the enumerator's visit.[20] In the endline surveys, we also attempted to over-sample those nominated as seeds (whether the village was in a seed treatment or not). Because the gossip survey was administered in the baseline, we can identify seeds and counterfactual seeds in all treatment cells. We present a test for baseline covariate balance across our four main experimental treatment arms in Table 3. Columns 1–4 present means by covariate in each treatment cell, while columns 5–10 present *p*-values for tests of pairwise differences in means across cells. Of the fifty-four pairwise comparisons, only 11% have a *p*-value below 0.1 and only 5.5% have a *p*-value below 0.05, indicating balance.

3.3. *Outcomes*

We have three main outcomes of interest at endline: engagement in social learning; general knowledge about facts surrounding demonetisation; and whether the respondent selected the demonetised Rs. 500 note as opposed to an IOU payable in 3–5 days for either Rs. 200 in non-demonetised notes or Rs. 200 in *dal*, a staple food.

First, we collected data on the volume of conversations about demonetisation, measured as the number of people each respondent spoke to about demonetisation in the prior three days. This allows us to see whether engagement in social learning increased or decreased based on the dissemination strategy.

19. Unfortunately, also due to the intense time pressure, in sixteen of the villages our field team administered the intervention and endline to the wrong hamlet. While these errors should be idiosyncratic and orthogonal to treatment, we collected outcome data in the intended hamlet and we redo our estimation using treatment assignment as instruments for treatment in Supplementary Material, Appendix K. All our results look nearly identical.

20. In these cases, the enumerators made at least two additional attempts to conduct surveys on the day of the visit. The biggest reason for locked doors was time of day—it was much easier to find respondents early in the morning or in the evening. Because surveyors were dispatched to villages in randomised order, we control for time of entry in the village in all of our main regression specifications.

TABLE 3
Baseline covariate balance

	Means				Pairwise differences p-values					
	(1) (S, NCK)	(2) (S, CK)	(3) (BC, NCK)	(4) (BC, CK)	(5) (S, NCK) – (S, CK)	(6) (S, NCK) – (BC, NCK)	(7) (S, NCK) – (BC, CK)	(8) (S, CK) – (BC, NCK)	(9) (S, CK) – (BC, CK)	(10) (BC, NCK) – (BC, CK)
Beyond 40 km of urban centre	0.14	0.21	0.1	0.22	0.39	0.53	0.35	0.13	0.93	0.11
Within 5 km of urban centre	0.31	0.4	0.35	0.31	0.41	0.73	1	0.63	0.39	0.72
Standardised entry time	−0.12	0.1	0.02	−0.21	0.23	0.49	0.65	0.71	0.13	0.3
Survey date	3.55	3.64	3.7	3.76	0.54	0.26	0.12	0.64	0.36	0.63
New strata	0.09	0.07	0.05	0	0.83	0.53	0.05	0.67	0.05	0.09
Female	0.32	0.25	0.33	0.39	0.25	0.91	0.29	0.17	0.02	0.29
Literate	0.8	0.8	0.82	0.78	0.89	0.75	0.6	0.66	0.74	0.41
Bank account holder	0.91	0.86	0.85	0.93	0.27	0.1	0.56	0.9	0.16	0.04
Age	40.01	40.06	38.27	38.24	0.97	0.12	0.15	0.14	0.16	0.98

Notes: Table compares (S, NCK), (S, CK), (BC, NCK), and (BC, CK) across whether the village is very rural, whether it is peri-urban, time of entry for endline survey, date of entry, whether the village was reassigned, gender of subject, literacy of subject, whether the subject has a bank account, and age of subject. Columns 1–4 present means by covariate in the four treatment cells aforementioned, in that order. Columns 5–10 present p-values of pairwise comparisons of differences in means across cells.

Second, we assessed knowledge of facts surrounding demonetisation. We surveyed the respondent on thirty-four facts and calculate the fraction of correct responses.[21]

Third, at endline, we offered subjects an unanticipated choice between: (a) a demonetised Rs. 500 note; (b) an IOU to be filled in 3–5 days for Rs. 200 in two Rs. 100 notes; (c) an IOU to be filled in 3–5 days for Rs. 200 worth of dal. With a probability of $1/6$, subjects actually received the item they chose. To implement the payment, we returned to each household in the sample before exiting the village, rolled the die, and provided either the Rs. 500 or the IOU notice.[22] The reason for using the IOU, which obviously relied on the villagers trusting us, was to make sure that the villagers did not go for the lower amount because they could get it right away, rather than after going to the bank. We nevertheless worried about the cost of going to the bank and depositing the 500 rupee note into an account. As noted already, 89% of respondents had bank accounts. We also collected data about the actual cost of going to the bank (see Table 4); based on the data we collected, the median wait time at banks was 10 min in the area and the median village in our sample was about 20 min from a bank by foot.[23] At the time of our experiment, depositing the bill required no documentation of the source of the cash. Thus, selecting Rs. 200 or the equivalent was giving up more than one day's wages, even accounting for the time spent travelling and at the bank. We argue that this is evidence of confusion and measures a willingness to pay to avoid holding on to the demonetised note in a period where it was both legal and easy to convert.[24] Furthermore, we asked respondents who did not choose the Rs. 500 to provide an open-ended justification for their choice at the end of the survey module. Figure 4 shows that most individuals who did not choose the Rs. 500 note believed, mistakenly, that the deposit deadline had already passed. The choice between 200 rupees and the equivalent in *dal* was intended to capture general trust in paper currency and confusion about whether the 100 rupee bills had also become demonetised. Taking the money offered more flexibility, since *dal* was easy to buy in village stores.[25]

4. MODEL

We present a simple framework to organise our analysis of how the treatments affect endogenous communication. The model plays two roles. First, it allows us to precisely articulate predictions based on the image concern frictions that motivated our study. Second, it provides a vocabulary on which we rely after presenting our results to consider a number of alternative theories of communication, both with and without other frictions, and assess how they line up with our empirical findings.

21. It is certainly the case that some of the facts in our index are more relevant for decision-making than others. Thus, our knowledge score should be viewed as a noisy measure of decision-relevant information.

22. In practice, we surprised the respondents by giving them the value in non-demonetised notes (Rs. 100 notes) even when they chose the Rs. 500 bill, saving them the cost of going to the bank. Note that this was our last action before we exited the village; it occurred after each subject had already locked in their response.

23. At this time, there were still news reports of very long queues at banks and ATMs in other, more urban parts of the country. In our study area, the waits had become much more manageable compared to the weeks following the policy announcement. Nevertheless, we were concerned that the villagers' perceived wait time could be very large. Our survey data showed that this was not the case—the median perceived wait time was 15 min, which was consistent with reality.

24. One issue is that given the tight time constraints, some households may not have been able to get to the bank on time. Thus, choosing the Rs. 500 option is likely an underestimate of the decision payoff that we would have seen for a less time-constrained decision.

25. We explore this further in Supplementary Material, Appendix H.

TABLE 4
Bank summary statistics

	Median	Mean	SD	Obs
Actual wait time at banks (min)	10.00	11.86	(7.87)	51
Perceived wait time at banks (min)	15.00	17.06	(22.13)	32
Nearest bank (min)	20.00	19.84	(9.88)	63

Notes: This table gives actual wait time at banks near our sample villages. On the last day on which SBNs were accepted, we surveyed as many banks as possible near the study villages. Our enumerators made it to fifty-one banks, where employees were surveyed. The table also gives perceived wait time and perceived time taken to reach the nearest bank according to a sub-sample of the endline respondents.

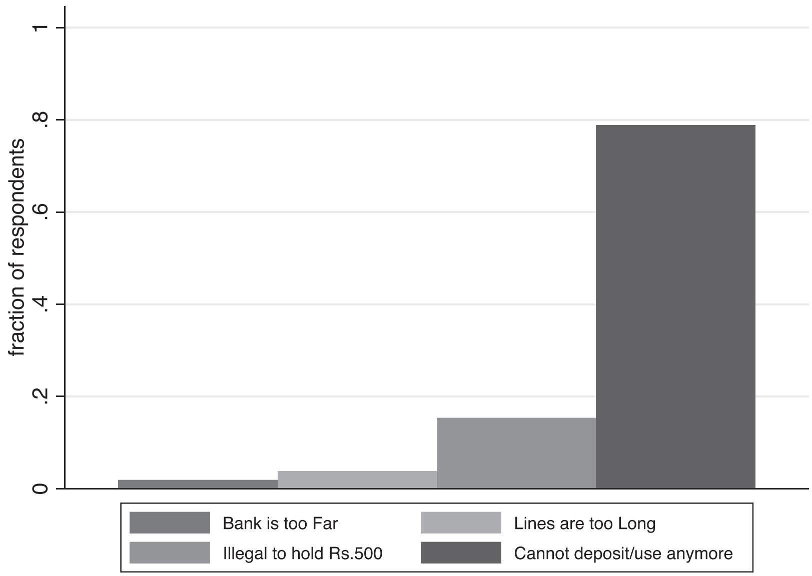

FIGURE 4
Why did you not choose 500?

4.1. *Basic framework*

This model is designed to study deliberate engagement in learning, with and without frictions.[26] An individual's information comes from three sources: pre-existing knowledge, information delivered by the experiment, and information acquired from social learning. We will define notation to keep track of these variables and analyse the engagement decision under different specifications of the individual's preferences.

Let us focus on one decision-maker, called D, and denote by $d \in \{0, 1\}$ his decision of whether or not to seek information by engaging in conversations. Let I_{D} be the indicator variable of whether this individual directly receives information in our experiment. Finally, let I_{S} be an indicator variable (potentially unknown to the decision-maker) of whether social information is available.

Let the random variable $V^{(I_{\mathrm{D}})}(0)$ denote the instrumental value to D if he chooses $d=0$ (does *not* engage in social learning); let $V^{(I_{\mathrm{D}}, I_{\mathrm{S}})}(1)$ be the instrumental value to D if he chooses

26. For the reasons behind this choice, as opposed to some alternative models where differences in proactive information sharing drive the effects, see Section 6.

$d = 1$ (*does* engage in social learning). Both random variables depend on I_{D}, whether D directly received information, but only the latter depends on I_{S}, the presence of social information. These values include all technological features of engaging in conversation—for example, the opportunity costs of time.

What is relevant to the individual's engagement decision is the instrumental payoff of engagement. This is a random variable

$$\Delta^{(I_{\mathrm{D}},I_{\mathrm{S}})} := V^{(I_{\mathrm{D}},I_{\mathrm{S}})}(1) - V^{(I_{\mathrm{D}})}(0). \tag{4.1}$$

This random variable has a cumulative distribution function $F_a^{(I_{\mathrm{D}},I_{\mathrm{S}})}$, which depends on the agent's *ability*, $a \in \{\mathrm{H}, \mathrm{L}\}$, in addition to $(I_{\mathrm{D}}, I_{\mathrm{S}})$.

The instrumental payoff of seeking depends on both I_{D} and I_{S}; D's beliefs about these are determined by what he knows in the given treatment. The individual always knows whether he directly received information, I_{D}, but he may be uncertain about whether social information is available, I_{S}.

The timing is:

(1) A dissemination treatment $\mathfrak{t}$ is exogenously selected (but not directly observed). The treatment has two dimensions: its breadth (BC, for broadcast, or S, for seed) and its publicness (CK, meaning that everyone is informed of the breadth, or NCK, where no information about breadth is delivered). The breadth determines the direct delivery indicator I_{D}.
(2) Individuals observe whether they are informed and, based on this, form beliefs about the presence of social information, I_{S}.
(3) Individuals observe the values $V^{(I_{\mathrm{D}})}(0)$ and $V^{(I_{\mathrm{D}},I_{\mathrm{S}})}(1)$ for their realised value of D and both possible values of S (*i.e.* conditional on both social information states). They may then engage in social learning.

The agent's total payoff from selecting $d = 1$ is

$$U = \Delta^{(I_{\mathrm{D}},I_{\mathrm{S}})} - f_{\mathfrak{t}}.$$

The important parameter $f_{\mathfrak{t}}$ is a treatment-dependent friction distorting the privately optimal engagement decision, such as an image cost of discussing the topic. For now, we leave this abstract, but our main hypothesis proposes a concrete form for this friction.

We assume that the distributions $F_a^{(I_{\mathrm{D}},I_{\mathrm{S}})}$ are exogenously given for each $(I_{\mathrm{D}}, I_{\mathrm{S}})$ and are not dependent on others' behaviour. Section 6 discusses extensions where others' engagement affects the returns to seeking information.

In our analysis throughout this section, we focus on the decision-making of a non-seed D, who receives information only in the broadcast treatments. Such individuals were the large majority of people in any village. Supplementary Material, Appendix D works out the details of seed D's as well.

4.2. *Frictionless model*

We first consider, as a benchmark, a frictionless model of engagement in learning. In this model, $f_{\mathfrak{t}} = 0$ for all $\mathfrak{t}$, and individuals engage in social learning if and only if $\mathbb{E}[\Delta^{(I_{\mathrm{D}},I_{\mathrm{S}})} \mid I_{\mathrm{D}}] \geq 0$, without any wedge distorting the decision.

We give two conditions, which are natural in our setting, jointly guaranteeing that, without frictions, (BC, CK) dominates (BC, NCK) in terms of volume of conversation as well as our other outcomes.

The first condition is that $\Delta^{(I_\mathrm{D},I_\mathrm{S})}$ is increasing in I_S: the incremental value of engaging in social learning is increasing in the availability of social information. The second condition is that D's subjective probability of the event that $I_\mathrm{S}=1$ is higher in (BC, CK) than under (BC, NCK). The rationale for this is that in the former case, the fact that everyone else has information is publicised along with the pamphlets themselves; the inequality is strict for our non-seed decision-maker D. We maintain the assumption from now on that both these conditions hold.

Under these two conditions, *there is strictly more engagement in social learning in (BC, CK) than in (BC, NCK)*: D places a higher probability on $I_\mathrm{S}=1$ and the distribution of values is shifted up in that case. These conditions imply that D places more probability on the events where Δ is large, and thus finds it worthwhile to seek in more states of the world. Turning to other outcomes, note that in (BC, CK), agents who choose to engage could have had a weakly greater payoff than in (BC, NCK) had they chosen not to seek, so they must receive a higher $\Delta^{(I_\mathrm{D},I_\mathrm{S})}$ by engaging in learning. This, in turn, implies that agents should be receiving greater informational benefits by engaging—*i.e.* greater knowledge and better choice outcomes. For similar reasons, (S, CK) and (BC, CK) should both dominate (S, NCK).

Adding common knowledge to broadcast treatments thus constitutes a key test of the frictionless model, and if this *hurts* outcomes, the frictionless hypothesis is rejected.

The effect of moving from (S, CK) to (BC, CK) under the frictionless model is more delicate. While it is natural to assume that clarification becomes more available in the second treatment ($V^{(I_\mathrm{D},I_\mathrm{S})}(1)$ stochastically increases), the endowment $V^{(I_\mathrm{D})}(0)$ also increases. The latter effect may dominate, reducing the net value of engagement and reducing overall engagement rates. That is, people may end up with less information in (BC, CK), precisely because they started with more and therefore have a weaker incentive to seek. Of course, this requires costs of seeking to be sufficiently high that the benefits of getting clarification are not worth it—in the introduction, we called this the high-cost hypothesis. Combined with the above-stated result about the comparison of (BC, CK) and (BC, NCK), the high-cost hypothesis predicts the ordering (BC, NCK) $<$ (BC, CK) $<$ (S, CK) in terms of seeking rates, and attributes the outcome to a certain ordering of instrumental costs and benefits.[27] We will probe this hypothesis in some of our more detailed empirical analysis.

4.3. *Seeking frictions from image concerns*

Now we posit that a nonzero friction term, f_t, comes from an image concern. We model image as being assessed by an observer, called O—a random person in the village who observes the seeking decision d that D makes and forms beliefs about D's ability. In turn, D cares about these beliefs. In particular, D values being perceived as more likely to be the high-ability type. Let $q_\mathrm{O}(d)$ be the observer's subjective probability that $a=\mathrm{H}$ (*i.e.* person D has high ability). Then the benefit of engaging in social learning is $\Delta^{(I_\mathrm{D},I_\mathrm{S})}$, while a potential cost is changing O's belief from $q_\mathrm{O}(0)$ to $q_\mathrm{O}(1)$, if the latter is lower. To incorporate both considerations, we posit that the net payoff of engaging in social learning is

$$U=\Delta^{(I_\mathrm{D},I_\mathrm{S})}-\underbrace{\lambda\mathbb{E}^\mathrm{D}[q_\mathrm{O}(0)-q_\mathrm{O}(1)]}_{\text{friction } f_\mathrm{t}},$$

27. Note that the first inequality depends on the cost of information, while the second one does not, as $\Delta^{(I_\mathrm{D}=0,I_\mathrm{S}=1)}$ always stochastically dominates $\Delta^{(I_\mathrm{D}=1,I_\mathrm{S}=0)}$.

where $\lambda > 0$ is a weight and the expectation is taken from D's perspective. D engages if and only if $U \geq 0$, *i.e.* if $\Delta^{(I_\mathrm{D}, I_\mathrm{S})} \geq f_\mathrm{t}$.

Solving this model involves solving for an equilibrium: D's seeking decision depends on the expected image payoff. In turn, the observer's $q_\mathrm{O}(d)$ is calculated using Bayes' rule, taking into account the different engagement rates of high and low types.

We assume that $\Delta^{(I_\mathrm{D}, I_\mathrm{S})}$ is stochastically higher for low-ability types,[28] so that a low-ability D always has at least as much to gain from seeking as a high-ability one, all else equal.[29] Then seeking is always weakly a signal of low ability, and thus $q_\mathrm{O}(1) \leq q_\mathrm{O}(0)$; it is strictly so when low-ability types have strictly more to gain. This implies that engagement in learning is distorted relative to the frictionless model.

The treatment affects both the informational benefits of engaging in social learning and the scope for image concerns to play a role. We do not develop the details here formally, but state the key ideas in a simple special case. (All omitted details can be found in Supplementary Material, Appendix D, where we make explicit D's beliefs about the informational benefits of seeking in the different treatments and consider weaker assumptions on payoffs.) Suppose that when they do not have information, both high- and low-ability types have equal distributions of $\Delta^{(I_\mathrm{D}=0, I_\mathrm{S}=1)}$: they stand to learn equally from others. But when they are informed, a low type needs clarification more than a high type. Then signalling concerns are strong when it is known that D has information, weaker when it is uncertain whether D has information, and weakest—indeed, in this special case, absent—when it is known that D does not have information.

4.3.1. Aggregate engagement rates. We now state predictions about the seeking rates in all four treatments, along with brief intuitions for the orderings.

(1) *(BC, NCK) dominates (BC, CK).* The expected informational benefits under (BC, CK) are shown to be only slightly greater. But under CK, it is known that the seeker received information and thus signalling concerns are fully active; in contrast, in NCK, it is not known whether D is informed, and therefore much less inference is drawn from his behaviour.

(2) *(S, CK) dominates (BC, CK).* In both cases, under our assumptions, a non-seed D assesses the same informational benefits from seeking, but has less initial information under (S, CK).[30] Moreover, a broadcast turns on signalling concerns (since it is known that D got information) whereas seeding makes it plain that D is uninformed, eliminating them.

(3) *(S, CK) dominates (S, NCK).* In the latter case, there is no particular reason to expect information to be available, whereas in the former, it is known that it can easily be found. Signalling concerns are small in either case, because others either know D is not informed or have no reason to believe that he is.

(4) *(BC, NCK) dominates (S, NCK).* In the former treatment, all individuals are in the same situation as seeds in the latter treatment: they have received information and do not know anything about who else has received it. Thus, to the extent that they engage in social learning, there should be much more of it happening in (BC, NCK), since there are at least ten times more non-seeds than seeds.

28. That is, the distribution $F_\mathrm{L}^{(I_\mathrm{D}, I_\mathrm{S})}$ first-order stochastically dominates $F_\mathrm{H}^{(I_\mathrm{D}, I_\mathrm{S})}$ for any value of the superscript.

29. Because the only thing that O observes is the decision d, while a and idiosyncratic value draws are private, the expected image penalty from D's perspective depends on d only.

30. If the information is easier to find in BC, then this could make the instrumental seeking benefits $V^{(I_\mathrm{D}, I_\mathrm{S})}(1)$ greater in (BC, CK). In that case, the prediction that we make here would hold under sufficiently strong image concerns.

In other words, in our model, more is not always more. In the NCK treatments, the broadcasting arm *increases* engagement by alerting people to at least the existence of information, without differentially activating signalling concerns, as explained above in (4). But under CK, broadcasting *decreases* engagement: with CK, people know about the social availability of information regardless, but broadcasting makes it clear that D is informed and activates ability-signalling concerns.

So far we have focused on volume of conversation, rather than knowledge or choice outcomes. We consider these other outcomes in the appendix. These comparisons are more delicate, because in some cases the information endowment is increased even as engagement in learning is decreased, making a comparison potentially ambiguous. For cases (1) and (3), the results noted above for volume of conversations extend straightforwardly, as endowments do not change in the comparisons. Moreover, we show that, under assumptions that are reasonable in our setting, the other comparisons also extend. For changes in endowments not to reverse our effects, we need that social learning is sufficiently important for enough of the population, relative to private processing of information. We believe that this is reasonable in our empirical context. The modelling details are in Supplementary Material, Appendix D.

The image concerns model has implications beyond the aggregate treatment effects discussed above. We next present predictions of the model concerning two dimensions of heterogeneity. First, we consider how low- versus high-ability agents respond to the mode of information delivery. Second, we look at how the complexity of the information itself affects the differences between treatment arms.

4.3.2. Differential predictions for high- versus low-ability agents. In our model, ability is the key dimension of heterogeneity across agents that drives the image concern. It is low ability—in the specific sense of low facility with understanding demonetisation facts—that agents are reluctant to reveal. Of course, whether engagement reveals low ability is endogenous.

We now elaborate on this in a bit more detail. Letting π_{H} be O's prior probability that $a = \mathrm{H}$, we have that, for either decision $d' \in \{0, 1\}$, by Bayes' rule,

$$\frac{1 - q_{\mathrm{O}}(d')}{q_{\mathrm{O}}(d')} = \frac{1 - \pi_{\mathrm{H}}}{\pi_{\mathrm{H}}} \frac{\mathbb{P}(d = d' \mid a = \mathrm{L})}{\mathbb{P}(d = d' \mid a = \mathrm{H})}.$$

Thus, seeking decisions are most informative about type (*i.e.* they change beliefs from the prior) when the rates of seeking are most different across types. The probabilities $q_{\mathrm{O}}(d')$ determine the seeking friction $f_{\mathrm{t}} = \lambda[q_{\mathrm{O}}(1) - q_{\mathrm{O}}(d)]$, and D seeks if and only if $\Delta^{(I_{\mathrm{D}}, I_{\mathrm{S}})} \geq f_{\mathrm{t}}$.

Thus, in the treatments where we predict that image concerns would drive down engagement, it should be that seeking is indeed done mostly by those of low ability. Equilibrium behaviour in our model entails that low aggregate engagement rates and over-representation of low-ability individuals among those engaging go hand in hand. Thus a signaling explanation predicts that:

(1) In (BC, CK), there should be less seeking than in (BC, NCK) and (S, CK) for both types.
(2) This reduction in seeking should be greater for high-ability types than for low-ability types. The appendix gives a natural assumption on the comparison of information values across ability types ensuring this holds.[31]

4.3.3. The difference between the long and short treatments. We now turn to the comparison of long and short pamphlets. Whether signalling concerns will be activated depends,

31. See Supplementary Material, Corollary 1 in Appendix D.

per the previous subsection, on the extent to which the instrumental payoffs to engaging differ between high- and low-ability types. That is, does the probability that $\Delta^{(I_{\mathrm{D}},I_{\mathrm{S}})} \geq f_{\mathrm{t}}$ differ significantly between H and L types? We will focus on the case (BC, CK), which, as we have argued above, creates the most potential for this difference.

Let us compare the short and long treatments (a two-pamphlet information delivery versus a lengthy pamphlet of twenty-four facts) through the lens of the model. One possibility is that long pamphlets confuse everyone, and nobody sees much value in speaking. In this case, we would expect engagement to go down in all treatments relative to short treatments, with no specific prediction for the relative seeking rates.

We now turn to the case where individuals do see value in discussing the pamphlets. In this case, a natural hypothesis is that high-ability people actually have a greater value of conversation in the long than in the short treatment. While a high-ability person is likely to be able to read and comprehend two facts, it is less likely that the person can make use of all twenty-four facts correctly without discussion. If this is the case, then relative to the short treatment, the random variable $V^{(I_{\mathrm{D}}=1,I_{\mathrm{S}}=1)}$ places more mass on high realisations even when D is a high-ability type. Suppose this happens in such a way that for all relevant values of the friction f_{t}

$$\frac{\mathbb{P}(d=1 \mid a=\mathrm{L})}{\mathbb{P}(d=1 \mid a=\mathrm{H})} = \frac{1-F_{\mathrm{L}}(f_{\mathrm{t}})}{1-F_{\mathrm{H}}(f_{\mathrm{t}})}$$

is smaller in long than in short treatments. This decreases the scope for signalling and increases seeking relative to short treatments, *ceteris paribus*.[32] Under this hypothesis, going from (BC, NCK, Long) to (BC, CK, Long) should generate less of a reduction in endogenous participation in social learning than going from (BC, NCK, Short) to (BC, CK, Short).

5. RESULTS

We now report the main empirical results and assess them in view of the model-based predictions (with and without frictions) that we have presented.

We begin with a visual inspection of the raw data for each of the primary outcomes in Figure 5. Panel A presents the number of people with whom the respondent had a conversation, panel B presents the knowledge index, and panel C presents the share choosing the Rs. 500 note. We can see the main results of the paper in the raw data itself. (S, CK) dominates (S, NCK) in all three measures. Similarly, (BC, NCK) dominates (BC, CK) in all outcomes, with most power in the raw data for volume of conversations. Finally, strikingly, (S, CK) and (BC, NCK) deliver similar, statistically indistinguishable results. Below we detail the results with more structured regression analysis and demonstrate that these findings carry through.

5.1. *Endogenous participation in social learning*

5.1.1. Volume of conversations. We begin by looking at which delivery mechanisms led to more or less engagement in social learning, measured by the number of people with whom the subject had spoken (henceforth "volume of conversations") over the prior three days about

32. Because the length of the pamphlet is common knowledge, both the decision-maker and observer know that the friction is lower, which makes things simple. See Supplementary Material, Appendix D for details.

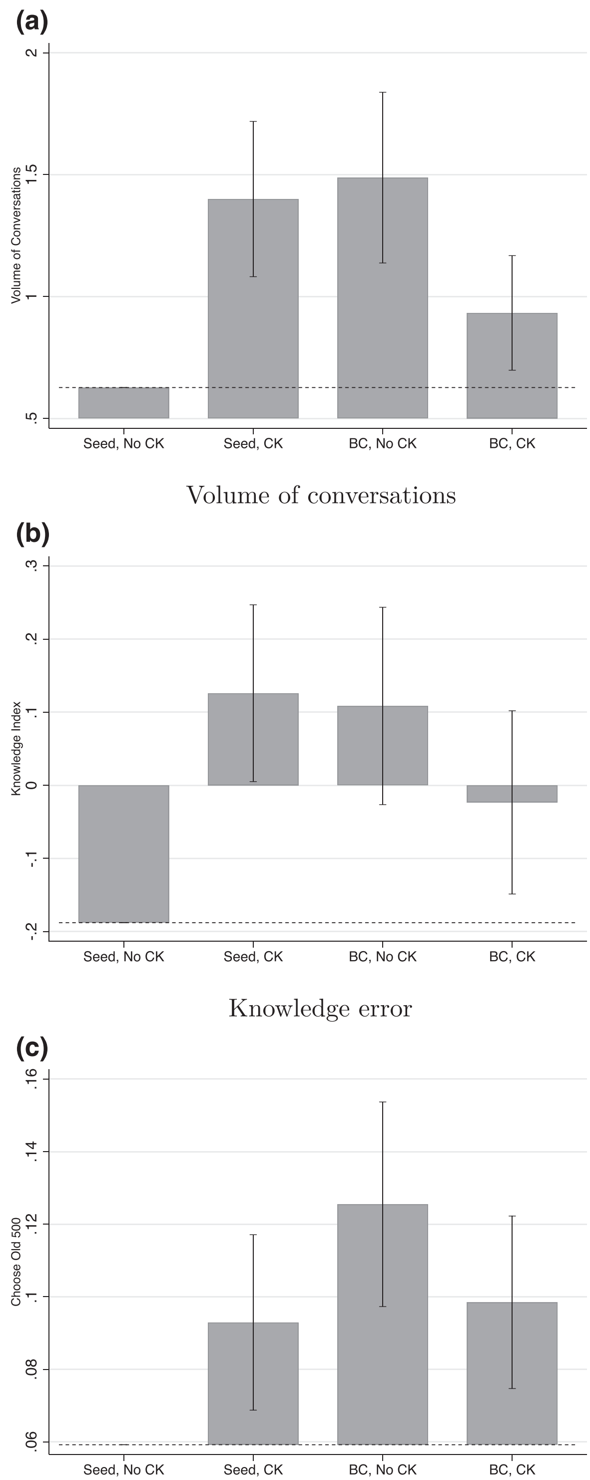

Volume of conversations

Knowledge error

Chose old 500

FIGURE 5

Raw data: core experimental outcomes

TABLE 5
Engagement in social learning, knowledge, and decision-making

Variables	(1) Volume of conversations	(2) # Secondary conversations	(3) # Primary conversations	(4) Knowledge	(5) Chose 500
CK	0.644	0.437	0.207	0.0313	0.0482
	(0.310)	(0.255)	(0.103)	(0.0126)	(0.0223)
	[0.0378]	[0.0871]	[0.0441]	[0.0129]	[0.0304]
BC	0.709	0.521	0.188	0.0280	0.0676
	(0.349)	(0.316)	(0.124)	(0.0140)	(0.0266)
	[0.0422]	[0.0986]	[0.130]	[0.0461]	[0.0109]
BC × CK	−1.493	−1.113	−0.380	−0.0505	−0.109
	(0.520)	(0.435)	(0.186)	(0.0189)	(0.0386)
	[0.00411]	[0.0105]	[0.0412]	[0.00764]	[0.00478]
Observations	1078	1078	1078	1082	1067
(S, NCK) mean	0.627	0.490	0.137	0.566	0.0592
CK + BC × CK = 0 p-value	0.0172	0.0262	0.239	0.154	0.0385
BC + BC × CK = 0 p-value	0.0254	0.0352	0.110	0.0572	0.104
CK = BC p-value	0.861	0.790	0.878	0.789	0.489

Notes: All columns control for randomisation strata (subdistrict) fixed effects. Other controls for each column selected with post-double selection Lasso from date and time of entry into the village, caste category of the treatment hamlet, distance from the village to an urban centre, and respondent-level controls such as age, gender, literacy, and potential seed status. Standard errors (clustered at the village level) are reported in parentheses and p-values are reported in brackets.

demonetisation.[33] Results are from regressions of the following form:

$$y_{ivd} = \alpha_d + \beta_1 CK_v + \beta_2 BC_v + \beta_3 BC_v \times CK_v + \gamma X_v + \lambda X_i + \varepsilon_{ivd}, \tag{5.1}$$

where i indexes the individual respondent, v indexes village, and d indexes the subdistrict, which was our unit of stratification. (S, NCK) is the omitted treatment arm. Village-level controls X_v include date and time of entry into the village, the caste category of the hamlet treated (and surveyed) in the village, and distance from the village to an urban centre. The respondent-level controls X_i include age, gender, literacy, and potential seed status. In all regressions, we use post-double-selection LASSO (see Belloni *et al.*, 2014) in order to efficiently select controls. Standard errors are clustered at the village level.

Table 5, columns 1–3, presents OLS regressions of the volume of conversations on the various treatments.[34] The coefficients are additive, so to compare (BC, CK) to the omitted category, it is necessary to add the coefficients: CK, BC, and BC × CK. In each regression specification, we present the p-values throughout, with standard errors clustered at the village level, and three additional key comparisons. The test (CK + BC × CK = 0) allows us to compare (BC, CK) to (BC, NCK), which, as argued above, represents a direct test of the frictionless model. The test "BC + BC × CK = 0" allows us to compare (BC, CK) with (S, CK), while the test "BC = CK" allows us to compare (BC, NCK) with (S, CK).

The outcome variable in column 1 is the volume of conversations about the demonetisation in which the respondent took part over the prior three days. Going from (S, NCK) to (S, CK)

33. Supplementary Material, Table P.1 of Appendix P presents results from the same regression where we look at whether an individual had any conversation as compared to the count of the number of conversation partners.

34. For all of our main results, we focus on our core 2 × 2 treatment design, pooling across the long and short lists of facts. Table 6 provides the analysis separately for long and short information and Section 5.4 discusses how one might interpret the length of the fact list through the lens of the model.

TABLE 6
Heterogeneity by length of information

Variables	(1) Volume of conversations	(2) # Secondary conversations	(3) # Primary conversations	(4) Knowledge	(5) Chose 500
CK	0.811	0.680	0.131	0.0208	0.0551
	(0.481)	(0.397)	(0.160)	(0.0158)	(0.0393)
	[0.0918]	[0.0865]	[0.414]	[0.186]	[0.161]
BC	0.973	0.667	0.306	0.0266	0.0789
	(0.534)	(0.474)	(0.213)	(0.0167)	(0.0350)
	[0.0687]	[0.160]	[0.152]	[0.110]	[0.0243]
Long	−0.0795	0.00630	−0.0858	−0.0127	−0.00757
	(0.361)	(0.320)	(0.127)	(0.0171)	(0.0292)
	[0.826]	[0.984]	[0.498]	[0.458]	[0.795]
BC × CK	−2.207	−1.608	−0.599	−0.0534	−0.144
	(0.719)	(0.613)	(0.258)	(0.0242)	(0.0545)
	[0.00215]	[0.00874]	[0.0202]	[0.0275]	[0.00833]
CK × long	−0.357	−0.479	0.122	0.0171	−0.0163
	(0.550)	(0.467)	(0.191)	(0.0249)	(0.0499)
	[0.515]	[0.304]	[0.522]	[0.490]	[0.744]
BC × long	−0.575	−0.317	−0.258	−0.000731	0.0261
	(0.665)	(0.603)	(0.227)	(0.0256)	(0.0534)
	[0.387]	[0.599]	[0.256]	[0.977]	[0.625]
BC × CK × long	1.428	0.991	0.437	0.00785	0.0713
	(0.790)	(0.716)	(0.276)	(0.0373)	(0.0769)
	[0.0706]	[0.167]	[0.114]	[0.833]	[0.354]
Observations	1078	1078	1078	1082	1067
(S, NCK, Short) mean	0.523	0.385	0.138	0.564	0.0374
CK + BC × CK = 0 *p*-value	0.00428	0.0233	0.0305	0.0803	0.0121
BC + BC × CK = 0 *p*-value	0.0138	0.0205	0.0589	0.118	0.117
CK = BC *p*-value	0.787	0.978	0.463	0.735	0.543
CK × long + BC × CK × long = 0 *p*-value	0.0783	0.339	0.0190	0.347	0.311
CK × long = BC × long *p*-value	0.776	0.800	0.170	0.468	0.858
CK + BC × CK + CK × long + BC × CK × long = 0 *p*-value	0.448	0.283	0.537	0.695	0.451
BC + BC × CK + BC × long + BC × CK × long = 0 *p*-value	0.330	0.408	0.475	0.254	0.503

Notes: All columns control for randomisation strata (subdistrict) fixed effects. Other controls for each column were selected with post-double selection Lasso from date and time of entry into the village, caste category of the treatment hamlet, distance from the village to an urban centre, and respondent-level controls such as age, gender, literacy, and potential seed status. Standard errors (clustered at the village level) are reported in parentheses and *p*-values are reported in brackets.

increases the number of conversations per capita by 103% (0.64 more conversations, $p = 0.04$). This result is consistent with the frictionless model detailed above—adding information about the identity of the seeds makes it easier to find someone to consult. It is clear in principle that this particular result could come from the fact that seeds have a stronger motivation to spread information under (S, CK). However, we do not think this kind of supply response fully drives our results for two reasons. First, in Supplementary Material, Appendix I, we show the same regression split by whether the household was a seed or not and demonstrate that the increase in conversations by seeds in (S, CK) cannot account for all the increase in conversations that non-seed households must have had in (S, CK). That is, conversations between non-seed households must have increased.[35] To see why, we can do a conservative back-of-the-envelope calculation. Supplementary Material, Table I.1 shows that a seed individual in (S, CK) participates in 1.8 extra conversations relative to (S, NCK). There are five such individuals and so there are nine more conversations generated. Looking at the non-seeded individuals, we see an increase of 0.6 conversations per respondent. In a village of fifty households, say with two eligible respondents per household (household head and spouse), this leaves ninety potential non-seed respondents out of whom we have an average increase of 0.6 conversations per respondent. That means that there are twenty-seven conversations that involve at least one non-seeded household; from before we know at most nine of these can be with seeds. So a minimum of eighteen, or 67%, of the conversations must be among non-seeds.[36] Second, we collected data about the nature of the conversations—whether they were the result of a directed question or statement about demonetisation (what we call primary conversations) or merely something that came up in a broader conversation (secondary conversations). These results are reported in Section 5.1. They make it clear that most of the increase came from secondary conversations, not from people going to ask questions of seeds or from seeds coming to deliver a message.

Next, we compare the common knowledge treatments. Going from (S, CK) to (BC, CK)—which typically corresponds to a tenfold increase in the number of households informed (from five households to all households)—leads to a 61% *decline* in the volume of conversations (0.78 fewer conversations per respondent, $p = 0.025$). This is a main prediction of the signalling model. It also could be consistent with the frictionless model if receiving the broadcast substantially lowers the net value of seeking. But then, as discussed in Section 4.2, we would see at least as much of a decline between (S, CK) versus (BC, NCK); in fact the corresponding two-point estimates are similar and not statistically distinguishable ($p = 0.861$). (The signalling theory does not predict a clear ranking between these two, but does imply that seeking in both should be high outside the high-cost regime.)

When we go to (BC, NCK) from (S, NCK), then, in sharp contrast to the previous result, we find an increase in the volume of conversations by 113% (0.709 more conversations, $p = 0.042$). This makes intuitive sense and is a prediction of any version of the model: essentially with (S, NCK) a typical household does not even know that there is something to converse about, whereas that is not true with (BC, NCK).

The move from (BC, NCK) to (BC, CK) leads to a 63% *decline* in the volume of conversations (0.84 fewer conversations, $p = 0.017$). Since it is easier to find people to consult when

35. Recall that every village had "seed" households selected by the same process *ex ante*, but in BC treatments all households were treated. In Supplementary Material, Appendix I, Table I.1 shows that all our main results hold for the households that are not seeds.

36. We can be even more conservative and imagine for some reason that every household has only one individual. Even in that case, the same calculation yields 13.5 new conversations involving at least one non-seed, and at the maximum nine of these could be between seed and non-seed, still leaving 33% purely among non-seeds.

there is common knowledge of who is informed, this is inconsistent, as we argued above, with the frictionless version of the model.

All of these observations are consistent with the version of the model with signalling frictions.[37]

5.1.2. Impacts on the types of conversations. We collected information on the number of conversations by type: primary (initiated with the purpose of talking about demonetisation) and secondary (the meeting was initiated for some other purpose but then one of the parties brought up demonetisation). Columns 2 and 3 of Table 5 break up the number of conversations that the subject participated in by whether they were secondary (column 2) or primary (column 3). Secondary conversations comprise the vast majority, 78%, of reported conversations. As columns 2 and 3 make clear, our core results broadly go through for each type of conversation, but significantly more of the impact of the interventions comes from the secondary conversations.[38] Consistent with that, column 3 of Supplementary Material, Appendix Table I.1 shows that the gap between the number of conversations in (S, CK) relative to (S, NCK) does not appear to be driven by the seed actively going out to explain the information to others, nor by others actively seeking out the seeds. The primary driver of the increase in conversations here is conversations among non-seeds, and we see no evidence of effort by seeds to coordinate conversations about the topic.

5.2. *Information aggregation and choice*

Columns 4 and 5 of Table 5 show how knowledge of the demonetisation rules and incentivised choice behaviour depend on the (randomised) information environment. Recall that the quality of the respondents' choices depended on their understanding of the demonetisation rules.

In column 4, we look at knowledge. It should be evident that more conversations need not imply greater knowledge—for example, even though there are fewer conversations happening in (BC, CK) as compared to (S, CK), ten times more households received information under broadcast treatments and it is entirely possible that they knew more.

Our main measurement of the knowledge outcome after our interventions is based on the answers to thirty-four questions about the demonetisation policy asked at the endline.[39] The mean knowledge index in the (S, NCK) group is 0.566. Going from (S, CK) to (BC, CK) leads to a 3.8% *reduction* in the knowledge index ($p = 0.057$). This shows that though 100% of households receive information instead of 10%, the amount of aggregated information that a random household has at the end of the day is actually less, not more. As we have discussed, while striking, this could happen in a frictionless world if endowments deter seeking enough (which happens under the high-cost hypothesis); still, by the same argument as before, the very similar performance of (BC, NCK) and (S, CK) refutes this theory.

Within broadcast strategies, adding common knowledge leads to a 3.2% *reduction* in knowledge, though the effect is not statistically significant ($p = 0.154$). In addition, going from (S, NCK) to (S, CK) increases the score on the knowledge index by 5.5% ($p = 0.0129$), and going from (S, NCK) to (BC, NCK) makes people better informed and improves knowledge by

37. Note that the same patterns emerge when using a binary indicator for having any conversations as the dependent variable. See Supplementary Material, Appendix Table P.1.

38. On the other hand, the *relative* increase in conversations is larger for the primary variety.

39. Recall that our treatment only gave information on a small subset of these thirty-four facts. We explore whether knowledge improvements are driven by the facts that were actually on the pamphlets in Supplementary Material, Appendix H.

4.9% ($p = 0.046$). Reassuringly, reductions in knowledge happen exactly where we see conversations declining, suggesting that conversations are an important source of information. In column 5, we turn to the impact of our experimental treatments on incentivised choice. We look at whether subjects choose the Rs. 500 note on the spot, which they could still deposit in their accounts, or an IOU worth Rs. 200 to be paid in 3–5 days, taking a loss of about 1.5 days' wages. The probability of selecting the Rs. 500 note in the omitted category (S, NCK) is only 5.92%. Within broadcasting strategies, adding common knowledge leads to a 48% *decline* in the probability of choosing the Rs. 500 note ($p = 0.0385$). This reversal is prima facie inconsistent with the frictionless model, as already argued. Going from seeding to broadcasting, conditional on common knowledge, leads to a 38.6% or 4.14pp *decline* in the probability of choosing the Rs. 500 note ($p = 0.104$), which is also striking. In addition, going from (S, NCK) to (S, CK) leads to a 4.8pp or an 81% increase in the probability of choosing the Rs. 500 note ($p = 0.03$); going from (S, NCK) to (BC, NCK) corresponds to a 6.76pp or 114% increase in the probability of choosing the Rs. 500 note ($p = 0.011$). These results are fully consistent with the results on conversations and knowledge. More conversations led to better knowledge, which in turn, allowed for improved decision-making.

In a world without common knowledge, the conventional wisdom holds: increasing the number informed encourages more conversations and better decision-making. However, under common knowledge, broadcasting information actually backfires, leading to worse outcomes across the board. One bottom-line result is that seeding just five households combined with CK makes the outcomes indistinguishable from (BC, NCK), where ten times as many people were seeded. Perhaps more strikingly, either holding CK fixed and moving from seed to BC or holding BC fixed and moving from NCK to CK actually reduces conversation volume, knowledge, and quality of choice. Less is more.

We also note that we find stark impacts of providing meta-knowledge despite the fact that our common knowledge treatments were most likely only partial. Some individuals in (BC, NCK) almost surely observed neighbours receiving pamphlets, and some in (BC, CK) may have never read the text indicating that everyone received the same pamphlet. This points to the power of meta-knowledge in our setting.

We now turn to some additional cuts of the data. First, we consider how low- versus high-ability agents respond to the mode of information delivery. Second, we look at how the complexity of the information itself affects the differences between treatment arms.

5.3. *Differential outcomes for high- versus low-ability agents*

In order to test the ability-based predictions laid out in Section 4.3, we need a pre-determined measure of agent ability pertaining to demonetisation, such as pre-intervention policy knowledge. Because we were not able to collect a household-level panel, we do not observe baseline knowledge for the same respondents that comprise our endline sample. Instead, we use information from our baseline surveys, conducted with a different set of households in the study villages, to construct a mapping from demographic covariates to predicted baseline policy knowledge. We then use this mapping to classify individuals in the endline survey sample into either high or low predicted ability categories.

Specifically, we first construct a random forest (RF) algorithm applied to our baseline sample to predict their baseline knowledge (raw score in our baseline knowledge survey). We allow as predictors all the demographic variables that were collected in both the baseline and endline surveys. This set of predictors includes age, gender, a coarse occupation category, subdistrict, distance to city, subcaste, and caste category. Our random forest model is calibrated to minimise the root mean squared error of the predictions relative to the true knowledge at baseline (which

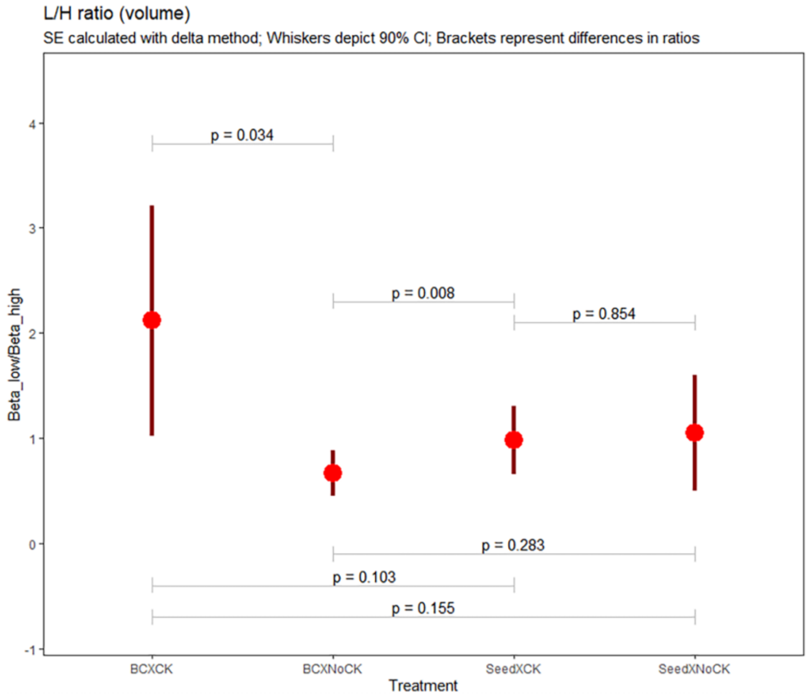

FIGURE 6

Treatment effect ratios of low-ability to high-ability. This figure illustrates relative seeking rates of low- versus high-ability agents for each treatment group. Ratios are constructed using coefficients from regressing an indicator for having any conversation on an indicator for high ability, treatment indicators, and their interactions, controlling only for subdistrict fixed effects (see Supplementary Material, Appendix Table F.2 for regression results)

we measured) and achieves a performance of 0.12 on out-of-sample data, which accounts for 17.9% of the variation in the data. Once we have the mapping from baseline characteristics to the knowledge score, we next apply it to the endline sample to generate a predicted baseline knowledge score for each household. We then classify endline respondents based on whether they are above- or below-mean within their village in this ability measure. This helps to ensure balance in ability across treatment assignment.

In Figure 6, we show that the model's ability-based predictions hold. Specifically, we construct and plot likelihood ratios of low-type to high-type seeking in each core treatment cell.[40] First, we see clearly that in (BC, CK), the likelihood ratio exceeds 1 ($p < 0.1$)—that is, more low types seek relative to high. Second, we find that when CK is added to BC, the relative seeking rate for low- versus high-ability individuals increases substantially ($p = 0.034$). That is, it is the high types whose conversations disproportionately decrease when moving from (BC, NCK) to (BC, CK).[41]

40. To do this, we regress an indicator for having any conversations on an indicator for high ability, treatment indicators, and their interactions, controlling only for subdistrict fixed effects. From the estimated coefficients, we construct the treatment-specific, relative seeking rates of low- versus high-ability agents. We use the delta method to estimate standard errors for the ratios and test the differences across treatments.

41. We would expect to see effects on both predicted types even if one is unaffected, since ability predictions are noisy and the separation between the types is imperfect.

In Supplementary Material, Appendix Table F.1, we show the results of expanding Equation (5.1) to incorporate heterogeneous treatment effects by predicted ability. Both high- and low-ability types decrease conversations in (BC, CK) relative to (BC, NCK) and (S, CK), although the reductions are not statistically significant for the low-ability types. Moreover, moving from (BC, NCK) to (BC, CK) leads to a greater differential reduction in communication for the high types ($p = 0.012$).

Our results are consistent with the predictions of our ability-based model, as outlined in Section 4.3, and support our preferred interpretation of the main experimental results. Nonetheless, we recognise that a few caveats are in order. Notably, our proxy for ability is a predictor based on demographic covariates, and so predicted ability could in principle be correlated with other individual-specific characteristics that are themselves relevant for signalling, but fall outside the scope of our model. For example, our model abstracts away from publicly observed attributes that might be correlated with how precisely others can assess someone's ability, or the activation of signalling concerns. We think that there is an interesting role for future research to assess where in the ability distribution signalling concerns are most acute.

5.4. *Differential outcomes for complex (long) versus simple (short) signals*

In Table 6, we test for the differential predictions in the long versus short treatments that were laid out in Section 4.3. Specifically, we look at how going from two to twenty-four facts differentially impacts the effects of interest. Focusing on volume first, in Table 6, we find that going from (BC, NCK) to (BC, CK) is less of a deterrent to primary conversations when there are many facts ($p = 0.078$).[42] This finding is consistent with our signalling story—it is natural that there would be less of an image cost involved in seeking clarification about a lengthy booklet of facts compared to seeking it about two facts. We find qualitatively similar patterns for our knowledge and choice outcomes, but, unsurprisingly, the results are noisier there.

For completeness, in Supplementary Material, Appendix Table G.1, we show the treatment effects from distributing long versus short pamphlets, pooling the data across the other two treatment cells. More information per pamphlet does not lead to more conversations or better outcomes. Providing a twelvefold increase in the number of facts leads to no statistically significant benefits in any of our primary outcomes. In all three cases, the coefficients are negative.

Again, these results are consistent with the model's predictions from Section 4.3, building confidence in our preferred interpretation. However, there could be alternative explanations for these results based on different channels, such as readers valuing clarification for different pamphlet lengths differently.

6. ASSESSING ALTERNATIVE MECHANISMS

We have presented evidence that our results are consistent with the image concerns mechanism that motivated our experimental design. We now turn to whether other mechanisms could be consistent with our findings. An important class of alternative models considers endogenous decisions of whether to *share* information. Indeed, sharing behaviour shaped by image concerns or other incentives can explain some of our findings. In this section, we examine both the parallels and the contrasts between models of endogenous seeking versus sharing.

42. We find a clearly negative effect of going from (BC, NCK) to (BC, CK) for short pamphlets, corresponding to the test on $\text{CK}+\text{BC}\times\text{CK}$ ($p = 0.0043$). The same effect is less pronounced when making the same move for long pamphlets.

We also discuss a number of other alternatives involving some well-studied social learning frictions to argue that the image frictions we emphasise are quite distinct in their predictions.

Our overall takeaway is that seeking with image concerns is a particularly parsimonious explanation of the facts. We do not insist that there cannot be other equally parsimonious stories, or that the signalling mechanism is the only one operating. However, we do give reasons why a number of natural alternatives cannot by themselves explain all the patterns we observe. In the process, we sketch out how to combine the image concerns model with other important mechanisms relevant to endogenous engagement.

6.1. *Active information sharing*

In our basic model, the main endogenous decision is whether to engage in social learning in order to acquire information, and the theoretical counterpart of conversation volume is $\bar{d}$, average seeking rate. Learning comes from active seeking. In this subsection, we introduce a distinct type of conversation, in which people spontaneously bring up the topic and actively share what they know. This would be a different contribution to volume, which we would also measure in our outcomes, but which is driven by different behaviour. Our goal is to examine which of our results can be explained by incorporating active information sharing that is endogenous to the treatments.

6.1.1. A simple active sharing model. The simplest sharing explanation is that people share information when they have information and believe that others may not have it. To tie this into the model, we can think of the observer O as the active sharer and posit that she is likelier to share when her counterparty is more likely to be uninformed, *i.e.* her sharing increases in $\mathbb{P}^{\mathrm{O}}(I_{\mathrm{D}}=0)$. For now, we study this as a mechanical rule: we simply posit that people like to offer helpful or interesting information for whatever reason. Can this type of theory by itself explain our results? To sharpen this question, consider a model with *only* active sharing. Each individual participates in a number of active sharing conversations. There is no seeking effort, and these conversations are accessed homogeneously throughout the village.

This hypothesis predicts the least active sharing in the (BC, CK) treatment: everyone is sure that others have heard the information. In (BC, NCK), everyone has information, but if they think it sufficiently likely that others might not have it, then we expect to see much more sharing. In (S, CK), there are also more conversations than under (BC, CK), because the seeds are aware that they know but others do not. Thus, this theory can explain at least some of the empirical patterns.

This simple model has a harder time accounting for heterogeneity across ability types. Recall from Section 5.3 that in (BC, NCK) and (S, CK), people of high and low predicted ability type (knowledge about demonetisation issues) report similar engagement in conversations about demonetisation. This is consistent with the active sharing story, where everyone is exposed to information. However, in (BC, CK), the agents of low predicted ability empirically have a much smaller reduction in conversation volume than those of high ability. This is inconsistent with the hypothesis that the overall reduction in this treatment is driven by a reduction in active sharing to which everyone is exposed.[43]

To account for these facts, one could layer on top of the active sharing behaviour a seeking decision—the frictionless seeking model from Section 4.2. Under this more elaborate alternative,

43. One could hypothesise that conversations are targeted by the active sharers to reach those in need, but this would not explain this happening only in the (BC, CK) treatment.

while there is little active sharing in (BC, CK) for the reasons discussed above, there is substantial active seeking, by low types only. This theory's explanation of the high types' not engaging is that their endowment strongly deters seeking in (BC, CK). The difficulty with this story is that we know that high types' informational benefits from conversation are actually considerable, based on knowledge and choice outcomes in (BC, NCK). So it would have to be that the *costs* of receiving information via others' active sharing are much lower than of seeking information, implying that high types are willing to do only the former—a version of the high-cost hypothesis we have mentioned above. These observations identify the conditions under which this theory could work.

However, in examining this explanation, it is worth recalling that in our data, a large majority of reported conversations are secondary to some other conversation (rather than sought out mainly for discussing demonetisation). In such secondary conversations, the physical costs (*e.g.* time and travel) have been paid regardless of whether a listener asks questions, and thus large incremental costs of seeking that are *not* dependent on treatment seem less plausible.

6.1.2. Active sharing with image concerns. We now return to a sharing-only theory and ask whether augmenting it with image concerns on the speaker's side can account for our findings.

In devising a theory of active sharing that can account for heterogeneity across types, it is worth noting that many of the concerns that encourage or deter seeking may also encourage or deter speaking. For example, suppose people are judged positively for being discerning about what information is interesting to others. Discerning types share novel information, whereas boring types share redundant information. Such image concerns reproduce the behaviour of the simple model of active sharing above (at least for the discerning types). Moreover, the heterogeneity by ability type that we observe in the data could be driven by this variant of a signalling story: the people we predict to have high ability are discerning, and they are the ones who refrain from sharing in (BC, CK), which explains why they are underrepresented among conversations there (even if they still benefit from and report hearing some active sharing).

This model is much closer to our main hypothesis, and so it could explain many of our comparisons. We thus certainly cannot reject that it plays a role. We will, however, argue that the theory has a hard time on its own accounting for the fact that (S, CK) and (BC, NCK) have comparable volume. Consider (S, CK) first. The active sharing story alone suggests that we should see a similar number of conversations in (S, CK) and (S, NCK)—in both cases seeds think they may have information that others do not, and thus start diffusions of information. It also predicts that we should see many fewer conversations in either of these two than in (BC, NCK)—a treatment where many more people are inclined to initiate an active sharing conversation. What we in fact see is that there are similar numbers of conversations in (S, CK) and (BC, NCK), and many more conversations in either of those two than in (S, NCK). Thus, the seeds in (S, CK) trigger a surprisingly large amount of active sharing, and this explanation can work only if the seeds' desire to reach out to others is stronger in (S, CK) than in (S, NCK).[44] This is not entirely implausible—perhaps the seeds feel pride or responsibility due to being known as one of a small number of informed people and, as a result, try harder to inform people. However, as discussed in Section 5.1, in the data, differences in conversations associated with

44. We spell out this argument in full detail in Supplementary Material, Appendix E.1.

seeds are only a small part of the observed difference between (S, NCK) and (S, CK), so this also cannot be the entire explanation.[45]

A final observation is that we would expect active sharing to play a larger role in villages where people start out being more informed at baseline. People are better at sharing information when they are more interested in and know more about the policy. In other words, better baseline information is a complement to sharing and, if anything, should reduce seeking. In Supplementary Material, Table Q.1, we split villages based on the average knowledge at baseline. We find that our main effects are much stronger in relatively *uninformed* villages. This supports an account based on demand for rather than supply of information.

6.1.3. Seeds being more motivated to provide public goods. A different kind of explanation focuses on the effort of informed people to understand, filter, and communicate the information in a useful way to others. Clearly, knowledge in our context is a public good. One could hypothesise that when a smaller group of people is publicly selected to provide a public good, as in (S, CK), they should provide more of it than in (BC, CK), where responsibility is diffuse. This, however, is at odds with standard models of public goods; as Banerjee *et al*. (2007) discuss, a fairly robust prediction of models of public goods is that while enlarging the set of people who are able to contribute to the public good often reduces per capita contributions, it should not markedly reduce aggregate provision in equilibrium. At a more basic level, theories based on intense effort by "deputised" seeds in (S, CK) are at odds with the fact that seeds in our data report few extra conversations in those treatments (recall Section 5.1). We flesh out these points in Supplementary Material, Appendix E.2.

6.2. *Standard social learning models*

Some standard social learning models are known to generate counterintuitive outcomes. In this subsection, we argue, nonetheless, that these mechanisms are not likely to explain our results.

A first observation is that many canonical "infection-type" models often used to study social learning, which have exogenous engagement in the learning process, share a basic monotonicity property (Bass, 1969; Bailey, 1975; Jackson, 2008; Jackson and Yariv, 2011; Aral and Walker, 2012): if more individuals are seeded with information, the number of people ultimately informed increases. In Supplementary Material, Appendix E.3, we discuss a version of this type of model that is most relevant to our setting, inspired by Möbius *et al*. (2015), in which people pass on everything they have heard, rather than summarizing it. We show that if initial endowments of information improve in the sense that they become Blackwell more informative about the state of interest, then the ultimate information of each individual also improves.

Thus, generating the kinds of reversals where adding information harms learning outcomes requires a different approach. One such approach is working with models where the focus is the quality of information aggregation rather than simply the extent of its diffusion. Another possibility is that an endogenous engagement margin (different from image concerns) plays a role. We consider several such models next.

6.2.1. Herding models. When agents aggregate their information into a coarse summary, we enter the world of herding or information cascades (Banerjee, 1992; Bikhchandani *et al*., 1992), where efficient information aggregation is not guaranteed. Is it possible that in this kind

45. A variant of this theory is that people share information that they are nearly sure that others do not have. This would explain why there are many conversations in (S, CK) which is the only case where people are sure that others did not also receive the same information. However, it cannot explain why there are almost as many conversations in (BC, NCK).

of setting, more information sometimes aggregates to worse outcomes? Unfortunately, characterising the extent of information aggregation and how it depends on parameters in general herding models tends to be very difficult. However, an approach of Lobel and Sadler (2015), which applies to sequential learning in arbitrary conversation networks, can be used to argue why strong "less is more" forces such as those our main model produces are unlikely to be produced by herding models. We flesh out the details of the argument in Supplementary Material, Appendix E.4, but the basic logic is something like this: consider, for simplicity, a binary decision—say, whether or not to accept certain denominations of currency. Individuals form opinions about this. Differences in private information and in messages received lead to heterogeneity in the strengths of their beliefs about the right decision. Lobel and Sadler (2015) show that in equilibrium, after social learning, most agents' decisions are at least as good as those decisions taken by "experts"—those agents who are very sure of the right answer based on private information (*i.e.* their own understanding) alone. The intuition can be seen most simply in a model where all predecessors are observed. If decisions were substantially worse than the expert benchmark for arbitrarily late movers, then the well-informed would act against the prevailing view, revealing their superior information, which would persuade others. Remarkably, the same remains true even when agents observe only some of their predecessors, under certain conditions. (The main substantive condition is that the network must be connected enough, with everyone having indirect access to many others.) An implication of this is that, in this type of model, improving information endowments always improves the expert lower bound, limiting the magnitude of any perverse effects. In other words, models of herding or information cascades will have difficulty explaining how adding information can lead to outcomes in which most people do worse than the individual decisions of "well-informed" individuals.[46]

6.2.2. Costs of remaining engaged in social learning. Acemoglu *et al.* (2014) elaborate a basic viral model of information diffusion by positing that people have the option of dropping out of the social learning process at any point of time, due to an opportunity cost of paying attention to it. When people drop out, it reduces what others can hope to learn, and thus precipitates further exit. In that model, under broadcasting, social learning is *improved* by making it common knowledge that many agents are informed. The reason is that this increases the amount of information that any one of them can expect to receive by a given time, and strictly increases incentives to stay engaged given a person's own level of informedness. Thus, such a model would predict an upward shift in equilibrium engagement in (BC, CK) relative to other treatments.

6.3. *Some other behavioural possibilities*

There are a number of more *ad hoc* behavioural assumptions that might account for some of our findings. We briefly review a few of them in light of the evidence.

6.3.1. Curiosity. We have so far assumed that the only reason for people to seek information is to be able to make better decisions. A potential alternative theory is that when something out of the ordinary happens and piques their interest, they investigate just to find out what is going on. In such a world, even absent signalling concerns, people's interests may be especially

46. Of course, one could simply posit that common knowledge of people being informed causes people to drop out of the learning process (thinning out the learning network enough that it hurts diffusion) but this simply begs the question of engagement incentives.

piqued in (S, CK) because they are told that there are some others who have information. In contrast, (BC, CK) creates no scope for such curiosity.

This argument by itself says nothing about why there is much more seeking in (BC, NCK) than in (BC, CK). To explain that, we could add the assumption that those who are informed in (BC, NCK) are trying to find out if others are informed and in the process have conversations that end up being informative. This still leaves unexplained the fact that in (BC, CK), the high-ability types seek less than low types, and that this is the only treatment where this is the case.

6.3.2. Mistaken perceptions and overconfidence. To explain the lack of conversations in (BC, CK)—in the data it is comparable to (S, NCK)—one could posit that participants mistakenly believed they understood the facts they were told (although, in fact, they had much room to understand them better). This runs counter to several different pieces of evidence. First, there is direct evidence from the knowledge surveys, in which many participants in (BC, CK) admitted ignorance even to us (panel B of Table 2). This evidence shows that substantial scope for learning remained and people knew that. Second and more fundamentally, such a theory does not predict less seeking in (BC, CK) than in (BC, NCK), which is what we observe. Indeed, insofar as subjects overcome overconfidence and ask others, the fact they know that others are informed should make it more, rather than less, appealing to ask them for clarifications.

6.3.3. Inferring the value of the information from the treatment. Finally, we consider the possibility that agents value the information differently across treatments, rather than having treatment-dependent seeking costs due to image concerns. One specific story that could match many of our key predictions is that people thought that the information was less valuable when it was distributed to more individuals. Under such a theory, they might not have even looked at the pamphlets in (BC, CK), throwing them away as "spam". We do not, however, find this possibility likely given what we know about our context. It goes against the fact that making public announcements—a small three-wheeler driving around the village with a loudspeaker attached to its top blaring out the message—is the most common way to get information to people in rural India about a possible tornado or other natural disaster.[47] This is (BC, CK) in our language, and people clearly do not assume all such messages are spam. Most people in our baseline also clearly wanted information and felt that neither they nor their neighbours knew enough about the post-demonetisation rules (Table 1).

6.4. *Taking stock*

We have presented a number of alternative frameworks in this section. In each case, we have argued that the alternative explanation is either incomplete (and requires additional *ad hoc* assumptions to fit all the facts) or inconsistent with what we know about the context. By contrast, the signalling model provides a fairly simple and unified account of all the rankings. Nevertheless, given the simplicity of our interventions, there may well be alternative behavioural mechanisms or combinations of those we mention above that could rationalise our findings. While we believe that signalling is an important component of our results, our main finding is there is a friction in seeking that varies not only with what people know but also what they think others know. A definitive decomposition of this friction into its ultimate constituents is beyond our scope and remains an important issue for further studies.

47. We also show evidence in Supplementary Material, Appendix Table H.2 that individuals in (BC, CK) did learn the facts *from their pamphlets*, but learned nothing more.

7. CONCLUSION

Whether social learning happens depends on whether people choose to ask questions. We show that, consistently with a model of endogenous engagement in social learning modulated by image concerns, the reach of an information delivery and the structure of meta-information matter *jointly* for the extent of participation in social learning—and in perhaps unexpected ways. In particular, we find some clear empirical rankings of policies; these rankings are consistent with image concerns deterring seeking, but not with a frictionless model of engagement in social learning. When looking at targeted seeding, going from no common knowledge to common knowledge increases conversations, but the opposite is true for broadcasting strategies. Moreover, conversations actually decline when, holding common knowledge fixed, more people are provided with information. Each increase or decline in conversation volume just mentioned is accompanied by a corresponding increase or decline in knowledge about the rules as well as in quality of choice.

Our model of image concerns provides a mechanism for meta-information to affect engagement that can explain both why the "more is more" ordering holds when it does, as well as the reversals in the data. The forces in the model are consistent with villagers' reports of their experiences in the context of the Indian demonetisation, as well as the more detailed cuts of the data by ability that we have discussed.

Of the full set of experimental interventions, two consistently perform well along all the dimensions—conversations, knowledge, and choice—and have comparable benefits to one another: seeding with common knowledge and broadcasting without common knowledge. While seeding with common knowledge is straightforward to implement, whether broadcasting without common knowledge is a feasible policy depends considerably on context. For example, posters in the village, loudspeakers on three-wheelers, radio, television, newspapers, or the village crier intrinsically entail a common knowledge component. Moreover, it may be difficult for the same entity (*e.g.* a local government or agricultural extension service) to consistently carry out a non-common knowledge broadcasting strategy without it eventually becoming (approximately) commonly known. At the same time, one can imagine contexts where flyers or SMS messages are a natural mode of communication, and in those cases it may take a while to become common knowledge.[48] Our results suggest that broadcasting without common knowledge can be quite effective when it is possible. But implementations may be blunted in their effectiveness if they are perceived as closer to broadcasting with common knowledge. On the other hand, seeding strategies benefit from common knowledge unambiguously.

Our empirical results, based on a three-day learning period, are most directly relevant for other settings where individuals need to act quickly on the information being disseminated. Often policy information is disseminated in advance of discrete decision points—*e.g.* information about new agricultural technologies disseminated near the time of annual planting decisions, information about specific job vacancies that may be filled quickly, or programmes advertised shortly before their application deadlines. Our results would also be relevant in settings where information stimulates bursts of conversation in the short run but stops being discussed as people move on to different topics.

Our results have implications for how researchers and policymakers should think about the use of broadcast media versus extension to educate individuals, and how extension strategies should be structured. Indeed, whenever contacting all households is feasible, the policymaker

48. This is perhaps most likely when the message is sent to everyone within some sub-group (for example, health messages about getting tested for diabetes or getting a flu shot may be targeted to specific age groups or risk categories) and people may not know the boundaries of the group.

may be able to do just as well by informing a few seeds and generating common knowledge of that as by directly informing all individuals (even without common knowledge). These lessons play out in related work. Banerjee *et al*. (2021) provide a policy-relevant example where seeding with common knowledge has been successful.[49]

Exploring other contrasts between seeding and broadcasting is a promising avenue for future work. Broadcasting strategies are inherently more democratic than seeding and may have different implications for information inequality. Moreover, in equilibrium, repeatedly calling upon the same individuals to act as seeds could concentrate power and change the social dynamic. These types of effects may affect the ultimate success of seeding strategies.

Finally, policymakers may be able to avoid engagement-related drawbacks of broadcast strategies by carefully curating the information shared in their campaigns. It is possible that in some applications, simplifying the information to the easy-to-process essential facts could remove the need for network-based aggregation after the official broadcast and neutralise the effects of endogenous social engagement frictions. Our results on message length show, however, that merely making messages brief may backfire, since that actually seems to intensify signalling concerns to the extent that communication is needed. In any case, the careful curation needed to make such a strategy successful may be costly and time-consuming, especially if communities have heterogeneous needs or if the policy information to be communicated varies across locations.

Acknowledgments. We thank Mohammad Akbarpour, George Beknazar-Yuzbashev, Leo Burzstyn, Chris Blattman, Pascaline Dupas, Matt Jackson, Rachel Kranton, Michael Kremer, Ted Miguel, Ben Olken, Evan Sadler, and many seminar participants for helpful comments and conversations. We are grateful to Ruben Enikolopov for an insightful discussion. We thank Pooja Suri, Arnesh Chowdhury, Piyush Tank, Tithee Mukhopadhyay, Sumhith Veda Aradhyula, Yann Calvó López, Anushka Chawla, Vasu Chaudhary, Yu-Chi Hsieh, Yixi Jiang, Devika Lakhote, Alokik Mishra, Shambhavi Priyam, Aneesha Parvathaneni, Awa Ambra Seck, Vedika Vyas, and Xuan Zhu for excellent research assistance. The paper was greatly improved by advice from four anonymous referees and the editor, Adam Szeidl. All errors are our own. Stanford IRB-35996.

Funding

We gratefully acknowledge financial support from The Pershing Square Fund for Research on the Foundations of Human Behavior and The National Science Foundation under SES-1658940.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

Code and data (all newly collected) used for analysis in this article are available with full documentation in the Zenodo online repository for *ReStud* at: https://doi.org/10.5281/zenodo.7860884 (Banerjee *et al*., 2023).

REFERENCES

ABALUCK, J. and GRUBER, J. (2011), "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program", *American Economic Review*, **101**, 1180–1210.

ACEMOGLU, D., BIMPIKIS, K. and OZDAGLAR, A. (2014), "Dynamics of Information Exchange in Endogenous Social Networks", *Theoretical Economics*, **9**, 41–97.

49. In a cross-randomised experiment with seventy-five policies covering 40% of the state of Haryana, Banerjee *et al*. (2021) examine information policies to encourage child immunisation. They study whether various incentives, reminders, and seeding policies (none, "information ambassadors", etc.) affect immunisation rates. The finding is that the most cost-effective policy is a combination of using seeds picked through the same "gossip survey" as described in this paper, together with light SMS reminders.

ARAL, S. and WALKER, D. (2012), "Identifying Influential and Susceptible Members of Social Networks", *Science*, **337**, 337–341.
BAILEY, N. (1975), *The Mathematical Theory of Infectious Diseases and Its Applications* (London: Griffin).
BANERJEE, A. V. (1992), "A Simple Model of Herd Behavior", *The Quarterly Journal of Economics*, **107**, 797–817.
BANERJEE, A., BREZA, E., CHANDRASEKHAR, A. G., *et al.* (2024), "When Less Is More: Experimental Evidence on Information Delivery During India's Demonetisation", *Review of Economic Studies*, **91**, 1884–1922.
BANERJEE, A., CHANDRASEKHAR, A. G., DALPATH, S., *et al.* (2021), "Selecting the Most Effective Nudge: Evidence from a Large-Scale Experiment on Immunization" (preprint arXiv:2104.09645).
BANERJEE, A., CHANDRASEKHAR, A. G., DUFLO, E., *et al.* (2013), "The Diffusion of Microfinance", *Science*, **341**, 1236498.
BANERJEE, A., CHANDRASEKHAR, A. G., DUFLO, E., *et al.* (2016), "Gossip: Identifying Central Individuals in a Social Network" (Working Paper).
BANERJEE, A., IYER, L. and SOMANATHAN, R. (2007), "*Public Action for Public Goods*", in SCHULTZ, T. P. and STRAUSS, J. A. (eds) *Handbook of Development Economics* (Vol. 4, Amsterdam: Elsevier Science), 3117–3154.
BASS, F. M. (1969), "A New Product Growth for Model Consumer Durables", *Management Science*, **15**, 215–227.
BEAMAN, L., BENYISHAY, A., MAGRUDER, J., *et al.* (2021), "Can Network Theory-Based Targeting Increase Technology Adoption?" *American Economic Review*, **111**, 1918–1943.
BELLONI, A., CHERNOZHUKOV, V. and HANSEN, C. (2014), "Inference on Treatment Effects After Selection Among High-Dimensional Controls", *The Review of Economic Studies*, **81**, 608–650.
BESHEARS, J., CHOI, J. J., LAIBSON, D., *et al.* (2013), "Simplification and Saving", *Journal of Economic Behavior & Organization*, **95**, 130–145.
BIKHCHANDANI, S., HIRSHLEIFER, D. and WELCH, I. (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, **100**, 992–1026.
BREZA, E., KAUR, S. and KRISHNASWAMY, N. (2020), "Propping Up the Wage Floor: Collective Labor Supply Without Unions" (Working Paper).
BREZA, E., KAUR, S. and SHAMDASANI, Y. (2017), "The Morale Effects of Pay Inequality", *The Quarterly Journal of Economics*, **133**, 611–663.
BREZA, E., KAUR, S. and SHAMDASANI, Y. (2021), "Labor Rationing", *American Economic Review*, **111**, 3184–3224.
BURSZTYN, L., EGOROV, G. and JENSEN, R. (2019), "Cool to Be Smart or Smart to Be Cool? Understanding Peer Pressure in Education", *The Review of Economic Studies*, **86**, 1487–1526.
BURSZTYN, L. and JENSEN, R. (2017), "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure", *Annual Review of Economics*, **9**, 131–153.
BUTERA, L., METCALFE, R., MORRISON, W., *et al.* (2019), "The Deadweight Loss of Social Recognition" (Tech. Rep., National Bureau of Economic Research).
CAI, J., DE JANVRY, A. and SADOULET, E. (2015), "Social Networks and the Decision to Insure", *American Economic Journal: Applied Economics*, **7**, 81–108.
CARVALHO, L. and SILVERMAN, D. (2019), "Complexity and Sophistication" (Working Paper No. 26036, NBER).
CHANDRASEKHAR, A. G., GOLUB, B. and YANG, H. (2018), "Signaling, Shame, and Silence in Social Learning" (Working Paper, Stanford University).
CHODOROW-REICH, G., GOPINATH, G., MISHRA, P., *et al.* (2020), "Cash and the Economy: Evidence from India's Demonetization", *The Quarterly Journal of Economics*, **135**, 57–103.
CONLEY, T. G. and UDRY, C. R. (2010), "Learning About a New Technology: Pineapple in Ghana", *American Economic Review*, **100**, 35–69.
DE HOOGE, I. E., ZEELENBERG, M. and BREUGELMANS, S. M. (2010), "Restore and Protect Motivations Following Shame", *Cognition and Emotion*, **24**, 111–127.
GERBER, A. S., GREEN, D. P. and LARIMER, C. W. (2008), "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment", *American Political Science Review*, **102**, 33–48.
GOLDFARB, A., MCDEVITT, R. C., SAMILA, S., *et al.* (2015), "The Effect of Social Interaction on Economic Transactions: Evidence from Changes in Two Retail Formats", *Management Science*, **61**, 2963–2981.
GRUENEWALD, T., DICKERSON, S. and KEMENY, M. (2007), "A Social Function for the Self-Conscious Emotions: Social-Self Preservation Theory", in TRACY, J. L. (ed) *The Self-Conscious Emotions: Theory and Research* (New York City: Guilford Press) 68–87.
JACKSON, M. O. (2008), *Social and Economic Networks* (Princeton, NJ: Princeton University Press).
JACKSON, M. O. and YARIV, L. (2011), "Diffusion, Strategic Interaction, and Social Structure", in BENHABIB, J., BISIN, A. and JACKSON, M. O. (eds) *Handbook of Social Economics* (Vol. 1, Amsterdam: Elsevier) 645–678.
KARING, A. (2018), "Social Signaling and Childhood Immunization: A Field Experiment in Sierra Leone" (Working Paper).
KAUR, S., MULLAINATHAN, S., OH, S., *et al.* (2019), "Does Financial Strain Lower Worker Productivity?" (Working Paper).
LESKOVEC, J., ADAMIC, L. A. and HUBERMAN, B. A. (2007), "The Dynamics of Viral Marketing", *ACM Transactions on the Web (TWEB)*, **1**, 5.
LOBEL, I. and SADLER, E. (2015), "Information Diffusion in Networks Through Social Searning", *Theoretical Economics*, **10**, 807–851.

MILLER, G. and MOBARAK, A. M. (2014), "Learning About New Technologies Through Social Networks: Experimental Evidence on Nontraditional Stoves in Bangladesh", *Marketing Science*, **34**, 480–499.

MÖBIUS, M., PHAN, T. and SZEIDL, A. (2015), "Treasure Hunt: Social Learning in the Field" (Working Paper No. 21014, NBER).

NIEHAUS, P. (2011), "Filtered Social Learning", *Journal of Political Economy*, **119**, 686–720.

PENNINGTON, C. R., HEIM, D., LEVY, A. R., *et al.* (2016), "Twenty Years of Stereotype Threat Research: A Review of Psychological Mediators", *PLoS ONE*, **11**, e0146487.

PEREZ-TRUGLIA, R. and TROIANO, U. (2018), "Shaming Tax Delinquents", *Journal of Public Economics*, **167**, 120–137.

RYAN, B. and GROSS, N. C. (1943), "The Diffusion of Hybrid Seed Corn in Two Iowa Communities", *Rural Sociology*, **8**, 15.

SIMONSOHN, U. (2021), "[99] Hyping Fisher: The Most Cited 2019 QJE Paper Relied on an Outdated STATA Default to Conclude Regression *p*-Values Are Inadequate", Data Colada. http://datacolada.org/99.

SPENCER, S. J., LOGEL, C. and DAVIES, P. G. (2016), "Stereotype Threat", *Annual Review of Psychology*, **67**, 415–437.

STEELE, C. and ARONSON, J. (1995), "Stereotype Threat and the Intellectual Test-Performance of African-Americans", *Journal of Personality and Social Psychology*, **69**, 797–811.