

A Report on “When Less Is More:  
Experimental Evidence on Information  
Delivery During India’s  
Demonetisation” by Banerjee et al.  
(2024)

Reviewer 2

February 07, 2026

v1



**isitcredible.com**

## Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made [here](#). Any errors in the report will be corrected in future revisions.

I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.

Plato, *The Apology of Socrates*, 21d

To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.

Karl Popper, 'Knowledge and the Shaping of Reality'

## Overview

**Citation:** Banerjee, A., Breza, E., Chandrasekhar, A. G., and Golub, B. (2024). When Less Is More: Experimental Evidence on Information Delivery During India’s Demonetisation. *Review of Economic Studies*, Vol. 91, No. 4, pp. 1884–1922.

**URL:** <https://doi.org/10.1093/restud/rdad068>

**Abstract Summary:** This paper offers theoretical and experimental evidence on the optimal strategy for policymakers to disseminate information, comparing broadcasting to everyone versus informing a small number of “seeds” who then spread the message. A field experiment during the 2016 Indian demonetisation found that seeding with common knowledge dominates broadcasting, and that adding common knowledge makes seeding more effective but broadcasting less so.

**Key Methodology:** Field experiment with a  $2 \times 2$  design (broadcasting vs. seeding, common knowledge vs. no common knowledge) and a third dimension (volume of information) during the 2016 Indian demonetisation, measuring conversation volume, policy knowledge, and incentivised choice quality.

**Research Question:** If the mode of dissemination affects engagement, which policy option ultimately generates the most knowledge?

## Summary

### Is It Credible?

This article presents a counterintuitive and policy-relevant thesis: that widely broadcasting information can sometimes be less effective than targeting a few individuals, specifically when the broadcasting creates “common knowledge” that everyone has been informed. Banerjee et al. argue that under a “broadcasting with common knowledge” (BC, CK) strategy, individuals are deterred from asking clarifying questions because doing so signals low ability or a lack of understanding—a reputational cost that outweighs the benefits of learning. Using a field experiment across roughly 225 villages during India’s 2016 demonetization, the authors claim that (BC, CK) leads to fewer conversations, lower knowledge, and worse decision-making compared to broadcasting without common knowledge (BC, NCK) or seeding information with common knowledge (S, CK). The headline claim is provocative: “adding common knowledge makes seeding more effective but broadcasting less so” (p. 1884).

The credibility of the “less is more” phenomenon—specifically the reversal where broadcasting with common knowledge performs worse than broadcasting without it—is supported by the data on social engagement, but the link to actual learning is statistically fragile. The article reports a significant 63% decline in conversations when moving from (BC, NCK) to (BC, CK). However, the corresponding decline in policy knowledge is not statistically significant ( $p = 0.154$ ), and the effect on the incentivized choice, while significant, is potentially confounded by factors other than information, such as trust in the researchers or transaction costs associated with banking. Furthermore, the operationalization of “common knowledge” was subtle—a line of text in a pamphlet—which raises questions about whether the treatment was strong enough to generate the complex higher-order beliefs required by

the game-theoretic model. The authors acknowledge that the treatment was likely “only partial,” yet the effects on conversation volume are stark. This suggests that while the suppression of dialogue is real, the precise causal chain leading from social friction to reduced knowledge is less definitively established than the headline implies.

The proposed mechanism—that image concerns deter high-ability individuals from seeking information—relies heavily on a proxy for “ability” that is methodologically weak. Because the authors could not collect panel data, they predicted baseline knowledge using a random forest model based on demographics like caste, gender, and occupation. This proxy explains less than 18% of the variation in actual knowledge and essentially conflates cognitive ability with social and economic status. Consequently, the finding that “high types” reduce seeking behavior in the common knowledge treatment might not reflect a signaling of intellect, but rather complex social dynamics related to caste or status hierarchies. Additionally, the article lacks a pure control group due to implementation challenges, meaning all effects are relative to a baseline of seeding without common knowledge. While the authors are transparent about these limitations, including the rushed nature of the fieldwork during a chaotic national crisis, the combination of a noisy mechanism proxy, a unique high-stigma context, and fragile statistical evidence for the knowledge outcome suggests that while the “less is more” effect occurred in this specific setting, the generalizability of the signaling mechanism remains an open question.

## **The Bottom Line**

The article provides compelling evidence that publicizing the breadth of an information campaign can paradoxically suppress social engagement, likely due to social frictions. However, the claim that this suppression leads to significantly worse knowledge outcomes is not supported by conventional statistical significance, and

the mechanism is obscured by a reliance on demographic proxies for ability. While the results powerfully illustrate how social context influences information diffusion, policymakers should be cautious about assuming these specific signaling dynamics will replicate outside the high-stakes, high-confusion context of India's demonetization.

## Potential Issues

**Weakness of the constructed proxy for ability:** The article’s analysis of heterogeneous effects, a key piece of evidence supporting its signaling model, relies on a constructed proxy for “ability” that has low predictive power and is potentially confounded. As the authors could not collect panel data, they used a random forest model on a separate baseline sample to predict policy knowledge from demographic covariates such as age, gender, caste, and occupation. This predicted score was then used as the ability proxy for the endline sample. The authors are transparent that this model accounts for only 17.9% of the variation in actual baseline knowledge, meaning it is a very noisy measure prone to misclassification (p. 1912). More importantly, because the proxy is built from demographic variables, it is not a clean measure of cognitive ability but rather a composite that likely includes social status and economic standing. The authors acknowledge this limitation, noting that “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” (p. 1913). This confounding makes it difficult to definitively attribute the observed heterogeneous responses to ability-based signaling, as opposed to signaling related to social position or other unobserved factors correlated with demographics. This weakness is particularly salient as the heterogeneity analysis is used to adjudicate between the article’s preferred model and alternative explanations (p. 1888).

**Execution flaws in the research design:** The study’s implementation suffered from several procedural issues that may compromise the integrity of the randomization and the precision of the results. The authors are transparent about these challenges, which arose from the rushed nature of the fieldwork. First, villages were randomized into treatments before it was verified that they met the study’s eligibility criteria, which could introduce bias if ineligibility correlates with treatment assignment (p. 1896). Second, in 16 villages, the intervention was administered to the “wrong

hamlet,” requiring the use of an instrumental variables strategy in a robustness check to account for the error (p. 1897, fn 19). Third, a set of 16 villages was added to the sample from a new subdistrict in a process that the authors state “was not randomly done” (p. L.1). While the authors conduct robustness checks that drop these villages (Appendix L), their inclusion in the main analysis introduces a non-random component to the sample. The cumulative effect of these deviations from the intended design introduces unquantifiable noise and potential for bias, reducing confidence in the study’s internal validity.

**Absence of a valid pure control group:** The study design lacks a valid “pure control” group that received no intervention, which limits the interpretation of the findings. The authors attempted to collect data from 30 “status quo” villages but ultimately excluded them from the main analysis because implementation failures, including non-random selection of replacement villages and geographic imbalances, rendered them “not entirely comparable to our core set” (p. 1895, fn 12; Appendix M). As a result, all treatment effects are estimated relative to the seeding of five individuals without a public announcement, the (S, NCK) arm. Without a “do nothing” benchmark, it is impossible to assess the absolute impact of any of the information delivery strategies. For instance, it remains unknown whether the best-performing strategies were genuinely beneficial compared to no intervention, or simply less detrimental than the worst-performing strategy. This omission restricts the article’s conclusions to the relative effectiveness of different interventions.

**Weak operationalization of the “common knowledge” treatment:** The study’s central claims hinge on the distinction between “common knowledge” (CK) and “no common knowledge” (NCK) treatments, but the manipulation used to induce common knowledge appears weak and its success was not validated. Common knowledge was operationalized via a single line of text in a pamphlet indicating that either all households received it or that specific “seed” individuals were informed (p. 1895). There is no manipulation check to confirm that participants read, under-

stood, and believed this text, or that it successfully generated common knowledge in the game-theoretic sense. The authors acknowledge this is a significant limitation, stating that their “common knowledge treatments were most likely only partial” and that some participants in the CK condition may not have read the text, while some in the NCK condition may have observed neighbors receiving pamphlets (p. 1911). While the authors interpret the fact that they still find effects as evidence for the “power of meta-knowledge,” this does not resolve the concern that the treatment was not cleanly implemented, making it difficult to attribute the observed differences solely to the presence or absence of common knowledge.

**Confounding of the incentivized choice outcome:** The article uses an incentivized choice—between a demonetized Rs. 500 note and an IOU for Rs. 200—as a key measure of “choice quality” and policy understanding. However, this measure is potentially confounded by factors other than knowledge. The authors acknowledge that the choice relied on villagers trusting the research team to return and honor the IOU, and that transaction costs (time and travel to a bank) or urgent liquidity needs could make the lower-value IOU a rational choice for some individuals (p. 1899). The experimental treatments themselves could have systematically affected trust or perceptions of transaction costs, thereby confounding the outcome. While the authors argue the choice still captures a “willingness to pay to avoid holding on to the demonetised note” (p. 1899), the conflation of policy knowledge with trust, liquidity preferences, and transaction costs means it is not a clean measure of choice quality derived from information alone.

**Debatable dismissal of the “active sharing” alternative mechanism:** The article considers an important class of alternative models where the results are driven by endogenous decisions to share information rather than to seek it. It argues against this mechanism, in part by showing it has difficulty explaining the observed heterogeneity across ability types (p. 1914). However, this argument relies on the weak and confounded ability proxy discussed previously. The article also argues that a

simple sharing model makes counterfactual predictions about the relative volume of conversations between certain treatment arms (p. 1915). While the authors engage substantially with this alternative (Section 6.1), their dismissal may not be fully conclusive. A more complex sharing model that also incorporates supply-side image concerns, such as a desire to signal knowledge by sharing novel information, could potentially explain some of the observed patterns. The article's analysis provides evidence against simpler sharing models but may not fully rule out this broader class of alternatives.

**Incomplete statistical support for a central claim:** The abstract and introduction make the strong claim that adding common knowledge to broadcasting leads to worse outcomes “on all three outcomes” (p. 1884). However, the statistical evidence for one of the three primary outcomes, policy knowledge, does not meet conventional standards of significance. The article reports that the negative effect of moving from broadcasting without common knowledge (BC, NCK) to broadcasting with it (BC, CK) on the knowledge index has a  $p$ -value of 0.154 (Table 5, p. 1907). The authors are transparent about this, stating in the text that the effect “is not statistically significant” (p. 1888). While the point estimate is in the hypothesized direction, the lack of statistical significance for a primary outcome variable weakens the claim of a uniformly negative effect and tempers the interpretation of the causal chain from reduced conversations to reduced knowledge.

**Limited generalizability due to extreme context:** The experiment's setting during India's 2016 demonetization—a sudden, chaotic, and highly salient national event—may limit the generalizability of its findings. The authors provide compelling evidence that in this specific context, seeking information was associated with negative traits like being “unintelligent” or “irresponsible,” making image concerns particularly acute (p. 1894). It is uncertain whether these strong image-based frictions would operate with the same intensity in more routine policy domains, such as agricultural extension or public health campaigns, where seeking clarification might be

viewed more neutrally or even positively. The authors acknowledge their results are “most directly relevant for other settings where individuals need to act quickly” (p. 1919), but the unique political and social intensity of the demonetization event raises questions about the external validity of the specific mechanisms identified.

**Presentation and transparency issues:** Several minor issues related to presentation and transparency appear in the article. First, the total number of villages in the study is reported inconsistently, with the introduction mentioning “200 villages” (p. 1885) while the methods section details a final analytical sample of 225 villages (p. 1897). Second, there is an unexplained discrepancy between the total endline sample of “1248 households” (p. 1896) and the smaller sample sizes of 1067–1082 used in the main regression tables, presumably due to item non-response that is not explicitly detailed (Table 5, p. 1907). Third, a critical detail about the implementation of the incentivized choice—that participants who chose the demonetized Rs. 500 note were “surprised” with its equivalent value in valid Rs. 100 notes, saving them the transaction cost—is disclosed in a footnote rather than the main methods section (p. 1899, fn 22). Finally, the qualitative evidence used to motivate the study’s mechanism was collected two years after the experiment in a different Indian state (Karnataka, not Odisha), a fact that is stated but could be easily missed by a casual reader (p. 1892).

## Future Research

**Validated measures of ability and confidence:** Future work should disentangle cognitive ability from social status to rigorously test the signaling mechanism. Rather than relying on demographic predictions, researchers could administer baseline cognitive assessments or confidence intervals regarding policy knowledge. This would allow for a precise test of whether it is truly the “high-ability” individuals who withdraw from social learning due to image concerns, or whether the effect is driven by status hierarchies where high-status individuals simply refuse to engage with peers they perceive as lower-status.

**Stronger manipulations of common knowledge:** To better understand the threshold at which meta-knowledge alters behavior, future experiments should vary the intensity of the common knowledge induction. Comparing subtle text-based notifications (as used here) against undeniable public signals—such as village-wide ceremonies, loudspeaker announcements, or public posting of recipient lists—would help validate whether the theoretical mechanism of higher-order beliefs is actually driving the observed behavioral changes, or if the effects are driven by simpler cues.

**Testing in routine policy environments:** Research should examine whether these image concerns persist in lower-stakes, routine policy settings. The demonetization context was uniquely chaotic and stigmatizing; it is unclear if the same “backfire” effect of broadcasting would occur in standard agricultural extension or public health campaigns. Replicating the design in a context where asking questions is normative or neutral would help establish the boundary conditions of the theory and determine if the “less is more” phenomenon is a universal feature of social learning or a pathology of crisis communication.

© 2026 The Catalogue of Errors Ltd

This work is licensed under a

**Creative Commons Attribution 4.0 International License**

(CC BY 4.0)

You are free to share and adapt this material for any purpose,  
provided you give appropriate attribution.

**isitcredible.com**