SIGNALING, SHAME, AND SILENCE IN SOCIAL LEARNING

ARUN G. CHANDRASEKHAR[‡], BENJAMIN GOLUB[§], AND HE YANG[§]

Abstract. We examine how a social stigma of seeking information can inhibit learning. Consider a Seeker of uncertain ability who can learn about a task from an Advisor. If higher-ability Seekers need information less, then a Seeker concerned about reputation may refrain from asking to avoid *signaling* low ability. Separately, low-ability individuals may feel inhibited even if their ability is known and there is nothing to signal, an effect we term *shame*. Signaling and shame constitute an overall *stigma* of seeking information. We distinguish the constituent parts of stigma in a simple model and then perform an experiment with treatments designed to detect both effects. Seekers have three days to retrieve information from paired Advisors in a field setting. The first arm varies whether needing information is correlated with a measure of cognitive ability; the second varies whether a Seeker's ability is revealed to the paired Advisor, irrespective of the seeking decision. We find that low-ability individuals do face large stigma inhibitions: there is a 55% decline in the probability of seeking when the need for information is correlated with ability. The second arm allows us to assess the contributions of signaling and shame, and, under structural assumptions, to estimate their relative magnitudes. We find signaling to be the dominant force overall. The shame effect is particularly pronounced among socially close pairs (in terms of network distance and caste co-membership) whereas signaling concerns dominate for more distant pairs.

Date: This Version: October 21, 2019.
We thank Ran Abramitzky, Nageeb Ali, Marcella Alsan, Abhijit Banerjee, Emily Breza, Leo Bursztyn, Gabriel Carroll, Esther Duflo, Pascaline Dupas, Georgy Egorov, Greg Fischer, Alex Frankel, Maitreesh Ghatak, Rema Hanna, Matthew O. Jackson, Navin Kartik, Asim Khwaja, Scott Kominers, Michael Kremer, Horacio Larreguy, Jeremy Magruder, Ted Miguel, Melanie Morten, Kaivan Munshi, Ben Olken, Juan Ortner, Gautam Rao, Andrei Shleifer, Frank Schilbach, Charlie Sprenger, Rob Townsend, Nils Wernerfelt, Juan Pablo Xandri, Leeat Yariv, as well as participants at the Berkeley Development Seminar, the MIT/Harvard Development Seminar, the Princeton Behavioral Seminar, the Chicago Booth Micro Seminar, the LSE CEP-STICERD Applications Seminar, the Stanford Junior Faculty Workshop, the Coalition Theory Network Workshop, and the Workshop on Information and Social Economics at Caltech. We are grateful to Linh Tô for a very helpful detailed reading of an earlier draft. Financial support from the NSF under grants SES-1156182 and SES-1559469, SCID at SIEPR, IRiSS, and The Pershing Square Fund for Research on the Foundations of Human Behavior is gratefully acknowledged. We thank Alokik Mishra, Devika Lakhote, Tithee Mukhopadhyay, Gowri Nagraj, Jimmy Narang, Aneesha Parvathaneni, Piyush Tank, and Vedika Vyas for their excellent research assistance.
[‡]Department of Economics, Stanford University; NBER; JPAL.
[§]Department of Economics, Harvard University.

1. Introduction

Even a fool, when he holdeth his peace, is counted wise; and he that shutteth his lips is esteemed as a man of understanding.

Proverbs 17:28

Information and advice from others can be crucial in making people aware of opportunities and guiding them to better choices—at work, in financial decisions, and in other domains (Duflo and Saez, 2003; Jensen, 2010; Cole and Fernando, 2014). Indeed, various interventions and policies—especially in environments where formal information channels are missing or unsuited to people's needs—seek to leverage social learning in order to spread information widely.[1] An important consideration is that to learn something, a person often needs to ask. The action of asking for or seeking information, in turn, is endogenous and may be inefficiently inhibited.[2]

We surveyed 122 villagers in rural Karnataka, India, asking about their social learning practices and the inhibitions they face. We focused on several topics: financial products, farming practices, and health decisions. First, the survey revealed that active learning was considerably less common than passive learning: 49% reported actively asking friends in their village for information on these topics; meanwhile, 95% reported hearing such information from friends without asking[3] and 90% from media sources (newspapers, TV, and Internet sources). Second, 88% of respondents reported that they felt a limit on the number of times they could approach a member of their community for information; in 64% of these cases, they reported that the cause of their reluctance was a desire not to appear uninformed or weak. These responses suggest that people trade off the value of seeking information against concerns about their image, and that there may be a social stigma of asking questions.

An essential feature of such an inhibition is that it involves others' beliefs about the asker—for instance, whether he is a competent, self-sufficient, or hardworking person. There are two mechanisms by which the prospect of bad assessments may inhibit people—a distinction that goes back at least to Goffman's (1963) seminal study of social stigma. One mechanism, a *signaling* concern, is about *managing others' beliefs*. For instance, a pupil concerned about how others perceive him may be reluctant to ask a teacher or a peer basic questions about an assignment, fearing that she could infer that the pupil is slow or lazy—a belief that would

[1]Examples in the context of economic development include financial and agricultural extension (Conley and Udry, 2010; Banerjee et al., 2013; Beaman et al., 2018). For an expanded discussion of inhibitions in social learning due to behavioral forces, and the significant welfare implications, see Kremer, Rao, and Schilbach (2019) Sections 2.4 and 6.3.

[2]For modeling of other endogenous decisions affecting information flow, see for instance Niehaus (2011), Acemoglu et al. (2014), Galeotti et al. (2013), and Calvo-Armengol et al. (2015).

[3]E.g., in conversations initiated by others, or by hearing others' conversations.

harm the student in the future.[4] But there is another type of inhibition that we consider part of stigma, which is not about managing beliefs but managing interactions in view of compromised beliefs. To illustrate, suppose a bad attribute—such as the pupil's ignorance—is apparent to the teacher irrespective of his decision about whether to ask questions. Even though seeking information is not a signal of low ability, the pupil may feel averse to seeking advice precisely when others have a negative assessment of him, as documented in a large qualitative literature in psychology. We call this type of inhibition *shame*.[5] Both forces can make the asking of questions by those who need help a stigmatized behavior.

In this paper, we focus on two questions. First, do signaling and shame concerns meaningfully inhibit learning, and how much? Second, what are the roles and magnitudes of these forces? Intuitively, signaling seems most important when there is (potentially) a lot to reveal—i.e., when there is considerable uncertainty on the part of an advisor about a seeker's ability. Thus, all else equal, it has a greater potential to shut down communication between people less familiar with each other, and about topics where cognitive ability matters more. On the other hand, signaling concerns are self-limiting, and naturally diminish as people learn about each other over time. Shame, in contrast, is most constraining when the asker's deficiencies are evident. It thus seems likelier to operate when the parties in the interaction know each other well, and can be a more durable obstacle to asking questions. Understanding the roles of these inhibitions, and their relative magnitudes, is essential for understanding how communication shuts down, in which relationships, and which policies are likely to reduce the frictions. As we will see, signaling concerns can be remedied with temporary interventions which move society to an equilibrium where seeking is done freely and doesn't signal much. Removing shame obstacles, on the other hand, requires more fundamental changes to preferences or social norms. Thus, which obstacles are operating, and in which social interactions, is important for policy. Because the manifestation of signaling and shame concerns depends both on the informational environment and people's preferences, identifying them presents an interesting problem.

[4]See Bénabou and Tirole (2006) for a seminal model of social signaling in a variety of contexts, with an emphasis on self-signaling. Recently, the effects of signaling in an educational context have been studied by Bursztyn et al. (2016), while Karing (2018) and Butera et al. (2019) have analyzed the details of signaling forces in the context of policy interventions.

[5]Goffman's (1963) discussion of the two forces distinguishes the situation of the "discreditable," who is inhibited by the possibility of revealing a bad attribute, and the "discredited," who is inhibited by the knowledge that is already possessed by those he may interact with; see also the analysis of self-presentation and "face" in Goffman (1959, 1967). In our choice of the term *shame*, we are guided by psychological studies of the self-conscious emotions, which define shame as involving "negative feelings about the stable, global self, ('I'm a dumb person')," (Lewis, 1971; Tracy and Robins, 2007), concern about being judged as deficient, and social withdrawal and concealment of the defective self (Lindsay-Hartz, 1984; Wong and Tsai, 2007). We discuss these connections further in Section 2.2.

In Section 2 we present a simple model of a decision about whether to seek advice, which is designed to distinguish the two components of stigma introduced above and study their distinct but overlapping effects. In the model, a Seeker (he) is deciding whether to ask a question of an Advisor (she), who may be uncertain about the Seeker's ability. The Seeker's decision is affected by (i) the *instrumental payoff* of getting advice, which may depend on his own ability; (ii) a *reputational payoff* based on how seeking *changes* the Advisor's belief about his ability; and, finally, (iii) an *interaction payoff* arising from the social interaction in which the advice is sought, which depends on the *level* of the Advisor's belief about the Seeker's ability. The reputational payoff creates signaling or reputation-management concerns, while the interaction payoff, as we will see, brings in the possibility of what we have called a shame effect. We use the model to derive predictions about the behavioral consequences of the signaling and shame effects. The model fleshes out our statements above on the differences between the two effects, and in their dependence on priors and information. The signaling effect will not be present if information about ability is symmetric—for instance, if a Seeker's low ability is evident regardless of his seeking decision. On the other hand, the shame effect—the disutility of an interaction because of the Advisor's adverse beliefs about the Seeker—is not about changes in the Advisor's beliefs but about their level, and is therefore strengthened when low ability is obvious. The two effects thus imply different responses to changes in the environment, which allows us to measure them. Importantly, if the model is misspecified and excludes the interaction payoff even though it is present, then estimates of the signaling effect can be seriously biased. The model formalizes these observations and sets the stage for our empirical analysis of signaling and shame effects in an experiment.

We conduct an experiment that allows us to assess the distinct contributions of signaling and shame in inhibiting information-seeking. We now describe the experimental setting, which is presented in full in Section 3. The experiment takes place in Karnataka, India, across 70 villages with 1247 total subjects. The main decision-maker in the experiment is a Seeker. On the first day of the experiment, the Seeker is told that on the third day, he or she will have the opportunity to win a prize by guessing which of two boxes contains it.[6] The Seeker has a choice about how to get clues about which box is correct. He or she is entitled to a certain number (k) of clues (independent hints that are correct with some probability). There is also an alternative option: the Seeker can opt to access a larger number of clues (some number $k' > k$, which the Seeker is told). This option, however, requires a physical visit to a named individual—the Advisor paired to that particular Seeker—to obtain a voucher for the clues. The Seeker has three days to do this, before we draw the clues and elicit the guess on the third day. By situating the experiment in the field, we make the seeking

[6] The prize, which is a lottery over a mobile phone and various cash amounts, has an expected value of over a day's wage.

interaction similar to one that might occur in subjects' daily lives, while the duration and nature of the task gives us experimental control of the setting.

Our experiment varies both whether there is scope for signaling in the task (whether ability matters), and what the Advisor knows about the Seeker. There are two treatment arms: $\{Random, Skill\}$ and $\{Private, Revealed\}$. Treatment is randomly assigned at the village level. The treatment arms activate signaling and shame concerns in different ways (and sometimes not at all), which is important for our identification of effects. The first arm varies whether k, which determines the Seeker's valuation of advice,[7] is *Random* or based on *Skill*. In the *Random* treatment, k is independent of any attribute of the Seeker. In the *Skill* treatment, it is proportional to the Seeker's performance on a cognitive ability test (a simple version of Raven's Progressive Matrices). The second arm varies whether the Seeker's ability score and identity are kept *Private* or *Revealed* by us to the Advisor, irrespective of the Seeker's actions. Importantly, in both arms, the Seeker and the Advisor have both been made familiar with the test, all the rules of the game, and how k is determined in that village—though the actual realized value of k is known only to the Seeker.

The outcome we focus on in our analysis is the Seeker's decision of whether to forgo his or her own clues in order to visit the paired Advisor over the course of the three days. In our main analysis, we hold fixed a low-ability Seeker, his or her need for information (k, own clue count), and the quality of advice to be received (k', Advisor's clue count). Thus, we compare behavior across individuals with the same attributes.[8] By comparing seeking rates across treatments, we measure the effects of stigma—both signaling and shame inhibitions—and analyze how these two forces depend on the circumstances. We now give some more detail on the measurement of the key effects. These are presented fully starting in Section 4.

Recall that we define the stigma effect to he be the reduction in the seeking of information that is caused by need being ability-dependent. We measure this as the reduction in seeking rates when we go from the (Random, Private) to the (Skill, Private) treatment. In the data, this comparison is associated with a large decline in the probability of seeking (55%). This is the effect we are most interested in practically, as it reflects the inhibition that arises when ability is implicated in a task. What accounts for the reduction? When a Seeker chooses to seek in (Skill, Private), the Advisor's belief that the Seeker has high ability moves from the prior to a lower posterior. (In our data, the imputed reduction in the belief turns out to be approximately 9pp, on average.) As we have discussed above, the reduction in seeking

[7]The other number, k', is independent of any attribute of the Seeker or the Advisor.

[8]We focus on low-ability Seekers because, in the Skill treatments, they are the ones who have few clues and therefore need to seek. If we had conditioned only on having a low clue count, a complication would be that the composition of this group varies across treatments: in the Skill treatment it consists only of those with a low ability score, but in the Random treatment there is a representative mix of ability scores.

could come from the signaling effect (to avoid this drop in beliefs), the shame effect (to avoid interaction given this drop), or a combination.

Thus, the next step is to examine the magnitude of the shame inhibition. Our basic measurement of shame is the comparison of (Random, Private) to (Random, Revealed)—a comparison in which there is no room for signaling to play a role. We see a 65% decline in the probability of seeking advice for a low-ability Seeker, holding everything else fixed. Note that in this case, the Advisor completely learns that a Seeker has a low ability score (when this is the case), so her beliefs that the Seeker is of high ability move from her prior to a zero probability. This is much greater than the movement that is typically caused by signaling that one has a lower ability than the Advisor thought—which ordinarily does *not* result in certainty ex post. Thus, while this difference allows us to identify the presence of a shame effect, we must be careful in scaling this number to assess how much of the stigma effect discussed in the last paragraph is due to shame.

By placing more structure on the problem—making parametric assumptions on payoffs and distributions of shocks—we can address these issues. There seems to be fairly little structural estimation of signaling in social behavior overall; Fang (2006) takes a structural approach in the classical context of job market signaling. Since in an experimental setup we are able to directly compare seeking rates by type across of all our treatment cells, we can estimate the model's key parameters and assess the relative contributions of signaling and shame. We find that both signaling and shame effects are sizable, and the signaling disutility is, on average, eight times the size of the shame disutility for the same change in beliefs under our structural assumptions.

In Section 5, we turn to how the effects on social structure, and find that the relative roles of signaling and shame become more nuanced. With detailed data on all respondents' network links, as well as their subcastes, we study whether stigma, signaling, and shame vary with the social distance between the Seeker and the Advisor. Most of the deterrent to seeking among those who are linked by friendship or of the same caste comes from shame. For instance, there is an 81% decline (16.3pp on a base of 20.1%) in the probability of seeking due to the shame effect among co-caste members. Building on our structural exercise, we find that signaling effects among the socially proximate are small and not statistically different from zero. This is consistent with the fact, which we document, that the socially proximate have stronger priors about each other's abilities. These stronger priors leave less room for a signaling effect. Meanwhile, among the socially distant (strangers or individuals of different castes), signaling appears to be the dominant force. The decomposition of these effects enabled by the structural exercise is consistent with our reduced-form results.

Taken together, our findings suggest that the endogenous formation of a communication network in a given social environment will depend in subtle ways on the kind of information being shared and the associated signaling and shame concerns. These can create sizable distortions in who talks to whom relative to what would be efficient in terms of the instrumental value of information. It can also account for considerable "missing conversations" between social groups.

Section 6 is a discussion that points out various implications of the frictions arising from signaling and shame. One of them is that reputational distortions are more transient, while interaction payoff frictions are potentially more persistent. This is because the updating of beliefs resolves uncertainty and limits future signaling concerns, whereas shame need not be self-limiting in this way.

Moreover, the forces we have identified can have considerable implications for the outcomes of information exchange and aggregation processes, and policies related to them. Consider, for example, how policymakers may attempt to spread information. One practical strategy is to offer a private learning environment that avoids the stigma concerns we have discussed. This has been implemented in experiment reported in Cole and Fernando (2014) via a "hotline"—a remote consulting service for production advice. If farmers can call in to request information without their actions being observed by peers, there is less room for strategic signaling concerns to operate. Thus, in addition to the other benefits of providing expert advice, such interventions may make endogenous seeking behavior more favorable for learning.[9]

Another aspect of information dissemination strategy is whether to broadcast information to everyone in a community (e.g., via a loudspeaker or the radio) or to provide information to a small set of seeds (e.g., via an extension program). Beyond their direct impact on who is informed, these policies can affect stigma concerns and thus seeking behavior: Broadcasts typically make it common knowledge that everyone was informed. Thus, in the broadcast case, needing clarification can be a signal of being unwilling or unable to learn on one's own. Signaling inhibitions can arise around this. Therefore, if social learning is an important part of comprehension, broadcasting can actually lead to less conversation and less knowledge. Banerjee et al. (2018) apply this perspective to study an information intervention during the 2016 Indian demonetization, where how information was delivered to villages was varied. We discuss this in Section 6.

Related Literature. One literature that this paper relates to, as we have already mentioned, focuses on image concerns in information/skill acquisition (Fryer and Austen-Smith, 2005; Bursztyn and Jensen, 2015; Bursztyn et al., 2016). Spence's (1973) seminal study of signaling focused on education. The perspective there was that seeking education is a signal

[9]On the other hand, in principle a shame inhibition can still operate even in an interaction with a remote consulting service, though may be lower.

(to potential employers, for instance) of a high ability, since education has lower costs for more able types. More recently, researchers have studied signaling concerns that can reduce demand for education. Fryer and Austen-Smith (2005) examine a dual-audience signaling model, where the signal that the labor market rewards may be one that an agent's local community penalizes: having a relatively high type from the perspective of the labor market can correspond to having a relatively low "social type," which is what peers care about. Of course, peers may also reward signaling high ability.

Bursztyn et al. (2016) (henceforth BEJ) perform carefully designed experiments in three schools designed to empirically examine whether peer perceptions inhibit the seeking of education. They focus on signaling two attributes, and find that the type of community matters a great deal for the nature of a signaling effect: In one type of school, students are concerned with avoiding signaling low ability, while in another, they are reluctant to reveal a disposition to exert high academic effort. We make a different distinction, focusing on just one type of cognitive ability, but separating the reputation-based and interaction-based distortions to asking. These are quite different in terms of their economics and policy implications (as we discuss throughout) but may be confounded. Another contrast is in terms of practical focus: in our setting the signaling is bilateral—the signaling occurs in a pairwise interaction (as opposed to public in BEJ) and thus quite sensitive to the nature of the relationship between the people involved—a dependence we explore in detail. Finally, in terms of empirical design, we study the signaling and shame effects across many communities (70 villages), with treatment being assigned at the village level.[10] This allows us to deal with the fact that within a cluster (a village or a school), there may be common shocks or other correlations. We conduct inference with suitable clustering to identify how properties of the task (e.g., whether ability is relevant) and the pairwise relationship in question (e.g., whether two people are friends) affect signaling and shame concerns.

More broadly, within economics, we relate to recent empirical literatures which focus on signaling, image, and status. The studies there focus on a number of domains: conspicuous consumption (Charles et al., 2009; Heffetz, 2011), commercial transactions (Goldfarb et al., 2015), productive effort (Besley and Ghatak, 2008; Ellingsen and Johannesson, 2008; Ashraf et al., 2014), professional education (Bursztyn et al., 2017), and political behavior (DellaVigna et al., 2016), among others. The survey of Bursztyn and Jensen (2017) offers a comprehensive overview. In many of these cases, the nature of the mechanism is open-ended. For example, DellaVigna et al. (2016) find that people vote, in part, to avoid having to report to their friends, when asked, that they haven't voted. This could occur due to reputational

[10]I.e., we tell the same rules to all individuals in the village (e.g., that skill is relevant to the task, if it is a Skill treatment village). At the same time, as we will explain later, every Seeker-Advisor pair has a guessing game with distinct labels, ensuring that every pair's interaction is separate.

concerns, or simply pride (closer to our interaction payoff). Since for our concerns the mechanism matters, we develop tools that could be applied to dig into the nature of the effect in other social signaling models. In Section 2.2 we also discuss how the payoffs we model relate to literatures in sociology and psychology on stigma, signaling, and shame.

Our paper relates to a separate literature on information flow and learning in networks (Conley and Udry (2010); Banerjee et al. (2013); Kremer and Miguel (2007); Foster and Rosenzweig (1995); Beaman et al. (2018)). Though exogenous transmission of information is a common baseline assumption, Niehaus (2011) argues that social learning endogenously filters information due to individuals' incentives to pass information or not. We focus on the other side of this interaction: Agents in our setting take an active decision to obtain information from a peer. Since this decision is shaped by agents' attributes and signaling/shame concerns, our analysis suggests new distortions and externalities that are relevant to deciding whether organic communication will spread information successfully. This ties into recent work on endogenous network formation that has studied networks formed in equilibrium (Galeotti et al., 2013; Calvo-Armengol et al., 2015) and conditions for learning to occur in them (Acemoglu et al., 2014).

2. Model

We formulate a simple model of decisions to seek information in the presence of reputational and interaction payoffs.

2.1. **Environment.** The model focuses on a decision taken by an agent called a Seeker. The Seeker either does or does not have a choice to make: for instance, an opportunity to invest in a new technology. Let the indicator of the choice being is available be C, and write $q := \mathbf{P}(C = 1)$ for the ex ante probability of this event. (The uncertainty here captures the fact that people do not always need information, and so there may not be much to infer from their not seeking it.) Conditional on having a demand for information, the Seeker (he) has a decision, $d \in \{0, 1\}$, to make: whether to seek ($d = 1$) or not to seek ($d = 0$) advice from an Advisor (she), who observes the decision and updates beliefs about the Seeker based on this observation.[11]

This is the key decision we focus on. The Seeker has an expected *instrumental payoff* of seeking: V is the increase in expected payoff (from later making the choice, e.g., an investment) due to seeking advice, net of all immediate material costs incurred to actually seek it (e.g., the opportunity cost of time). The payoff V is privately known to the Seeker.[12] From the perspective of others, V is random. Its distribution depends on the Seeker's *ability*

[11] For a formal description of the game and equilibrium, see the beginning of Appendix A.

[12] Of course, the realization of V need not resolve all uncertainty regarding the value of information or the cost of getting it: it is simply an expectation based on the Seeker's private information of his need for advice and other relevant facts.

(or *skill*) type, $a \in \{\mathrm{H}, \mathrm{L}\}$ (high or low).[13] Let F_a be the c.d.f. of V for a Seeker of ability a, and let G_a be the complementary c.d.f.

The Advisor has a prior about the ability, a, of the Seeker: $\pi := \mathbf{P}(a = \mathrm{H})$.[14] The Advisor observes the decision, d, of whether to seek. If there is asymmetric information about the Seeker's ability (i.e., if the Advisor does not know a), then the Advisor will update her beliefs about it based on the Seeker's decision. The Seeker cares about this updating, potentially in two ways, which we now discuss.

The Seeker's utility, as a function of whether he chooses to seek is

$$U(d) = \underbrace{V\mathbf{1}_{d=1}}_{\text{instrumental}} + \underbrace{\varphi\left(\mathbf{P}(a = \mathrm{H} \mid d)\right)}_{\text{reputational}} + \underbrace{d \cdot \psi\left(\mathbf{P}(a = \mathrm{H} \mid d)\right)}_{\text{interaction}}, \tag{2.1}$$

and involves three kinds of payoffs. From left to right we label these terms the (i) *instrumental* payoff; (ii) the *reputational* payoff; and (iii) the *interaction* payoff. We have discussed the instrumental payoff. The reputational payoff corresponds to the the Seeker caring about $\mathbf{P}(a = \mathrm{H} \mid d)$, the Advisor's assessment[15] of the Seeker's ability, which is made in view of the Seeking decision d. Finally, there is the interaction component, which reflects that the Seeker may derive more or less utility from the seeking interaction with the Advisor that occurs when $d = 1$ depending on what the Advisor thinks of him during that interaction. We discuss interpretations and foundations of these payoffs below in Section 2.2.

We assume that both $\varphi(\cdot)$ and $\psi(\cdot)$ are continuous, increasing, and bounded functions $[0, 1] \to \mathbb{R}$. We also make the assumption that $\psi(\pi) = 0$—the shame term makes no contribution when the assessment of the Seeker is the same as the prior. This is a normalization in the sense that if we did not have it, we could achieve it by making a suitable adjustment to the distribution of V.

Rewrite (2.1) as

$$\Delta U = V + \Delta\varphi + \psi\left(\mathbf{P}(a = \mathrm{H} \mid d = 1)\right), \tag{2.2}$$

where $\Delta U = U(1) - U(0)$ and $\Delta\varphi = \varphi(\mathbf{P}(a = \mathrm{H} \mid 1)) - \varphi(\mathbf{P}(a = \mathrm{H} \mid 0))$.

This clarifies the difference between the reputational and interaction parts of the payoff: reputational payoffs are about $\Delta\varphi$, i.e., about changing the beliefs of the Advisor, $\mathbf{P}(a = \mathrm{H} \mid d)$, by changing his information from $d = 0$ to $d = 1$. The interaction payoff only depends on the final level of the Seeker's belief, irrespective of whether or how that level is changed by the Seeker's decision.

---

[13] The binary-type simplification makes for a simpler and more intuitive exposition. The theory extends readily to more types.

[14] For simplicity we assume that C and a are independent, though it is immediate to relax this assumption; indeed, our formulas below remain correct if we take π to be the probability of high ability conditional on having the chance to seek information.

[15] Here we follow a standard and tractable parameterization—see, e.g., Bernheim (1994), Bénabou and Tirole (2006), and Ali and Lin (2013).

The change in seeking caused by the $\Delta\varphi$ term in (2.2) is called the *signaling* effect on seeking,[16] while the change caused by the final ψ term is the *shame* effect. The assumption we made earlier, that $\psi(\pi) = 0$, amounts to saying that when the Seeker does something (e.g., seeking) that lowers beliefs below the prior, the contribution to his payoff is negative.

The total of the two terms (on payoffs or seeking rates, depending on context) will be called the *stigma* effect. We discuss these terms and their interpretations further in Section 2.2 below.

A tuple of primitives $\left(\pi, q, (F_a)_{a\in\{\mathrm{H,L}\}}, \varphi, \psi\right)$ is called an *environment*.

2.2. **Interpretation.** An especially clear distinction between the signaling and shame parts of stigma is drawn in Goffman (1963):

> "The term stigma and its synonyms conceal a double perspective: does the stigmatized individual assume his differentness is known about already [or not]? In the first case one deals with the plight of the *discredited*, in the second with that of the *discreditable*. This is an important difference, even though a particular stigmatized individual is likely to have experience with both situations." (Goffman, 1963)

Goffman and the literature that has followed him (see, e.g., Scambler (1998); Page (2015)) have emphasized that both kinds of concerns prompt the deliberate management of interactions, but in different ways and with different consequences. Modeling the preferences underlying such behavior leads to the payoff terms we have described above—in particular, an interaction term in addition to a signaling term. The specification (2.1) highlights that the payoffs underlying these concerns may differ in their magnitudes, and these differences will be important in understanding both behavior and reaction to policy. We now discuss these terms in turn.

2.2.1. *Reputational payoff.* The reputational payoff reflects the Seeker's concern with managing the Advisor's beliefs. One simple foundation for this, and one that we believe is relevant in many applications, is that the Seeker will have access to materially valuable future opportunities if the Advisor thinks well of him. (Appendix D fleshes this out and relates it to our functional forms.) Alternatively, the reputational term may come simply from a hedonic valuation of others' esteem, aggregated over the future occasions when their assessment of the Seeker will be relevant. In the introduction, we have mentioned some entry points into the extensive literature on reputational concerns and signaling effects, so we will not reiterate that here.[17]

[16]The implicit counterfactual is that this term is set to zero: the signaling effect is the difference between the level of seeking in the true model and this counterfactual with reputational concerns turned off.

[17]As is often the case, reputational concerns can be reinterpreted through the lens of self-signaling (insofar as the act of seeking serves as a memorable reminder of one's ability). Within our experiment, because

2.2.2. *Interaction payoff.* The interaction payoff, in contrast, pertains only to the *particular act of seeking that the decision d controls*; it describes the utility of this particular information-seeking engagement as a function of how the Advisor will be assessing the Seeker. The associated payoffs have several psychological interpretations. Earlier, we labeled the inhibition associated with a negative interaction payoff (for low-ability types) as shame. This is likely to be the right interpretation when the question at hand or the knowledge required is basic in the relevant context: For example, few people in a literate community take special pride in the ability to read directions, but illiterate people are often ashamed that they cannot. However, in other cases the more natural description is that the interaction term captures pride in being seen to have a trait that is especially good.

The interaction payoff may be caused by external events or may come from within. In the external case, the payoff may arise from behavior of the Advisor (to take an extreme example, explicit shaming or praising the Seeker as a result of the seeking interaction). In the internal case, it comes from how the Seeker feels about the situation of asking, even if the interaction is not visibly affected by the Advisor's beliefs. The interaction payoff in practice is likely to be a combination of these.[18] Note, however, that this part of the payoff must be related *somehow* to this particular act of seeking. If the Seeker feels bad about his low ability irrespective whether he seeks or not, that is not reflected in his interaction payoff. It would, instead, be a separate, d-independent term in (2.1), which we omit from our analysis as it does not affect the decision to seek.[19]

The shame interpretation ties into a literature in psychology. This literature views shame as a negative emotion prompted by one's own perceived inadequacies, especially persistent attributes (such as intelligence) subject to external evaluation and judgment. The characteristic manifestsation of shame is social withdrawal (Tracy and Robins, 2007; Wong and Tsai, 2007). Thus even aside from concerns about future reputation, has the potential to deter the seeking interaction.

2.3. **Basic analysis.** We will study the Bayes-Nash equilibria of our model. This is a specification of behavior in which the Seeker is best-responding to a given belief-updating function of the Advisor, and this belief updating is consistent (in the sense of Bayesian inference) with the actual behavior of the Seeker.[20]

ability score is always clearly revealed to the Seeker, our interpretation is that external reputation concerns are likely the dominant ones there.

[18]Adam Smith's discussion of image concerns (*Theory of Moral Sentiments*, I.iii.2) is at least partly about interaction payoffs, as opposed to instrumental signaling. He argues that the *main* motive of "bettering our condition" is "to be observed. . . to be taken notice of with sympathy. . . and approbation." He stresses the value of satisfying one's "vanity" by enjoying the positive assessment of others, beyond securing their good opinion.

[19]Our experiments always make comparisons within ability type, differencing this variation out.

[20]One interpretation of this two-player game is as an Advisor facing a population of Seekers, with a distribution of valuations and attributes. The formal definition of the solution concept is in Appendix A.1.

We make an assumption about how the privately known value of seeking, V, compares across the two types. Recall that the c.d.f. of this random variable for a Seeker of ability a is F_a.

Assumption **1.** F_{L} first-order stochastically dominates F_{H}, and it does so strictly, in the sense that $F_{\mathrm{L}}(v) < F_{\mathrm{H}}(v)$ for every v.

This says that the value to low-ability types from seeking exceeds (in distribution) the value to high-ability types; they need help less. This assumption corresponds to our focus on "basic" questions, and holds by construction in the experiments we conduct.[21]

In terms of practical examples: a less able farmer may have more difficulty using a new fertilizer without help. A less intelligent or experienced student may need help making sense of instructions. In addition, high-ability types may face higher opportunity costs of time, which is part of V (the net material value of seeking).

We also make a technical assumption to simplify the analysis.

Assumption **2.** For both ability types $a \in \{\mathrm{H}, \mathrm{L}\}$, the random variable V has an atomless distribution (i.e., its c.d.f. F_a is continuous) whose support contains the positive reals (i.e., $F_a(v) \in (0, 1)$ for all $v > 0$).

The first basic result is:

Proposition **1** (Equilibrium Existence and Characterization). *Under Assumption 2, an equilibrium exists and every equilibrium is in cutoff strategies, where an agent seeks if and only if $V \geq v$ for some v.*[22] *An equilibrium is characterized by a cutoff v (used by all Seekers, independent of the value of a) which satisfies*

$$v = \varphi\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 0)\right) - \varphi\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right) - \psi\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right), \tag{2.3}$$

where $\mathbf{P}_v(a = \mathrm{H} \mid d = \cdot)$ are uniquely determined by the posterior odds ratios

$$\frac{\mathbf{P}_v(a = H \mid d = 0)}{1 - \mathbf{P}_v(a = H \mid d = 0)} = \frac{\pi}{1 - \pi} \frac{1 - qG_H(v)}{1 - qG_L(v)} \qquad \frac{\mathbf{P}_v(a = \mathrm{H} \mid d = 1)}{1 - \mathbf{P}_v(a = \mathrm{H} \mid d = 1)} = \frac{\pi}{1 - \pi} \frac{G_H(v)}{G_L(v)}. \tag{2.4}$$

Under Assumption 1, the cutoff v is positive, meaning that $V > 0$ is necessary (but not sufficient) for an agent to seek.

We start with some reminders on the notation in this result and a few remarks on its interpretation. First, v is the cutoff in a particular equilibrium, whereas V is a random

[21] In general, asking questions (or some other social behavior) could signal high ability rather than low ability, or it could have more complex informational content. For instance, some pupils are eager to ask certain kinds of questions in class to signal how much they understand. Though the assumptions we work with in the sequel are specific to our setting and experiments, the basic point that information-acquisition is distorted by signaling concerns would hold—with suitable modifications—in other environments.

[22] We identify two strategy profiles if outcomes are different across them with probability zero.

variable drawn for a given Seeker, which may or may not exceed that cutoff. Second, recall that the model considers a Seeker about whom the prior, $\mathbb{P}(a = H)$, is fixed at π. Once we apply this model to the real world, different Seekers will be associated with different priors π, and the cutoff v will depend on one's identity and covariates—we omit notation that would make this dependence explicit. Third, recall that G_a is the complementary cdf of F_a.

The argument is straightforward. (Details of all proofs are in Appendix A.) First, since the payoffs that depend on d (reputational and interaction) enter additively, it is clear that the best-response d (taking Advisor's beliefs as given) is increasing in the instrumental payoff V. Under the assumption of no atoms, in any best response the Seeker uses a cutoff strategy, choosing $d = 1$ if and only if his value is high enough: $V \geq v$.[23] Actual ability type does not enter an agent's utility function at the time of the seeking decision, and is not observed by the Advisor; these facts imply that a does not matter for the cutoff v. Given the cutoff, one can compute the probability that the agent's type is high conditional on not seeking, and conditional on seeking. The quantity on the right-hand side of (2.3) is the decrease induced by seeking in the reputational and interaction payoffs. The cutoff type v is indifferent between this and the expected value v he will get from the information. It only remains to verify that a positive cutoff v solving (2.3) exists, which follows by a standard argument.[24]

We have noted that ability does not play a role in cutoffs. Indeed, it plays a role only in determining the ex ante distribution of the gains to seeking V, whose realization the Seeker knows when deciding. This will, of course, affect signaling implications.

Proposition 2 (Inferences in Equilibrium). *Under Assumptions and 1 and 2, in any equilibrium of the signaling game, the mass of high-ability agents choosing to seek ($d = 1$) is strictly smaller than the mass of low-ability agents choosing to seek ($d = 0$). Therefore, not seeking signals high ability: if v is the equilibrium cutoff, then*

$$\mathbf{P}_v(a = \mathrm{H} \mid d = 0) > \mathbf{P}_v(a = \mathrm{H} \mid d = 1).$$

The proof is simply that with *any* cutoff, as long as it is the same for both types (as must be the case in equilibrium), Assumption 1 guarantees that the mass of high-ability agents choosing to seek is strictly smaller than the corresponding mass of low-ability agents. The conclusion then follows by Bayes' rule.

To make predictions corresponding to some of our treatment arms, as well as to study the effect of pre-existing familiarity between the Seeker and the Advisor, we look at some extreme cases of our problem. In particular, we wish to explore how equilibria behave as we

[23]Here we are using a tie-breaking rule, which does not matter as the proposition ignores zero-probability events.

[24]Note that at $v = 0$, the left-hand side of (2.3) is zero while right-hand side is positive; as $v \to +\infty$, the left-hand side tends to $+\infty$ while the right-hand side remains bounded. Since the left-hand side of (2.3) is continuous, we can apply the intermediate value theorem.

make the Advisor's priors precise, or make the marginal value of seeking unrelated to ability type. In both cases, the signaling effect disappears and both ability types have the same behavior.

PROPOSITION **3** (Known Ability or Ability-Irrelevance). *Suppose that Assumptions 1 and 2 hold. Take a sequence of environments satisfying either*

(1) $\pi \to 0$ *or* $\pi \to 1$, *fixing all other parameters; or*
(2) *the distribution* F_H *converges to* F_L *in the total variation norm, fixing all other parameters.*

For any sequence of equilibria corresponding to those environments, the cutoff v *converges to* 0 *and the seeking decision becomes uninformative:*

$$\mathbf{P}_v(a = \mathrm{H} \mid d = 0) \to \mathbf{P}_v(a = \mathrm{H} \mid d = 1).$$

2.4. **A parameterized family.** We introduce a parametric family of environments to illustrate the mechanics of the model; this setup will also be useful in our structural exercise. Fix "basic" functions $\widehat{\varphi}$ and $\widehat{\psi}$, and for some real numbers $\lambda, \gamma > 0$, let $\varphi = \lambda\widehat{\varphi}$ and $\psi = \gamma\widehat{\psi}$.[25] Thus (2.1) becomes:

$$U(d) \quad = \quad V\mathbf{1}_{d=1} \quad + \quad \lambda\widehat{\varphi}\left(\mathbf{P}(a = \mathrm{H} \mid d)\right) \quad + \quad \gamma d \cdot \widehat{\psi}\left(\mathbf{P}(a = \mathrm{H} \mid d)\right).$$

Then (2.3) becomes

$$v = \lambda\left[\widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 0)\right) - \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right)\right] - \gamma\widehat{\psi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right).$$

To see how the equilibria are affected by parameters, it is convenient to divide through by λ and obtain

$$\frac{v}{\lambda} = \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 0)\right) - \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right) - \frac{\gamma}{\lambda}\widehat{\psi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right).$$

Then, fixing the ratio γ/λ, we can visualize the dependence of equilibria on the magnitude of stigma by looking at intersections between the scaled return to seeking information, v/λ, and the right-hand side, which we define as $R(v)$.

Figure 1 presents plots where we vary λ and γ and study the resulting equilibria. In Panel A with $\lambda \in \{3, 10\}$ and $\gamma = 1$, we see that there is a unique equilibrium in each case. Further, when image matters considerably ($\lambda = 10$), the equilibrium is one in which the threshold to seek is very high and therefore the incurred reputation loss is high as well. Meanwhile the equilibrium when $\lambda = 3$ has the threshold for seeking quite low, and accordingly the reputational cost is low. Panel B repeats the exercise but with $\lambda = 5.5$ to demonstrate that

[25]In our illustrations, we take the foundation for $\widehat{\varphi}$ from Appendix D, with the c.d.f. $H(x)$ given by $x/(1+x)$, and set $\widehat{\psi} = \widehat{\varphi}$.

multiple equilibria are possible. Here A is an equilibrium where the threshold to seeking is low and there is essentially no reputation cost. In contrast, in equilibrium B, the threshold to seek is considerably higher and therefore an individual incurs a larger reputational cost. (An equilibrium is stable when $R(v)$ has a smaller slope, in absolute value, than $1/\lambda$ at the point of intersection.) Finally in Panel C we look at the same case as in Panel B, except that we remove the shame effect. This flattens the reputational cost functions and will tend to move the equilibria to the left.

Note that if the model is misspecified, and we estimate a signaling-only model when shame is present, then this can lead to seriously biased estimates of λ and wrong predictions (e.g., about the impact of reducing information asymmetries).

2.5. **Uniqueness.** Finally, we briefly turn to the issue of uniqueness. In general, our model may feature multiple stable[26] equilibria. This reflects the realistic feature that, despite fundamental parameters being the same in two communities, the same actions may have different meanings due to culture or custom. This is manifested in equation (2.3) potentially having multiple solutions v. (See Figure 1 for an illustration.)

However, we can give conditions for this not to happen. For simplicity, focus on a signaling-only model. Let $\varphi = \lambda\widehat{\varphi}$, where $\lambda > 0$ is a parameter and $\widehat{\varphi}$ is a continuous, increasing, and bounded function $[0,1] \to \mathbb{R}$. Let μ be the supremum of the derivative (in v) of

$$\widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 0)\right) - \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right)$$

over all $v \geq 0$.

Proposition 4 (Equilibrium Uniqueness). *Fix all aspects of the environment except λ. Fix a constant $\varepsilon > 0$. If Assumptions 1 and 2 hold, then: (i) If $\lambda < \mu^{-1}$ there is a unique equilibrium cutoff. (ii) There is a constant $\overline{\lambda}$ such that, if $\lambda > \overline{\lambda}$, the equilibrium probability of seeking is at most ε (so there is an essentially unique equilibrium with no seeking).*

Thus, when signaling concerns are low or high, there is a unique or essentially unique equilibrium. However, with signaling concerns of intermediate magnitude, it is possible that there are multiple equilibria. Temporary interventions that shift the equilibrium can therefore have considerable impact on seeking rates and welfare, as discussed in Section 6.3.2. The restriction to a reputational payoff only (i.e. φ, not ψ) is not essential to these statements. On the other hand, without signaling concerns, multiple equilibria cannot arise.

2.6. **Predictions and implications.** We now discuss some practical implications of the model that follow directly from model and results above.[27]

[26]A *stable* solution is one for which best-response dynamics—between the strategy (i.e., cutoff) of the Seeker and the belief-updating of the Advisor—converge back to it after a slight perturbation, and corresponds to the right-hand side of (2.3) having a slope between -1 and 1 at the equilibrium.

[27]Connections to our specific experimental treatments in Section 3.2.1 below.

(1) When the Seeker's need for information is negatively associated with his ability, (i) high-ability Seekers seek less than low-ability Seekers on average and (ii) Advisors update their beliefs, upon observing seeking, in the direction of low ability.

(2) Suppose the Seeker's low ability is known to the Advisor. Then low-ability types seek less than high-ability types (even though beliefs are not updated conditional on seeking) if and only if the final term of (2.2) is nonzero. In other words, a test of whether shame is present is whether seeking rates are ranked in this way.

(3) Suppose the net valuation of information (V) is independent of ability. Then (i) the seeking rate does not depend on ability given the Advisor's prior π about the Seeker and (ii) there is no belief updating about ability based on seeking.

Now that we have formalized the model, we can also flesh out some remarks we made earlier about contrasts between shame and signaling. The two are very different in how they operate. Signaling, because it is about the *change* in beliefs that an interaction induces, is inherently bounded in how many times it can deter Seeking: once ability is revealed, it plays no role via the signaling term. Similarly, if the task in question is one where ability is easy to diagnose, signaling effects will be limited. But shame can be a severe obstacle in both situations.

Signaling and shame effects also depend on social context in different ways. Signaling is most powerful when there is substantial asymmetric information. In particular, between strangers, the act of asking could update beliefs from the population prior to a fairly precise posterior, and so signaling is likely to be a stronger deterrent in such interactions compared to ones in which the Seeker's type is known. On the other hand, shame can be strong even without updating of beliefs, and in particular can be very strong between close associates, who are likely to know that ability is low when this is the case.

Sections 6 and 7 extend this discussion and explore further implications.

3. Setting and experimental design

3.1. **Setting and background.** We conducted surveys and experiments with 1247 subject pairs in 70 villages in Karnataka, India. The majority of villagers have occupations in agriculture, sericulture, and dairy production. This is a setting where villagers rely on word-of-mouth learning to obtain information useful for production, so understanding what obstacles impede social learning is important.

Prior to our experiment, we conducted surveys with 122 respondents in four villages not from our sample, but in the same region. Our goal was to learn about the frictions villagers face in social learning. Among other questions, we asked them about how they get information from their networks on a number of topics (financial products, farming inputs, and health practices). We found that passive learning was considerably more common than active learning: 95% of respondents reported hearing information passively from friends and 90% from broadcast media (newspapers, TV, and the Internet). In contrast, only 49% reported actively asking friends in their own village for information about the above-mentioned topics.

Our further questions were aimed at understanding what factors might deter active seeking. In particular, we probed whether the respondents felt constrained in seeking information. We did this by asking whether they felt there were limits on the number of times one could approach another villager for information. In 88% of surveys, they reported feeling a limit on the number of times they could approach a member of their community to ask for farming, health, or financial advice. In 64% of these cases, our respondents reported that they chose not to seek out information to avoid appearing weak or uninformed. This is consistent with some sort of stigma (signaling, shame, or a mix) being an important force in this context.

Finally, respondents also told us that the constraints keep them from the information that would be most valuable. About 70% of respondents told us they consider it important to talk to others before making consequential decisions about financial products, farming, and health practices. We then asked respondents to consider (i) the best person in the village whom they felt comfortable approaching for advice about the topic as well as (ii) the best person in the village to ask for advice if, hypothetically, they did not feel constrained. About 60% of the respondents said that they already knew everything relevant known by the approachable people. However, 70% of the respondents said that the unconstrained-best person to ask for advice *did* know more information. Overall, our survey evidence points to stigma as a potentially important force that constrains the seeking of information.

Turning now specifically to the 70 villages in which we conducted our experiment: We had previously collected data on networks among households to help us identify whether two subjects are socially connected, in what ways, and to what extent.[28] Our measures of the frequency or intensity of social interaction allow us to examine how frictions in social learning vary with social distance—a crucial variable in combination with demographic covariates.

3.2. **Design.** In every village, we selected approximately 18 people as Seekers and 18 as Advisors, drawing randomly from our village census, and formed Seeker-Advisor pairs. The roles of Seeker and Advisor were randomly assigned. The experiment is at the pair—that is, Seeker-Advisor—level (though, as we describe below, treatment is randomized at the village level). We stratified the sample into friend pairs, co-caste pairs, and pairs belonging to neither class.[29]

[28]This data was collected and analyzed in Banerjee et al. (2016). Specifically, we collected data on (i) social relationships (whose house the subject visits and who visits their house to socialize), (ii) financial relationships (with whom they engage in borrowing/lending of small amounts of money or material goods such as kerosene or rice), and (iii) advising relationships (whom they talk to for advice regarding a financial, health, or farming decision).

[29]We reweight observations appropriately when we conduct the pooled analysis.

To situate the act of getting information in a context similar to typical informational interactions in the subjects' lives, the experiment was designed to take place over three days in each village. At a high level, on Day 1 we collected baseline data and introduced the guessing game Seekers would play on Day 3. Between Day 1 and Day 3, each Seeker had the opportunity to increase the number of clues he would have access to in the guessing game by visiting his paired Advisor. On Day 3, the guessing game was played.

In more detail: On Day 1, we approach both Seekers and Advisors and collect demographic data in a baseline survey. We also administer—separately to all participants—a simple test of cognitive ability based on the Raven's Progressive Matrices.[30] The ability test occurs before subjects know which treatment they are in, at the same time as we obtain informed consent. The participants are informed of the possible treatments. Therefore, importantly, the subjects know that their score and identity may be revealed to another member of their community, irrespective of their seeking decisions. Since they are informed of this before randomization, we do not need to worry about endogenous variation in effort on the Raven's test depending on treatment. The test score maps to the ability variable in our theory. We inform the Seekers that on Day 3, each will play a game for a prize. The game is to guess which of the two boxes contains a cellphone. The prize is a cellphone worth Rs. 1350, or one of several cash prizes. The exact prize is determined by a dice roll, but everyone is guaranteed to win something for guessing correctly; the expected cash value of the random prize is Rs. 180, which is over a day's wage in the area.

In addition to being told about the game, on Day 1 the Seeker is also given a choice about clues for the guessing game. Each clue is an independent, identically distributed draw of a chit from a bag which has the label of the box with the prize with probability 0.55, and the wrong box otherwise. The Seeker has the right to get k clues without doing anything else. (We discuss the determination of k shortly.) He can, instead, decide to have a voucher for k' clues, typically more than k, delivered to his paired Advisor. The paired Advisor's clues are usable only by the paired Seeker.[31] If the Seeker chooses this option, he has to physically visit the Advisor between Day 1 and Day 3 to obtain the voucher. Our main outcome of interest is which of these two options the Seeker chooses. On Days 1-3, the Seekers are able to seek information from their paired Advisors, if they chose to do so. On Day 3, we revisit the village to solicit the Seekers' guesses and distribute prizes.

[30] For the instrument, see Appendix B.1.

[31] This is because every Seeker–Advisor pair has a uniquely labeled game. For example, one Seeker has to guess whether the "cat" or "dog" box has the prize and another has to decide whether the "tree" or "river" box has the prize, and so on. Therefore, each Advisor had unique clues that were only usable by her Seeker, and further, the Advisor had no intrinsic use for the clues herself.

Our experiment has a simple 2×2 design, randomized at the village level.[32] The treatment arms are as follows:

(1) What determines quality of information:
   - *Random:* The number of clues that the Seeker is endowed with, k, is drawn uniformly at random from $\{1, 2, 3, 4, 5\}$.
   - *Skill:* k is increasing in the Seeker's score on an ability test—the Raven's Matrices Test (see Figure B.2 in Appendix B for the exact schedule).

   In either case, if the seeker decides to forego his k clues, he can obtain, by visiting the Advisor, an entitlement to k' clues. The number k' is drawn uniformly at random from $\{4, 5, 6, 7, 8\}$. This number is known to the Seeker at the time of his decision. Typically $k' > k$.[33]

(2) Whether the Seeker's ability is known:
   - *Private:* we do not reveal a Seeker's test score to the paired Advisor.
   - *Revealed:* irrespective of the seeking decision, the Seeker's identity, test score, and percentile are immediately revealed to the paired Advisor.[34]

We emphasize the following to all participants, making it clear that they all receive the same briefing: (i) the treatment that the subjects are assigned to; (ii) the fact that all subjects in our study have taken a Raven's Matrices Test, and therefore both Seekers and Advisors are familiar with it.

3.2.1. *Implications of the theory for the treatments.* We now sketch how the arms correspond to the theoretical primitives. One arm—Skill versus Random—varies whether F_{L} first-order stochastically dominates F_{H} or is identical to it (where the "ability" variable comes from the cognitive test we conduct). The second arm—Private versus Reveal—varies the Advisor's prior about the score on our ability test; in the latter case, it is set to a degenerate prior that reveals all the information about the ability test that could be implicit in the Seeker's decision.

We now summarize the model's most immediate qualitative implications for the comparisons of these arms; later, we turn to more detailed structural estimation.

(1) In (Skill, Private), high-ability Seekers, on average, will seek less than low-ability Seekers; Advisors will update their beliefs, upon observing seeking, in the direction of low ability.

[32]We randomize at the village level, as opposed to the pair level, to guard against the scenario that players in the same village would discuss the rules of the game and get confusing information.

[33]In about 14% of cases, we had $k' \leq k$; very little seeking occurred (about 3% within this group) when k' was equal to k and no seeking occurred when k' was strictly less than k.

[34]We implemented this with one surveyor stationed with the Seeker and another with the Advisor, in all treatments, so that the revelation would be immediate. This avoided any doubt as to when the Advisor would learn the Seeker's attributes.

(2) If shame is present, in Revealed treatments, low-ability types seek less than high-ability types even though beliefs are not updated conditional on seeking. The shame effect is strictly larger under the Revealed treatment, because the posterior is extreme (as opposed to the interior values that it takes due to Bayesian updating in equilibrium under Private).
(3) In the Random treatments, high- and low-ability types have similar seeking rates.
(4) In the (Skill, Private) treatment, both types seek less than in (Random, Private).[35]

3.2.2. *Measurement of effects.* The design schematic is presented in Figure 2. The outcome we focus on is the Seeker's decision of whether to forgo his or her own clues in order to visit the paired Advisor over the course of the three days. In looking at this outcome variable, our main effects can be measured most directly by comparing the outcome across treatment cells, holding fixed the ability of the individual (test score), the need for clues (k), and the number of clues that the Advisor has (k'). In our preferred specifications, we operationalize this by using Seeker-by-Advisor clue count fixed effects (i.e., (k, k') fixed effects) and conditioning on the sample of Seekers with low test scores. In discussing comparisons next, we will take for granted that we are conditioning on all these, unless noted otherwise.

First, by comparing (Random, Private) to (Skill, Private), we can estimate a total *stigma effect* that arises when ability is relevant, compared to a benchmark situation where ability is not relevant (with ability being held private in both cases). Recall from (2.1) that stigma here consists of the effects of both signaling (from the reputational payoff) and shame (from the interaction payoff).

Second, by comparing (Random, Private) to (Random, Revealed), we can estimate the *shame effect*. Recall that in the latter treatment, we reveal the ability type of the a low-score individual, irrespective of the seeking decision d. Thus, the signaling term cannot contribute to difference and payoffs, and the incremental utility of seeking in (2.2) is only affected by the shame term, $d \cdot \psi(0)$. Note that this is a pure shame effect, and it is maximal in that beliefs are fully updated to reflect the true ability of the seeker. While comparisons to (Skill, Revealed) are not part of our core design, we leave open the possibility that there may be a Skill-Revealed interaction.[36] After presenting our results, we discuss the interpretation of such a term in Section 6.3.1.

Third, by putting some additional structure on the problem, we can estimate the relative magnitudes of the signaling and shame effects as components of the stigma detected in our first measurement. Suppose the utility of seeking decision d is given by

$$U(d) = \alpha d + \beta P(k_d) + \lambda \mathbf{P}(a = \mathrm{H} \mid d) + \gamma d \left[\mathbf{P}(a = \mathrm{H} \mid d) - \frac{1}{2} \right] - u,$$

[35]This follows from the fact that, under (Skill, Private), given (1), seeking moves the belief of the Advisor down and incurs a positive reputation cost.
[36]"Interaction" here is distinct, of course, from the sense in "interaction payoff."

where $P(k_d)$ is the probability of guessing correctly given k_d signals under decision d, and where u is drawn from a Type I extreme value distribution. Then the incremental value of seeking is

$$\Delta_d U(d) = \alpha d + \beta \Delta_d \mathbf{P}(k_d) + \lambda \Delta_d \mathbf{P}(a = \mathrm{H} \mid d) + \gamma \left[\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2}\right] - \epsilon$$

where ϵ has a logistic distribution and is independent across individuals. The idea, informally, is that by comparing (Random, Private) to (Random, Revealed), we identify γ, since $\mathbf{P}(a = \mathrm{H} \mid d) = 0$ for low-ability types.[37] With this in hand, we can use the comparison of (Random, Private) and (Skill, Private) to back out λ. To do this, we use the differential seeking rates in our treatment to back out how much information seeking conveys, which gives us values of $\Delta_d \mathbf{P}(a = \mathrm{H} \mid d)$.

Fourth, we are particularly interested in how the signaling and shame effects, and their total impact, vary with aspects of the social context. We will thus estimate these effects by interacting them with measures of social proximity (whether individuals are friends, of whether they are of the same caste). We also estimate our structural model separately for the socially proximate and socially distant to tease out the relative magnitudes γ and λ in each case, and explain how the results may reflect differences in the affective strength of these relationships (relevant to shame) and the precision of prior beliefs (relevant to signaling).

It goes without saying that the usual caveats of an experiment conducted at scale—decisions and actions occurring over three days, at the scale of half a day's wage and a maximal payoff of four day's wage—certainly does not tell us about the magnitude of effects in an environment of much larger scales such as agricultural technology adoption. Nonetheless, the decisions made are at a meaningful scale and require real-world interactions. So, not only does the design allows for an internally valid, clean identification of the effects of interest, but also measures the effect at a relevant day-to-day scale.

3.3. **Sample statistics.** Table 1 presents the sample statistics. We have comparable numbers of male and female subjects: 54% of the Seekers and 53% of the Advisors are female.

Though we do have rich *jati* (subcaste) data, because major divisions occur at broad categories, as is standard we look at caste blocks of General, Other Backward Class (OBC), Scheduled Caste (SC) and Scheduled Tribe (ST). In our analysis we treat General and OBC as upper caste and SC/ST as lower caste. 61% of the Seekers and 56% of the Advisors are General or OBC Caste—the remainder are SC or ST.

We also have a wide distribution of ability as measured by the Raven's Matrix test, scored out of 15. The mean Seeker score is 9.5 with a standard deviation of 3.2, and similarly the mean is 9.2 (3.3) for Advisors.

[37]Assume for this discussion α and β have been identified, just so we can focus on the intuition. We explain parameter identification formally in Appendix C.

Turning to the network data, the subjects have on average 8.6 links overall.[38] Within-caste links far outnumber across-caste links both in social and informational relationships.

4. Results

This section presents our results. We start with some preliminaries on the updating of perceptions based on the test score and the raw seeking rates. Then we report the effects of our treatments on seeking, controlling for all relevant heterogeneity. Next, we conduct a more detailed examination of how people make inferences. We conclude with a structural estimation of the parameters of our model.

4.1. **Preliminaries.**

4.1.1. *Do people update perceptions from the test score?* We first show that the score on the Raven's Matrix Test can effectively update people's beliefs about the test-taker's ability. Working with a separate survey sample of 399 subjects (drawn from the same population as the subjects in our main experiment), we asked them to score the intelligence of several other subjects (an average of three per respondent) on a 0-100 scale. Three days later, we revealed to them these subjects' Raven's matrices scores and asked them for a revised score. We regress the outcome of whether there was an increase in this intelligence rating on the Raven's matrix score that we revealed to the subject. Table 2 presents the results. Column 1 only uses surveyor and caste controls, Column 2 adds village fixed effects, and Column 3 uses respondent fixed effects. Note that the the regression identifies the effect of seeing a higher Raven's matrix score on the probability of increasing the score given to the test-taker. We find that a one standard deviation increase in a person's reported score corresponds to a 3-4pp increase in this probability, on a base of 39% (roughly a 10% effect size).

4.1.2. *A first look: Raw probabilities of seeking.* The mean seeking rate across all treatments is 14%, while among those with a low skill and a low clue count (those we will focus on in our main analysis), it is 21%. In Figure 3, we graphically compare the probability of seeking for high-skill and low-skill Seekers, without conditioning on the number of clues obtained by the Seekers. Here, high-skill Seekers are defined as those having an ability score above the median[39] and the low-skilled are defined to be the rest. We see that in the (Random, Private) treatment, the seeking rate is roughly the same for the two types of Seekers. This is because the distribution of the clues is random and independent of Seekers' skill, and so the need for seeking is expected to be similar.[40] In the (Skill, Private) treatment, however,

[38] Recall footnote 28. Subjects have 6.7 friendship links, 5.7 informational links, 6.6 within-caste, and 2.1 across-caste links on average.

[39] In our entire sample, which corresponds to a score of at least 10 out of 15.

[40] Of course, this needn't have been the case: the taste for seeking advice could have turned out to be correlated with ability, and so even absent any signaling concerns, one ability type would have done more seeking.

the high-skilled Seekers seek significantly less often—they have less need for seeking as they obtain more clues by performing better on the skill test.

4.2. **Treatment effects on seeking rates.** Our main sample conditions on low-skilled Seekers (below the median score on the ability assessment) who received low clue counts (a below-median number of clues). We focus on these because it is the Seekers with few clues that have a need to seek. Focusing on the low-ability types avoids confounding variation in the composition of the sample: note that in Skill *only* low-ability types can have low clues, while in Random, the low-clue population is a mix of ability types.

Given this sample, we estimate, for Seeker i with Advisor j in village v:

$$\begin{aligned}\text{Seek}_{ijv} = \alpha &+ \beta_1 \,\text{SkillTreatment}_v + \beta_2 \,\text{RevealTreatment}_v \\ &+ \beta_3 \,\text{SkillTreatment}_v \times \text{RevealTreatment}_v + \epsilon_{ijv}.\end{aligned} \tag{4.1}$$

Table 3 presents the results of the regression analysis. We add Seeker score as well as Seeker-by-Advisor clue count fixed effects across specifications.[41]

4.2.1. *Stigma: Do low-skill Seekers seek less in the Skill treatment?*

We first measure the total stigma effect, if any, making low-skill Seekers in the (Skill, Private) treatment seek less than they do in the (Random, Private) treatment. Recall that we are comparing low-skill Seekers across treatments who are offered the same number of clues (either by random chance in control, or deterministically in treatment) in order to hold the need to ask constant in the comparison.

Table 3 presents the results. We focus on β_1 in equation (4.1). We find that moving to the (Skill, Private) treatment from (Random, Private) leads to a 11.5pp drop (column 3, $p = 0.04$) in the probability that the Seeker seeks on a base of 20.9% for the (Random, Private) treatment. The effect remains robust as we add various fixed effects. Relative to column 1, column 2 adds fixed effects for the Seeker's score on the skill test, and column 3 additionally adds fixed effects for the Seeker-by-Advisor clue counts. This demonstrates that in fact there is a 55% decline in the probability of seeking advice for a low ability individual simply due to stigma, holding everything else fixed.

4.2.2. *Shame: Revealing Seeker ability to Advisor.* We now look at what happens when we reveal the Seeker's ability score to the Advisor. Without a shame inhibition, this would allow the low-skilled Seekers to resume seeking at the same rate as in control, since there would be nothing left to signal. On the other hand, if there is an interaction payoff and a shame effect, then revealing the ability score will discourage seeking relative to (Random, Private).

[41]We also control for surveyor fixed effects and phase of experiment and we reweight in order to compensate for the stratification and have the estimates representative of the population distribution.

To examine whether there is a shame effect, we first focus on the estimates in Table 3 of β_2 from equation (4.1). We find evidence of such an effect: if in the Random treatment one reveals the Seeker's score to the Advisor, then the Seeker is 13.6pp less likely to seek on a base of 20.9% (column 3, $p = 0.08$). This effect is large: the shame effect leads to a 65% decline in the probability of seeking advice for a low ability individual, holding everything else fixed.

It is important to note that the shame effect measured here by comparing (Random, Private) and (Random, Revealed) comes from moving the Advisor's belief that the Seeker's ability is high from the prior to full revelation—i.e., $\mathbf{P}(a = \mathrm{H} \mid d) = 0$ for a low-ability individual. In contrast, when we move from (Random, Private) to (Skill, Private), the corresponding shift is smaller: from π to the equilibrium assessment of the Seeker conditional on his seeking, $\mathbf{P}(a = \mathrm{H} \mid d) \in (0, 1)$. In other words, because in equilibrium a mix of high and low types will seek, the conditional probability of being a high type given a seeking decision will be some non-extreme value, which we examine in more detail below. So when shame is present, it will have less than the full impact we have measured above.

4.3. **Inferences in equilibrium and structural estimates.** Now that we have an estimate of the "full" shame effect, we can return to understanding the components of stigma in the more realistic case we are primarily interested in, where stigma consists of both signaling and shame effects. The first step is to study the size of the shift in beliefs driving both these effects. Thus, in this subsection we begin by examining the information content of seeking decisions. That is, we compute the conditional probabilities of seeking for both types of agent and the implied likelihood ratio[42] that is relevant for updating beliefs about Seeker ability. We then place more structure on the problem, making parametric assumptions about $\varphi(\cdot)$ and $\psi(\cdot)$. Putting these pieces together allows us to decompose the stigma effect into its constituent parts.

4.3.1. *Inferences in equilibrium.* Recall the predictions of Section 3.2.1. The second prediction, based on Proposition 2, is essential to the mechanics of the signaling equilibrium: low-skilled types seek more, so that seeking is a signal of low ability. To check whether this holds in the data, we then compare the seeking behavior of low-skill and high-skill types in the (Skill, Private) treatment. In our sample, due to ties, 53.2% of Seekers had high skill. We examine the behavior of both types of Seekers in the skill treatment. We find that 12.9% of the low-skill Seekers chose to seek (averaging across all clue counts, etc.), whereas only 7.9% of high-skill Seekers chose to seek. The difference is statistically significant.

Given the base rate (53.2% have high skill) and these seeking rates, we can compute the Bayesian posteriors conditional on seeking. Given that someone chooses to seek, the

[42]That is, $\frac{\mathbf{P}(a=\mathrm{H}|d=1)}{\mathbf{P}(a=\mathrm{L}|d=1)}$.

posterior odds ratio that he has a high type is

$$\frac{\mathbf{P}(a = \mathrm{H} \mid d = 1)}{\mathbf{P}(a = \mathrm{L} \mid d = 1)} = \frac{\pi}{1-\pi} \frac{\mathbf{P}(d = 1 \mid a = \mathrm{H})}{\mathbf{P}(d = 1 \mid a = \mathrm{L})} = \frac{\pi}{1-\pi} \times 0.612.$$

Using the population base rate $\pi = 53.2\%$, $\mathbf{P}(a = \mathrm{H} \mid d = 1)$ in this case is 41%. Thus, conditional on a random individual seeking, the probability assigned to high ability falls from the base rate of 53.2% to 41%.

It is also instructive to do a similar computation in the (Random, Private) treatment, where there ought to be no signaling motive. Both types seek more frequently in (Random, Private) as compared to (Skill, Private). Figure 3 presents this graphically. We can compute analogously

$$\frac{\mathbf{P}(a = \mathrm{H} \mid d = 1)}{\mathbf{P}(a = \mathrm{L} \mid d = 1)} = \frac{\pi}{1-\pi} \frac{\mathbf{P}(d = 1 \mid a = \mathrm{H})}{\mathbf{P}(d = 1 \mid a = \mathrm{L})} = 0.98.$$

Here we have $\mathbf{P}(a = \mathrm{H} \mid d = 1) = 49.5\%$: conditional on seeking, the probability that one is high-skill essentially stays at 1/2.

This suggests that the basic force operating in the theory is present in the data, and gives a sense of the magnitude of the belief updating induced by the signaling.

4.3.2. *A structural estimate.* Our structural exercise allows us to decompose the stigma effect into its constitutent parts. To do this, we make the parametric assumptions that

$$\varphi\left(\mathbf{P}(a = \mathrm{H} \mid d)\right) = \lambda \mathbf{P}\left(a = \mathrm{H} \mid d\right) \text{ and } \psi\left(\mathbf{P}(a = \mathrm{H} \mid d)\right) = \gamma\left(\mathbf{P}\left(a = \mathrm{H} \mid d\right) - \frac{1}{2}\right).$$

Therefore we parameterize utility by

$$U(d) = \alpha d + \beta \mathbf{P}(k_d) + \lambda \mathbf{P}(a = \mathrm{H} \mid d) + \gamma d \left[\mathbf{P}(a = \mathrm{H} \mid d) - \frac{1}{2}\right] - u,$$

where u is Type I extreme value, so the marginal utility is given by

$$\Delta U_a(d) = \alpha + \beta \Delta_d \mathbf{P}(k_d) + \lambda \Delta_d \mathbf{P}(a = \mathrm{H} \mid d) + \gamma \left[\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2}\right] - \epsilon$$

with ϵ drawn from a logistic distribution. This expresses the marginal utility of seeking for an individual of ability a in terms of the changes in (i) the instrumental payoff $(\beta \Delta_d \mathbf{P}(k_d))$, (ii) the reputational payoff $(\lambda \Delta_d \mathbf{P}(a = \mathrm{H} \mid d))$, and (iii) the interaction payoff $(\gamma \left[\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2}\right])$.

Appendix C describes identification and estimation of the parameters α, β, γ, and λ and Panel A of Table 4 presents the results. Our parameters of interest are λ (signaling) and γ (shame). We estimate $\widehat{\lambda} = 0.643$ and $\widehat{\gamma} = 0.0846$, so that $\frac{\widehat{\lambda}}{\widehat{\gamma}} \approx 7.6$. Thus, the signaling effect is about eight times as large as the shame effect in our sample. The fact that the signaling parameter is much larger than the shame parameter is consistent with our reduced form effects.

We repeat this exercise below with separate estimates for socially proximate pairs (friends and same caste) and those that are socially distant.

4.4. **Further reduced-form results.**

4.4.1. *High skill and high clue counts.* So far we have presented results for the low-skill population with low clue count to keep the sample constant across treatment and identify the signaling effect of interest. In Table 5, we give the analog of Table 3 for those with high ability and a high number of clues.[43]

The results are largely consistent with our story. The level of seeking is considerably lower—5.8% in (Random, Private)—as compared to 20.9% for the low ability and low clue count sample. This is unsurprising: the need for signals is smaller if one has a higher clue count. Moreover, we see neither signaling nor shame effects. A Seeker is just as likely to seek in (Skill, Private) as (Random, Private). He is also just as likely to seek in (Random, Private) as (Random, Revealed).

Taken together, the results from the high skill and high clue count sample suggest that because these individuals receive more clues, they have less of a need for information and do not seek very often. If they do, it is likely for idiosyncratic reasons.

4.4.2. *The interaction of the skill treatment and ability revelation.* There is one cell we have not yet discussed: (Skill, Revealed). By design, our main effects are identified off other comparisons. We now discuss the comparisons associated with this cell.

Consider the effect of going from (Skill, Private) to (Skill, Revealed), given by $\beta_2 + \beta_3$ in equation (4.1). By revealing ability, the signaling effect faced in (Skill, Private) is removed, which should encourage seeking: we expect $\beta_3 > 0$. Consistent with this, in Table 3 we find that the total effect of going to (Skill, Revealed) from (Skill, Private) corresponds to an 8.8pp increase (column 3). However, this effect is "too large": the resulting seeking rate is nearly the same as (Random, Private), meaning that the seeking rate more than compensates for the shame effect faced in (Random, Revealed).

In interpreting this effect, note that there is an important contrast between the (Skill, Revealed) cell and all of the others. For the population we focus on (low-skill, low-clue-count Seekers) the *need for information is evident to the Advisor*. This is because the Advisor in this case is told that her paired Seeker has a low score, and also knows the rules of the game. This eliminates some plausible explanations for not seeking, such as having no need to (which is a possibility, as far as the Advisor knows, in (Random, Revealed)). Indeed, when we condition instead on *high*-skill Seekers in this treatment, the extra seeking due to the interaction of the Skill and Revealed arms is not significantly different from zero.

[43]Notice that presenting these tables separately is equivalent to a saturated model where we put a dummy for high ability and high clue interacted with every regressor in our main regressions (on the low ability and low number of clues sample).

Though a full analysis of the motives that become relevant in this case is beyond the scope of our study, we offer some hypotheses in a discussion in Section 6.3.1 below.

5. The role of social structure

We have seen that introducing an ability-signaling motivation in social learning induces a stigma effect, inhibiting seeking. We have also studied how this effect can be decomposed into signaling and shame effects. In this section we explore how heterogeneity in these effects makes seeking outcomes dependent on the social context. Specifically, in what relationships is each effect strong?

5.1. **Beliefs and social distance.** Our theoretical analysis makes clear that, holding the signaling parameter λ constant, the strength of the signaling effect depends on the strength of the prior: when the Advisor has a stronger prior that ability is low, beliefs will move less following seeking, creating less room for a signaling effect. This motivates a look at individuals' confidence about others' abilities.

We begin by looking at how confidence in the assessments of others' ability (intelligence) varies with social distance. Panel A of Figure 4 shows that there is a steady decline in confidence as a function of network distance (i.e., graph distance in the friendship graph). Panel C repeats the exercise with caste. When a high caste respondent rates a fellow high-caste individual's ability, she is more confident in her assessment than when she rates a low-caste individual. There is no detectable difference however among low-caste respondents. Overall these findings are consistent with individuals having, on average, stronger priors about those that are socially close.

Next we ask whether respondents update less when they have stronger priors. Panels B and D present the results conditioned on network distance and relative caste, respectively. We define the "upward revision probability" to be the empirical frequency of respondents increasing their assessment of the assessed individual's skill. We compute the upward revision probability conditional on learning that an individual has a test score that is at least one standard deviation above the mean. We see that the upward revision probability when network distance is 1 is about 0.35. That is, the majority do not update their assessment when they are socially close. But at distance 3, the upward revision probability is 0.55. A similar pattern holds looking within caste versus across caste group.

Taken together, the evidence suggests that individuals feel more confident about their assessments of those who are socially proximate to them. Accordingly, they are less responsive to a signal about the individual's ability. This suggests that we should expect signaling to be a greater concern when the Advisor is socially distant rather than socially proximate to the Seeker.

5.2. **Stigma, signaling, and shame by social distance.** We now study the heterogeneous effects of our experiment by social proximity between Seeker and Advisor. We run the regression

$$\begin{aligned}\text{Seek}_{ijv} = \alpha &+ \beta_1 \,\text{SkillTreatment}_v + \beta_2 \,\text{RevealTreatment}_v \\ &+ \beta_3 \,\text{SkillTreatment}_v \times \text{RevealTreatment}_v \\ &+ \delta_0 \,\text{SocialProx}_{ijv} + \delta_1 \,\text{SkillTreatment}_v \times \text{SocialProx}_{ijv} \\ &+ \delta_2 \,\text{RevealTreatment}_v \times \text{SocialProx}_{ijv} \\ &+ \delta_3 \,\text{SkillTreatment}_v \times \text{RevealTreatment}_v \times \text{SocialProx}_{ijv} + \epsilon_{ijv}\end{aligned} \tag{5.1}$$

for Seeker i with Advisor j in village v. SocialProx indicates the social proximity between the Seeker and Advisor, here measured either by whether they are friends or by whether they are of the same subcaste. Again we include Seeker score fixed effects and Seeker-by-Advisor clue count fixed effects.

5.2.1. *Signaling with the socially distant.* We find that among the socially distant, signaling likely drives more of the stigma effect than shame. Figures 5, Panels A and B present our results graphically, simply plotting the raw data. We see that the seeking rate declines when going to (Skill, Private) from (Random, Private) for both non-friends and different-caste pairs, consistent with a stigma effect. However, in the raw data there is no evidence of a shame effect, indicating that most of the stigma effect must be driven by signaling.

Tables 6 and 7 present regression results, which show large stigma effects (β_1 in equation (5.1)) but shame effects (β_2) that are noisy and not statistically different from zero. When we look at subject pairs that are not friends and compare (Skill, Private) to (Random, Private), we identify a 13.9pp decrease in the probability of seeking, on a base of 18.2%, ($p = 0.018$). Similarly, when we look at pairs that differ in caste, we see a 15.4pp decline in the probability of seeking, on a base of 18.2% in (Random, Private). We cannot statistically reject zero shame, though the point estimates are 11.4pp and 8.36pp for non-friends and different-caste pairs, respectively.

Our structural exercise gives us a different lens for assessing the relative magnitudes of the contributing effects. Panels D and E of Table 4 repeat our structural exercise looking only at socially distant pairs. When we estimate the model on strangers or those of different castes, we find that $\widehat{\lambda} = 1.1$ or 1.57, respectively, while $\widehat{\gamma} = 0.009$ or -0.045, respectively. That is, signaling is the primary concern, and the parameter is much larger than the value of the shame parameter.

5.2.2. *Shame with the socially proximate.* Panels C and D of Figures 5 present raw data for the socially proximate. We see small declines in seeking probability when we compare (Skill, Private) to (Random, Private). This indicates that the stigma effect is small among the socially proximate. However, we see sizable shame effects in the comparison of (Random, Revealed) to (Random, Private). Therefore, the raw data suggests that stigma is driven entirely by shame, but the equilibrium shifts in the conditional probability of seeking by type in (Skill, Private) are insufficient for shame to have a meaningful impact on the seeking rate. On the other hand, when ability is revealed, low ability is revealed with certainty, and therefore the shame effect is substantial.

Turning to the regression results, we see that the point estimate for the stigma effect is quite small—a decline of 2.9pp on a base of 25.9% and not statistically different from zero ($p = 0.824$). However, moving from (Random, Private) to (Random, Revealed) results in a large decline in the seeking rate—by about 17.4pp. This indicates a potentially large, but noisily estimated, shame effect amongst friends ($p = 0.263$). Similarly, among same-caste pairs under (Skill, Private), we see no detectable stigma effect (6pp decline on a base of 20.1%, $p = 0.385$). But we do see a significant shame effect among same-caste pairs, with a 16.3pp decline in seeking relative to (Random,Private) ($p = 0.057$).

This is consistent with the idea that the socially proximate may know more about their friends and co-caste members, so with tighter priors there is less scope for updating and therefore signaling. At the same time, people tends to interact more with socially proximate individuals, and the emotional stakes in these interactions can be higher, so shame may be a greater issue, compared to interactions with strangers.

Panels B and C of Table 4 repeat our structural exercise with the socially proximate. We see that when we estimate the model on friend or same-caste pairs, we obtain $\widehat{\lambda} = -0.061$ or 0.285, respectively (neither statistically different from zero), while $\widehat{\gamma} = 0.355$ or 0.163, respectively. That is, signaling is essentially of no concern, though the caste estimates are noisy, whereas the shame parameters are both large and more precisely estimated with confidence intervals excluding zero.

6. Discussion

This section discusses some implications, both for social network theory and for policy; supporting evidence from qualitative work; and further hypotheses suggested by our results.

6.1. **Implications for social structure.** This section moves beyond the two-agent model that has been our main focus throughout. We informally discuss how the stigma of signaling low skill can depend on and feed back into patterns of *homophily* (the tendency of similar types to connect) and how this can contribute to inequality.

6.1.1. *Homophily.* For simplicity, suppose that there are two types of agents in the network, in group[44] $g \in \{r, b\}$, and that an agent's category and skill are independent. Consider the following two line networks depicted in Figure 6: (A) shows a homophilous network, where people of the same group tend to be neighbors within the network, and (B) shows a heterophilous network, where group membership alternates across nodes on the line. The former pattern is more commonly observed in a residential setting, but the latter may be induced, for example, in a workplace.

Assume that each node has a prior belief about the skill level of each of her neighbors: π_{L} and π_{R} for the neighbor to the left and right, respectively, drawn from a full-support distribution. We further assume that a node has a more accurate prior (i.e., one closer to 0 or 1) about a neighbor of the same demographic type.[45]

Assume for simplicity that there is no shame effect, in order to focus on signaling. In the homophilous network, there will be chains of agents who all have strong priors about each other's type. Thus, they are more likely to opt into seeking information from each other. On the other hand, in the heterophilious network, neighbors do not have strong priors about each other's skill, and there is greater scope for signaling across every link in the network. As a result, individuals in a heterophilous network are less likely to engage in social learning. This phenomenon, driven by the forces we have elucidated, feeds on itself: less talking means less familiarity, which in turn means less talking.

Bringing the shame effect back in, assume (again motivated by the data) that the shame effect is stronger among homophilous pairs. In this case, the shame effect would deter learning in the homophilous network and encourage it in the heterophilous network.

Which effect dominates is an empirical question. Both our reduced-form and structural analyses suggest that while the shame effect is stronger among the socially proximate than the socially distant, the signaling effect among the socially distant is so large that it serves as the dominant deterrent. Therefore, we expect that in real-world settings, homophilous networks should allow for greater social learning since they ought to have a lower overall stigma effect.

6.1.2. *Inequality.* Suppose there are some agents who are known to be of high ability, with $\pi \approx 1$ when they are Seekers. This may happen, for instance, because of prior interactions in which they have convincingly signaled high ability. Suppose also that there are other, "ordinary" agents whose abilities are not known with high confidence to be either high or low—that is, for whom the prior π is closer to $1/2$.

[44]A group can be interpreted as caste or other demographic type.

[45]This can capture a number of effects: for instance, those of the same type may interact more frequently (be friends socially) or perhaps be members of the same caste. What is key is that those of the same identity may carry more information about each other, for instance simply because they have opportunity to draw more inferences from a wider range of interactions.

In this case, "ordinary" individuals are deterred from seeking by stigma, whereas our results in case of a precise prior show that known-high-ability individuals are not. This can create a multiplier effect and exacerbate inequalities in information. Those considered very intelligent are permitted to ask questions and are immune to the stigma, while those who need information more protect their reputations.

On the other hand, consider prior beliefs that have a different structure that might be called "bad news": some people are known to be of low ability ($\pi \approx 0$) but, as before, the abilities of "ordinary" individuals are less confidently known. Then those known to be of low ability face no signaling inhibition, whereas those whose abilities are uncertain are more reluctant. On the other hand, in this case the shame inhibition may still deter those of low-ability, depending on whom they are interacting with.

This sketch gives a sense of how the forces we have described may be relevant in exacerbating inequality, paralleling our discussion of homophily above.

6.2. **Policy implications.**

6.2.1. *Information delivery strategies.* Two information dissemination strategies are often used by policymakers: broadcasting information to all in a community (e.g., radio or television) versus seeding information with a sparse set of individuals (e.g., through an extension program). Broadcasting information to all in a community typically makes the broadcast *public*. In particular, community members know that information has been broadcast in the newspaper or on the radio, so it is likely that others have heard the information as well. On the other hand, seeding information with a small set of individuals, such as the introduction of new agricultural technology via an extension program to a village, often makes it known to the entire community the identity of the initially informed, and that others did not receive any instruction.

Our model implies that these differences will matter not only for who is initially informed, but also for signaling concerns and seeking incentives. In particular, if it is public that everyone had a chance to learn, and the information is reasonably "straightforward," then those who need further clarification may signal low ability or laziness if they need to seek clarification. On the other hand, if it is common knowledge that information is delivered only to a few seeds, all others can freely seek information without a signaling deterrent: it is known they had none to begin with.

In Banerjee et al. (2018), our model is adapted to explore how stigma concerns affect information seeking during the 2016 Indian demonetization, in which there was widespread confusion but also a reluctance to appear uninformed, which we document in that paper. The experiment varied whether information about the demonetization was broadcast or seeded with five households, as well as whether the information delivery method was common knowledge. The results show that individuals' asking and learning behaviors follow the patterns suggested by our model: when there is common knowledge about how information is delivered, providing more people with information deters communication and can actually reduce knowledge.

6.2.2. *Learning environments: Pros and cons of easy access to information.* The findings have implications for the consequences of making information available to those who may need it. On the one hand, the presence of stigma clearly highlights the value of anonymous query protocols, such as *Avaaj Otalo* in Gujurat (Cole and Fernando, 2014). These offer individuals a way to access information without exposing themselves to judgment by peers or shame concerns that interact with the rest of their lives.[46]

These implications interact in interesting ways with access to technology. Widespread use of the Internet provides an outside option for access to information (Cole and Fernando, 2014), and the different modes of communicating and interacting may also mitigate shame.

On the other hand, the fact that information is easy to access may make it especially compromising to still need help with it. If social learning is in fact still essential for a substantial number of people to make good use of information, then the increase in signaling or shame concerns can in fact outweigh the value of access (Banerjee et al., 2018). Thus, in the spirit of Gagnon and Goyal (2017), it is important to investigate when a learning environment amplifies the frictions we have been emphasizing, and when it alleviates them.

6.3. **Further hypotheses.** This section offers hypotheses suggested by our study. One comes from an experimental finding, and the other is an implication of the theory.

6.3.1. *Revealing need and pridefulness.* Recall from Section 4.4.2 that in the (Skill, Reveal) treatment when the Advisor is aware of the Seeker's need for information, the seeking rate increases. As we noted there, in this treatment, for the low-ability types we focus on, the Advisor learns that the Seeker has low ability. Therefore, she learns with certainty that the Seeker needs advice, and further, the need for advice is generated by his low skill. Finally, these facts are all commonly known (at least in an approximate sense) between the Seeker and the Advisor. This feature of the treatment within this subgroup is somewhat peculiar. But the surprising outcome we observe there—a strong positive effect on seeking, especially among friends—does suggest some hypotheses for future study.

In particular, given the common knowledge we have described, if the Seeker chooses not to seek advice, he may be signaling an unwillingness to ask for help, which might be called *pridefulness*. Our model can be readily extended to posit that agents care about how others assess attributes of theirs beyond the main one—ability—that we have been concentrating on. In view of such concerns, agents could be reluctant to signal such an attribute. Another

[46] Of course, such concerns can be present—though probably to a lesser degree—in relatively anonymous interactions (Goldfarb et al., 2015).

way of looking at the relevant attribute is that it can be considered good to seek the help of a friend, and refusing to do so when it's clearly needed is a negative signal about the relationship. Note that in all other treatments, there are benign reasons not to be seeking advice, even for the low-skill group—the simplest one being not having a need for it.

This hypothesis is formulated after making our observations, and is necessarily tentative. But the relevance of perceived need on seeking behavior is potentially important, and could even be leveraged to encourage learning.

6.3.2. *Welfare and removing stigma.* In our model, when signaling concerns are present, there may be a unique equilibrium or multiple equilibira. This occurs for fundamental reasons rather than due to peculiar off-path beliefs, which often drive severe multiplicity in signaling models. Recall equation 2.3 in Section 2. A stable equilibrium is a v at which the left-hand side and right-hand side intersect, with $R(v)$ having a slope of less than 1 in absolute value at the point of intersection. Whether there is just one such intersection or several depends on the shapes of the functions φ, ψ, G_{L}, and G_{H}.

When there are several stable equilibria, shifting between equilibria can have substantial welfare consequences. Considering only the preferences in the model (i.e., those of Seekers), the equilibria will have different levels of welfare in utilitarian terms, though in general a shift may be better for some ability types and worse for others. In a richer model, the gains to the Seekers would have to be traded off against the value to Advisor of learning more about the Seeker's type.

A shift between equilibria gives a meaning to the intuitive notion of "removing the stigma" of asking questions. For example, an intervention may involve inducing people to ask questions by providing idiosyncratic, one-time rewards for doing so, essentially adjusting the shocks—ideally in a way that is not known to everyone (to avoid an adjustment in the equilibrium). This can change the interpretation of the act of seeking. Once these interpretations change, the new equilibrium can persist under the old fundamentals. Note that our approach need *not* entail changing the shape of φ or ψ, which are fundamental parameters describing how agents value reputation and interactions. On the other hand, some interventions could have their effect by changing fundamentals temporarily. Analyzing these in detail would require a fuller theory of learning and equilibrium adjustment.

The economic fundamentals of the model involved in (2.3) determine whether temporary policies of this sort can improve welfare. Consider an example illustrated in Figure 1 (B), where points A and B represent two stable equilibria. Suppose a society starts at point B, which represents a higher cutoff v for seeking information. It is possible that a one-shot incentive to seeking, which induces more people to seek information, can "remove the stigma" and lead to the other stable equilibrium represented by point A. In this case, the cutoff v is lower, and there is a greater degree of information flow within the society.

Next, note that some of the interventions that reduce signaling deterrents also mitigate shame. In particular, if seeking can somehow be made less informative about one's type, whether by shifting among equilibria or by changing the fundamentals of the model, then shame deters seeking less. On the other hand, some other interventions that eliminate signaling severely worsen shame concerns and dramatically deter seeking by low types—our (Revealed, Random) treatment being the most prominent example.

Finally, we note that the multiplicity comes from the signaling dimension of the model. If $\varphi = 0$, then it can readily be seen that the model has a unique prediction of behavior. Thus, it is the presence of signaling concerns that creates the opportunity for short-run interventions with considerable effects. If the main impediment to seeking comes from shame, then freeing people to seek requires more fundamental changes in social norms or people's preferences.

7. Conclusion

This paper offers theory and evidence on how information-seeking can be dampened by an endogenous stigma of asking questions. In particular, the act of seeking advice may signal that one is of low ability. Further, there may be an aversion to interact with someone in view of an unfavorable assessment, aside from signaling concerns—a phenomenon we have called shame. We develop a simple model to think about how stigma, through signaling and shame, can affect social learning and then conduct an experiment to explore this mechanism, identifying which effects are operating and how strongly. We find evidence in favor of the view that individuals do worry about both signaling and shame in information-seeking, and this can deter social learning. The signaling and shame effects are different in terms of when and how they operate: signaling matters more when agents need to exchange information with strangers or acquaintances, whereas shame matters more with friends.

In closing, it is useful to benchmark the frictions we have identified against other, familiar sorts of social boundaries. The decline in seeking rate due to stigma when we move from (Random, Private) to (Skill, Private) is 1.5 (respectively, 5) times as large as the decline in seeking rate due to the Advisor not being a friend (respectively, not being of the same caste as the Seeker). Similarly, the decline in seeking rate due to shame in going from (Random, Private) to (Skill, Revealed) is 1.75 (respectively, 6) times as large as the decline in seeking rate due to the Advisor not being a friend (respectively, not being of the same caste). To the extent that social and caste connections are significant in communication frictions, this suggests that seeking frictions due to stigma can have comparable effects.

We have highlighted (in Section 6.1) that perception concerns play an important role in mediating the activation of information networks. These concerns can distort network formation outcomes even relative to what is efficient for the people involved (Niehaus, 2011). The nature of the distortion will vary across settings. But broadly, this observation suggests a practical connection between the behavioral literature on signaling and related phenomena (Bénabou and Tirole, 2006; Bursztyn and Jensen, 2017) and the network literature on endogenous network formation (Galeotti and Goyal, 2010; Acemoglu et al., 2014), as well as work on inequality and social structure more broadly. On the policy side, we have emphasized the implications of signaling and shame concerns for the design of information interventions (such as public announcements) and various opportunities (hotlines, extension services, etc.) designed to facilitate learning.

References

Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2014): "Dynamics of Information Exchange in Endogenous Social Networks," *Theoretical Economics*.

Ali, S. N. and C. Lin (2013): "Why people vote: Ethical motives and social incentives," *American economic Journal: microeconomics*, 5, 73–98.

Ashraf, N., O. Bandiera, and S. S. Lee (2014): "Awards unbundled: Evidence from a natural field experiment," *Journal of Economic Behavior & Organization*, 100, 44–63.

Banerjee, A., E. Breza, A. G. Chandrasekhar, and B. Golub (2018): "When Less is More: Experimental Evidence on Information Delivery during India's Demonetization," NBER Working Paper No. 24679.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013): "The Diffusion of Microfinance," *Science*, 341.

——— (2016): "Gossip: Identifying Central Individuals in a Social Network," .

Beaman, L., A. BenYishay, J. Magruder, and A. M. Mobarak (2018): "Can Network Theory-based Targeting Increase Technology Adoption?" Tech. rep., National Bureau of Economic Research.

Bénabou, R. and J. Tirole (2006): "Incentives and prosocial behavior," *American economic review*, 96, 1652–1678.

Bernheim, B. D. (1994): "A theory of conformity," *Journal of political Economy*, 102, 841–877.

Besley, T. and M. Ghatak (2008): "Status incentives," *American Economic Review*, 98, 206–11.

Bursztyn, L., G. Egorov, and R. Jensen (2016): "Cool to Be Smart or Smart to Be Cool? Understanding Peer Pressure in Education," .

Bursztyn, L., T. Fujiwara, and A. Pallais (2017): "'Acting Wife': Marriage Market Incentives and Labor Market Investments," *American Economic Review*, 107, 3288–3319.

Bursztyn, L. and R. Jensen (2015): "How Does Peer Pressure Affect Educational Investments?" *Quarterly Journal of Economics*, 130, 1329–1367.

——— (2017): "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure," *Annual Review of Economics*, 9, 131–153.

Butera, L., R. Metcalfe, W. Morrison, and D. Taubinsky (2019): "The Deadweight Loss of Social Recognition," Tech. rep., National Bureau of Economic Research.

Calvo-Armengol, A., J. de Marti, and A. Prat (2015): "Communication and Influence," *Theoretical Economics*, 10, 649–690.

Charles, K. K., E. Hurst, and N. Roussanov (2009): "Conspicuous consumption and race," *The Quarterly Journal of Economics*, 124, 425–467.

Cole, S. A. and A. N. Fernando (2014): "The value of advice: Evidence from the adoption of agricultural practices," *HBS Working Group Paper*.

Conley, T. G. and C. R. Udry (2010): "Learning about a New Technology: Pineapple in Ghana," *The American Economic Review*, 35–69.

DellaVigna, S., J. A. List, U. Malmendier, and G. Rao (2016): "Voting to tell others," *The Review of Economic Studies*, 84, 143–181.

Duflo, E. and E. Saez (2003): "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 118, 815–842.

Ellingsen, T. and M. Johannesson (2008): "Pride and prejudice: The human side of incentive theory," *American economic review*, 98, 990–1008.

Fang, H. (2006): "Disentangling the College Wage Premium: Estimating a Model with Endogenous Education Choices," *International Economic Review*, 47, 1151–1185.

Foster, A. and M. Rosenzweig (1995): "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy.*, 103, 1176–1209.

Fryer, R. and D. Austen-Smith (2005): "An Economic Analysis of 'Acting White'," .

Gagnon, J. and S. Goyal (2017): "Networks, markets, and inequality," *American Economic Review*, 107, 1–30.

Galeotti, A., C. Ghiglino, and F. Squintani (2013): "Strategic information transmission networks," *Journal of Economic Theory*, 148, 1751–1769.

Galeotti, A. and S. Goyal (2010): "The Law of the Few," *American Economic Review*, 100, 1468–1492.

Goffman, E. (1959): "The Presentation of Self in Everyday Life," *Garden City, NY.*

——— (1963): *Stigma: Notes on the Management of Spoiled Identity*, Englewood Cliffs, N.J.: Prentice-Hall.

——— (1967): *Interaction ritual: Essays in face-to-face behavior*, Garden City, N.Y.: Doubleday.

Goldfarb, A., R. C. McDevitt, S. Samila, and B. S. Silverman (2015): "The effect of social interaction on economic transactions: Evidence from changes in two retail formats," *Management Science*, 61, 2963–2981.

Heffetz, O. (2011): "A test of conspicuous consumption: Visibility and income elasticities," *Review of Economics and Statistics*, 93, 1101–1117.

Jensen, R. (2010): "The (Perceived) Returns to Education and the Demand for Schooling," *The Quarterly Journal of Economics*, 125, 515–548.

Karing, A. (2018): "Social Signaling and Childhood Immunization: A Field Experiment in Sierra Leone," *Working Paper*.

Kremer, M. and E. Miguel (2007): "The Illusion of Sustainability," *The Quarterly Journal of Economics.*, 122, 1007–1065.

Kremer, M., G. Rao, and F. Schilbach (2019): "Behavioral development economics," *Handbook of Behavioral Economics*, 2.

Lewis, H. B. (1971): "Shame and guilt in neurosis," *Psychoanalytic review*, 58, 419.

Lindsay-Hartz, J. (1984): "Contrasting experiences of shame and guilt," *American behavioral scientist*, 27, 689–704.

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

Page, R. (2015): *Stigma*, London: Psychology Press.

Scambler, G. (1998): "Stigma and disease: changing paradigms," *The Lancet*, 352, 1054–1055.

Spence, M. (1973): "Job market signaling," *The quarterly journal of Economics*, 355–374.

Tracy, J. L. and R. W. Robins (2007): "The self in self-conscious emotions: A cognitive appraisal approach," in *The self-conscious emotions: Theory and research*, ed. by J. L. Tracy, R. W. Robbins, and J. P. Tangney, New York: Guilford Press, 3–20.

Wong, Y. and J. Tsai (2007): "Cultural models of shame and guilt," in *The self-conscious emotions: Theory and research*, ed. by J. L. Tracy, R. W. Robbins, and J. P. Tangney, New York: Guilford Press, 209–223.

Figures

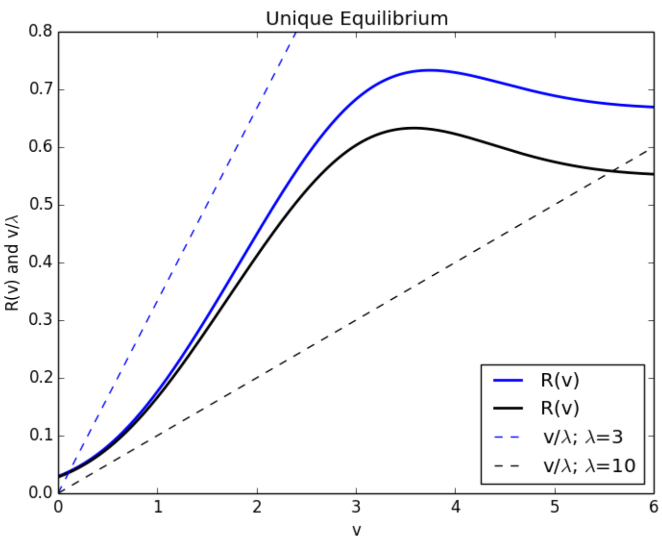

(A) Unique Equilibrium

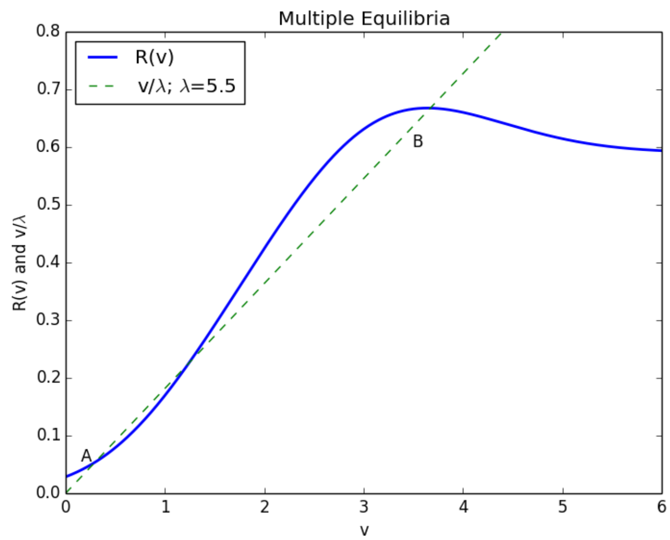

(B) Multiple Equilibria

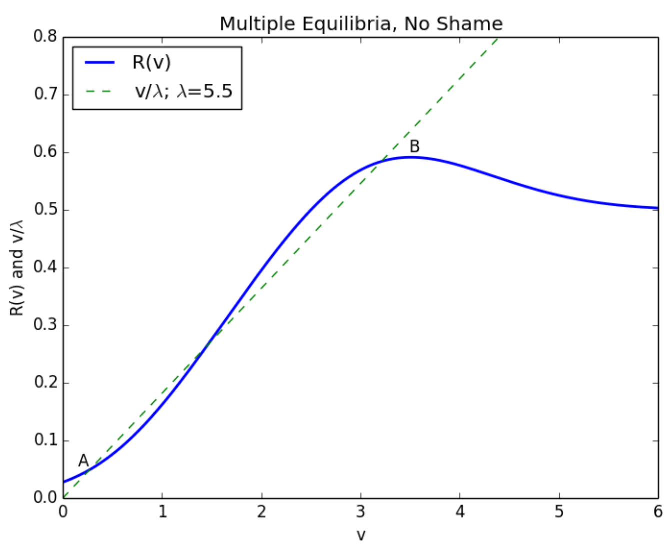

(C) Multiple Equilibria without Shame

Figure 1. A graphical depiction of finding equilibria by intersecting v/λ and $R(v) := \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 0)\right) - \widehat{\varphi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right) - \frac{\gamma}{\lambda}\widehat{\psi}\left(\mathbf{P}_v(a = \mathrm{H} \mid d = 1)\right)$. In Panel A we plot two different configurations where the equilibrium is unique. Panels A and B set $\gamma = 1$ (shame) and Panel C sets $\gamma = 0$. In this example we take the foundation for $\widehat{\varphi}$ from Appendix D, with the c.d.f. $H(x)$ given by $x/(1 + x)$ and then we set $\widehat{\psi} = \widehat{\varphi}$.

	Private	Revealed
Random	value of info − seeking cost	value of info − seeking cost − shame
Skill	value of info − seeking cost − stigma	value of info − seeking cost − stigma + skill-revealed interaction

FIGURE 2. Experimental design schematic. Stigma in the Skill treatment consists of signaling and shame effects.

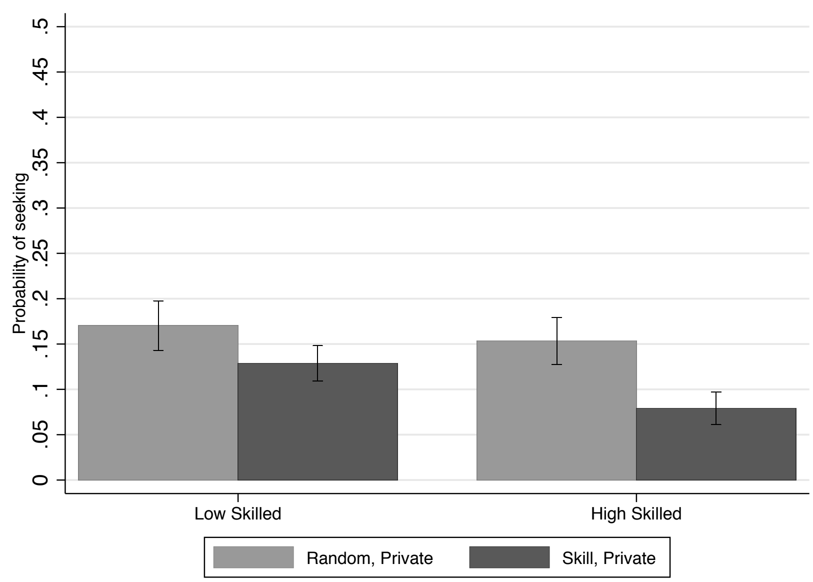

FIGURE 3. Probability of seeking by low skilled (below median score) or high skilled (above median score), by treatment. This does not condition on the clue count in any way.

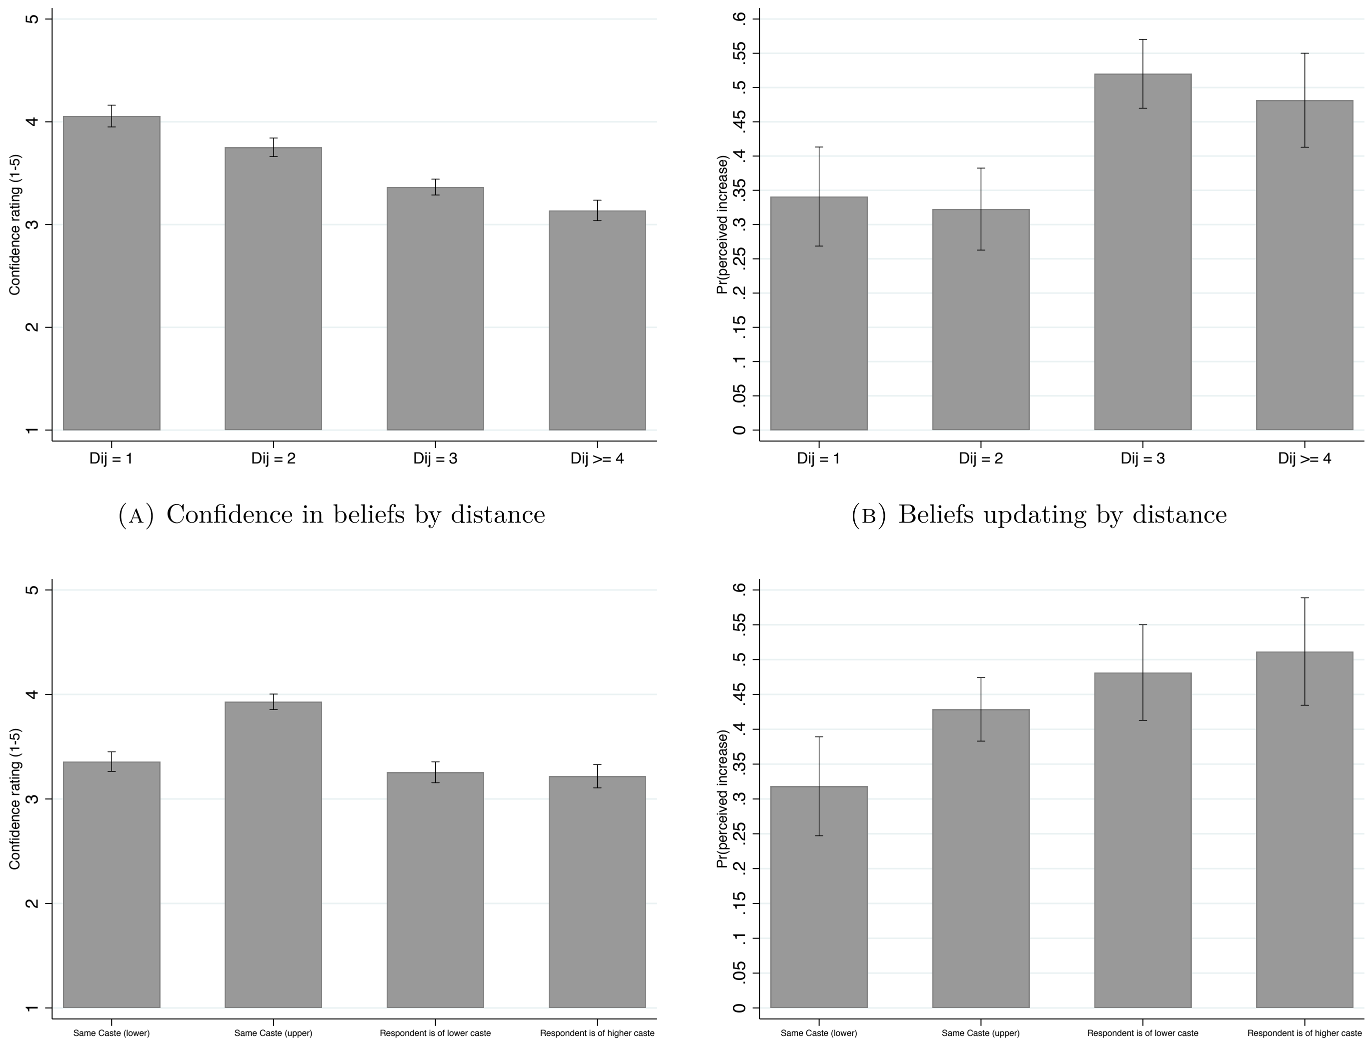

(A) Confidence in beliefs by distance

(B) Beliefs updating by distance

(C) Confidence in beliefs by caste

(D) Beliefs updating by caste

Figure 4. A respondent's stated level of confidence in her belief about a given subject's intelligence at baseline. In (A) and (C) these are averaged conditioning on distance in the friendship network (D_{ij}) and relative caste relationship. In (B) and (D), we plot the probability that the respondent increases her evaluation of a given subject's intelligence after learning that this person got a Raven Matrix score at least one standard deviation above the mean. Again we plot conditioning on distance in the network (D_{ij}) and relative caste relationship. Distance in the friendship network is defined as the length of the shortest path connecting the two individuals in the graph where links reflect friendship (recall footnote 28). We categorize caste relationship into three categories: same, higher, or lower caste (relative to the subject).

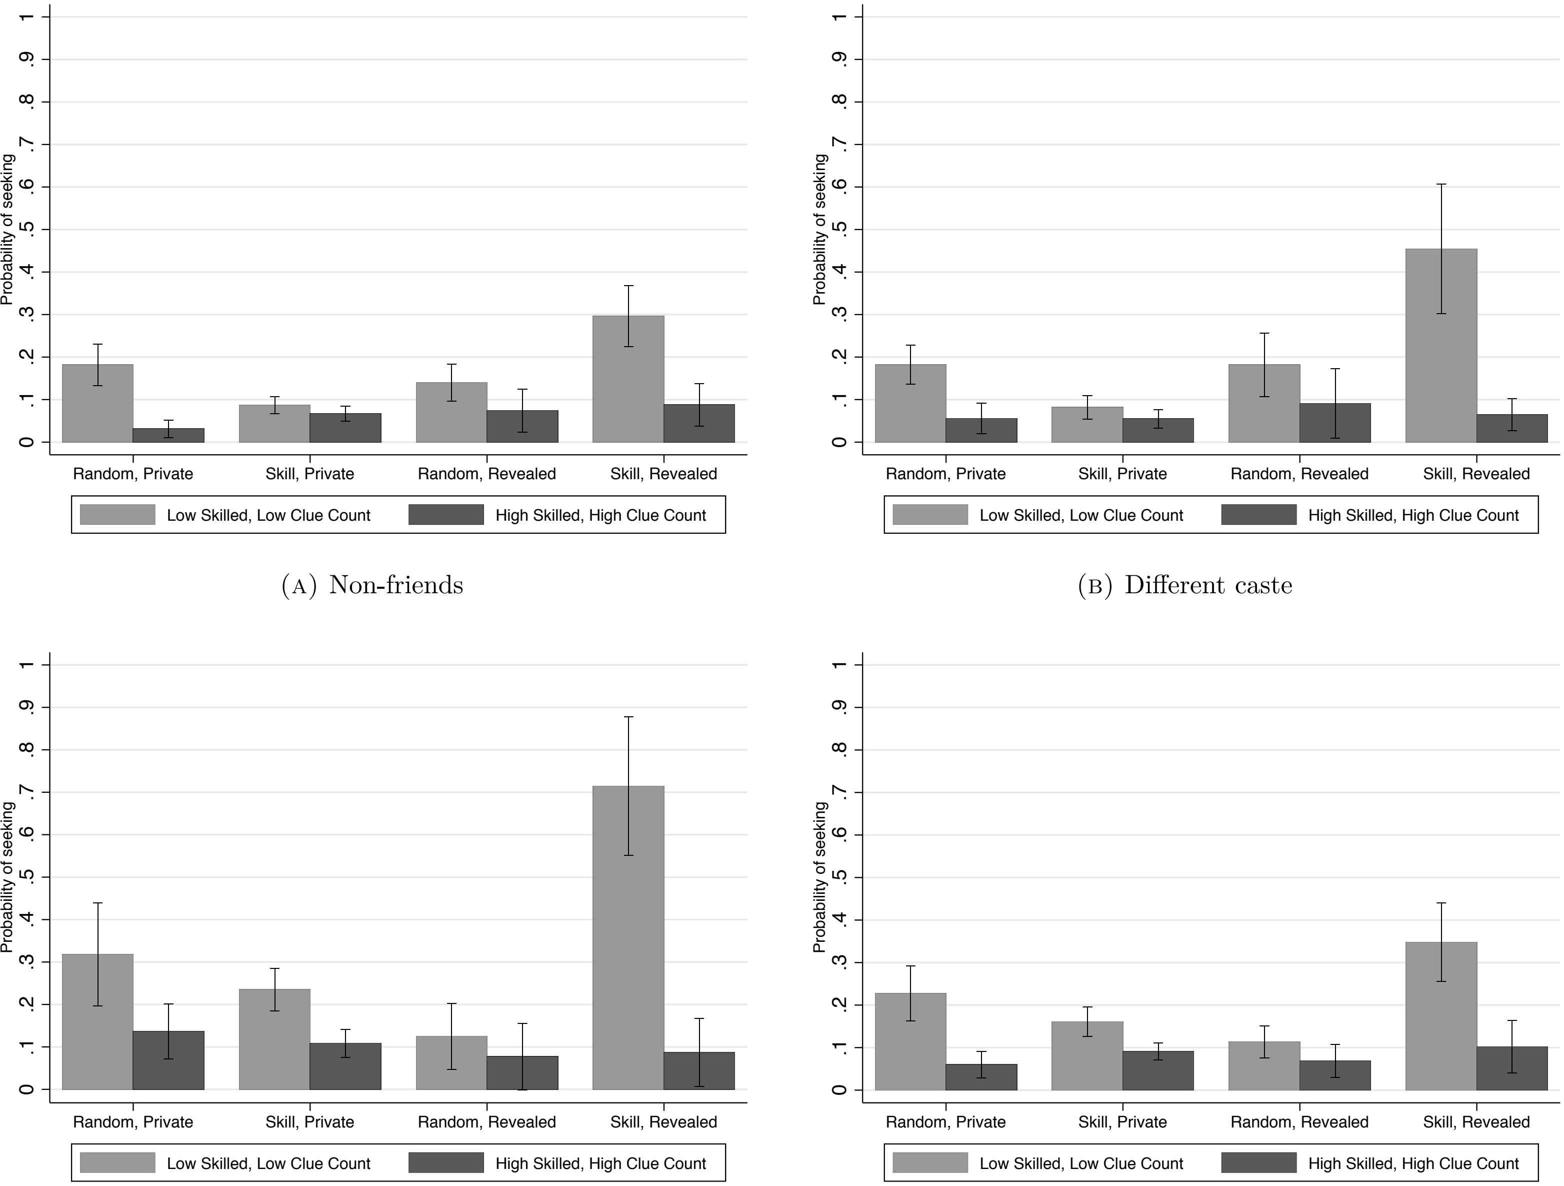

(A) Non-friends

(B) Different caste

(C) Friends

(D) Same caste

Figure 5. Probability of seeking plotted by treatment, with standard errors. Two samples are plotted: low-skill with low clue count and high-skill with high clue count, in order to hold ability and incentive to seek fixed across treatments.

(a) Homophilous Network

(b) Heterophilous Network

Figure 6

Tables

Table 1. Summary statistics

Panel A: Experimental sample

	mean	sd
probability of seeking	0.14	(0.35)
test score of advisor	9.17	(3.29)
test score of seeker	9.46	(3.20)
female advisors	0.53	(0.50)
female seekers	0.54	(0.50)
upper caste advisors	0.56	(0.50)
upper caste seekers	0.61	(0.49)

Panel B: Links of various types within/across caste

	mean	sd
all links	8.62	(5.28)
within-caste all links	6.57	(4.60)
across-caste all links	2.13	(2.73)
social links	6.66	(3.64)
within-caste social links	5.13	(3.33)
across-caste social links	1.60	(2.13)
information links	5.65	(3.99)
within-caste info links	4.30	(3.41)
across-caste info links	1.42	(2.02)

Table 2. Belief updating

VARIABLES	(1) Upward revision prob.	(2) Upward revision prob.	(3) Upward revision prob.
Std. Score	0.0416	0.0447	0.0341
	(0.0142)	(0.0146)	(0.0181)
Observations	1,306	1,306	1,306
Depvar Mean	0.388	0.388	0.388
Village FE		✓	✓
Respondent FE			✓

Notes: Standard errors (clustered at the village level) are reported in parentheses. All columns control for caste and surveyor fixed effects.

Table 3. Main Results: Low IQ, low clue count

VARIABLES	(1) Seeking	(2) Seeking	(3) Seeking
Skill	-0.145	-0.142	-0.115
	(0.0559)	(0.0559)	(0.0551)
Reveal score to Advisor	-0.129	-0.123	-0.136
	(0.0696)	(0.0688)	(0.0765)
Skill × Reveal score to Advisor	0.219	0.203	0.224
	(0.0940)	(0.0886)	(0.0899)
Observations	452	452	452
Random, Private Mean	0.209	0.209	0.209
Seeker Score FE		✓	✓
Advisor-Seeker clue count FE			✓

Notes: Standard errors (clustered at the village level) are reported in parentheses. All columns control for surveyor fixed effects and phase of experiment.

Table 4. Structural estimates of signaling and shame

	(1)	(2)
Panel A: Full Sample	λ	γ
	0.6431	0.0846
	[0.0209, 1.5854]	[-0.0119, 0.1810]
Panel B: Friends	λ	γ
	-0.0606	0.3553
	[-1.0459, 1.6908]	[0.2364, 0.4542]
Panel C: Same Caste	λ	γ
	0.2854	0.1632
	[-0.8420, 1.9681]	[0.0551, 0.2668]
Panel D: Strangers	λ	γ
	1.1002	0.0092
	[0.0717, 3.0203]	[-0.0398, 0.1639]
Panel E: Different Caste	λ	γ
	1.5711	-0.0453
	[0.5801, 3.4712]	[-0.1276, 0.1233]

Notes: Confidence intervals computed from block bootstrap by resampling with replacement at the village level are reported in brackets.

Table 5. High IQ, high clue count

VARIABLES	(1) Seeking	(2) Seeking	(3) Seeking
Skill	0.0181	0.0184	0.00505
	(0.0248)	(0.0252)	(0.0289)
Reveal score to Advisor	-0.00975	-0.00640	-0.0212
	(0.0503)	(0.0501)	(0.0556)
Skill × Reveal score to Advisor	-0.0398	-0.0433	-0.0276
	(0.0778)	(0.0785)	(0.0787)
Observations	484	484	484
Random, Private Mean	0.0581	0.0581	0.0581
Seeker Score FE		✓	✓
Advisor-Seeker clue count FE			✓

Notes: Standard errors (clustered at the village level) are reported in parentheses. All columns control for surveyor fixed effects and phase of experiment.

Table 6. By social relationship

VARIABLES	(1) Seeking	(2) Seeking	(3) Seeking
Skill	-0.149	-0.155	-0.139
	(0.0561)	(0.0559)	(0.0574)
Reveal score to Advisor	-0.112	-0.106	-0.114
	(0.0694)	(0.0704)	(0.0752)
Skill × Reveal score to Advisor	0.217	0.203	0.217
	(0.100)	(0.0973)	(0.0991)
Friend	0.124	0.113	0.0775
	(0.138)	(0.134)	(0.139)
Skill × Friend	0.0387	0.0688	0.110
	(0.148)	(0.146)	(0.150)
Reveal score to Advisor × Friend	-0.131	-0.119	-0.0610
	(0.170)	(0.171)	(0.178)
Skill × Reveal score to Advisor × Friend	0.397	0.336	0.275
	(0.251)	(0.251)	(0.260)
Observations	452	452	452
Random, Private, Non-Friend Mean	0.182	0.182	0.182
Seeker Score FE		✓	✓
Advisor-Seeker clue count FE			✓

Notes: Standard errors (clustered at the village level) are reported in parentheses. All columns control for surveyor fixed effects and phase of experiment.

Table 7. By caste

VARIABLES	(1) Seeking	(2) Seeking	(3) Seeking
Skill	-0.155	-0.164	-0.154
	(0.0629)	(0.0595)	(0.0618)
Reveal score to Advisor	-0.0841	-0.0709	-0.0836
	(0.0881)	(0.0922)	(0.0982)
Skill × Reveal score to Advisor	0.341	0.349	0.355
	(0.190)	(0.185)	(0.176)
Seeker Caste same as Advisor Caste	0.0451	0.0499	0.0192
	(0.0765)	(0.0722)	(0.0719)
Skill × Seeker Caste same as Advisor Caste	0.0402	0.0583	0.0935
	(0.0927)	(0.0863)	(0.0829)
Reveal × Seeker Caste same as Advisor Caste	-0.108	-0.117	-0.0802
	(0.118)	(0.118)	(0.119)
Skill × Reveal × Seeker Caste same as Advisor Caste	-0.0690	-0.122	-0.141
	(0.215)	(0.211)	(0.207)
Observations	452	452	452
Random, Private, Diff Caste Mean	0.182	0.182	0.182
Seeker Score FE		✓	✓
Advisor-Seeker clue count FE			✓

Notes: Standard errors (clustered at the village level) are reported in parentheses. All columns control for surveyor fixed effects and phase of experiment.

Appendix A. Theoretical details and proofs

A.1. **Definition of equilibrium.** A (mixed) strategy for the Seeker is a map $\sigma : \{\mathrm{H}, \mathrm{L}\} \times \mathbb{R} \to [0,1]$ where $\sigma(a, w)$ is the probability of $d = 1$ given ability a and a realization of V equal to w. A belief function for the Advisor is a map $b : \{0,1\} \to [0,1]$ giving the posterior belief that $a = \mathrm{H}$ conditional on observing d. The Seeker's payoff is

$$U(d) = V\mathbf{1}_{d=1} + \varphi\left(b(d)\right) + d \cdot \psi\left(b(d)\right). \tag{A.1}$$

A Bayes-Nash equilibrium is defined to be a pair (σ, b) where b is consistent with Bayesian updating assuming the Seeker plays according to σ, and σ maximizes the Seeker's payoff (among all σ) taking b as given.

A.2. **Proof of Proposition 1.** Equation (A.1) above implies

$$U(1) - U(0) = V - \left[\varphi(b(0)) - \varphi(b(1)) - \psi(b(1))\right]. \tag{A.2}$$

Given V, this difference does not depend on a. Thus (up to probability-zero events) by the no-atoms assumption, in any best-response profile the Seeker will choose $d = 1$ if and only if the difference is nonnegative, i.e. if

$$V \geq \varphi(b(0)) - \varphi(b(1)) - \psi(b(1)). \tag{A.3}$$

A.2.1. *Inferences.* Consider any strategy profile in which $d = 1$ is chosen if and only if $V \geq v$. First we will use Bayes' rule to compute the Advisor's inferences based on the seeking behavior for any candidate cutoff v. As a preliminary calculation, note that in such a strategy profile,

$$\begin{aligned}
\mathbf{P}_v(d = 1 \mid a) &= \mathbf{P}(C = 1)G_a(v) \\
&= qG_a(v) \\
\mathbf{P}_v(d = 0 \mid a) &= 1 - \mathbf{P}_v(d = 1 \mid a) \\
&= 1 - qG_a(v).
\end{aligned}$$

Letting $b_v(d)$ denote the beliefs induced by this cutoff rule, we compute by Bayes' rule that

$$\frac{b_v(0)}{1 - b_v(0)} = \frac{\mathbf{P}_v(a = \mathrm{H} \mid d = 0)}{\mathbf{P}_v(a = \mathrm{L} \mid d = 0)} = \frac{\pi}{1-\pi}\frac{1 - qG_{\mathrm{H}}(v)}{1 - qG_{\mathrm{L}}(v)} \tag{A.4}$$

and

$$\frac{b_v(1)}{1 - b_v(1)} = \frac{\mathbf{P}_v(a = \mathrm{H} \mid d = 1)}{\mathbf{P}_v(a = \mathrm{L} \mid d = 1)} = \frac{\pi}{1-\pi}\frac{G_{\mathrm{H}}(v)}{G_{\mathrm{L}}(v)}. \tag{A.5}$$

A.2.2. *Equilibrium condition.* In any equilibrium the Seeker with $V = v$ has to be indifferent between seeking and not seeking. This means (A.3) holds with equality, where $b_v(0)$ and

$b_v(1)$ are determined by (A.4) and (A.5), respectively. Thus for the cutoff type we have

$$v = \varphi(b_v(0)) - \varphi(b_v(1)) - \psi(b_v(1)), \tag{A.6}$$

with the b determined as we have said.

A.2.3. *Existence of cutoff.* A v solving (A.6) exists because the right-hand side of (A.3) is a continuous function of v (by Assumption 2) bounded in absolute value (by boundedness of the functions φ and ψ), and thus taking values in some bounded interval $\mathfrak{I}$. The function $v \mapsto v$ is a continuous function crossing the interval $\mathfrak{I}$, so there must be a solution of (A.6) by the Intermediate Value Theorem.

A.2.4. *Positivity of cutoff.* Now we argue that any v solving (A.6) is positive. At $v \leq 0$, the left-hand side is nonpositive. By Lemma 1 in Section A.6 below, $b_v(1) < b_v(0)$. By the Law of Iterated Expectations, since $b_v(1)$ and $b_v(0)$ must average to the prior π, it must in fact be the case that $b_v(1) < \pi < b_v(0)$. Putting this together with our assumptions that φ is increasing and that ψ in increasing with $\psi(\pi) = 0$, it follows that the right-hand side of (A.6) is positive for any v. Thus there are no solutions with negative v.

A.3. **Proof of Proposition 2.** Proposition 1 in Section A.6 establishes that the cutoff is the same for both types. Lemma 1 establishes that $b_v(1) < b_v(0)$.

A.4. **Proof of Proposition 3.** If it is hypothesis (1) that holds, and the prior becomes precise, then using (A.4) and (A.5) we can verify that $b_v(d)$, for any v and any d, tends to π. Thus, the right-hand side of (A.6) tends pointwise to 0, and so the v solving (A.6) must be very close to 0. Using the formulas in Section A.2.1 above, it follows that

$$\frac{\mathbf{P}_v(a = \mathrm{H} \mid d = 1)}{\mathbf{P}_v(a = \mathrm{L} \mid d = 1)} - \frac{\mathbf{P}_v(a = \mathrm{H} \mid d = 0)}{\mathbf{P}_v(a = \mathrm{L} \mid d = 0)} \to 0.$$

In other words,

$$\frac{\mathbf{P}(a = \mathrm{H} \mid d = 1)}{1 - \mathbf{P}(a = \mathrm{H} \mid d = 1)} - \frac{\mathbf{P}(a = \mathrm{H} \mid d = 0)}{1 - \mathbf{P}(a = \mathrm{H} \mid d = 0)} \to 0.$$

Since the function $x \mapsto \frac{x}{1-x}$ is increasing and continuous, this shows that

$$\mathbf{P}(a = \mathrm{H} \mid d = 1) - \mathbf{P}(a = \mathrm{H} \mid d = 0) \to 0.$$

Indeed, both must be arbitrarily close to the prior. Then $\psi(\mathbf{P}(a = \mathrm{H} \mid d))$ converges to zero as well (by the assumption that $\psi(\pi) = 0$) for both values of d, and so the net benefit of seeking becomes arbitrarily close to V.

The argument assuming that (2) holds is very similar.

A.5. **Proof of Proposition 4.** Write condition (A.6) defining the right-hand side as $R(v)$, so that

$$\frac{v}{\lambda} = R(v) := \widehat{\varphi}(b_v(0)) - \widehat{\varphi}(b_v(1)). \tag{A.7}$$

If $1/\lambda$ is larger than the maximum slope of the right-hand side, it is clear (from the Mean Value Theorem) that there cannot be two values of v satisfying the equation.

Now for the "large λ" result: first note that $R(v) \neq 0$ for any v by the "strict" aspect of Assumption 1 on first-order stochastic dominance. Choose an arbitrarily large v_0. Let R_0 be the minimum of R over $[0, v_0]$, which is strictly positive. Choose λ so large that v/λ remands less than R_0 over the same interval. Then any intersection must satisfy $v_0 > 0$. For v_0 chosen large enough, the probability of seeking will be arbitrarily small.

A.6. **Auxiliary results.**

LEMMA **1.** *For any $v \in \mathbb{R}$,*

$$\frac{\pi}{1-\pi}\frac{G_H(v)}{G_L(v)} < \frac{\pi}{1-\pi}\frac{1-qG_H(v)}{1-qG_L(v)}. \tag{A.8}$$

and consequently the $b_v(0)$ and $b_v(1)$ defined by (A.4) and (A.5) satisfy $b_v(1) < b_v(0)$.

Proof. The ratios are well-defined by Assumption 2. By Assumption 1 on first-order stochastic dominance, we have $\frac{G_{\mathrm{H}}(v)}{F_{\mathrm{H}}(v)} < \frac{G_{\mathrm{L}}(v)}{F_{\mathrm{L}}(v)}$, and therefore

$$\frac{\pi}{1-\pi}\frac{G_{\mathrm{H}}(v)}{G_{\mathrm{L}}(v)} < \frac{\pi}{1-\pi}\frac{F_{\mathrm{H}}(v)}{F_{\mathrm{L}}(v)}.$$

Also by Assumption 1 on first-order stochastic dominance,

$$\frac{\pi}{1-\pi}\frac{G_{\mathrm{H}}(v)}{G_{\mathrm{L}}(v)} < \frac{1}{1}.$$

Now, for any positive reals x, y, z, y', z', if we have $x < y/z$ and $x < y'/z'$ then it follows that $x < \frac{qy+(1-q)y'}{qz+(1-q)z'}$. Thus,

$$\frac{\pi}{1-\pi}\frac{G_{\mathrm{H}}(v)}{G_{\mathrm{L}}(v)} < \frac{\pi}{1-\pi}\frac{qF_{\mathrm{H}}(v)+(1-q)}{qF_{\mathrm{L}}(v)+(1-q)}.$$

To deduce the desired equation, use the identity $G_a(v) = 1 - F_a(v)$ to show

$$\frac{\pi}{1-\pi}\frac{qF_{\mathrm{H}}(v)+(1-q)}{qF_{\mathrm{L}}(v)+(1-q)} = \frac{\pi}{1-\pi}\frac{1-qG_{\mathrm{H}}(v)}{1-qG_{\mathrm{L}}(v)}. \tag{A.9}$$

This completes the proof of the first statement. The ordering of $b_v(1)$ and $b_v(0)$ follows from the fact that the function $x \mapsto \frac{x}{1-x}$ is strictly increasing for $x \in (0, 1)$. ∎

Online Appendix
Not for Publication

Appendix B. Experimental materials

B.1. **Skill test.** We present sample questions from the skill test.

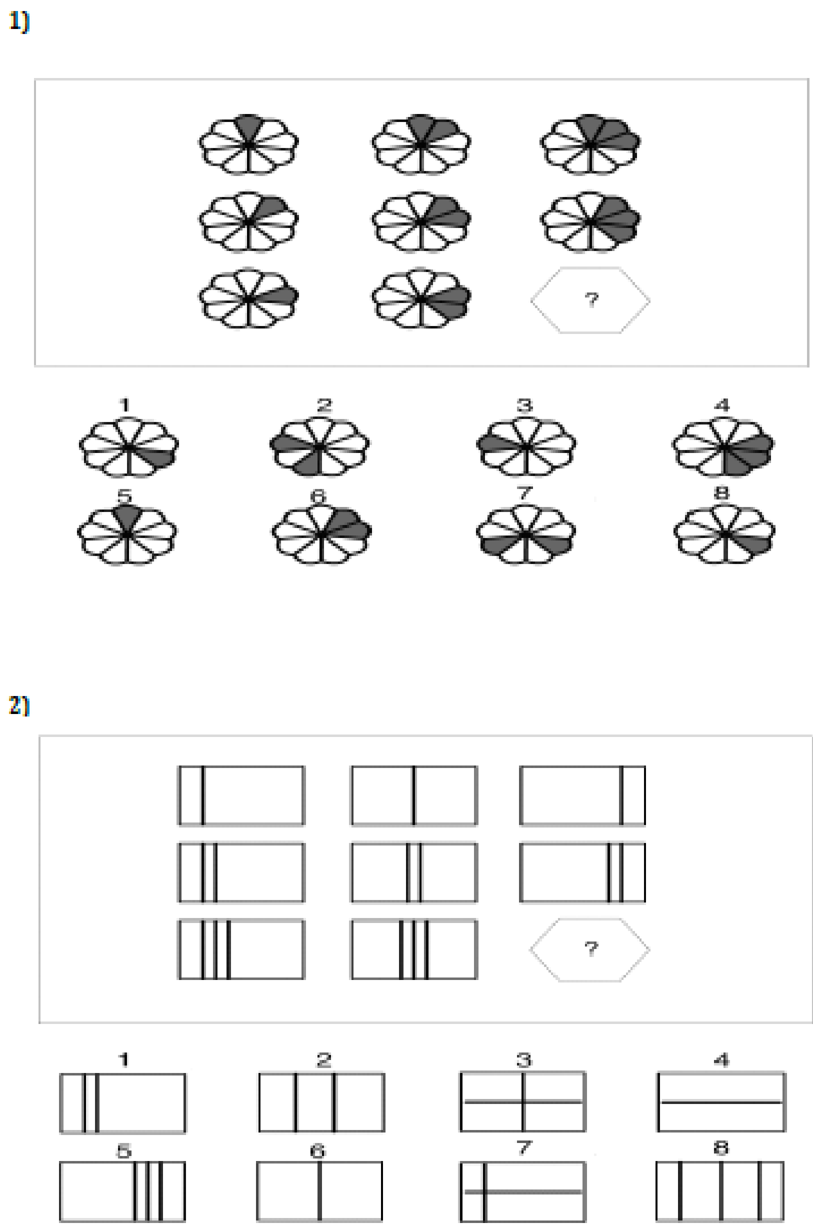

Figure B.1. Examples from Skill Test.

B.2. **Score to clues.** We present the mapping from the skill test score to the number of clues drawn in the Skill treatment.

Test Score	Clues
0	1
1	
2	
3	
4	2
5	
6	
7	3
8	
9	
10	4
11	
12	
13	5
14	
15	

FIGURE B.2. Score to Clues Schedule.

Appendix C. Structural estimation

We take a simple stand on preferences where the reputational and interaction payoffs are linear functions of posterior beliefs:

$$U_a(d) = \alpha d + \beta \mathbf{P}(k_d) + \lambda \mathbf{P}(a = \mathrm{H} \mid d) + \gamma d \left[\mathbf{P}(a = \mathrm{H} \mid d) - \frac{1}{2} \right] - u,$$

where u is Type I extreme value. In that case the marginal utility of seeking is

$$\Delta_d U_a(d) = \alpha + \beta \Delta_d \mathbf{P}(k_d) + \lambda \Delta_d \mathbf{P}(a = \mathrm{H} \mid d) + \gamma \left[\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2} \right] - \epsilon$$

where ϵ independent across individuals and its distribution is logistic, with a c.d.f. given by $\Lambda(\cdot)$.

The role of priors can be thought as follows. Though the presentation is for the case of $\pi = 1/2$, when priors deviate from this, they are, in a first-order sense, absorbed into λ. Suppose, for instance, that the Advisor thinks the paired Seeker has an ability less than 1/2. Then a first-order approximation yields that the estimated λ decomposes into the underlying utility parameter multiplied by an attenuation factor reflecting that beliefs are updated less when the prior is stronger. It is in this sense that priors that vary by whether the pair are friends versus strangers can greatly change the effect of signaling on decision making through a smaller λ. The shame term, on the other hand, depends only on the posterior.[47]

Thus the probability that the individual chooses to seek is simply

$$\mathbf{P}(d = 1) = \Lambda \left(\alpha + \beta \Delta_d \mathbf{P}(k_d) + \lambda \Delta_d \mathbf{P}(a = \mathrm{H} \mid d) + \gamma \left[\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2} \right] \right).$$

The experimental variation allows us to identify parameters of the model in a rather straightforward way. Because we observe the share of a-type individuals who choose to seek as well as the change in information, $\Delta_d \mathbf{P}(k_d)$, and the type share conditional on a seeking decision, $\mathbf{P}(a = \mathrm{H} \mid d)$, we can estimate all pieces of the model.

Specifically, we can obtain estimates of α and β directly from a logistic regression with the (Random, Private) data:

$$d_i = \Lambda\left(\alpha + \beta \mathbf{P}(k_d)_i\right)$$

Note that under (Random, Private), we have $\Delta_d \mathbf{P}(a = \mathrm{H} \mid d) = 0$ (since ability is unrelated to the need to seek) and, averaging across the population, $\mathbf{P}(a = \mathrm{H} \mid d = 1) = \frac{1}{2}$.[48]

With $\left(\widehat{\alpha}, \widehat{\beta}\right)$, we then estimate for $\theta = (\lambda, \gamma)$,

$$\widehat{\theta} = \mathrm{argmin}_{\theta \in \mathbb{R}^2} \widehat{m}(\theta)' \, \widehat{m}(\theta)$$

[47]We could change the reference point from 1/2 to another value, even one that is relationship-dependent, but for this simple exercise we aim to keep the number of parameters small.

[48]This uses less information than jointly estimating all moments together, but is sufficient for our purposes.

where the moments are given by simply matching the empirical share of those who seek in each treatment with the probabilities when evaluated at the appropriate parameters. The first is simply the share of low types who seek in the (Random, Revealed) treatment minus the quantity that, in equilibrium, it is equal to: the c.d.f. at the estimated parameters. The second moment is the share of low types who seek in the (Skill, Private) minus the prediction of this at the estimated parameters.

We use the following moments by treatment, where the sums and normalizations are over the respective samples:

(1) (Random, Revealed, Low type):

$$\widehat{m}_1(\theta) := \frac{1}{n}\sum_i d_i - \Lambda\left(\widehat{\alpha} + \widehat{\beta}\mathbf{P}(k_d)_i + \gamma\left(0 - \frac{1}{2}\right)\right).$$

(2) (Skill, Private, Low type):

$$\widehat{m}_2(\theta) = \frac{1}{n}\sum_i d_i - \Lambda\left(\widehat{\alpha} + \widehat{\beta}\mathbf{P}(k_d)_i + \lambda\Delta_d\mathbf{P}(a = \mathrm{H} \mid d) + \gamma\left(\mathbf{P}(a = \mathrm{H} \mid d = 1) - \frac{1}{2}\right)\right).$$

Recall that all of $\mathbf{P}(k_d)_i$, $\Delta_d\mathbf{P}(a = \mathrm{H} \mid d)$, and $\mathbf{P}(a = \mathrm{H} \mid d = 1)$ are observed quantities in the data.

We conduct inference by block-bootstrap. Specifically, we sample with replacement 70 villages and then recompute $\left(\widehat{\alpha}_b, \widehat{\beta}_b, \widehat{\lambda}_b, \widehat{\gamma}_b\right)$ for $b = 1, \ldots, B$, here $B = 500$, to compute standard errors. We also bootstrap bias correct our parameter estimates.

Appendix D. Foundations for the reputational payoff

For example, suppose that the the Advisor receives utility W_a from collaborating with the Seeker on some later project if the true skill is a. If the Advisor chooses to collaborate, the Seeker receives a deterministic payoff of $\lambda > 0$ which enters his utility additively. The Advisor wants to collaborate if and only if

$$\frac{\mathbf{P}(a = \mathrm{H} \mid y)}{\mathbf{P}(a = \mathrm{L} \mid y)} \geq \frac{-W_{\mathrm{L}}}{W_{\mathrm{H}}}. \tag{D.1}$$

The ratio $\frac{-W_{\mathrm{L}}}{W_{\mathrm{H}}}$ corresponds to the relative value of working with a high-skill Seeker compared to the loss of working with a low-skill Seeker, and corresponds to the odds—in the likelihood-ratio sense—required to make working with this person worthwhile. Let H be the c.d.f. of $\frac{-W_{\mathrm{L}}}{W_{\mathrm{H}}}$ and assume it has a positive density supported on the positive reals. The utility of a Seeker of skill a is then

$$U_a(d) = V(d) + \lambda H\left(\frac{\mathbf{P}(a = \mathrm{H} \mid d)}{\mathbf{P}(a = \mathrm{L} \mid d)}\right).$$

Appendix E. High stakes

We reran our experiment in a new sample of subjects all of low score in two treatments (Random, Private) and (Skill, Private). The goal was to make the expected payoffs considerably larger conditional on guessing correctly, now Rs. 330 along with the same chance ti win the Rs. 1350 phone, but now changing the odds to make the value of seeking almost miniscule for a high type and enormous for someone with low clue count.

In this case, $k \in \{2, 6\}$, so the number of clues is binary. Further, while the probability of a clue being correct in the low case is 0.6, it is 0.8 in the high case. This corresponds to an individual with 2 clues of quality 0.6 being about 60% likely to make the right decision whereas the number is 94% for someone with 6 clues of quality 0.8. Finally, Advisors always receive 8 clues of quality 0.8. This implies that getting the clues from the Advisor leads to getting the right choice 96% of the time. Therefore, the expected gains in monetary terms for someone who received a high number of clues here is Rs. 7, whereas it is Rs. 110 (just under a day's wage) for someone who received a low number of clues.

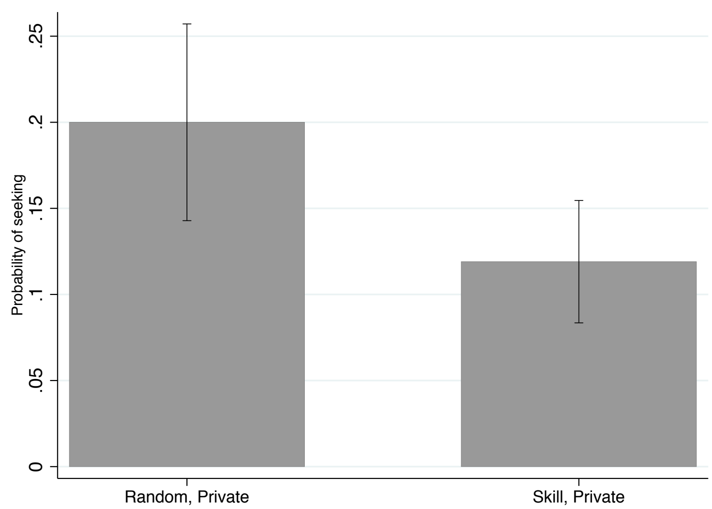

Figure E.1. Seeking rates by treatment across 134 subjects who have low IQ and low clue count.

Figure E.1 presents the results graphically.[49] As in our main experiment, (Random, Private) has a mean seeking rate of 0.2. This drops by 9.31pp when we move to (Skill, Private). This suggest that the sizeable reduction in seeking due to the signaling effect persists even when we raise the stakes considerably.

[49]A regression with heteroskedastic-robust standard errors yields significant differences between the two treatments.

Table E.1. High Stakes Results

VARIABLES	(1) Seek
Skill	-0.0931
	(0.0624)
Observations	134
Random, Private Mean	0.200

Notes: Robust standard errors are reported in parentheses.