Whom do we learn from? Beliefs, preferences, and identity in learning

Akshay Moorthy*

May 3, 2024

Abstract

This paper reports from a series of large-scale survey experiments aimed at identifying whether people have preferences over the social identity of information sources. I examine both naturally occurring identities (caste and religion in India), and experimentally assigned identities (in a EU/US sample). The results show that when the quality of information is known, the identity of the messenger does not influence learning. Further, participants react strongly to signals of information quality in all settings and the evidence suggests that people may rely on pre-existing beliefs about the abilities of different identity groups in the absence of quality signals. Finally, I show that people prefer to learn from non-social sources (a computer algorithm) rather than from other people. Taken together, the results suggest that experts and policymakers should prioritise emphasising the quality of information.

*I gratefully acknowledge guidance and support from Bertil Tungodden, Erik Sørensen, and Benjamin Enke. Kata Urban provided excellent administrative support. I further thank Kai Barron, Luigi Butera, Alexander Cappelen, Daniel Carvajal, Fabio Galeotti, Sam Hirshman, Chris Roth, Mikko Silliman, and participants at various seminars and conferences for helpful comments. This study was funded by ERC grant 788433.
Moorthy: Department of Economics and FAIR, Norwegian School of Economics; akshay.moorthy@nhh.no.

1 Introduction

Social learning is a key component of human evolutionary success (Boyd et al., 2011; Henrich, 2016) and influences many economic decisions (Mobius and Rosenblat, 2014). In social learning, individuals observe others' actions, extract relevant information from those actions, and decide whether to use that information. Determining when to use this information is challenging because the quality of information is often unobservable (Stigler, 1961). In such situations, individuals may use other observable attributes like an information provider's social identity to gauge information quality. The focus of this paper is on how social identity affects social learning. While a vast literature in the social sciences has explored whether and when people learn, the question of whom they learn from has received relatively little attention.

Research has shown that social identity influences economic decisions through both beliefs and preferences (Akerlof and Kranton, 2000; Shayo, 2020).[1] In the context of social learning, individuals may use a person's identity to evaluate the quality of the information that they provide, as in statistical discrimination (Phelps, 1972; Arrow, 1973). At the same time, work on taste-based discrimination (Becker, 1957; Guryan and Charles, 2013) suggests the possibility of a preference channel which might lead individuals to ignore information not because of its intrinsic quality but because of the identity of its source. For example, people may be more (less) likely to trust information that comes from their in- (out-) group, or they may incur psychological costs from listening to their out-groups.

Understanding whether the preference channel exists is important as inefficient information aggregation could lead people to make poorer decisions. At scale, this could even lead to adverse effects such as the formation of information echo chambers (Levy and Razin, 2019). However, identifying the preference channel is challenging as many factors that affect information quality (such as education or experience) could be correlated with social identity. Disentangling beliefs about the quality of information from preferences for the identity of the information source calls for an experimental approach.

In this paper, I conduct large-scale online survey experiments studying how identity affects how people learn from others. The primary experiments study this question within the context of religion and caste in India, which matter in a variety of important social and economic contexts. I complement these experiments with additional treatments using experimentally assigned minimal identity (with a US/EU sample). Finally, to assess whether there is a more fundamental aversion to learning from another human being, I study whether people prefer to get information from a social source (another human) or a non-social source (a computer algorithm).

The experiment consists of a two-period task that combines the widely used "balls-and-

[1] Social psychologists have conceptualised identity as people's sense of who they are based on membership to relevant social groups such as race, gender, or religion (Tajfel and Turner, 1978).

urns" and "judge-advisor" paradigms. In the first period, participants are given some information and have to make an incentivised estimate. In the second period, participants are shown an estimate made by another person (the *source*). Importantly, sources and participants make estimates on the same tasks and with the same information. Participants then decide whether they want to *stick* to their initial estimate, or *switch* to the shown estimate, which is the outcome of interest.

In addition to the estimate, participants observe two attributes associated with the information source. First, they receive a signal of the *quality* of the source's estimate, which is the probability that the estimate is incentive maximising. By design, the quality of the source is independent of the identity of the source and controls participants' beliefs. Second, they observe the *identity* group of the person who made that estimate. This is exogenously assigned and varies between treatment conditions. In experiment *Caste*, the sources belong to a high-status caste group or a low-status caste group. In experiment *Religion*, the sources are either Hindu or Muslim. In experiment *Minimal*, participants from a EU/US online sample are randomly assigned to one of two minimal identity groups and see sources from either the same group or the other group. In experiment *Computer*, participants receive information from either a human source or a computer algorithm. Across experiments, the focus of analysis is on comparing how participants switch when receiving the same information from sources belonging to different identity groups, holding the quality of information fixed.

The design focuses on observational learning in situations where everyone (i) has identical information, (ii) wants to find the same (objectively correct) solution, and (iii) finds obtaining the correct solution cognitively demanding. In such a setting, a Bayesian decision-maker without processing costs would not need to learn from others. However, descriptive results from my experiments show that a substantial fraction of participants choose to switch to the source (between 25% and 37%). These decisions are not driven by a belief that others have better or different information, as both participants and sources have the same information set; rather, participants learn from others because they believe that others are better at processing the same information. Such situations are relatively unexplored in the literature on social learning which has largely focused on situations where people possess private information. Studying how people learn when interpretations differ is important because many learning problems involve a search for both better information and better interpretations. For example, many investment decisions involve sifting through documents to find the best options, but can also be made through advice from experts with experience or specialised knowledge.

The main finding from these experiments is that across contexts, I find no evidence that participants have preferences over the identity of an information source when the quality of information is the same. Participants are, on average, equally likely to switch to a source from a high- or low-status caste group in experiment *Caste*, or a Hindu or a Muslim source in experiment *Religion*. Looking at heterogeneity by group status in these experiments – for instance, do Hindus prefer to get information from other Hindus rather than from Muslims. I find that participants are just as likely to use information that comes from their in-group than from their out-group. These results are robust to a wide range of tests, and analyses looking at different combinations of identity in- and out-groups are largely consistent with the main result.

Experiment *Minimal* provides further support for the absence of a preference channel. In this experiment (conducted with a US/EU sample), participants are assigned to one of two groups based on choices in an unrelated task (following the design in Tajfel et al. (1971); Chen and Li (2009)). They see estimates made either by a person from the same (in-) group or the other (out-) group. The results from this experiment replicate the earlier patterns and do not find any evidence that participants prefer to get information from in-group sources.

Turning to the role of beliefs, I find that participants are sensitive to the quality of information in all treatment conditions. Participants switch to the source's estimate more often when they see a high-quality source than when they see a low-quality source. I also look at how participants learn when the quality of the source's estimate is unknown. I find suggestive evidence that when participants are not given a signal about the quality of the source's estimate, they rely on the source's identity to gauge quality. Participants' prior beliefs about the task-specific ability of people from a caste group are correlated with whether they switch when seeing a source from that group. When participants are given a signal about the quality of the source's estimate, these correlations disappear. In conjunction with the other results, this suggests that when the quality of information is precisely known, the identity of the information source does not play a role in whom people choose to learn from.

Finally, experiment *Computer* compares social learning, where participants get information from another human to non-social learning, where participants get an estimate from a computer algorithm. As before, the quality of information is precisely known to participants. Despite the formal similarity of these two choice environments, participants are much more likely to switch to the estimate when the information is provided by an algorithm than when it is provided by another human. Previous work on social versus non-social learning has found mixed results. Some have shown that people may have an aversion to learning from algorithms relative to other humans (Goeree and Yariv, 2015; Dietvorst et al., 2018, 2015). Others find that people display algorithm appreciation, and favour algorithms over humans in situations (Logg et al., 2019). Attempts to reconcile these conflicting findings have revolved around the accuracy of algorithms (Hou and Jung, 2021). My results show evidence for appreciation at all levels of information quality (ranging from a coin toss to a near-certainty).

Summing up, the results show that when deciding whether to learn from others, people respond strongly to beliefs about the quality of information. However, there is no strong evidence that people have preferences for the identity of whom they get information from through social interactions. However, I do find that people learn more in non-social settings than in social settings. The near-absence of evidence for the preference channel stands in contrast to the vast literature on discrimination which shows that preference-based discrimination occurs in many settings. The choice of the naturally occurring identities, religion and caste in India, was driven by their importance in Indian social and economic life (Munshi, 2019; Mosse, 2019; Jaffrelot, 2021) and previous research has documented the existence of discriminatory behaviour (both belief- and preference-based) because of differences in religion and caste. The lack of strong evidence for preferences for the identity of information sources suggests that policymakers can focus on highlighting information quality and on making information transmission as impersonal as possible.

Contribution and related literature. The paper contributes to the literature on social learning by showing how beliefs and preferences driven by social identity affect learning. The importance of, and various frictions in, social learning have been documented in a variety of lab and field settings: Weizsäcker (2010) finds that individuals learn less than they should from others in a meta-study. Conlon et al. (2022) find that individuals underweight the information that they get from others relative to information gathered by themselves, and go on to show differing levels of this barrier within couples. Mobius and Rosenblat (2014) reviews this large literature, in which the role of social identity is conspicuous by its absence. The links between identity and learning have been explored as secondary outcomes in a few field experiments: for instance BenYishay and Mobarak (2018) provide suggestive evidence that farmers learn more about new agricultural technologies from others like themselves. Macours (2014) shows that farmers learn more from other high-skilled farmers. In the laboratory, Berger et al. (2018), Parys and Ash (2018), and Zou and Xu (2022) find that people learn more from their in-groups in an experimentally assigned identity setting, but are not able to speak to specific mechanisms.

To the best of my knowledge, this paper is the first to identify the effect of identity-related beliefs and preferences in social learning and to experimentally demonstrate how these can be disentangled. The novel setting of the experiment – everyone involved has the same information and the same objectives – studies a crucial component of social learning that has been relatively ignored. The growing literature studying the role of narratives in economic decision-making highlights the relevance of this approach (Shiller, 2017; Graeber et al., 2022; Barron and Fries, 2023). The paper also contributes by providing evidence from a non-western sample and extends the validity of some results in the belief-updating literature that have mostly come from studies in the global North or through convenience samples.

Experiments based on the influential design developed by Anderson and Holt (1997) have long been used to study social learning. The abstract simplicity of this paradigm has yielded a rich set of results. However, this very simplicity limits the ability of this paradigm to study more detailed contexts such as social identity. Previous studies using this design to study in-group effects on learning have found significant and interesting results but were not able to link them to specific mechanisms (Berger et al., 2018; Zou and Xu, 2022). Other approaches have tried using a large relative increase in incentives to argue against the existence of preferences for political in-groups (Zhang and Rand, 2023). The experimental approach developed in this paper allows for the identification of beliefs and preferences, and can be a useful tool to study the effect of various contextual factors on learning.

The results further speak to a growing literature that studies whether discrimination is driven by tastes or inaccurately specified beliefs (Guryan and Charles, 2013; Bohren et al., 2023). Bohren et al. (2019), Ayalew et al. (2021), Chan (2022), and Gallen and Wasserman (2023) find that indicating expertise reduces discrimination in seeking advice from mentors or selecting experts such as doctors. Relative to these studies, the experiment in this paper controls information quality directly related to the decision task, which strengthens identification. My results add to this strand of literature by showing that controlling beliefs to minimise the possibility of identity-based discrimination can be effective even in low-stakes settings where information providers are not necessarily experts.

This paper also relates to a nascent literature on political polarisation and information demand. Dekel and Shayo (2023), Robbett et al. (2023), and Zhang and Rand (2023) find that people show an inclination to evaluate members of their political in-groups more favourably than out-group members. This occurs even in settings where the decision criteria should not be affected by political affiliation. Zhang and Rand (2023) attempt to distinguish between preferences and beliefs by using high-stakes treatments. My paper provides an improved approach to identifying preferences and supports the results of these papers by showing that much of the partisan bias could be driven by beliefs or motivated reasoning concerns related to identity, rather than by preferences for the identity of information sources.

Finally, this paper contributes to the literature on how social identity shapes economic decision-making, with direct evidence on the role of religion and caste on social learning in India. A large body of work has documented the influence of religious and caste identity on consumption (Atkin et al., 2021), hiring in labour markets Siddique (2011), labour supply (Cassan et al., 2019; Oh, 2023), marriage markets (Banerjee et al., 2013), teamwork (Ghosh, 2022) and many others (Munshi (2019) reviews the literature on caste and economics in India, Mosse (2019) provides a more general overview of caste and modern India). This paper adds to this literature by showing that preferences for the religious or caste identity of information sources do not affect how people learn.

Next, Section 2 presents a conceptual discussion and provides background information on religion and caste in India. Section 3 presents the experiment design and describes the various experiments and treatment conditions. Sections 4–8 present descriptive and causal evidence from the experiments. Section 9 concludes the paper.

2 Conceptual discussion and background information

This section provides a brief conceptual discussion of the channels through which identity can affect learning. It also provides details on the specific identity contexts that are used in the experiments.

2.1 Conceptual discussion

Consider a situation where a decision maker (DM) has to correctly estimate the probability y_T that a particular state of the world realises. First, the DM observes some information I and forms an independent estimate y_1. This estimate is formed by attempting to apply Bayes' Rule (*BR*) to the information. The DM makes this estimate with some error in applying *BR*, and they have beliefs about the accuracy of their estimate. This is their subjective certainty or confidence in their estimate, which is their belief about the probability of the accuracy of their estimate (as in Enke and Graeber (2023)).

Next, the DM is given additional information in the form of an estimate y_s made by another person; the other person's estimate y_s is formed based on the same information I. After observing y_s, the DM must make a decision y_2. The DM can either choose to stick to their initial estimate, i.e. $y_2 = y_1$, or they can switch to the other person's estimate, $y_2 = y_s$.

Belief channel. The DM seeks information to form beliefs about the accuracy of y_s, Q. In the absence of contextual features or preferences for particular states, the DM compares their subjective certainty (or confidence) c in their estimate against the quality Q of the other person's estimate.

$$y_2 = \begin{cases} y_S & \text{if } Q > c \\ y_1 & \text{if } Q \leq c \end{cases}$$

First, suppose that only the estimate y_s is observable and participants do not know anything about the quality Q of y_s. It is therefore difficult to compare Q and c when making the decision y_2.

Now suppose that the social group g that s belongs to is observable. g can be a person's race, ethnicity, or gender. Models of statistical discrimination show that people make inferences about individuals from beliefs about the characteristics of their social groups. In this case, if the DM has some prior beliefs about the ability of a group g they can evaluate the quality of the estimate y_s based on these beliefs. Thus, differences in how DMs react to sources from different groups will be driven by differences in the beliefs that individuals hold about these groups.

Next, suppose that the DM receives a (possibly noisy) signal about the quality Q of the specific estimate y_s made by the source s. The DM's posterior belief about the quality of the source will be a variance-weighted function of the prior beliefs about the group g and the noisy signal of Q. If this signal is sufficiently precise, beliefs about the quality of the source will be equal to the signal Q, and underlying beliefs about the ability of groups will not play a role in decision-making. In other words, providing a precise signal of the quality of an estimate made by another person will eliminate the role of any group-specific beliefs in this situation.

Preference channel. Suppose that DMs find switching to y_s inherently costly. This assumption is supported by the literature on belief formation which documents "conservatism bias" (Benjamin, 2019). Suppose additionally that this cost depends upon the social group g of the source s. In this case, both preferences for groups and beliefs about groups will play a role in determining whether the DM chooses to switch to y_s. If a DM has an intense dislike for a particular group, or if there is some prescription against interacting with individuals from some groups, this makes switching costlier than when observing another source where these costs are smaller or do not exist.

If the signal about quality is noisy, then this leads to an identification problem as behaviour conditional on observing identical signals can be explained both by differences in underlying beliefs about groups and by differences in preferences for groups (potentially compounded by misspecified beliefs, as in Bohren et al. (2023)). However, if the signal is precise, then any differential switching conditional on other observable characteristics (other than the identity of the source) being equal can be attributed to preferences for (or against) a particular identity. This is the linchpin of the experimental design – the quality of information must be independent of the identity of the source and must be precisely known to separate the belief and preference channels.

Next, I provide some intuition on why social identity may give rise to such preferences, along with some context on religion and caste in India.

2.2 Social Identity and Preferences

Social psychologists conceptualise identity as an individual's sense of self based on their membership in social groups. Such groups may be driven by demographic factors like race, gender, nationality, or through more subjective considerations such as hobbies, personal interests, and affiliations with specific organizations (Tajfel and Turner, 1978). Economic models of the channels through which identity influences behaviour broadly assert that membership to a group carries prescriptions for appropriate behaviour. Individuals incur costs from deviating from these prescriptions, and this motivates individuals to adhere to these prescriptions (Akerlof and Kranton, 2000). In the context of social learning, the act of learning from other people involves associating with others in situations where some aspects of their identity may be visible. This could invoke various identity-related prescriptions that make associating with specific identity groups a transgression of their group's norms. If these considerations are sufficiently strong, this may prevent people from learning from individuals belonging to certain identity groups.

Identity effects have been documented with both naturally occurring identities, and with experimentally assigned identities. In this paper, I look at both of these categories. Next, I provide some relevant information about the natural identity contexts studied in this paper – religion and caste in India.

Caste in India. The caste system is a system of social stratification in India that goes back thousands of years (Munshi, 2019). The system consists of thousands of caste groups called *jatis*. Membership to caste groups is largely determined by birth. Additionally, many population groups were excluded from the caste system and were regarded as "untouchables". The caste system is deep-rooted, complex, and has evolved over the centuries. Many prominent features of the caste system such as endogamy, social hierarchy, segregation, and ritual purity continue to exert a considerable influence on modern Indian society. While the caste system is largely associated with Hinduism, most of the other religions are also a part of it, and people can possess caste identities even if they are not Hindu.

To effectively target welfare programs and affirmative action policies, the Government of India classifies castes into four "categories". The "General" category (also known as the Forward Castes, $\approx$ 30% of the population) consists of *jatis* that are considered socially and economically advanced (in relative terms). People from the General category are more likely to be wealthier and better educated than people from the other categories. The "Scheduled Castes" (SC, $\approx$ 20%) and "Scheduled Tribes" (ST, $\approx$ 9%) categories are formed of castes that are the most economically and socially disadvantaged. This category also contains the erstwhile "untouchables" or *Dalits*, and people from indigenous tribes. Last, the "Other Backward Classes" (OBC, $\approx$ 40%) category consists of many *jatis* that are economically and socially disadvantaged relative to the General category.

Caste and social learning. Individuals are likely to possess beliefs about the education or cognitive ability of people belonging to particular caste groups. This can be driven by economic realities such as disparities in education and income between caste groups, and by caste-specific stereotypes. For example, some specific castes are perceived as intellectuals, while others are perceived as entrepreneurial or business-oriented. Thus, if people believe that individuals from one caste group are likely to be intelligent or well-educated, they may prefer to learn from them. Conversely, people may prefer not to learn from people belonging to caste groups that are perceived negatively.

Additionally, people may have strong preferences for associating with others from their caste groups and may have very strong preferences for not associating with caste groups considered lower in status. Many caste groups are organised around occupational identity, with many occupations associated with ritual pollution (such as tanning or sanitation). On the other hand, people from downtrodden castes may bear a sense of pride or resentment against the status hierarchy. They may prefer to be self-reliant rather than accept help or advice from people belonging to castes above them in the hierarchy. Survey data from Pew Research Center (2021) and data from my experiments indicate that $\approx$ 70% of people's friendship networks are restricted to their caste group. A reluctance to associate or work with someone from a different caste based on the hierarchy may lead to not using information provided by them.

Religion in India – Hinduism and Islam. Religion is an important part of daily life in India. According to the 2011 census, $\approx$ 80% of the country is Hindu, $\approx$ 14% is Muslim, and the rest of the population follows other religions such as Christianity, Sikhism, Buddhism, Jainism, etc. The tensions between Hindus and Muslims in India have deep historical roots (often traced back to the Mughal era) and were exacerbated during the British colonial period and the eventual Partition of India. Contemporary conflicts arise from various socio-political factors such as territorial disputes, religious nationalism, and competition for resources. Such tensions are often heightened by political polarisation along communal lines.

Like caste, religion can affect social learning through both belief and preference channels. Since Muslims are, on average, poorer and less well-educated than Hindus (especially those from the higher status castes), people may have beliefs that Muslims are less likely to be intelligent or well-educated than upper-caste Hindus. In terms of preferences, purity norms associated with caste differentiation also extend to Muslims (especially in terms of eating beef). However, a history of communal tension coupled with rising mistrust and animosity in recent times (Jaffrelot, 2021) are possibly more influential in driving preferences.

3 Experiment Design

Design goals. Studying the role of identity-driven beliefs and preferences for information requires a setting where these channels can be cleanly separated. The priority is to precisely control beliefs about the quality of information as any ambiguity in these beliefs leaves room for various confounding factors to affect participant choices. Controlling beliefs requires a decision task that has an objective truth, and that all individuals possess the same information. If not, participants may always have beliefs that other people know something different from themselves, or that they have different skills and experiences which could influence performance in a given domain. It is also important to minimise the role of image concerns and other forms of motivated reasoning that may influence learning choices.

Below, I describe the experimental framework, the treatment conditions, and the identity contexts. A full set of screenshots of the experiment along with additional details are presented in the appendix.

3.1 Experimental task

The experimental task builds on the "balls-and-urns" and "judge-advisor" paradigms, which are widely used in experiments studying belief updating and social learning. The task takes place in two periods. Participants make an incentivised decision in each period.

First period. Participants are shown two urns, urn A and urn B which contain 100 red or black coloured balls. Urn A contains $\theta \in \{70, 90\}$ red balls, and urn B contains $100 - \theta$ red balls. Urn A is randomly selected with probability $p \in \{0.5, 0.7, 0.9\}$ else urn B is selected. $k \in \{3, 5\}$ balls are drawn (with replacement) from the selected urn. The participant does not know which urn is selected. The drawn balls are then shown to the participant.

After seeing the drawn balls, participants make their first incentivised decision. They estimate $y_1 \in (0, 100)$, the posterior probability that urn A was selected using a slider. Participants also state their subjective certainty that y_1 lies in a 5% interval centred around the correct estimate.

Importantly, the true posterior can be calculated accurately by applying Bayes' rule using the provided information. In other words, there is an objective truth that a costless Bayesian will be able to compute by applying Bayes' rule. Details of Bayes' rule and how it can be applied to calculate the posterior are provided to all participants.[2]

Second period and main outcome. After making this decision, participants are shown a value between 0-100, y_s, which is associated with a source s. In addition to the estimate y_s, participants see one or more of (i) the group *identity* g of the source, and (ii) the *quality* of the estimate. Figure 1 shows the representation of these elements in the experiment. Participants are then asked to make a second decision y_2 which is the result of a binary choice: whether to *stick* to their first guess $y_2 = y_1$, or *switch* to the shown value $y_2 = y_s$. This is the main outcome variable, *Switch*, which is 1 when $y_2 = y_s$ and 0 when $y_2 = y_1$.

Participants complete six of these tasks, with different configurations of balls, urns, and probabilities. These tasks are deployed in a series of experiments, listed in Table 1. Each experiment studies a specific identity context and consists of two treatment groups. Participants are assigned to one of these experiments and are randomly assigned upon entry (between-subjects) into one of the two treatment groups. The treatment interventions are

[2] In open text responses, many participants indicated that they tried to apply Bayes' rule while working on the task.

Table 1: Experiment overview

Experiment	Identity Treatment	Sample	Sample size
Religion	H – Hindu	Panel survey	431
	M – Muslim		422
Caste	G – General	Panel survey	564
	O – SC/ST/OBC		575
Caste – No Quality	G – General	Panel survey	186
	O – SC/ST/OBC		190
Minimal identity	In-group	Prolific	142
	Out-group		136
Computer	Computer	Prolific	163
	Human		430

Notes. Participants are assigned to one experiment, and to one identity condition within each experiment. General – Source is from the General caste category. SC/ST/OBC – Source is from the Scheduled Castes, Scheduled Tribes, or Other Backward Classes categories. In treatment *Human*, the sample from the *Minimal* experiments is also used.

based on manipulating the identity of the source between-subjects, while varying a participant's knowledge and beliefs about the source s within a treatment. Below, I describe how identities are made salient in each of the identity contexts.

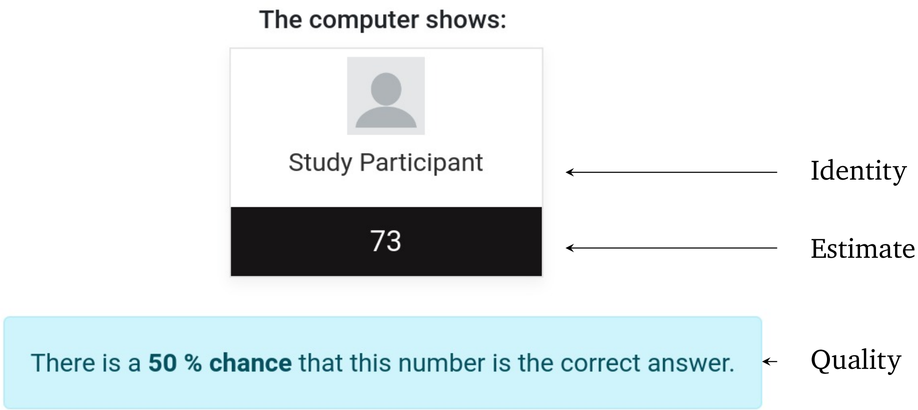

Figure 1: Source – Attributes

3.2 Identity contexts

Caste. In this experiment, the caste identity of the source is made salient through surnames that are informative of the caste category of the individual. These surnames are shown along with an arbitrarily chosen initial (eg. Mr A. Moorthy) in place of "Study Participant" in Figure 1. There are two treatments in this experiment, G and O. In treatment G, surnames belonging to the General (or "Forward Castes") category are used. In treatment O, surnames belonging to the Scheduled Castes, Scheduled Tribes, or Other Backward Classes

(SC/ST/OBC) are used. The surnames used for *G* and *O* were validated for recognisability through a separate survey that was conducted through the same survey provider. 200 individuals were incentivised to correctly classify a list of surnames into one of the four caste categories.

The *Caste* experiment and the validation study were conducted through online surveys with a sample of participants from India provided by Norstat. The effective sample size for the main experiment is 851 Hindu participants – $\approx$ 63% of the participants belonged to the General (*G*) caste, and the remaining $\approx$ 37% belonged to one of the other caste groups. Relative to the population, the *G* category is overrepresented in the sample.

Religion. The religious identity of the source is made salient using surnames that are informative of the religion of the individual. In treatment *Hindu*, the names are the same as those used in the *G* treatment of the *Caste* experiment. In treatment *Muslim*, the surnames are common Muslim surnames. The Muslim surnames were not validated separately as these are generally easily identifiable by Indians.

This experiment was conducted through online surveys with a sample of participants from India provided by Faktum Research. The effective sample size is 853 participants of whom 647 were Hindus ($\approx$ 75%), 67 Muslims ($\approx$ 8%), and the rest ($\approx$ 17%) of other religions. Muslims are under-represented in the sample, relative to the population.

Minimal Identity. In this experiment, identities are created through the minimal identity paradigm that has been used for decades by social scientists to study identity effects. Originally conceived by Tajfel et al. (1971), the method involves asking participants to first participate in an innocuous task and then using their responses to classify them into arbitrarily named groups. In experiment *Minimal*, I follow the method used in Chen and Li (2009) to assign minimal identities to participants. First, participants are asked to examine a few pairs of paintings and indicate which painting they preferred in each pair.[3] Next, participants are classified into either the Orange or Purple group based on whether they preferred more Klee (Orange) or Kandinsky (Purple) paintings. Participants are informed that they have been assigned to a group based on their choices, and that others in their group also liked similar paintings. Participants are also reminded of their assigned group identity after the third of the six tasks.

Participants in this experiment were recruited through Prolific (a popular survey platform for social science experiments) and came from the US, UK, and EU.

[3] One of the paintings in each pair was created by the artist Paul Klee, and the other by the artist Wassily Kandinsky. Paintings by these artists have been used in many studies that use this method because of their similarity (at least, to the untrained eye).

Computer. This experiment is a comparison of whether people switch differently when seeing information from a non-social source relative to seeing information from a social source. In this experiment, participants see either an anonymous “Study Participant” or a “Computer”. In treatment *Computer*, participants are told that the estimate is generated by a computer algorithm and that this is correct with a given probability. The phrasing of this text was carefully designed to be as close as possible to the phrasing for the other treatment groups. The comparison group for this treatment are participants pooled from a separate treatment within this experiment (*Choice*, where sources are labelled as “Study Participant”, with everything else identical to the other experiments) and responses from the minimal group experiment.[4]

Participants in this experiment were recruited through Prolific (a popular survey platform for social science experiments) and came from the US, UK, and EU.

3.3 Source Quality

In all experiments except *Caste – No Quality*, participants are given a signal about the quality Q of the source’s estimate. This is the probability with which the source’s estimate is incentive maximising, implying that the estimate lies within ±2% points of the Bayesian posterior. This probability is either $Q \in \{50, 90\}$ in *Caste*, or is $Q \in \{50, 60, 70, 80, 90\}$ in the other experiments.[5] Participants see an inaccurate answer with a $100 - Q$ percentage chance. The quality varies within-subject – the probability is chosen randomly at the task level.

In experiment *Caste*, the following information is shown to participants:

> There is a Q % chance that this number is the correct answer.
>
> The computer has access to a pool of participants who made correct guesses on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess. The computer randomly chooses one of these people and shows you a number:
>
> - With Q% probability, the shown number is the chosen person’s correct guess.
> - Otherwise, the shown number is incorrect.

In all other experiments, the following information is shown to participants:

> There is a Q % chance that this guess is within ±2% points of the correct answer.

[4]This was pre-registered, under the condition that the treatment effects in *Minimal* were minimal.

[5]This difference in the set of options was pre-registered, following an iteration on the design after the first set of studies.

> A computer randomly chooses this guess from the guesses made on this exact task by participants in a previous study - they had the same information, and saw the same ball colours when making their guess.

In the *Computer* treatment, the following information is shown to participants:

> There is a Q % chance that this value is within ±2% points of the correct answer.
>
> Otherwise, the computer chooses a random number between 0 and 100.

Experiment *Caste – No Quality*

To provide a baseline, participants assigned to experiment *No Quality* do not see the quality of the estimate. They only see the caste identity of the source and the estimate. The following information is presented to participants in this treatment instead of the quality information:

> The computer has access to a pool of participants who made guesses on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess. The computer randomly chooses one of these people and shows you the chosen person's guess.

3.4 Study details and procedures

Choice of parameters. Individuals complete six of these tasks, which vary in the parameters. The tasks are pre-defined in the sense that the values of the base rates, urn composition, and the signals (colours of drawn balls) are the same for all participants.[6] The tasks, corresponding parameters, and other details are listed in Table 2. Participants see tasks in a randomised order.[7]

Source Estimates. The estimates shown to participants in the different tasks of the main experiments are displayed in the last two columns of Table 2. The source estimates are generated from the first period responses to the tasks in Table 2 in a separate survey with an Indian sample. The estimates made in the source survey were also incentivised. The estimates that are shown to participants in the main study are chosen from a pool of correct and incorrect responses made by people from different identity groups. I curate the estimates in the following manner:

- For each task, two estimates are selected such that at least one participant from each relevant identity group made that estimate. For example, 51 is chosen on task 1, and at least one participant from the G, O, and Muslim categories made this estimate.

[6]The motivation behind this design choice was to control the task difficulty and distance between shown estimates and true posteriors across tasks and treatment conditions, while staying within budgetary constraints.

[7]In experiment *Caste*, the last of the six tasks is repeated – it is the same as one of the first two tasks faced by a participant to facilitate a test-retest analysis. This is described in more detail in Section 6.

- For each task, one of the estimates is within ±2% of the Bayesian posterior. This is the "correct" answer.
- The second estimate ("incorrect") is chosen such that it is at least 15% points away from the true value.

Table 2: List of balls and urns tasks

#	Base Rate	Red Balls	Total (Red) Draws	True Value	Correct	Incorrect
1	0.7	70	3 (1)	50	51	25
2	0.7	70	5 (2)	50	48	26
3	0.9	70	3 (2)	95.5	95	76
4	0.9	90	3 (0)	1.2	3	23
5	0.9	70	5 (4)	99.1	97	82
6	0.5	70	5 (4)	92.7	92	77

Notes. Base rate is the probability with which the red bag is selected. *Red balls* is the number of red balls in the red bag. *Total (red) draws* is the number of balls that are drawn from the selected bag, with the number of red balls drawn in brackets. *True value* is the Bayesian posterior probability that the red bag was selected. *Correct* and *Incorrect* values for the source in each task are selected from responses on the same tasks in a previous study using the procedure described in Section 3.4.

Incentives – Participation and main tasks. Participants complete six tasks, and they make two decisions in each task. One of the 12 decisions is randomly chosen for a bonus reward. Participants get a $ 3 bonus payment if the chosen decision is within ±2% of the true posterior, else they get no bonus. All participants receive a $ 2 fee for participating in the study. Participants are paid in cash or in "channel points" with a value equivalent to the stated dollar amounts.

Attention and comprehension checks. Participants are only allowed to participate in the incentivised tasks if they can successfully pass multiple attention checks and provide correct answers to four comprehension questions. In the *Religion* and *Caste* experiments, we screen out about 75% of all participants who enter the experiment. In the experiments conducted on Prolific, the pass rate is about 45%.

Additional details. The experiments were programmed using OTree (Chen et al., 2016). The experiments on Prolific were conducted in March 2023. The *Religion* experiment was conducted in collaboration with Faktum Research in April 2023. The *Caste* experiments were conducted in September 2023 in collaboration with Norstat. The experiments were pre-registered at the AEA RCT registry (#0011066 and #0011924). All experiments were reviewed and approved by the IRB at the Norwegian School of Economics. Appendix Table B.1 provides more details on the demographics of the different samples.

3.5 Survey module

Demographics. In all experiments, participants provide their age, sex, education level, and employment status. In the *Caste* and *Religion* experiments, participants also state their religion, caste group, religiosity, and their favourite religious festival (free text). These questions also serve as a mild priming device, drawing attention to these characteristics before the main decision tasks (as in Chen et al. (2014)). Participants are only asked these questions after clearing the attention and comprehension tests.

Reflection question. We ask respondents to type free-form responses to:

> Please tell us how you used the recommendation when making the second decision in the tasks. How did you think about the choice of using your own decision or the shown number?

The belief and reflection questions are presented after participants complete the decision tasks.

Exposure and Attitudes. Participants in the *Caste* experiments respond to three additional survey questions before the end of the study. The first two elicit the extent to which respondents have close associations with people who belong to their religious or caste out-groups. The questions are:

> How many of your friends belong to the same [Religion/Caste category] as you?

The final question elicits people's attitudes towards caste-based affirmative action policies.

> Do you support reservations in jobs and educational institutions based on caste?

Beliefs – Group performance. Participants in the *Caste* experiments state their beliefs about the probability that an anonymous person belonging to a caste or religious group will answer the experimental task correctly. These beliefs are incentivised for accuracy – participants earn an additional \$0.50 if they guess the number within ±5% points of the true probability, which is calculated based on a previous study.

> If a randomly selected individual belonging to the [Caste group] category attempted the Decision task (the task that you just completed). What do you believe is the probability (0% to 100%) that they will answer it correctly? 0% means that they will never get it correct. 100% means that they will always get it correct.

These questions are asked for the "General" and "Scheduled Castes" caste groups.

4 Learning: Aggregate patterns

I start by looking at the descriptive patterns of how people learn from others in the experiments. The main outcome variable is the participants' second-period decision: whether they stick to their first-period decision or switch to the source.

The bars in the Panel (a) of Figure 2 show the fraction of decisions in each of the experiment groups in which individuals choose to switch to the source in the second period. Results are shown by pooling the different treatment conditions within each study (Panel (a) of Appendix Figure A.1 shows this separately for each treatment). The figure shows that participants switch in a substantial fraction of decisions, ranging from $\approx$ 24% in experiment *Religion* to $\approx$ 37% in the experiments conducted with the US/EU sample on Prolific. The triangles in the figure show the fraction of decisions where participants would have earned the incentive if they had switched, instead of sticking to their first-period guess. The triangles are much higher than the fraction of individuals switching, which means that participants leave a lot of money on the table by not switching to the source when it would have benefited them to do so.

Across experiment conditions, participants' task performance is fairly low. Panel (b) of Appendix Figure A.1 shows the fraction of first-period decisions which are incentive maximising, which ranges from $\approx$ 4% in the Religion experiment, to $\approx$ 12% in the Caste experiment. These success rates highlight that participants find the task challenging. At the same time, participant's self-reported confidence in the accuracy of their first estimates is very high which suggests that participants vastly underestimate the difficulty of the experimental task.

Appendix Figure A.2 shows the distribution of the number of switches made by participants in each experiment. In the experiments conducted with an Indian sample, the mode of switching decisions is 0 – about 30% of participants never switch in any of the tasks. In the Prolific experiments, the fraction of participants who never switch is smaller than in the India sample. The patterns are quite different in the *Computer* treatment – here, participants are more likely to switch on the intensive margin as well as the extensive margin.

Subjective certainty. Intuitively, the more confident individuals are about their performance on a given task, the less likely they would be to switch to the source. In all experiments, participants report their certainty in the accuracy of their independent estimate. This is a participant's belief about the probability that their estimate is within ±2% points of the incentive maximising answer.[8] Panel (a) of Figure 3 is a binned scatter plot of certainty in a decision and the magnitude of the error in that decision. Panel (b) of Figure 3 shows a binned scatter plot of certainty and the likelihood of switching to the source. Across exper-

[8]A caveat here is that subjective certainty elicitations could be noisy. However, they provide a sense of how confident a participant is in their own estimate.

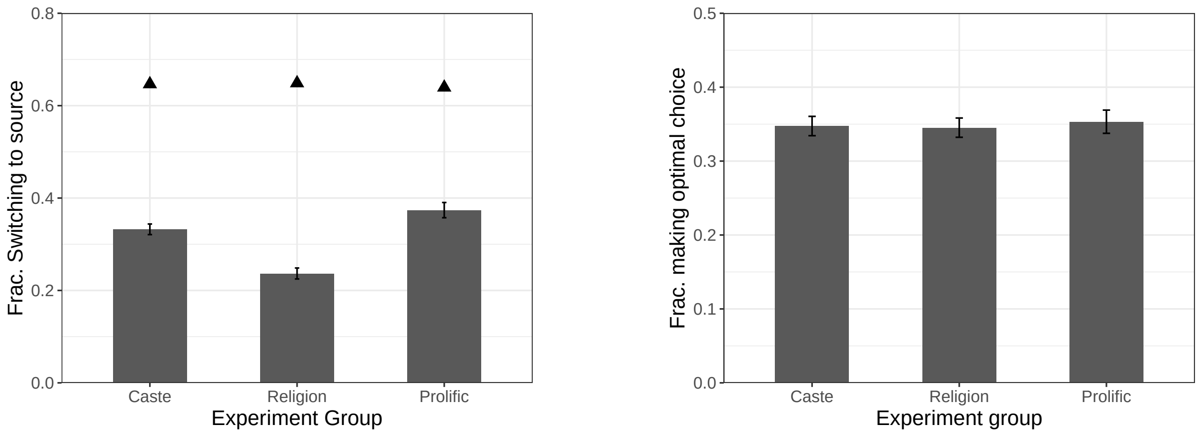

(a) Switching behaviour

(b) Decision optimality

Figure 2: (a) Bars: fraction of decisions where participants choose to switch to the source in each experiment. Triangles: fraction of decisions where switching to y_s would have been incentive maximising. Error bars indicate 95% confidence intervals. (b): Fraction of participants who make optimal decisions, in the sense of comparing their subjective certainty against the quality of the source. Error bars indicate 95% confidence intervals. Results are shown pooling treatment conditions in each study (*G* and *O* in *Caste*, Hindu and Muslim in *Religion*, and *Choice, Minimal* and *Computer* in Prolific).

iments, higher subjective certainty is correlated with more accurate first-period decisions, and with a decreased likelihood of switching to the source in the second period. These patterns indicate that while participants are, on average, over-confident, their behaviour is in line with their subjective confidence.

As motivated in Section 2, a relevant benchmark to evaluate the optimality of participants' decisions is comparing a participant's confidence in the accuracy of their estimate to the quality of the shown estimate. With this definition, an optimal decision is one where participants stick to their own decision if their certainty is higher than the quality of the source, and switch to the source if their certainty is lower than the quality of the source. The right-hand panel of Figure 2 shows the fraction of participants who make such optimal decisions in the different experiments. Across experiments, about $\approx$ 35% of decisions are optimal. The remaining 65% either stick when they should have switched, or switch when they should not have.

5 The role of beliefs

The main goal of the experiments is to identify preferences for the identity of information sources by controlling the role of beliefs in participants' decisions. This section provides evidence that the exogenous manipulation of beliefs about the quality of information influences learning. It also shows that controlling beliefs is important because people believe that identity groups differ in task ability.

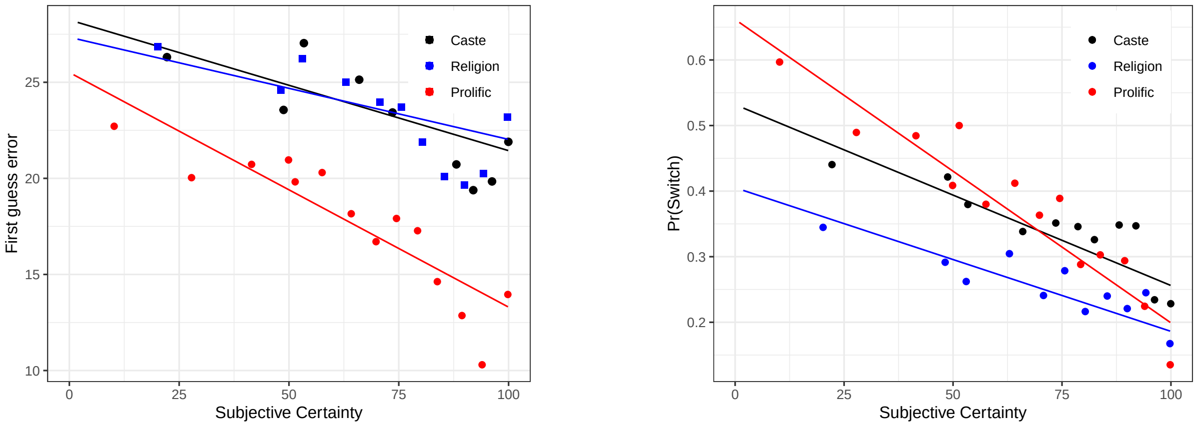

(a)

(b)

Figure 3: Binned scatter plots showing the relationship between (a): the error made on a decision and the participant's subjective certainty in the accuracy of their decision. (b): the likelihood of switching to the source's estimate and subjective certainty. Results are shown pooling treatment conditions in each study (G and O in *Caste*, Hindu and Muslim in *Religion*, and *Choice, Minimal* and *Computer* in Prolific).

5.1 Exogenous beliefs about the quality of information

In most of the experiments in this paper, participants receive a signal about the quality of the estimate that they are shown while making the switching decision. This signal is the exact probability with which switching to the shown estimate will yield the incentive. This is either $\in \{50, 90\}$ in experiment *Caste*, or one of $\in \{50, 60, 70, 80, 90\}$ in the other experiments. The quality varies within treatments and is randomly chosen at the task level (all participants do 6 tasks).

The standard view in economics that information demand is primarily a search for quality would imply that an increase in the quality of information ought to lead to an increase in switching to the source. Figure 4 shows that this is indeed the prevalent pattern in each of the experiments where the quality of information is provided to participants. Participants are much more likely to switch to the shown estimate y_s when the information provided by s is of higher quality. The pattern varies across samples and is more pronounced in the experiments conducted on the US/EU samples on Prolific.

The patterns also show conservatism in learning from others even at high levels of information quality. Switching to the source when the source's estimate has a 90% probability of being correct is quite low – ranging from $\approx 36\%$ to $\approx 53\%$.

5.2 Underlying beliefs about identity groups

People may use the social identity of an information source to form beliefs about the quality of information provided by them. In a separate survey ($N = 327$), I elicit incentivised beliefs about the likelihood that a randomly chosen individual from the G, O, or Muslim categories would make an accurate independent estimate on the decision task. Figure 5

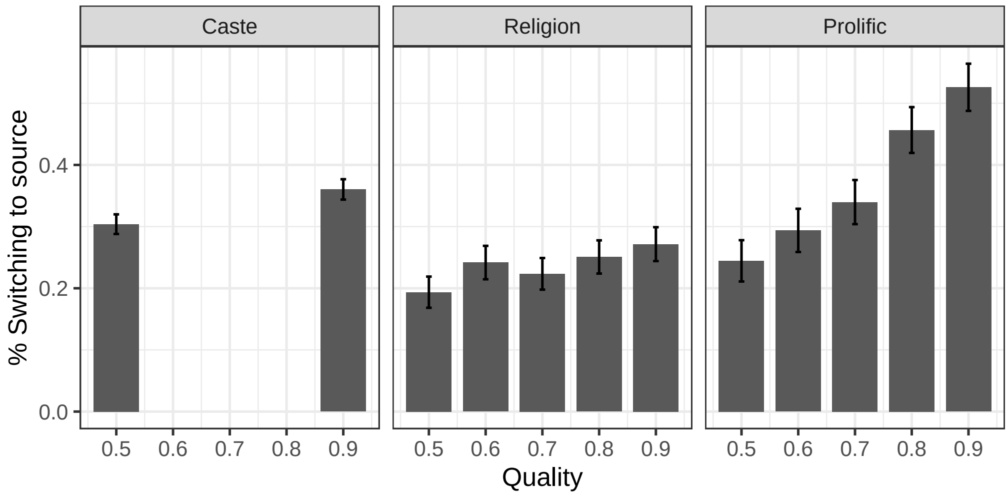

Figure 4: The graph shows the fraction of decisions where participants choose to switch to the source in each experiment, split by the quality of the information. The quality $Q \in \{50, 90\}$ in *Caste*, and $Q \in \{50, 60, 70, 80, 90\}$ in the other experiments. Error bars indicate 95% confidence intervals.

shows the cumulative distribution of these beliefs and indicates that people do indeed have different beliefs about these groups. Panel (a) shows that survey respondents believe that the G category outperforms the O category, and Panel (b) shows that respondents believe that Hindus outperform Muslims on the decision task.

Apart from these differences, another striking fact is that beliefs are substantially misspecified – most participants vastly overestimate the likelihood of success of all identity groups. Further, the levels of these beliefs are very similar across the various groups. For example, high caste status Hindus believe that the likelihood of success is 76.6% for fellow high caste Hindus, 71% for low caste status Hindus, and 61.7% for Muslims. The actual success rates are much lower – $\approx$ 9% for Hindus, and $\approx$ 11% for Muslims.

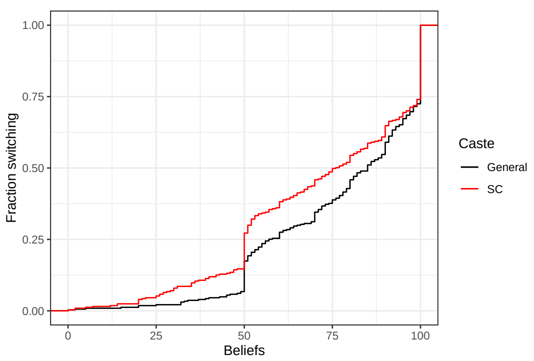

(a) Caste: High vs. low status

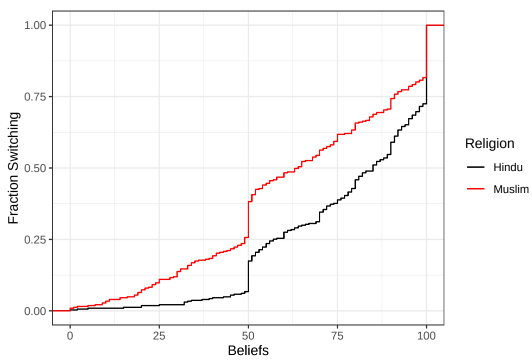

(b) Religion: Hindus vs. Muslims

Figure 5: CDF of beliefs about the likelihood that a person from the G or O caste category succeeds at the task. Panel (a) – Red line: O category, Black line: G category. Panel (b) – Red line: Muslims, Black line: Hindus.

To summarise, these results highlight that (i) beliefs about the quality of information influence switching, (ii) switching is fairly conservative, even when information quality is high, and (iii) people believe that identity groups differ on task ability. The results show the importance of controlling beliefs to identify preferences for the identity of information sources.

6 The role of preferences

This section presents results on whether people have preferences for the identity of information sources. The section begins with an analysis of the average treatment effects, proceeds to examine heterogeneity by the religious or caste identity of the participant, and concludes with a battery of robustness tests and additional heterogeneity analyses.

6.1 Preferences for the identity of information sources

The experiments are designed to identify whether people prefer to learn more from information sources belonging to one social group over another. I estimate the following specification using OLS regressions to identify preferences for the identity of information sources:

$$Switch_{i,n} = \beta_0 + \beta_1 T_i + Q_{i,n} + \gamma_i + \nu_n + \epsilon_i, \tag{1}$$

i is a participant who decides whether to switch in task n. T_i is an indicator variable for the identity treatment group to which participant i is assigned to in a given experiment. In the *Caste* experiment, $T = 1$ if the participant is assigned to treatment group G, and $T = 0$ if the participant is assigned to treatment group O. In the *Religion* experiment, $T = 1$ if the participant is assigned to treatment group H, and $T = 0$ if the participant is assigned to treatment group M. In the *Minimal* experiment, $T = 1$ if the participant is assigned to the treatment group where they see in-group sources, and 0 if they see out-group sources. In the *Computer* treatment, $T = 1$ if the participant sees information from a Computer, and $T = 0$ if the source is an anonymous human. $Q_{i,n}$ is the quality of the information source seen by participant i in task n. γ_i is a set of demographic controls (age, gender, tertiary education, and employment status). ν_n are controls for task order and individual tasks. Standard errors are clustered at the participant level. The regressions also control for participants' subjective certainty of the accuracy of their first guess.

Figure 6 presents the results from estimating this equation for each of the four main experiments separately (the regression results are presented in Appendix Table B.2, and detailed regression analyses are presented in Appendix Tables B.4–B.7). The figure shows the estimated coefficients and 95% confidence intervals for the treatment dummy variable T for each of the experiments. The estimated average treatment effect in the *Caste*, *Religion*, and *Minimal* experiment groups are statistically indistinguishable from zero. In other words, participants do not switch differently whether the information comes from a G caste or *O* caste group individual in the *Caste* experiment, a Hindu or a Muslim in the *Religion* experiment, or their experimentally assigned in-group member or an out-group member in the *Minimal* experiment.

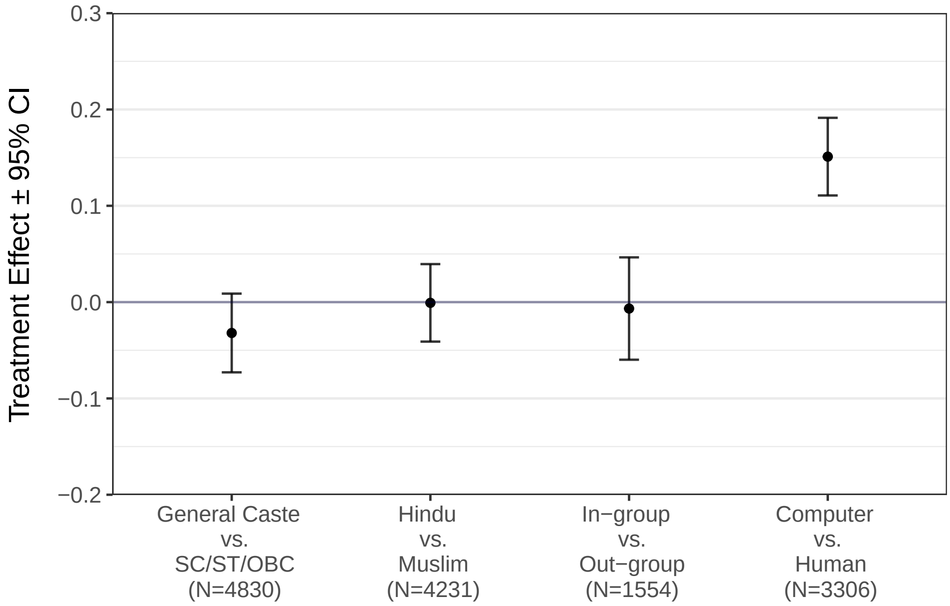

Figure 6: Estimated average treatment effects in each experiment. Each point is the estimated coefficient for the treatment dummy variable. The estimated specification is Equation 1, and standard errors are clustered at the participant level. Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source. Table B.2 presents the corresponding regression tables.

Robustness tests and alternate mechanisms – presented in Section 6.3 – show that these null results cannot be explained by overconfidence, experimenter demand, motivated reasoning, or beliefs. The experiments are also well-powered and the 95% confidence intervals are small compared to effect sizes found in other experiments on learning and group or social identity.

Taken together, the lack of an economically or statistically significant treatment effect in the experiments where people have the option of learning from other humans supports the interpretation that people may not have preferences for whom they get information from when the quality of information is precisely known. On the other hand, the estimated treatment effect is large and positive for the *Human vs. Computer* experiment (the right-most estimate presented in Figure 6), indicating that participants respond much more strongly to information that comes from a computer or algorithm, relative to an anonymous individual. A detailed investigation of this result is presented in Section 8.

6.2 Heterogeneity: Decision maker identity

6.2.1 In-group preferences

Next, I study whether there are in-group preferences for the identity of information sources. A person's identity relative to the identity of the source of information could be an important determinant of such preferences. Indeed, a large body of work has found that in-group favouritism affects social preferences in many contexts (Shayo, 2020; Charness and Chen, 2020). The data allows for an investigation of whether the overlap between a participant's identity and the (exogenously assigned) identity of the information source matters. Since we elicit people's religion and caste, we can study whether people prefer to get information from their caste or religious in-groups than from their out-groups. For example, G caste participants may prefer to get information from G caste sources rather than from O caste sources.

I estimate the following specification using OLS regressions to identify the existence of in-group preferences for the identity of information sources:

$$Switch_{i,n} = \beta_0 + \beta_1 T_i + \beta_2 C_i + \beta_3 T_i \times C_i + +Q_{i,n} + \gamma_i + \mu_n + \epsilon_i, \tag{2}$$

Here, $C = 1$ if the identity of the participant is the same as the identity for which $T = 1$ in the relevant experiment. In experiment *Caste*, $C = 1$ if the participant belongs to the general category (G) and $C = 0$ if the participant belongs to the SC/ST/OBC categories. In experiment *Hindu*, $C = 1$ if the participant is Hindu and $C = 0$ if the participant is Muslim.

The coefficient β_1 is the treatment effect of seeing a $T = 1$ source relative to a $T = 0$ source for $C = 0$ participants. $\beta_1 + \beta_3$ is the treatment effect of seeing a $T = 1$ source relative to a $T = 0$ source for $T = 1$ group participants. The interaction coefficient β_3 is the difference in the treatment effects between $T = 0$ and $T = 1$ identity group participants. The main (pre-registered) hypothesis here is that $\beta_3 > 0$, which would identify that participants prefer to get information from their in-groups than from their out-groups.

Table 3 presents the results from estimating equation 2 on different samples. Columns (i)-(ii) show results for the *Caste* experiment, and Columns (iii)–(iv) show results for the *Religion* experiment. The estimated coefficients on the interaction terms are not statistically significant (either individually, or jointly). Therefore, there is no evidence that participants belonging to different caste categories or religions have in-group preferences for information from their caste or religion in-groups. The results are robust to the inclusion of a host of demographic controls, and to controlling for participant's stated subjective certainty.

Table 3: Regression analysis: In-group preferences

	Switch to source			
	Caste		Religion	
	(i)	(ii)	(iii)	(iv)
G Source × G Participant	-0.045	-0.034		
	(0.045)	(0.044)		
G Source	0.003	-0.010		
	(0.036)	(0.036)		
Hindu Source × Hindu Participant			-0.088	-0.092
			(0.071)	(0.070)
Hindu Source			0.075	0.081
			(0.068)	(0.067)
G Participant	0.008	0.001		
	(0.032)	(0.032)		
Hindu Participant			-0.005	0.003
			(0.043)	(0.042)
Quality	0.061***	0.060***	0.138***	0.134***
	(0.015)	(0.014)	(0.047)	(0.047)
Certainty		-0.003***		-0.002***
		(0.000)		(0.000)
Constant	0.300***	0.484***	0.140***	0.329***
	(0.027)	(0.069)	(0.052)	(0.083)
Task controls		✓		✓
Demog. controls		✓		✓
R^2	0.006	0.035	0.004	0.021
Observations	4,830	4,830	4,231	4,231
Clusters	851	851	733	733

Notes. The table presents estimates from OLS regressions of Equation 2. The dependent variable is 1 if the participant switches to the source's estimate y_s. Columns (i)-(ii) present estimates from experiment *Caste*, (ii)-(iv) from experiment *Religion*. $G = 1$ if the participant/source is from the General caste category, and 0 if the participant/source is from the SC/ST/OBC categories. Hindu source/Hindu participant indicates that the source or participant is Hindu, else they are Muslim. The samples are restricted to the identity groups mentioned above. Task controls include order and problem fixed effects. Demographic controls are age, education, employment, and gender. Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

6.2.2 Multiple identities – Religion and caste

The *Religion* experiment allows for the study of the intersection of religious and caste identity. The surnames of the sources in treatment group *Hindu* indicate both their religion and their caste category. All Hindu surnames used in the experiment belong to the General caste category (these are the same surnames as used in the *G* treatment group in the *Caste* experiment). Therefore, Hindus belonging to the *G* category and the *O* category may respond differently to *G - Hindu* sources, relative to Muslim sources.

I estimate the specification in Equation 2 using OLS on the full sample, as well as on caste category sub-samples for Hindu participants in the experiment. Table 4 presents the results and shows that participants belonging to the General category do not switch differently based on the religious identity of the source. However, participants belonging to the SC/ST/OBC categories switch less often when they see a *G* category Hindu source than when they see a Muslim source. The effect is substantial and is roughly 40% of the mean. This heterogeneity result suggests that preferences for the identity of information sources may exist for (or between) certain identity groups. In light of the relatively low statistical power and limited causal interpretability of this result, I interpret this as weak evidence for the existence of preferences for the identity of information sources.

Table 4: Regression analysis: Heterogeneous effects by caste category

	Switch to source			
	All Hindus	SC/ST/OBC	General	Muslims
	(i)	(ii)	(iii)	(iv)
Hindu Source	-0.104***	-0.106***	0.040	0.067
	(0.038)	(0.039)	(0.026)	(0.063)
Hindu Source × General caste	0.144***			
	(0.046)			
Quality	0.136***	0.164**	0.124**	0.154
	(0.048)	(0.081)	(0.060)	(0.159)
Certainty	-0.002***	-0.001*	-0.002***	-0.003*
	(0.000)	(0.001)	(0.000)	(0.002)
Constant	0.391***	0.332**	0.329***	0.181
	(0.083)	(0.132)	(0.100)	(0.236)
General caste	-0.103***			0.016
	(0.035)			(0.066)
Task controls	✓	✓	✓	✓
Demog. controls	✓	✓	✓	✓
R^2	0.027	0.029	0.032	0.059
Dependent variable mean	0.226	0.226	0.226	0.273
Observations	3,743	1,334	2,409	488
Clusters	647	231	416	86

Notes. The table presents estimates from OLS regressions of Equation 2. The dependent variable is 1 if the participant switches to the source's estimate y_s. *G* source/participant = 1 when the source/participant belongs to the General caste category and is 0 otherwise. Hindu source/participant = 1 when the source/participant is Hindu. Quality is a continuous variable which indicates the percentage chance with which the source's estimate is correct. Demographic controls are the participant's age and dummies for employment, college education and gender. Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

6.3 Robustness of the main results

Noisy responses. Given the nature of the task and the fact that the India samples are not typically exposed to incentivised economic experiments (as is the case with participants from Prolific), a relevant concern is that the results may be driven by confusion or inattention to the main tasks or instructions. By design, the experiment minimises the possibility of a lack of understanding of the requirements of the task, the incentive structure, or the nature of the decision context in driving the results. All participants are required to pass an extensive training module and screening tests. These tests screen out nearly 80% of the participants in the India samples. In the Prolific samples, the screening rate is around 40%. The nature of the experiment task allows for further tests of whether the main results are driven by noisy responses.

Figure 7 shows the results from a battery of robustness tests for the *Caste* and *Religion* experiments. These tests use features of participant's responses that are indicative of reduced levels of noisiness, fatigue, or inattention. First, participants may be more motivated, energetic, or likely to remember instructions at the beginning of the experiment. Restricting the sample to the first two (of six) tasks does not make any difference to the estimated treatment effect.

Next, completing tasks very quickly or very slowly may reflect inattention or distraction. I find that excluding decisions that are in the top and bottom 20% of time taken while making the switching decision does not affect the estimates.

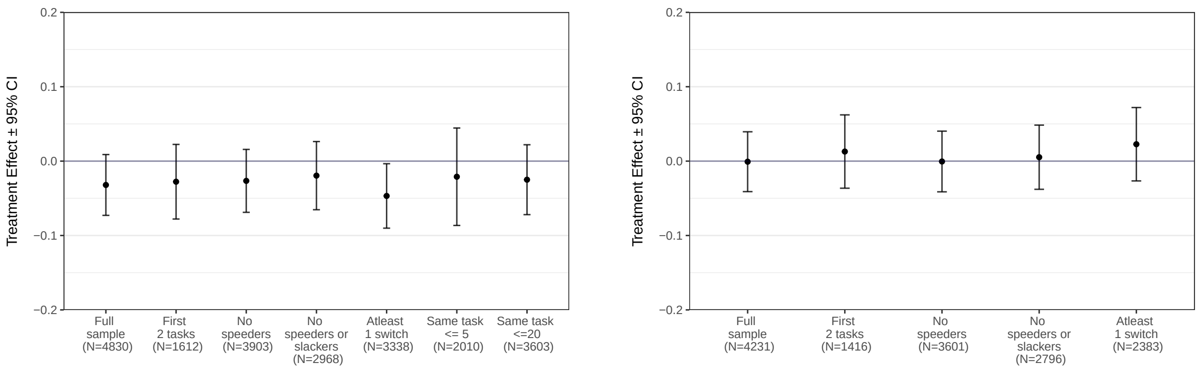

(a) Caste experiment

(b) Religion experiment

Figure 7: Robustness tests – Each point is the estimated coefficient on the treatment dummy variable. Error bars indicate 95% confidence intervals. The estimated specification is Equation 1, and standard errors are clustered at the participant level. Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source. Table B.9-B.10 presents the regression tables used to construct this plot.

Next, participants who do not switch in any of the tasks may be less likely to be engaged when completing the tasks. I restrict the sample to participants who switch at least once, and I find that the treatment effect is just about statistically significant at the 5% level. However, this is the only robustness test that finds a statistically significant result, and the similarity of the effect size on this particular sub-sample does not provide strong support for the existence of a general preference for learning from a particular caste group.

Finally, I use an additional feature of experiment *Caste* to test whether attention or engagement drives the results. Participants complete six tasks, and the last task that each participant faces is the same as one of the first 3 tasks. Two analyses based on this feature lend support to the robustness of the main results. First, the correlation between the responses in the identical tasks is 0.62, which is quite strong. This indicates that participants behave similarly – by making similar independent estimates – in the same task, which occur at the beginning and end of the task sequence. A low correlation in the independent guesses (y_1) between two iterations of the same task would indicate noisiness. Second, the estimated treatment effect is not distinguishable from zero for participants whose estimates in identical tasks are more similar. It is plausible that participants who are more consistent in their choices may behave differently than those who are less focused and hence less likely to make consistent estimates. However, the results show that there is no difference between these types of participants – the last two points in Figure 7 show that the estimates for subsamples where this difference between the two estimates on the same task is ≤ 5 and ≤ 20 respectively. These estimates are very similar to those with the full sample.

Figure A.3 presents similar robustness tests for in-group preferences. The figure shows the coefficient on the interaction term when estimating Equation 2 within experiment *Caste*. The pattern is consistent and replicates the main result that the estimated treatment effect does not differ by the caste identity of the participant.

Experimenter demand effects. Experimenter demand is often an issue in online experiments, and I implement several measures to minimise these effects. First, the within-subjects design means that participants are unaware of the other treatment variations. Participants in the caste and religion experiments would not be surprised upon seeing G caste and $Hindu$ names in an online survey environment as these populations are overrepresented, relative to the population, in online settings. Therefore, they are unlikely to anticipate a treatment where the sources are exclusively O caste or $Muslim$.

Second, if experimenter demand were a factor, it would be far more surprising to see six consecutive O caste category names than it would be to see six G category names. This would mean that the decisions in the O treatments would be more susceptible to experimenter demand – participants would be more likely to switch more when seeing an O type source, relative to G type sources. However, the data shows that this is not the case: Appendix Figure A.5 shows the fraction of individuals switching within each treatment (the identity of the source, G or O) by the respondent's caste group (G or O). From the figure, we can see that G type participants switch slightly *less* when seeing a G source, than when they see an O source. These patterns are consistent with a limited impact of experimenter demand.

Biases in probabilistic reasoning. A large literature (reviewed in Benjamin (2019)) has shown that belief updating in the balls-and-urns paradigm is subject to a variety of systematic biases. People update differently if they receive information that confirms (or disconfirms) their original estimate, if they are more (or less) confident, or if they are more (or less) rational (in terms of applying Bayes' rule). Given this body of evidence, a pertinent concern is that the null treatment effect might conceal significant systematic variation along some of these dimensions. For example, preferences for the identity of information sources may be different for confirmation-seeking or overconfident individuals.

Figure 8 presents estimates of the average treatment effect from regressions that slice the data along various dimensions that are the source of systematic biases in belief updating. The first panel splits the sample into quartiles based on their reported certainty in their independent guess (with higher quartiles being more confident). The second panel splits the sample by the distance of the source estimate from the independent guess (lower quartiles are more likely to exhibit confirmation seeking behaviour). The third panel splits the sample by the error made on the first guess, with higher quartiles representing more inaccurate decisions. The figure shows that there is no evidence of heterogeneous treatment effects along (i) participants' subjective certainty in their first period guess, (ii) the distance between their first period guess and the source's estimate, and (iii) the error made on the first guess. This set of results supports the evidence pointing towards the lack of a preference channel.

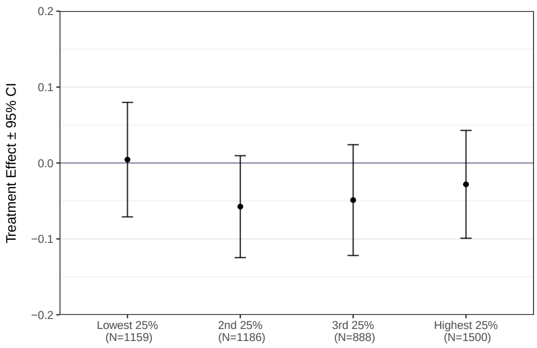

(a) Subjective Certainty (quantiles)

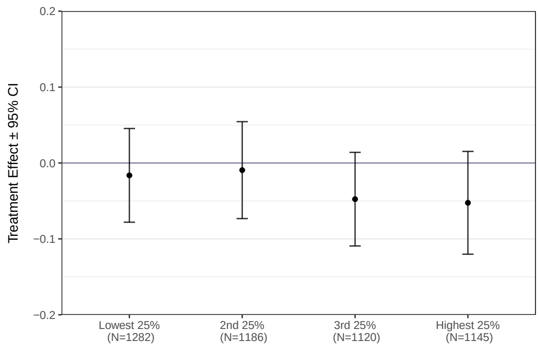

(b) Distance to y_s (quantiles)

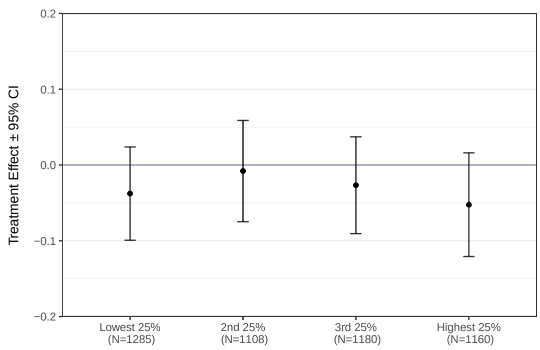

(c) First guess error (quantiles)

Figure 8: Heterogeneity analyses in the *Caste* experiment. Panels show regressions on different subsamples. Figures (a)-(c) show regressions for each quantile of first guess certainty, the absolute distance between the y_1 and y_s, and first guess error ($|y_1 - y_T|$). Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source. Standard errors are clustered at the participant level in all regressions.

Further heterogeneity analyses studying behaviour by (a) exposure to different caste and religious groups, (b) religiosity, and (c) individualism are presented in Appendix C.

6.4 Learning when quality is unknown

The main results show that while there are perceived differences in the abilities of different identity groups, there is no evidence for preferences for particular identity groups when the quality of information is exogenously provided. I conduct an additional experiment to study behaviour when the identity of the information source is salient but the quality of information is unknown. In experiment *Caste – No Quality* (as in experiment *Caste*) participants are assigned to see sources belonging to either the G or O caste category. The key difference in this experiment is that participants are not provided any information about the quality of y_s. In other words, the decision to switch will be influenced by participants' group-level beliefs, which as we have already seen implies that they may be more likely to listen to G than O sources. On the other hand, the demonstrated conservatism in switching implies that the effect of belief differences is likely to be relatively small.

The binned scatter plots in Figure 9 show the correlations between whether participants switch to a source belonging to a particular caste group against their beliefs about the likelihood that a person from that caste group succeeds at the task. The left panel of Figure 9 shows the correlations in experiment *Caste – No Quality*, and shows that beliefs about the group that a source belongs to are correlated with the decision to switch to a source belonging to the same group. The right-hand panel of Figure 9 shows these correlations in experiment *Caste* where participants are additionally given information about the quality of the source's estimate. Appendix Table B.12 shows this in the form of regressions.

The raw correlation between beliefs about the performance of a given caste group and switching to a source from that caste group is ≈ 0.09. The correlation increases to ≈ 0.14 when looking at participants who report being subjectively uncertain in the accuracy of their first guess. This suggests that participants do indeed rely on underlying beliefs about groups to make the stick-or-switch decision. Additionally, when the quality of the information is provided (experiment *Caste*), these correlations mostly vanish – ≈ 0.02 unconditional, and ≈ 0.04 for the low certainty participants. Strikingly, the increase in switching when beliefs about the success of a particular caste group moves from 50% to 90% is approximately the same as the difference in switching when beliefs about quality are exogenously varied by the same amount.

Treatment effects. Figure 10 shows the results from estimating the specification in Equation 1 for this experiment. The figure shows the estimated coefficients of the treatment dummy variable using the full sample, and sub-samples based on the participant's caste group. The results show that the estimated treatment effect is not statistically significant – participants do not appear to switch differently whether they see a G source or an O source. The results can be explained by the difference in magnitude between the difference in beliefs about the performance of different groups and the responsiveness to signals about the

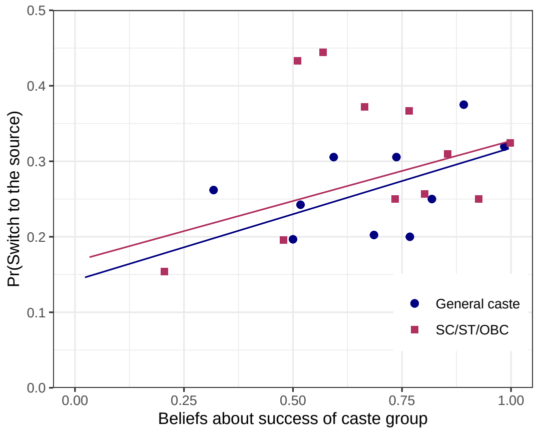

(a) Quality of estimate is unknown

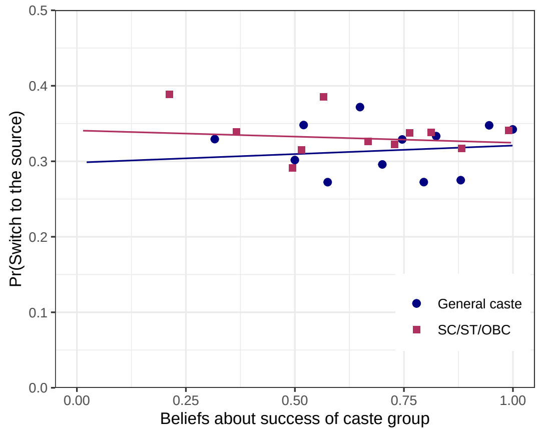

(b) Quality of estimate is known

Figure 9: Binned scatter plots of the likelihood of switching to a source belonging to the G or O caste categories against incentivised beliefs about the likelihood that a person from the G or O caste group is likely to be correct on the task. (a): Participants are not informed about the quality of the source's estimate. (b): Participants are given precise information about the quality of the source's estimate.

quality of information. Participants' responsiveness to the quality of information is relatively muted – a 40% increase in quality leads to ≈ 10% increase in switching in the *Caste* experiment. In comparison, the magnitude of participants' beliefs about the differences in ability between the caste groups is much smaller (about $5-10\%$).[9]

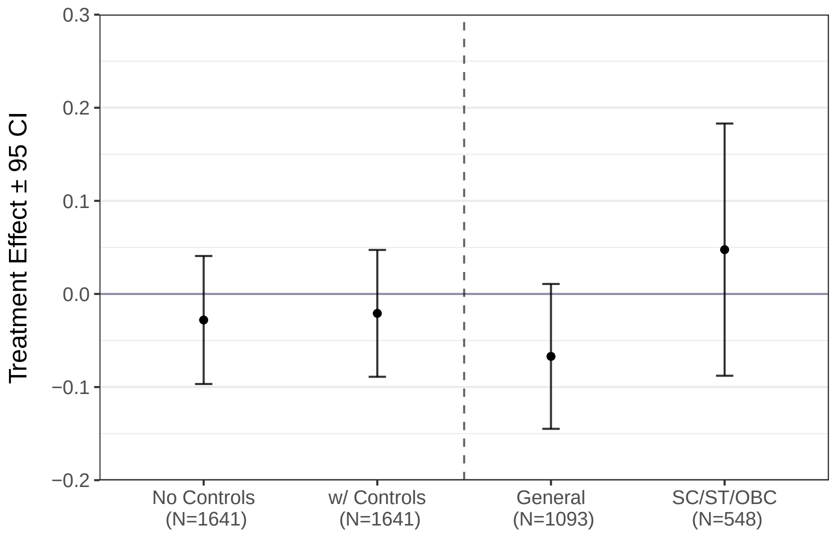

Figure 10: Average treatment effects in experiment *Caste – No Quality*. Each point is the estimated coefficient for the treatment dummy variable, and indicates the difference in the likelihood of switching to the shown estimate when the caste identity of the source is G, relative to O. The first two points show estimates from the full sample, and the next two show estimates restricting the sample by the participant's caste category. Standard errors are clustered at the participant level. Error bars indicate 95% confidence intervals. Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source.

[9]Experimenter demand effects in this experiment can be ruled out – Appendix Figure A.4 shows that both G and O category participants are equally likely to switch when seeing an O source. However, when seeing a G source, G category participants are less likely to switch to it than O category participants.

7 Identity and sensitivity to quality

Section 5 presented evidence that participants switch more when the quality of the information provided by the source increases. A pertinent question is whether the extent to which people are sensitive to quality depends upon the identity of the information source. Since the quality of information is randomly assigned within subjects (at the task level), it is possible to causally estimate whether the sensitivity to the quality of information depends on the identity of the information source using the following specification:

$$Switch_{i,n} = \beta_0 + \beta_1 T_i + \beta_2 Q_{i,n} + \beta_3 T_i \times Q_{i,n} + \gamma_i + \nu_n + \epsilon_i, \tag{3}$$

Where $Q_{i,n}$ is the quality of the information y_s seen by participant i in task n. The interaction term β_3 can be interpreted as the difference in the responsiveness to information quality between identity groups in a given treatment. A non-zero coefficient would mean that preferences for the identity of information sources affect their sensitivity to an increase in the quality of the information source.

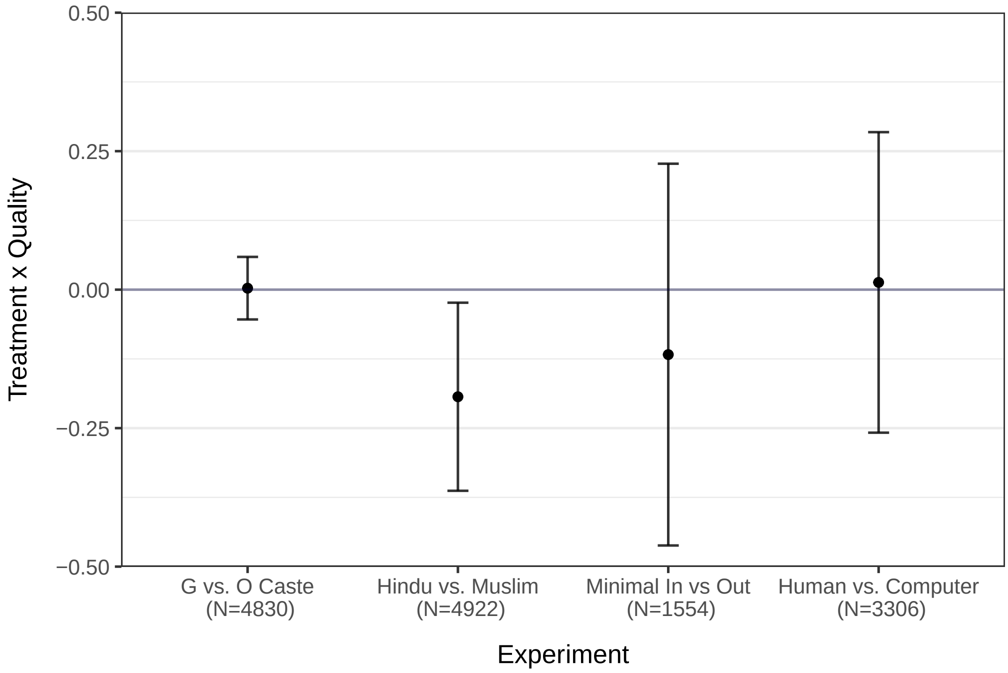

Figure 11: Estimated coefficients of the interaction term Treatment × Quality. Each point indicates whether sensitivity to the quality of the source's estimate varies depending on the identity of the source and is estimated separately for each experiment group. Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source. Error bars indicate 95% confidence intervals. Standard errors are clustered at the participant level

Figure 11 presents the results from estimating this equation for the experiments where quality and identity are both manipulated. In experiments *Caste, Minimal*, and *Computer*, there is no evidence that participants are differentially sensitive to the quality of information based on the identity of the source. In experiment *Religion*, participants appear to be more sensitive to the quality of Muslim sources than of Hindu sources. Figure A.6 further shows that the difference in sensitivity based on the religion of the source is driven by General category Hindus – SC/ST/OBC Hindus do not appear to be differentially sensitive based on the source's religion.

To summarise, these results provide only weak evidence of a conditional preference for the identity of information sources and support the interpretation of a null result for the existence of preferences for the identity of whom people learn from.

8 Social vs. non-social learning

Across experiments, the results provide only weak evidence for preferences for the identity of information sources. However, the results also show that people do not learn all that often even when high-quality information is available. Is there something particular about learning from another individual that drives the reluctance to learn from others' decisions? Experiment *Computer* investigates this question by reframing the information source as the output of a computer algorithm that displays the correct answer with a known probability else it shows a random number. The probability of the answer being correct is the same as the quality of the sources in the other experiments.

Figure 6 in Section 6 shows a substantial treatment effect in the *Computer* experiment. The magnitude of the effect (0.151% points) is quantitatively large, about a third of the fraction switching in the *Human* condition. Figure 12 further shows that this is a level effect as the fraction of people who switch is higher in the *Computer* treatment at all levels of source quality. The figure also shows that while participants switch a lot more when getting information from a non-social source, most participants could benefit from switching more. Using a simple benchmark of decision optimality – where participants who receive a signal that the quality of the source's estimate is higher than their subjective certainty in their independent estimate should switch – shows that while participants are more optimal when information is provided by a non-social source, there is still a substantial gap between participant's behaviour and optimal switching. Appendix Figure A.8 shows similar patterns when conditioning on decisions where the realisation of the shown estimate was objectively correct.

Taken together, the results indicate that participants (i) switch more often, (ii) more optimally, and (iii) at all levels of information quality when they get information through a non-social source rather than from a social source. These results are surprising since participants know the quality of information in both conditions. They provide direct causal evidence of algorithmic appreciation from a relatively abstract setting where participants only care about maximising their earnings.

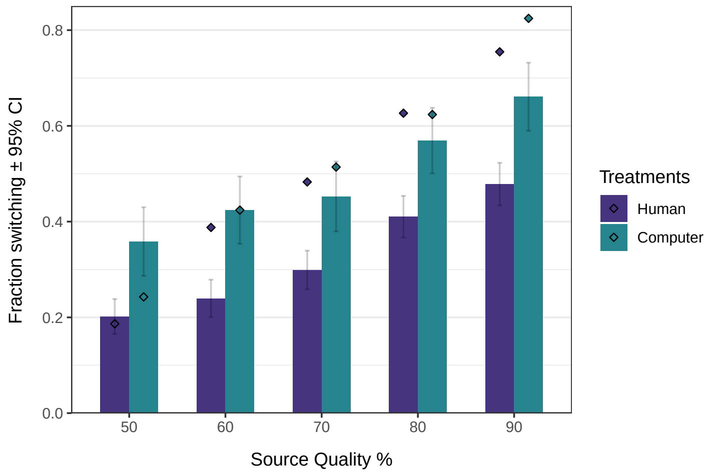

Figure 12: Fraction of decisions where participants choose to switch to the source in treatments *Human* and *Computer*, by the quality of information. The diamonds indicate the fraction of decisions where switching would have been the optimal choice (based on a comparison of information quality and subjective certainty). Error bars indicate 95% confidence intervals.

9 Conclusion

Social learning influences how people make consumption and investment decisions, and how they learn about social norms. Understanding how social learning works is important, and the main contribution of this paper lies in providing rich evidence on how people choose whom to learn from. The results presented in this paper show that while beliefs about the quality of information or the abilities of the provider of information matter, identity-based preferences do not appear to play a significant role when people learn from others.

A limitation of this paper is the abstract nature of the experiment. This is an important consideration as there are many other channels through which identity can affect information acquisition. The experiment shuts down many of these channels to isolate the role of preferences for identity in a stylised setting where there is no possibility of a prolonged association with out-group members. Yet, the experiment design nests a wide variety of decisions that people make every day – we read newspaper articles, read online reviews, or watch news on the TV which may be delivered by people belonging to different social groups. Future research can adapt the experiment introduced in this paper to study learning in settings where channels such as motivated reasoning and image concerns play a role.

The absence of evidence for the preference channel stands in contrast with the existence of preference-based discrimination along caste and religious lines. This has been documented in a variety of settings: In consumption behaviour (Atkin et al., 2021), health care choices (Islam et al., 2023), and in labour markets (Siddique, 2011). Therefore, being able to rule out a role for this channel in how people learn from others has important implications for how governments and organisations structure their communications and outreach efforts. People are often unresponsive to information provision in many settings (Haaland et al., 2023). Focusing on the quality of information rather than on the identity of the messenger, and perhaps delivering information through non-social channels may have better results in encouraging people to take up social welfare programs, adopt better health practices, and drive changes in social norms.

References

Akerlof, George A. and Rachel E. Kranton, "Economics and Identity," *Quarterly Journal of Economics*, aug 2000, *115* (3), 715–753.

Anderson, Lisa R and Charles A Holt, "Information cascades in the laboratory," *The American economic review*, 1997, pp. 847–862.

Arrow, Kenneth J, "The Theory of Discrimination, Discrimination in Labor Markets, Ashenfelter, O. and A. Rees eds., 3-33," 1973.

Atkin, David, Eve Colson-Sihra, and Moses Shayo, "How Do We Choose Our Identity? A Revealed Preference Approach Using Food Consumption," *Journal of Political Economy*, apr 2021, *129* (4), 1193–1251.

Ayalew, Shibiru, Shanthi Manian, and Ketki Sheth, "Discrimination from below: Experimental evidence from Ethiopia," *Journal of Development Economics*, jun 2021, *151*, 102653.

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune, "Marry for What? Caste and Mate Selection in Modern India," *American Economic Journal: Microeconomics*, may 2013, *5* (2), 33–72.

Barron, Kai and Tilman Fries, "Narrative Persuasion," *SSRN Electronic Journal*, 2023.

Becker, Gary S., "The economics of discrimination," *University of Chicago Press*, 1957.

Benjamin, Daniel J., "Errors in probabilistic reasoning and judgment biases," in "Handbook of Behavioral Economics - Foundations and Applications 2," Elsevier, 2019, pp. 69–186.

BenYishay, Ariel and A Mushfiq Mobarak, "Social Learning and Incentives for Experimentation and Communication," *The Review of Economic Studies*, jul 2018, *86* (3), 976–1009.

Berger, Sebastian, Christoph Feldhaus, and Axel Ockenfels, "A shared identity promotes herding in an information cascade game," *Journal of the Economic Science Association*, may 2018, *4* (1), 63–72.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg, "The Dynamics of Discrimination: Theory and Evidence," *American Economic Review*, oct 2019, *109* (10), 3395–3436.

__ , Kareem Haggag, Alex Imas, and Devin G. Pope, "Inaccurate Statistical Discrimination: An Identification Problem," *Review of Economics and Statistics*, sep 2023, pp. 1–45.

Boyd, Robert, Peter J. Richerson, and Joseph Henrich, "The cultural niche: Why social learning is essential for human adaptation," *Proceedings of the National Academy of Sciences*, jun 2011, *108* (supplement_2), 10918–10925.

Cassan, G, D Keniston, and T Kleineberg, "A division of laborers: Identity and efficiency in India," Technical Report feb 2019.

Chan, Alex, "Discrimination and Quality Signals: A Field Experiment with Healthcare Shoppers," *Unpublished manuscript*, 2022.

Charness, Gary and Yan Chen, "Social Identity, Group Behavior, and Teams," *Annual Review of Economics*, aug 2020, *12* (1), 691–713.

Chen, Daniel L., Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, mar 2016, *9*, 88–97.

Chen, Yan and Sherry Xin Li, "Group identity and social preferences," *American Economic Review*, 2009, *99* (1), 431–57.

_ , _ , Tracy Xiao Liu, and Margaret Shih, "Which hat to wear? Impact of natural identities on coordination and cooperation," *Games and Economic Behavior*, 2014, *84*, 58–86.

Conlon, John J, Malavika Mani, Gautam Rao, Matthew Ridley, and Frank Schilbach, "Failing to Learn from Others," 2022.

Dekel, Inbal and Moses Shayo, "Follow the Crowd: But Who Follows, Who Counteracts, and Which Crowd?," *Who Counteracts, and Which Crowd*, 2023.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err.," *Journal of Experimental Psychology: General*, 2015, *144* (1), 114–126.

_ , _ , and _ , "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science*, mar 2018, *64* (3), 1155–1170.

Enke, Benjamin and Thomas Graeber, "Cognitive Uncertainty," *The Quarterly Journal of Economics*, may 2023, *138* (4), 2021–2067.

Gallen, Yana and Melanie Wasserman, "Does information affect homophily?," *Journal of Public Economics*, jun 2023, *222*, 104876.

Ghosh, Arkadev, "Religious Divisions and Production Technology: Experimental Evidence from India," *SSRN Electronic Journal*, 2022.

Goeree, Jacob K. and Leeat Yariv, "Conformity in the lab," *Journal of the Economic Science Association*, feb 2015, *1* (1), 15–28.

Graeber, Thomas, Florian Zimmermann, and Christopher Roth, "Stories, Statistics, and Memory," *SSRN Electronic Journal*, 2022.

Guryan, Jonathan and Kerwin Kofi Charles, "Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots," *The Economic Journal*, nov 2013, *123* (572), F417–F432.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart, "Designing Information Provision Experiments," *Journal of Economic Literature*, mar 2023, *61* (1), 3–40.

Henrich, Joseph, *The Secret of Our Success*, Princeton University Press, dec 2016.

Hou, Yoyo Tsung-Yu and Malte F. Jung, "Who is the Expert? Reconciling Algorithm Aversion and Algorithm Appreciation in AI-Supported Decision Making," *Proceedings of the ACM on Human-Computer Interaction*, oct 2021, *5* (CSCW2), 1–25.

Islam, Asad, Debayan Pakrashi, Liang Choon Wang, and Yves Zenou, "Determining the Extent of Taste-Based and Statistical Discrimination: Evidence from a Field Experiment in India," 2023.

Jaffrelot, Christophe, *Modi's India: Hindu nationalism and the rise of ethnic democracy*, Princeton University Press, 2021.

Levy, Gilat and Ronny Razin, "Echo Chambers and Their Effects on Economic and Political Outcomes," *Annual Review of Economics*, aug 2019, *11* (1), 303–328.

Logg, Jennifer M, Julia A Minson, and Don A Moore, "Algorithm appreciation: People prefer algorithmic to human judgment," *Organizational Behavior and Human Decision Processes*, 2019, *151*, 90–103.

Macours, Karen, "Heterogeneous quality of agricultural commercial inputs and learning through experimentation," feb 2014.

Mobius, Markus and Tanya Rosenblat, "Social Learning in Economics," *Annual Review of Economics*, aug 2014, *6* (1), 827–847.

Mosse, David, "The Modernity of Caste and the Market Economy," *Modern Asian Studies*, oct 2019, *54* (4), 1225–1271.

Munshi, Kaivan, "Caste and the Indian Economy," *Journal of Economic Literature*, dec 2019, *57* (4), 781–834.

Oh, Suanna, "Does Identity Affect Labor Supply?," *American Economic Review*, aug 2023, *113* (8), 2055–2083.

Parys, Jessica Van and Elliott Ash, "Sequential decision-making with group identity," *Journal of Economic Psychology*, dec 2018, *69*, 1–18.

Pew Research Center, "Religion in India: tolerance and segregation," in "Pew Forum" 2021.

Phelps, Edmund S, "The statistical theory of racism and sexism," *The american economic review*, 1972, *62* (4), 659–661.

Robbett, Andrea, Lily Colón, and Peter Hans Matthews, "Partisan political beliefs and social learning," *Journal of Public Economics*, apr 2023, *220*, 104834.

Shayo, Moses, "Social Identity and Economic Policy," *Annual Review of Economics*, aug 2020, *12* (1), 355–389.

Shiller, Robert J., "Narrative Economics," *American Economic Review*, apr 2017, *107* (4), 967–1004.

Siddique, Zahra, "Evidence on Caste Based Discrimination," *Labour Economics*, dec 2011, *18*, S146–S159.

Stigler, George J, "The economics of information," *Journal of political economy*, 1961, *69* (3), 213–225.

Tajfel, Henri and John C Turner, "Intergroup behavior," *Introducing social psychology*, 1978, *401*, 466.

_, M. G. Billig, R. P. Bundy, and Claude Flament, "Social categorization and intergroup behaviour," *European Journal of Social Psychology*, 1971, *1* (2), 149–178.

Triandis, Harry C and Michele J Gelfand, "Converging measurement of horizontal and vertical individualism and collectivism.," *Journal of personality and social psychology*, 1998, *74* (1), 118.

Weizsäcker, Georg, "Do We Follow Others when We Should? A Simple Test of Rational Expectations," *American Economic Review*, dec 2010, *100* (5), 2340–2360.

Zhang, Yunhao and David G. Rand, "Sincere or motivated? Partisan bias in advice-taking," *Judgment and Decision Making*, 2023, *18*.

Zou, Wenbo and Xue Xu, "Ingroup bias in a social learning experiment," *Experimental Economics*, dec 2022, *26* (1), 27–54.

APPENDIX

A Additional Figures

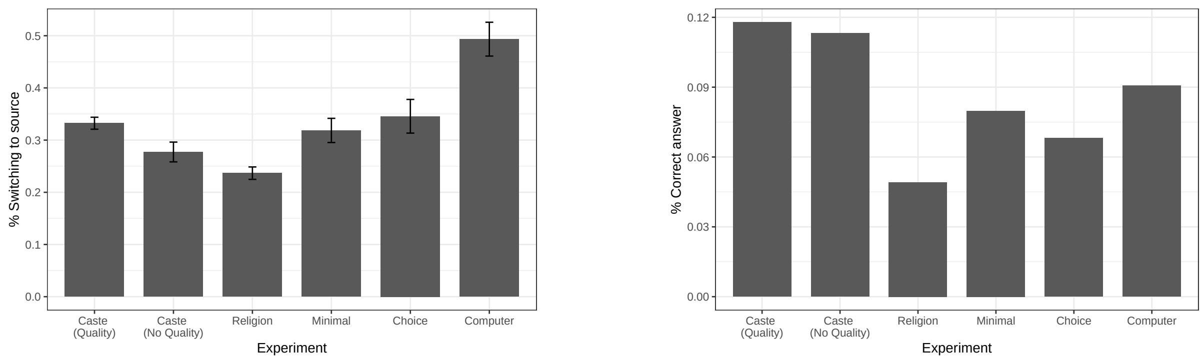

(a) (b)

Figure A.1: (a): fraction of decisions where participants choose to switch to the source in each experiment. Error bars indicate 95% confidence intervals. (b): fraction of decisions where participants' first estimate y_1 is within ±2% of the correct answer. Decisions where the absolute difference between the shown guess and a participant's first guess is ≤ 1 are excluded from these plots.

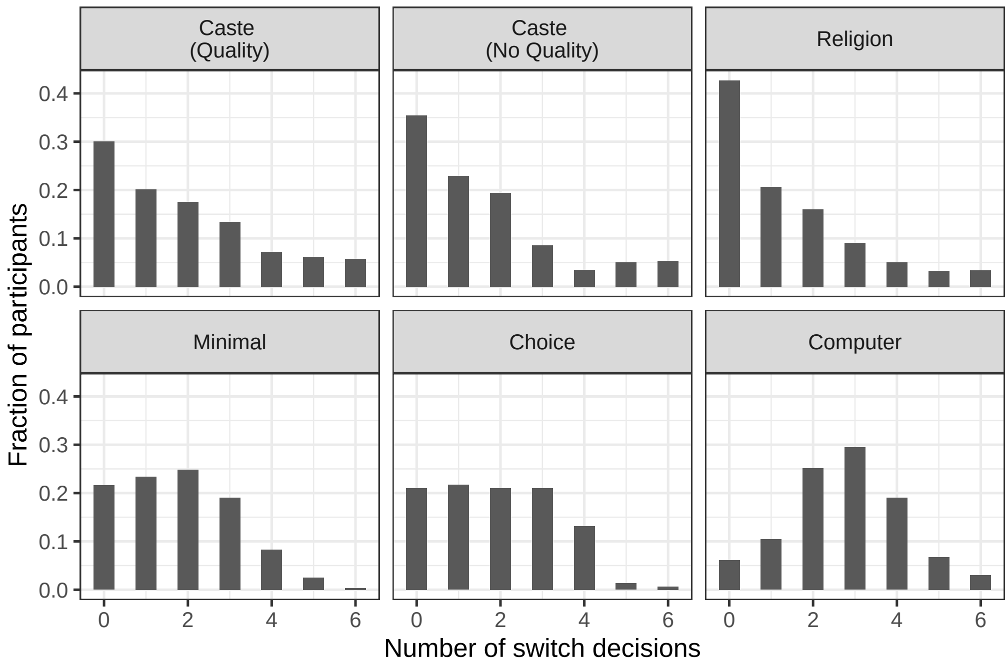

Figure A.2: The graph shows the distribution of the number of switching decisions made by participants in each experiment.

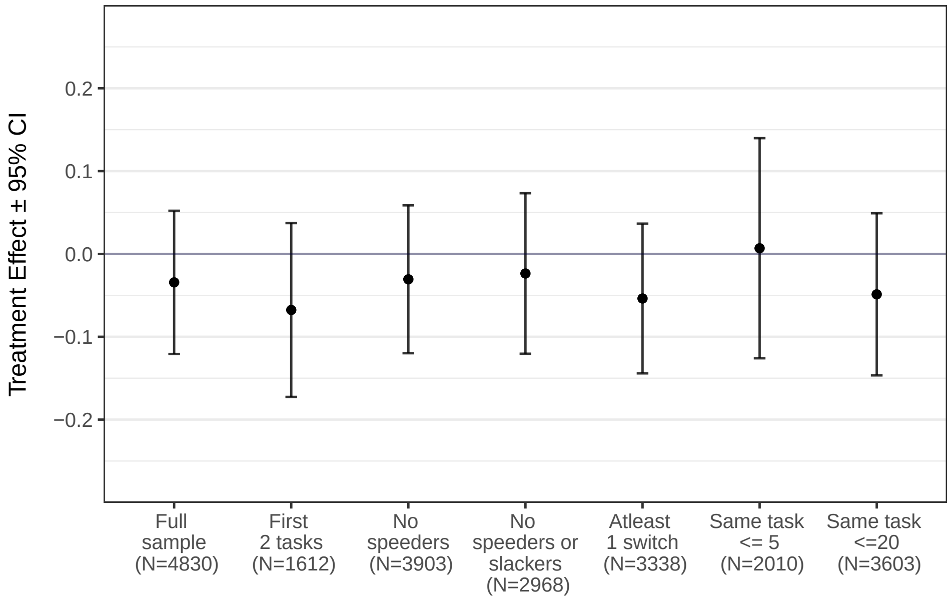

Figure A.3: Robustness tests – Difference in average treatment effect by the identity of the participant in experiment *Caste*. Each point is the coefficient on the treatment dummy variable, and indicate the preference for one identity group relative to the other in the *Caste*. The estimated specification is Equation 2, and standard errors are clustered at the participant level. Controls are: A set of demographic controls, task and order fixed effects, stated subjective certainty, and the quality of the information source. Table B.11 presents the regression tables used to construct this plot.

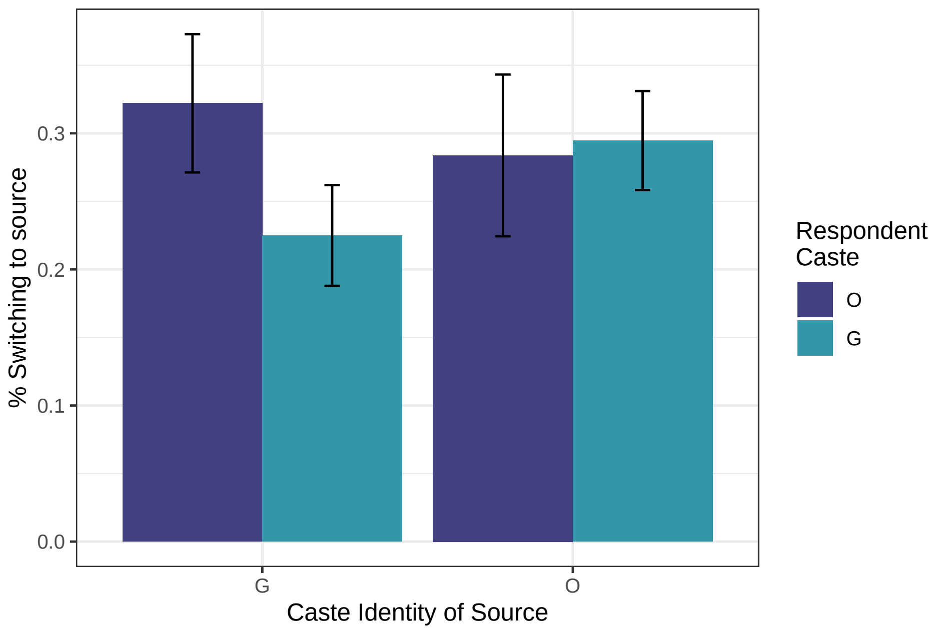

Figure A.4: Fraction of the decisions in which people switch in experiment *Caste – No Quality*. Each panel is a treatment within the experiment, (*G* or *O* source), and the bars within each panel show the fraction of people switching to y_s for each respondent caste group. Error bars show 95% confidence intervals.

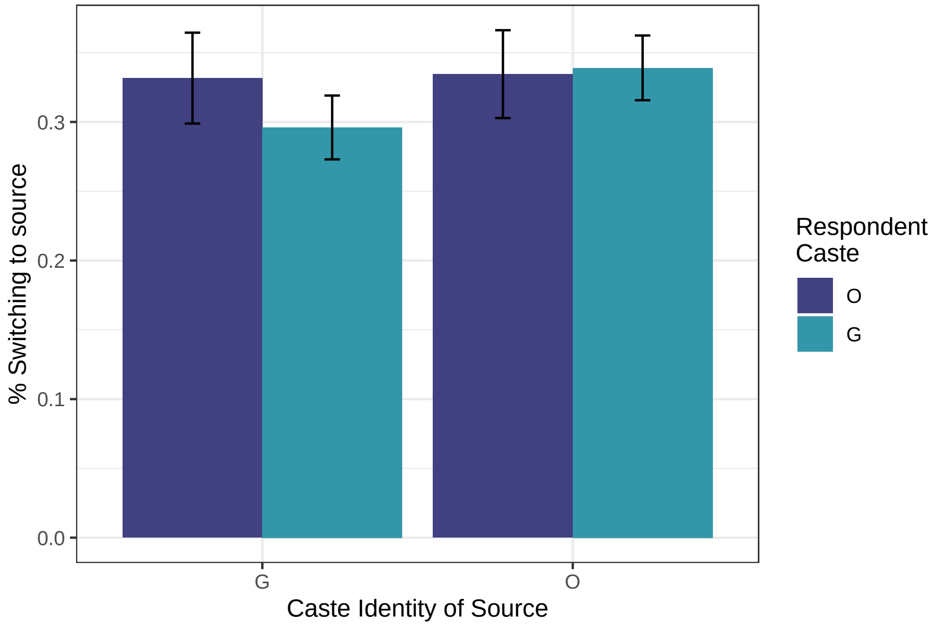

Figure A.5: Fraction of the decisions in which people switch in experiment *Caste*. Each panel is a treatment within the experiment, (*G* or *O* source), and the bars within each panel show the fraction of people switching to y_s for each respondent caste group. Error bars show 95% confidence intervals.

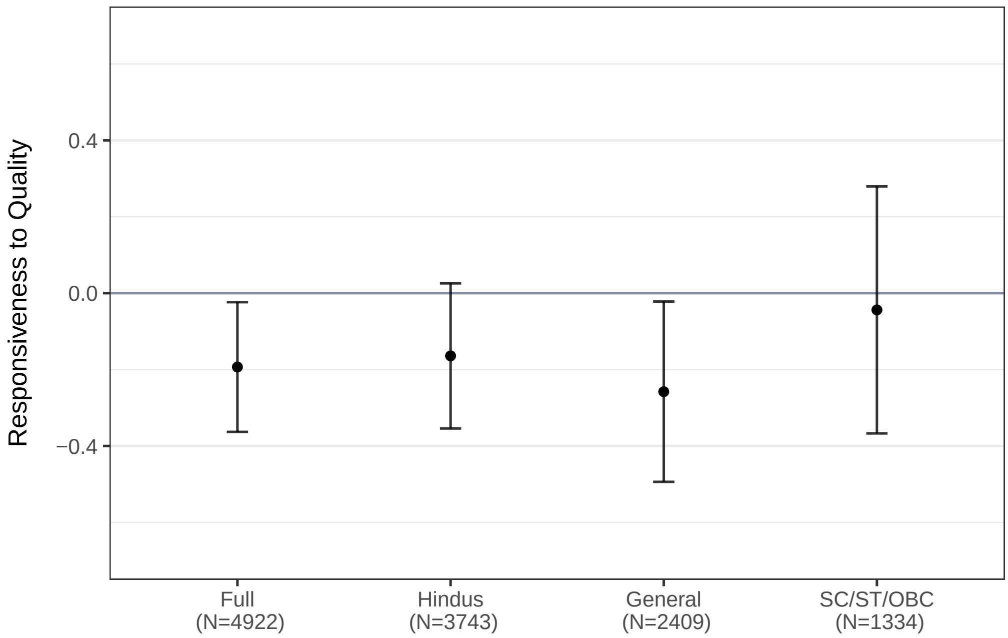

Figure A.6: The graph shows the responsiveness to source quality in experiment *Religion*. Each point is the coefficient on the interaction term *Treatment* × *Quality*. The estimates shown are for the full sample, and sub-samples of Hindus, General caste Hindus, and SC/ST/OBC Hindus respectively. Error bars indicate 95% confidence intervals.

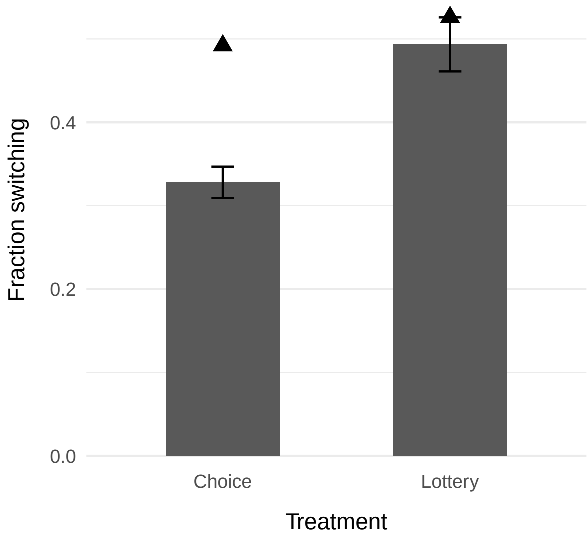

Figure A.7: The bars show the fraction switching to the source estimate in treatments *Choice* and *Computer*. The triangle indicates the number of decisions where switching would have been optimal, based on subjective certainty and the quality of information on a particular decision. Error bars indicate 95% confidence intervals.

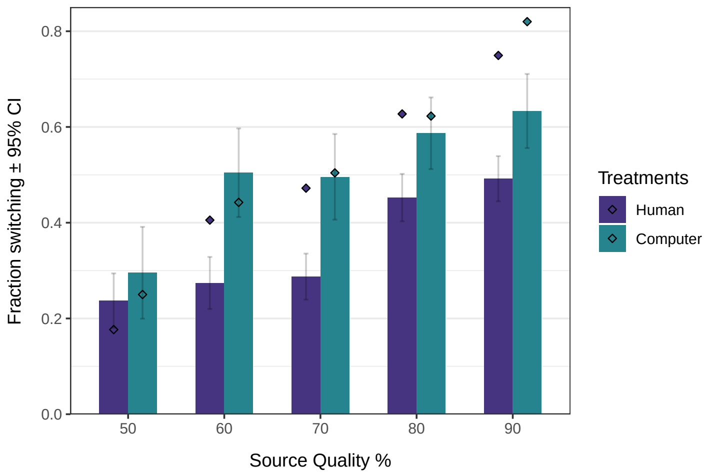

Figure A.8: Fraction of decisions where participants choose to switch to the source in treatments *Choice* and *Computer*, by the quality of information. The sample is restricted to decisions where the realisation of the source was accurate (the correct answer). The diamonds indicate the fraction of decisions where switching would have been the optimal choice (based on a comparison of information quality and subjective certainty). Error bars indicate 95% confidence intervals.

B Additional Tables

Table B.1: Sample Descriptives

	Main Experiments			Robustness
	Caste	Religion	Prolific	No quality
Age	32	31	32	32
% Men	52	60	57	52
% College	90	89	76	93
% Employed	84	83	70	85
Decision Time	25	20	17	
Participants	851	853	593	295
Decisions	5106	5118	3558	1770

Notes. Sample descriptives for the experiments presented in the paper. College and Employment are indicator variables = 1 if the participant has tertiary education or is not unemployed. Response time is the average time taken by participants on the main decision screen (decision to switch or stick).

Table B.2: Regression analysis: Average treatment effects

	Switch to source			
	Caste (i)	Religion (ii)	Minimal (iii)	Computer (iv)
G Source	-0.032 (0.021)			
Hindu Source		-0.001 (0.021)		
In-group Source			-0.007 (0.027)	
Computer Source				0.151*** (0.021)
Quality	0.060*** (0.014)	0.129*** (0.047)	0.703*** (0.089)	0.724*** (0.061)
Certainty	-0.003*** (0.000)	-0.002*** (0.000)	-0.005*** (0.001)	-0.005*** (0.000)
Constant	0.493*** (0.066)	0.339*** (0.075)	0.176* (0.102)	0.150** (0.071)
Task controls	✓	✓	✓	✓
Demog. controls	✓	✓	✓	✓
R^2	0.034	0.019	0.114	0.130
Observations	4,830	4,231	1,554	3,306
Clusters	851	733	278	593

Notes. Average treatment effects in the different experiments. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* \, p < 0.10, ** \, p < 0.05, *** p < 0.01$.

Table B.3: Regression analysis: Demographics – Caste experiment

	Switch to source			
	(i)	(ii)	(iii)	(iv)
G Source	-0.033	-0.033	-0.033	-0.033
	(0.021)	(0.021)	(0.021)	(0.021)
Quality	0.061***	0.060***	0.061***	0.061***
	(0.014)	(0.014)	(0.014)	(0.014)
Certainty	-0.003***	-0.003***	-0.003***	-0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.518***	0.462***	0.522***	0.514***
	(0.050)	(0.046)	(0.051)	(0.052)
Age	0.000			
	(0.001)			
Employed		0.075***		
		(0.026)		
College			-0.003	
			(0.034)	
Male				0.004
				(0.021)
Task controls	✓	✓	✓	✓
R^2	0.031	0.034	0.031	0.031
Observations	4,830	4,830	4,830	4,830
Clusters	851	851	851	851

Notes. OLS analyses of demographic characteristics and switching in experiment *Caste*. The demographic characteristics are a participant's age, employment status, whether college education or not, and their gender. Controls are task and order fixed effects. Standard errors are clustered at the participant level. $* \, p < 0.10, ** \, p < 0.05, *** p < 0.01$.

Table B.4: Regression analysis: Average treatment effects – Caste

	Switch to source				
	(i)	(ii)	(iii)	(iv)	(v)
G Source	-0.027	-0.028	-0.027	-0.026	-0.032
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)
Quality				0.059***	0.060***
				(0.015)	(0.014)
Certainty					-0.003***
					(0.000)
Constant	0.337***	0.313***	0.309***	0.280***	0.491***
	(0.016)	(0.024)	(0.050)	(0.050)	(0.058)
Age			-0.001	-0.001	0.000
			(0.001)	(0.001)	(0.001)
Task controls		✓	✓	✓	✓
Demog. controls			✓	✓	✓
R^2	0.001	0.004	0.007	0.011	0.035
Observations	4,830	4,830	4,830	4,830	4,830
Clusters	851	851	851	851	851

Notes. Average treatment effects in experiment *Caste*. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

Table B.5: Regression analysis: Average treatment effects – Religion

	Switch to source				
	(i)	(ii)	(iii)	(iv)	(v)
Hindu Source	-0.006	-0.007	-0.005	-0.005	-0.003
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)
Quality				0.165***	0.161***
				(0.044)	(0.043)
Certainty					-0.002***
					(0.000)
Constant	0.240***	0.227***	0.296***	0.181***	0.351***
	(0.014)	(0.021)	(0.053)	(0.062)	(0.068)
Task controls		✓	✓	✓	✓
Demog. controls			✓	✓	✓
R^2	0.000	0.003	0.006	0.009	0.024
Observations	4,922	4,922	4,922	4,922	4,922
Clusters	853	853	853	853	853

Notes. Average treatment effects in experiment *Religion*. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

Table B.6: Regression analysis: Average treatment effects – Minimal identity

	Switch to source				
	(i)	(ii)	(iii)	(iv)	(v)
In-group Source	-0.015	-0.016	-0.020	-0.026	-0.007
	(0.029)	(0.029)	(0.029)	(0.029)	(0.027)
Quality				0.696***	0.703***
				(0.089)	(0.089)
Certainty					-0.005***
					(0.001)
Constant	0.326***	0.374***	0.467***	-0.020	0.176*
	(0.020)	(0.040)	(0.082)	(0.106)	(0.102)
Task controls		✓	✓	✓	✓
Demog. controls			✓	✓	✓
R^2	0.000	0.007	0.012	0.055	0.114
Observations	1,554	1,554	1,554	1,554	1,554
Clusters	278	278	278	278	278

Notes. Average treatment effects in experiment *Minimal*. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* \, p < 0.10, ** \, p < 0.05, *** p < 0.01$.

Table B.7: Regression analysis: Average treatment effects – Human vs. Computer

	Switch to source				
	(i)	(ii)	(iii)	(iv)	(v)
Computer Source	0.165***	0.165***	0.168***	0.170***	0.151***
	(0.022)	(0.023)	(0.022)	(0.022)	(0.021)
Quality				0.729***	0.724***
				(0.062)	(0.061)
Certainty					-0.005***
					(0.000)
Constant	0.328***	0.388***	0.456***	-0.056	0.150**
	(0.012)	(0.028)	(0.057)	(0.073)	(0.071)
Task controls		✓	✓	✓	✓
Demog. controls			✓	✓	✓
R^2	0.023	0.030	0.032	0.076	0.130
Observations	3,306	3,306	3,306	3,306	3,306
Clusters	593	593	593	593	593

Notes. Average treatment effects in experiment *Computer*. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* \, p < 0.10, ** \, p < 0.05, *** p < 0.01$.

Table B.8: Regression analysis: Subsample analysis, experiment *Caste*

	Switch to source			
	Unconditional	w/ Controls		
Sample	Full		G	O
	(i)	(ii)	(iii)	(iv)
G Source	-0.029	-0.034	-0.042	-0.013
	(0.022)	(0.021)	(0.026)	(0.038)
High Quality		0.060***	0.063***	0.055**
		(0.015)	(0.018)	(0.026)
Certainty		-0.003***	-0.003***	-0.002***
		(0.000)	(0.001)	(0.001)
Constant	0.337***	0.480***	0.514***	0.440***
	(0.016)	(0.066)	(0.084)	(0.112)
Task controls		✓	✓	✓
Demog. controls		✓	✓	✓
R^2	0.001	0.033	0.044	0.028
Dependent variable mean	0.323	0.323	0.318	0.333
Observations	4,733	4,733	3,088	1,645
Clusters	834	834	542	292

Notes. Average treatment effects in the *Caste* experiment. (i)–(ii): Full sample with and without controls. (iii)–(iv): Subsamples by the case of the participant. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

Table B.9: Regression analysis: Robustness, *Caste* experiment

	Switch to source						
	Full sample	First 2 tasks	No Speeders	No speeders or slackers	Atleast 1 switch	Same task <= 5	Same task <= 20
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
G Source	-0.032	-0.028	-0.027	-0.020	-0.047**	-0.021	-0.025
	(0.021)	(0.026)	(0.022)	(0.023)	(0.022)	(0.033)	(0.024)
Quality	0.060***	0.070***	0.070***	0.061***	0.167***	0.158***	0.128***
	(0.014)	(0.024)	(0.016)	(0.018)	(0.047)	(0.056)	(0.042)
Certainty	-0.003***	-0.004***	-0.003***	-0.003***	-0.003***	-0.004***	-0.003***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Constant	0.493***	0.514***	0.488***	0.545***	0.550***	0.475***	0.489***
	(0.066)	(0.082)	(0.070)	(0.080)	(0.079)	(0.109)	(0.080)
Task controls	✓	✓	✓	✓	✓	✓	✓
Demog. controls	✓	✓	✓	✓	✓	✓	✓
R^2	0.034	0.047	0.037	0.039	0.028	0.043	0.035
Observations	4,830	1,612	3,903	2,968	3,338	2,010	3,603
Clusters	851	851	840	822	586	354	635

Notes. OLS estimates for the robustness exercises presented in Figure 7. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

Table B.10: Regression analysis: Robustness, *Religion* experiment

	Switch to source				
	Full sample (i)	First 2 tasks (ii)	No Speeders (iii)	No speeders or slackers (iv)	Atleast 1 switch (v)
Hindu Source	-0.001	0.013	0.000	0.005	0.023
	(0.021)	(0.025)	(0.021)	(0.022)	(0.025)
Quality	0.129***	0.139*	0.171***	0.165***	0.219***
	(0.047)	(0.077)	(0.050)	(0.056)	(0.074)
Certainty	-0.002***	-0.002***	-0.002***	-0.002***	-0.002***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)
Constant	0.339***	0.347***	0.341***	0.378***	0.343***
	(0.075)	(0.103)	(0.078)	(0.085)	(0.097)
Task controls	✓	✓	✓	✓	✓
Demog. controls	✓	✓	✓	✓	✓
R^2	0.019	0.020	0.025	0.031	0.019
Observations	4,231	1,416	3,601	2,796	2,383
Clusters	733	732	732	724	413

Notes. OLS estimates for the robustness exercises presented in Figure 7. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $*\,p < 0.10, **p < 0.05, ***p < 0.01$.

Table B.11: Regression analysis: Robustness – Interaction effect, *Caste* experiment

	Switch to source						
	Full sample (i)	First 2 tasks (ii)	No Speeders (iii)	No speeders or slackers (iv)	Atleast 1 switch (v)	Same task <= 5 (vi)	Same task <= 20 (vii)
G Source × G Participant	-0.034	-0.068	-0.031	-0.024	-0.054	0.007	-0.049
	(0.044)	(0.054)	(0.046)	(0.049)	(0.046)	(0.068)	(0.050)
G Source	-0.010	0.016	-0.006	-0.004	-0.012	-0.025	0.006
	(0.036)	(0.043)	(0.036)	(0.040)	(0.037)	(0.052)	(0.039)
G Participant	0.001	0.005	-0.006	-0.025	-0.002	0.017	0.021
	(0.032)	(0.037)	(0.032)	(0.035)	(0.033)	(0.050)	(0.035)
Quality	0.060***	0.069***	0.070***	0.061***	0.167***	0.159***	0.129***
	(0.014)	(0.024)	(0.016)	(0.018)	(0.047)	(0.056)	(0.042)
Certainty	-0.003***	-0.004***	-0.003***	-0.003***	-0.003***	-0.004***	-0.003***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Constant	0.484***	0.496***	0.483***	0.546***	0.540***	0.474***	0.469***
	(0.069)	(0.084)	(0.073)	(0.083)	(0.080)	(0.115)	(0.083)
Task controls	✓	✓	✓	✓	✓	✓	✓
Demog. controls	✓	✓	✓	✓	✓	✓	✓
R^2	0.035	0.049	0.038	0.041	0.029	0.043	0.036
Observations	4,830	1,612	3,903	2,968	3,338	2,010	3,603
Clusters	851	851	840	822	586	354	635

Notes. OLS estimates for the robustness exercises presented in Figure 7. Controls are task and order fixed effects, and demographic characteristics(the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $*\,p < 0.10, **\,p < 0.05, ***p < 0.01$.

Table B.12: Regression analysis: Role of relevant beliefs, experiment *Caste – No Quality*

	Switch to source			
	No Quality			Quality
	(i)	(ii)	(iii)	(iv)
Beliefs (relevant)	0.188**	0.235***	0.189*	0.068
	(0.078)	(0.077)	(0.100)	(0.050)
G Source		-0.010	-0.081	
		(0.034)	(0.108)	
G Source × Beliefs (relevant)			0.104	
			(0.159)	
Certainty		-0.002***	-0.002***	-0.003***
		(0.001)	(0.001)	(0.000)
Constant	0.145***	0.322***	0.350***	0.461***
	(0.052)	(0.113)	(0.123)	(0.073)
Task controls		✓	✓	✓
Demog. controls		✓	✓	✓
R^2	0.009	0.042	0.043	0.030
Observations	1,676	1,676	1,676	4,830
Clusters	295	295	295	851

Notes. OLS regressions of whether a participant switches to the source's estimate on their beliefs about the likelihood that a person sharing the source's identity is correct on a task. Columns (i)–(iii): No signal about the quality of the source's estimate, (iv): Receive a signal about the quality of the source's estimate. Controls are task and order fixed effects, and demographic characteristics (the participant's age and dummies for employment, college education and gender). Standard errors are clustered at the participant level. $* p < 0.10, ** p < 0.05, *** p < 0.01$.

C Experiment details and additional analyses

C.1 Surnames used in the experiments

The recognisability of the Hindu surnames used in the experiments were validated in a separate survey ($N = 350$). These individuals completed much of the training module and comprehension tests that will be used in the main treatments. The source of truth for the incentivisation of these classifications came from official classifications of individuals belonging to these communities, or the common nature of these surnames. The recognisability ranged from $\approx$ 60% to 85%, with names from the SC/ST/OBC castes being, on average, more recognisable than General caste surnames. The Muslim names were not validated, given the common nature and clear identifiability of these names.

Hindu and General caste names. Iyer, Banerjee, Chaturvedi, Tiwari, Bharadwaj, Mishra.

Hindu and SC/ST/OBC caste names. Paraiyar, Bhil, Jatav, Manjhi, Mahar, Chamar.

Muslim names. Khan, Shaikh, Abdullah, Syed, Moinuddin, Ali.

C.2 Learning from others in India.

Evidence on how much people rely on social learning is scarce. Figure C.1 presents some indicative figures from a recent survey (Pew Research Center, 2021), showing that a substantial fraction of individuals do indeed learn from others. There is also some evidence of heterogeneity along caste and religious dimensions – Hindus are more likely to get their news from others than Muslims, and General caste individuals are more like to get news from others than SC/ST/OBC individuals. While this does not help answer our research questions, Figure C.1 illustrates that social learning is indeed a significant channel through which people get information in India.

C.3 Exposure, attitudes, and religiosity.

Participants' attitudes towards caste-based affirmative action, exposure to people from different caste groups, or religiosity may influence whom they learn from. It is plausible that people who are against the idea of benefits for certain caste groups, are not exposed to people from different caste groups, or are religious may have stronger in-group preferences for their own caste-group. Figure C.2 shows analyses exploring these various dimensions. There is no evidence of any heterogeneous treatment effects along these dimensions. It must be noted that these analyses do not easily lend themselves to a causal interpretation as many of

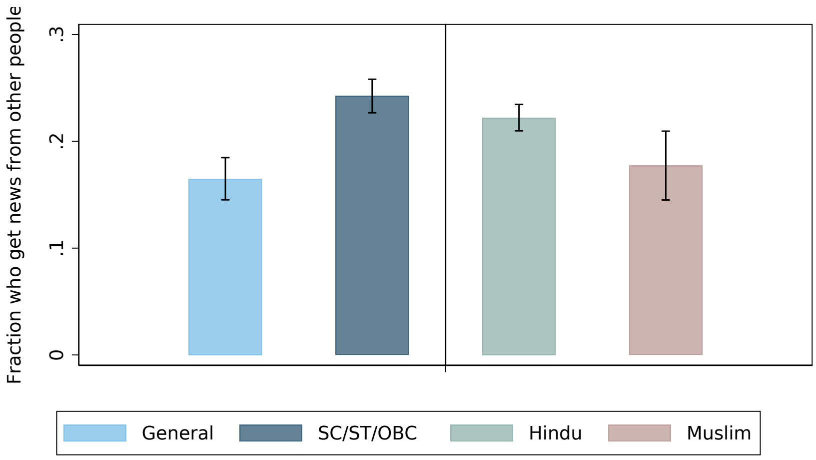

Figure C.1: Each bar represents the fraction of respondents who select “Other people” as one of their top three sources of getting news, controlling for income, education, and urbanicity. Author’s calculations based on data from a survey conducted by Pew Research Center (2021).

these factors are endogenous to the experiment context, and (except religiosity) are elicited after the main tasks.

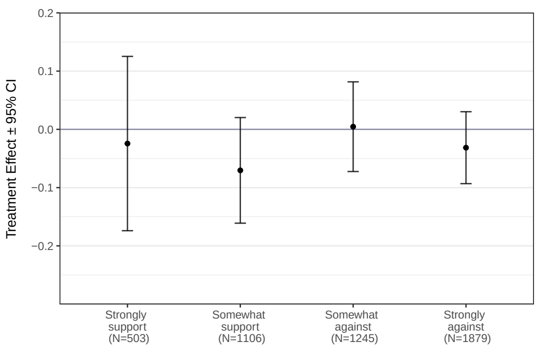

(a) Support for caste-based affirmative action (levels)

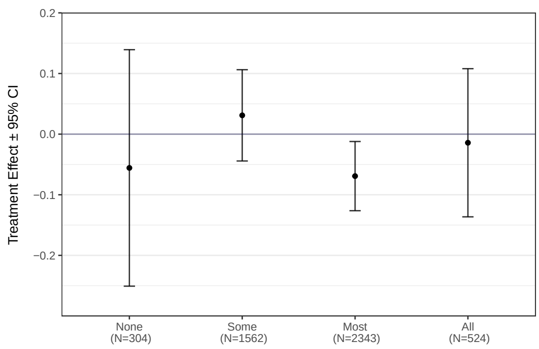

(b) Exposure to other castes (levels)

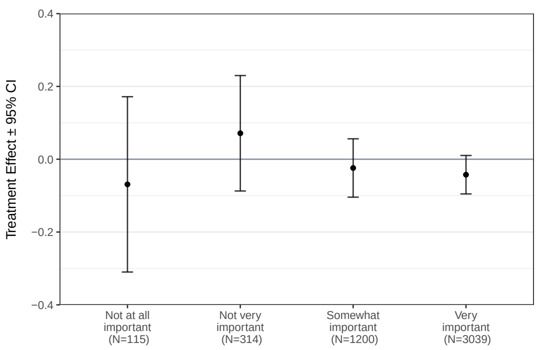

(c) Religiosity – Caste experiment

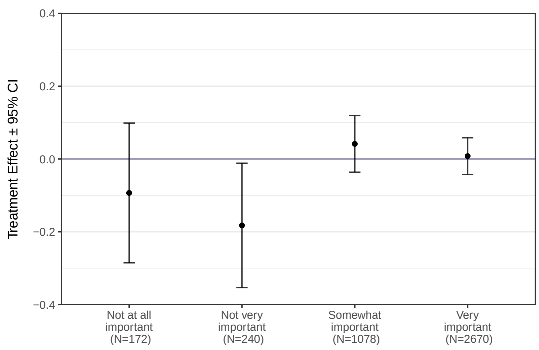

(d) Religiosity – Religion experiment

Figure C.2: Heterogeneity analyses in the *Caste* and *Religion* experiment. Panels show estimated coefficients from regressions on different sub-samples. Figures (a)–(d) show regressions for each level of self-reported support for caste-based affirmative action policies, exposure to other castes, and religiosity (in the *Caste* and *Religion* experiments). Standard errors are clustered at the participant level in all regressions.

C.4 Individualism and learning

Individualism is a measure of the extent to which individuals strive to be distinct from others, and nests two concepts Triandis and Gelfand (1998). Horizontal individualism captures the extent to which individuals strive to be different without desiring special status. Vertical individualism captures the extent to which individuals strive to be distinct and acquire special status. In the context of social learning, highly individualistic people may choose to ignore social information more than less individualistic people. This may also influence the extent to which they use information from groups that they are members of. To study this potential channel, I elicit participants' individualism in experiments *Religion, Minimal, Choice,* and *Computer* using a subset of questions from the Individualism - Collectivism scale (Triandis and Gelfand, 1998).

Individualism measure. Experiments *Religion, Minimal, Choice, and Computer* include 8 additional questions that measure Horizontal and Vertical Individualism. These measures are taken from the scale developed by Triandis and Gelfand (1998).[10] Answers are coded on a scale from 1 to 5, with 1 indicating "Strongly disagree" and 5 indicating "Strongly agree".

The questions are:

I would rather depend on myself than others.
I rely on myself most of the time. I rarely rely on others.
I often do my own thing.
My personal identity, independent of others, is very important to me.
It is important that I do my job better than others.
Winning is everything.
Competition is the law of nature.
When another person does better than I do, I get tense.

Summing up the responses to these questions results in an "Individualism" score for each participant.

Results. The results indicate that individualism is correlated with the decision to switch only in settings where group identity is present in the context. Individualism is negatively correlated with the decision to switch to the source's decision in the experiments *Minimal* and *Religion*. In treatments *Choice* and *Computer*, the correlation is much lower. Figure C.3 graphically shows these correlations. These correlation suggest that individualism may play

[10]The full battery comprises 24 questions, half of which measure individualism and the other half collectivism. In the experiment, the questions regarding collectivism are dropped to avoid making the survey module too lengthy, focusing on the role that individualism may play on an individual's potential preference for consistency or autonomy, which was a strong pattern in the pilot experiments conducted before the main study.

a role in social learning, but only when identity concerns are activated. However, individualism does not appear to be correlated with any differences in how people switch when seeing an in-group relative to an out-group source. Thus, while individualism may be related to whether people learn from others in situations where identity is salient, there is no evidence that it leads people to treat groups differently.

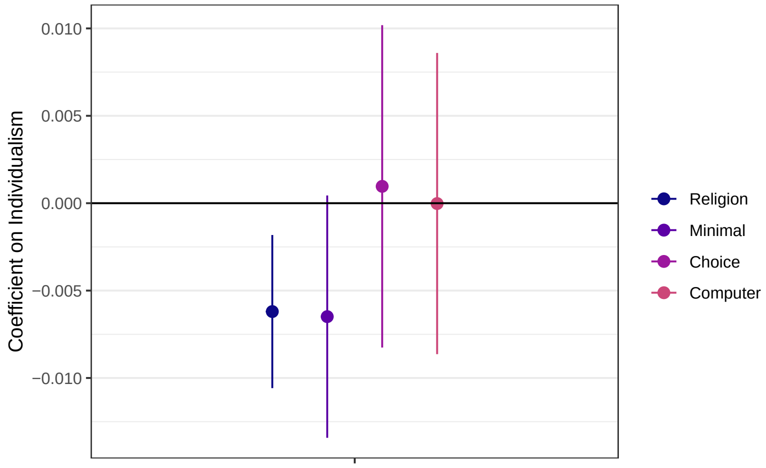

Figure C.3: The graph shows the coefficients on the individualism measure from regressions of the outcome variable *Switch* on individualism for experiments *Religion, Minimal, Choice, and Computer*. Error bars indicate 95% confidence intervals.

D Experiment instructions

Experiment *Caste – Quality*

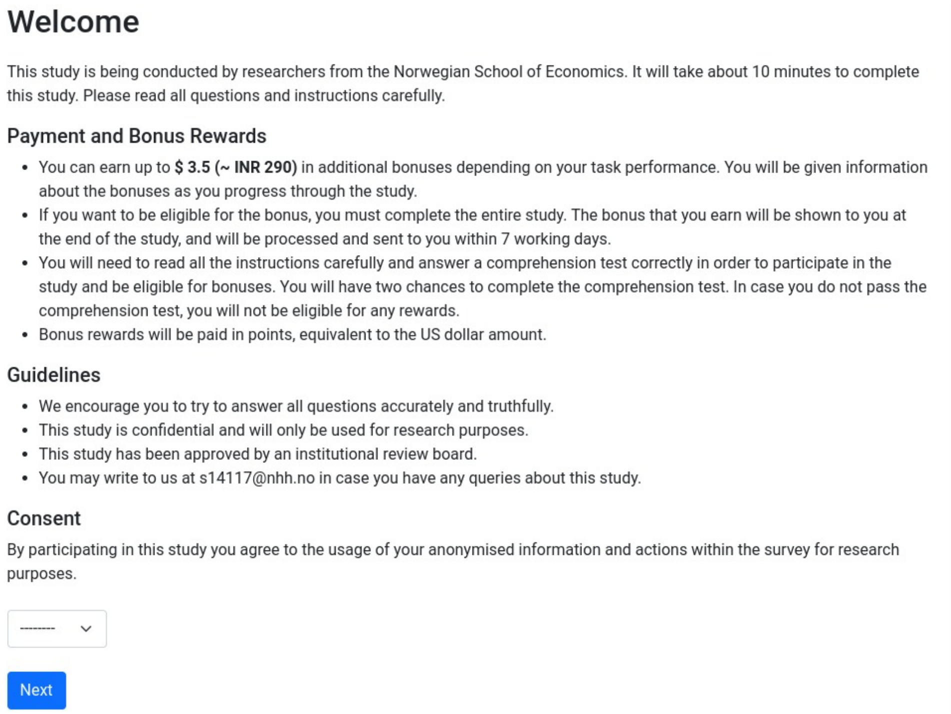

Welcome

This study is being conducted by researchers from the Norwegian School of Economics. It will take about 10 minutes to complete this study. Please read all questions and instructions carefully.

Payment and Bonus Rewards

- You can earn up to **$ 3.5 (~ INR 290)** in additional bonuses depending on your task performance. You will be given information about the bonuses as you progress through the study.
- If you want to be eligible for the bonus, you must complete the entire study. The bonus that you earn will be shown to you at the end of the study, and will be processed and sent to you within 7 working days.
- You will need to read all the instructions carefully and answer a comprehension test correctly in order to participate in the study and be eligible for bonuses. You will have two chances to complete the comprehension test. In case you do not pass the comprehension test, you will not be eligible for any rewards.
- Bonus rewards will be paid in points, equivalent to the US dollar amount.

Guidelines

- We encourage you to try to answer all questions accurately and truthfully.
- This study is confidential and will only be used for research purposes.
- This study has been approved by an institutional review board.
- You may write to us at s14117@nhh.no in case you have any queries about this study.

Consent

By participating in this study you agree to the usage of your anonymised information and actions within the survey for research purposes.

Next

Figure D.1: Welcome and consent

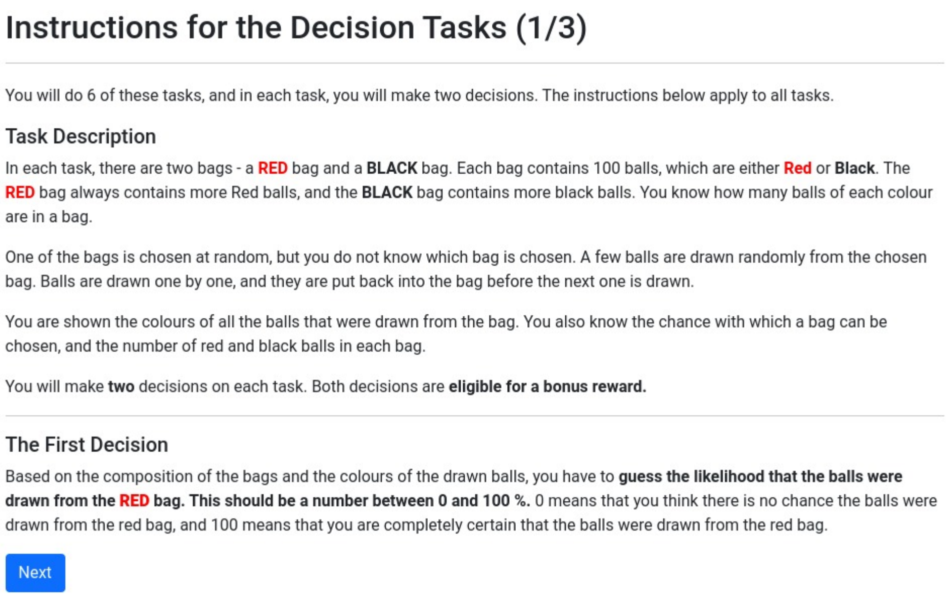

Instructions for the Decision Tasks (1/3)

You will do 6 of these tasks, and in each task, you will make two decisions. The instructions below apply to all tasks.

Task Description

In each task, there are two bags - a **RED** bag and a **BLACK** bag. Each bag contains 100 balls, which are either **Red** or **Black**. The **RED** bag always contains more Red balls, and the **BLACK** bag contains more black balls. You know how many balls of each colour are in a bag.

One of the bags is chosen at random, but you do not know which bag is chosen. A few balls are drawn randomly from the chosen bag. Balls are drawn one by one, and they are put back into the bag before the next one is drawn.

You are shown the colours of all the balls that were drawn from the bag. You also know the chance with which a bag can be chosen, and the number of red and black balls in each bag.

You will make **two** decisions on each task. Both decisions are **eligible for a bonus reward.**

The First Decision

Based on the composition of the bags and the colours of the drawn balls, you have to **guess the likelihood that the balls were drawn from the RED bag. This should be a number between 0 and 100 %.** 0 means that you think there is no chance the balls were drawn from the red bag, and 100 means that you are completely certain that the balls were drawn from the red bag.

Next

Figure D.2: Instructions 1/3

Instructions for the Decision Tasks (2/3)

After making the first decision, you will be shown a number.

How is this number generated?

- The computer has access to a pool of participants who made **correct guesses** on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess.
- The computer randomly chooses one of these people, and shows you a number:
  - With some probability, the shown number is the chosen person's **correct** guess.
  - Otherwise, the shown number is **incorrect**.

The probability that the correct guess is shown can change from task to task.

The Second Decision

You have to make a choice:

1. Stick to your first guess, OR
2. Switch to the shown number.

Next, you will learn about how you can earn bonuses on these tasks.

Next

Figure D.3: Instructions 2/3

Instructions for the Decision Tasks (3/3)

Bonuses

- You will do 6 of these tasks, making 2 decisions in each task. **One** of these 12 decisions will be randomly selected for a bonus.
- Each task has a mathematically correct answer.
- If the chosen guess is **within ± 2 % points of the correct answer**, your bonus will be $ 3.
- If your guess is more than 2 % away from the correct answer, then you will not earn a bonus.

It is therefore in your best interest to try and make all your guesses as accurately as possible.

Note: The correct answer to each task can be calculated using a probability formula. Click to see formula

Next, you will see an example task.

Next

Figure D.4: Instructions 3/3

Example Task (1/3)

This is an example of how the actual task will work.

Example Task

One of these bags is selected.
The **RED** bag is selected with a **50%** chance.

RED bag	BLACK bag
100 **Red**, 0 **Black** balls	0 **Red**, 100 **Black** balls
50% chance of being selected	50% chance of being selected

When you click the button, **3 balls** will be drawn from the selected bag one by one, with replacement. You will be allowed to make your guess after drawing the balls.

Click to draw balls

Figure D.5: Example 1/3

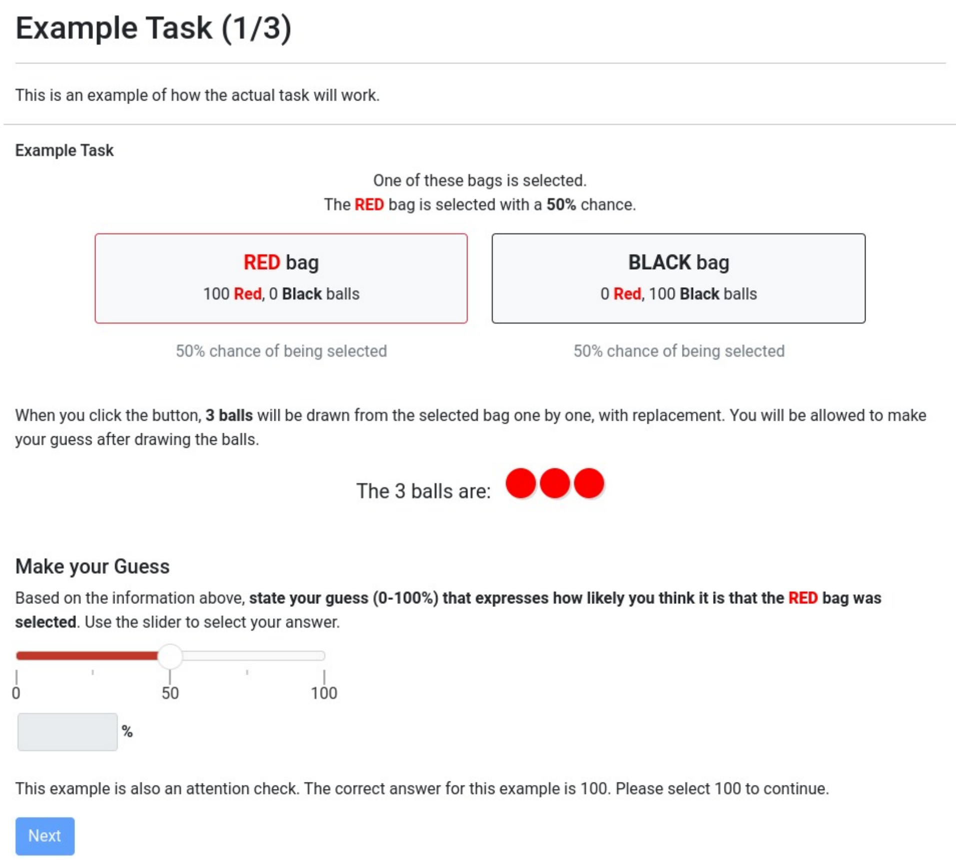

Figure D.6: Example 1/3

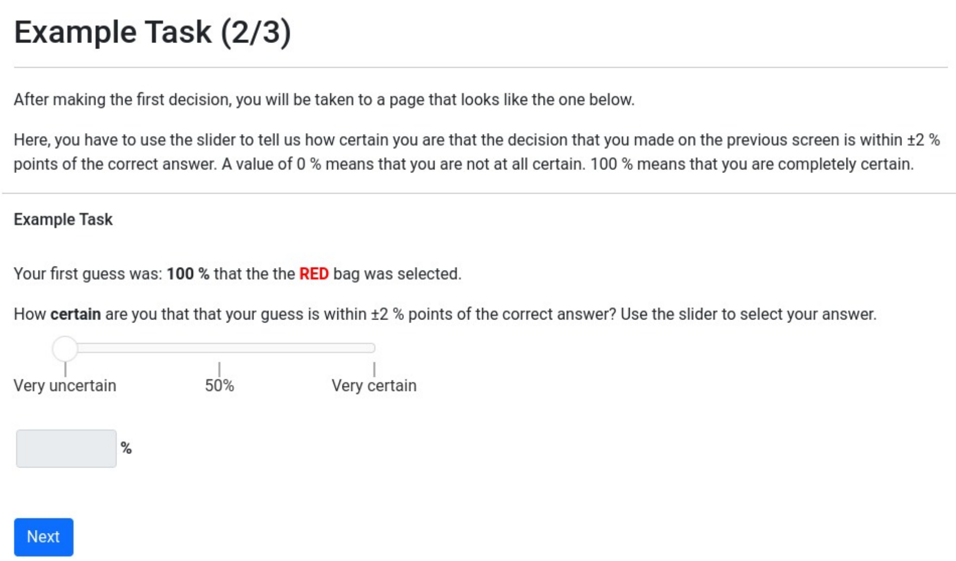

Figure D.7: Example 2/3

Example Task (3/3)

The example continues on this page - no bonuses will be awarded for this task, but you have to provide an answer to continue to the next page.

Example Task - The Second Decision

The computer shows:

Study Participant

73

There is a **50 % chance** that this number is the correct answer.

More info:

- The computer has access to a pool of participants who made **correct guesses** on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess.
- The computer randomly chooses one of these people, and shows you a number:
    - With some probability, the shown number is the chosen person's **correct** guess.
    - Otherwise, the shown number is **incorrect**.

Make your Second Decision

Your first guess was: **100 %**.

Based on the information above, **choose one of these options to make your decision (0-100%) that expresses how likely you think it is that the RED bag was selected**.

- **100 %** - Stick to your first guess
- **73 %** - Switch to the shown number

Next

Figure D.8: Example 3/3

Summary

Congratulations! You have successfully completed the example task.

Next, you will answer a short comprehension test. You must pass the test in order to be eligible for survey completion and bonus payments. Below is a short summary of the task, please familiarise yourself with them before moving on to the test.

Task Description

There are two bags - a **RED** bag and a **BLACK** bag. Each bag contains 100 balls, which are either **Red** or **Black**. The **RED** bag always contains more Red balls, and the **BLACK** bag contains more black balls. You know how many balls of each colour are in a bag.

One of the bags is chosen at random, and a few balls are drawn randomly from the chosen bag. Balls are drawn one by one, and they are put back into the bag before the next one is drawn. You are shown the colours of all the balls that were picked.

First decision: Based on this information, you have to **guess the likelihood that the balls were drawn from the RED bag. This should be a number between 0 and 100 %.**

Second decision: The second decision is made after you see what someone else has guessed on the exact same task. They have seen the same bags and drawn balls as you. A computer will show you a number. You have to decide whether to stick to your first guess, or switch to the shown number.

Important points

1. You make two decisions on each task, and there are 6 tasks.
2. Each of the Decision Tasks has a mathematically correct answer.
3. One of the decisions will be randomly chosen for a bonus. You will get a $ 3 bonus if the chosen decision is within 2% points of the correct answer.

Click the button to proceed to the comprehension test.

Next

Figure D.9: Summary

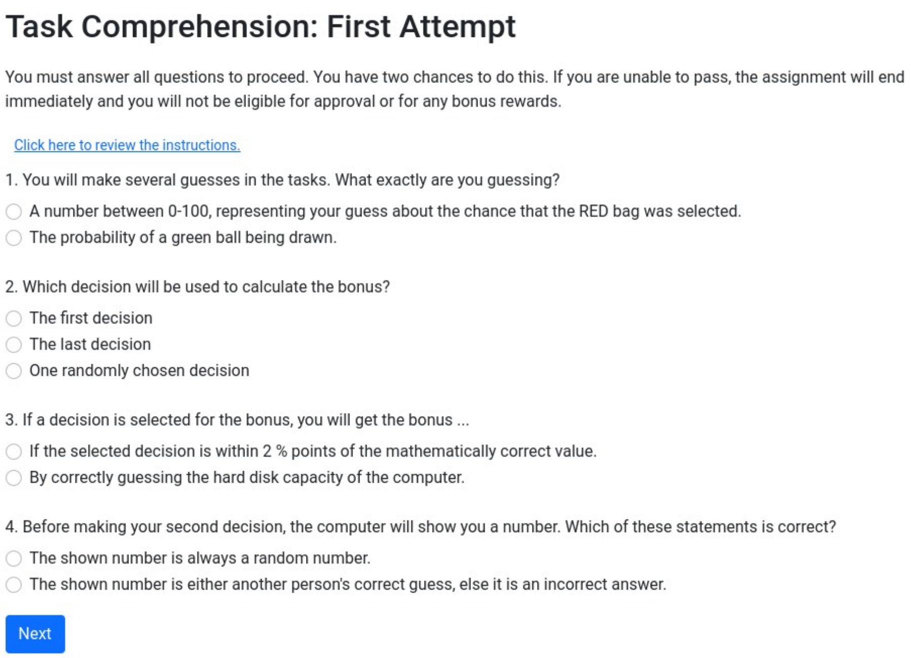

Task Comprehension: First Attempt

You must answer all questions to proceed. You have two chances to do this. If you are unable to pass, the assignment will end immediately and you will not be eligible for approval or for any bonus rewards.

Click here to review the instructions.

1. You will make several guesses in the tasks. What exactly are you guessing?

- ○ A number between 0-100, representing your guess about the chance that the RED bag was selected.
- ○ The probability of a green ball being drawn.

2. Which decision will be used to calculate the bonus?

- ○ The first decision
- ○ The last decision
- ○ One randomly chosen decision

3. If a decision is selected for the bonus, you will get the bonus ...

- ○ If the selected decision is within 2 % points of the mathematically correct value.
- ○ By correctly guessing the hard disk capacity of the computer.

4. Before making your second decision, the computer will show you a number. Which of these statements is correct?

- ○ The shown number is always a random number.
- ○ The shown number is either another person's correct guess, else it is an incorrect answer.

Next

Figure D.10: Comprehension (two attempts)

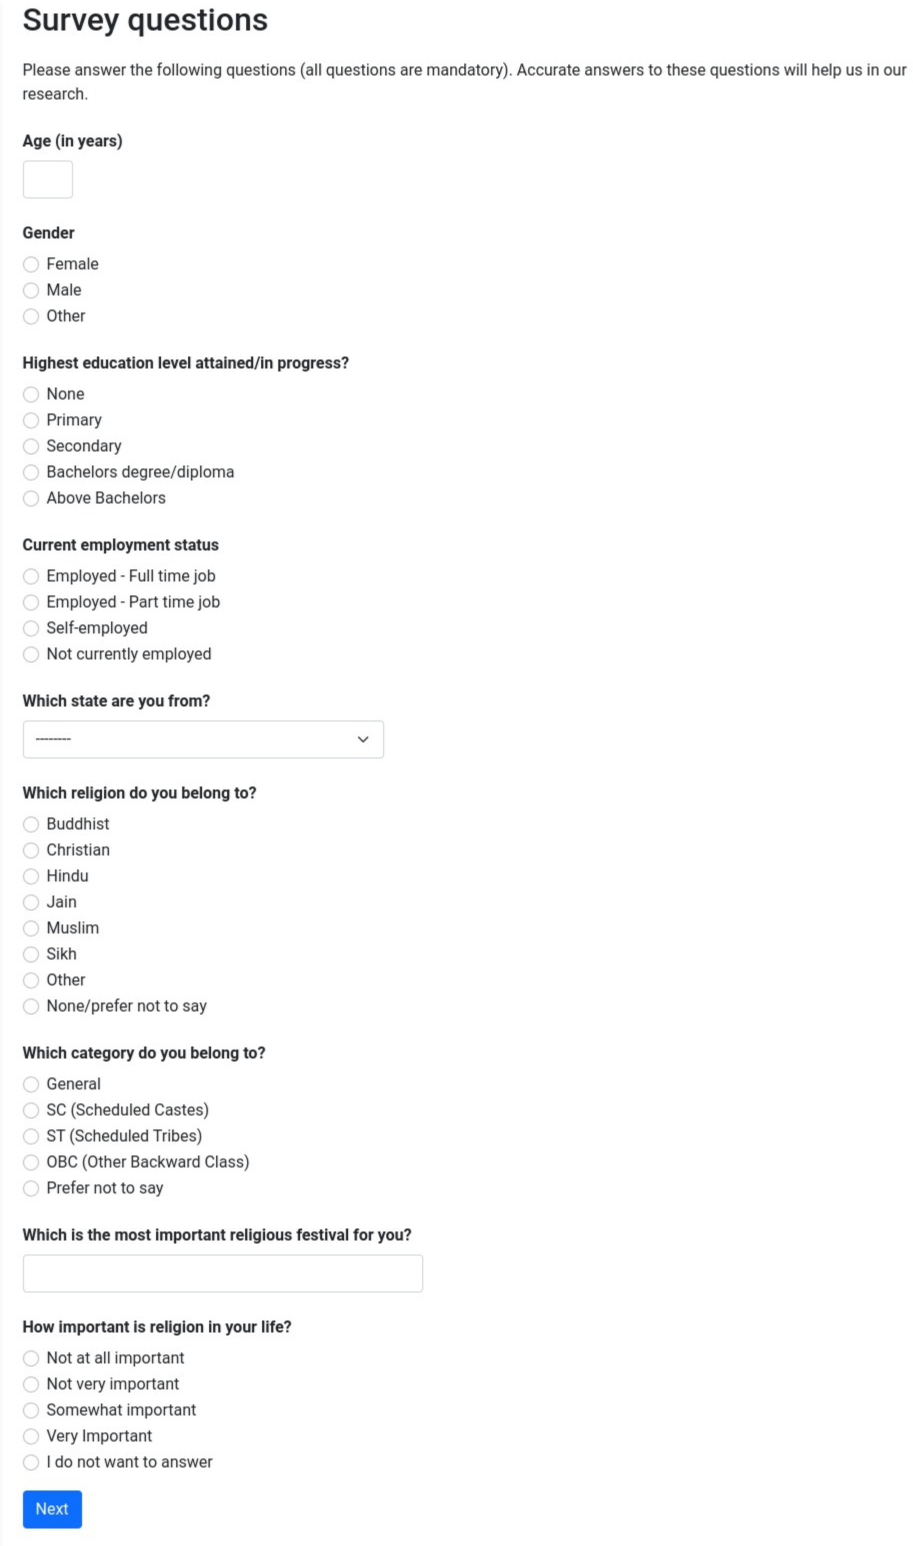

Survey questions

Please answer the following questions (all questions are mandatory). Accurate answers to these questions will help us in our research.

Age (in years)

Gender

- Female
- Male
- Other

Highest education level attained/in progress?

- None
- Primary
- Secondary
- Bachelors degree/diploma
- Above Bachelors

Current employment status

- Employed - Full time job
- Employed - Part time job
- Self-employed
- Not currently employed

Which state are you from?

Which religion do you belong to?

- Buddhist
- Christian
- Hindu
- Jain
- Muslim
- Sikh
- Other
- None/prefer not to say

Which category do you belong to?

- General
- SC (Scheduled Castes)
- ST (Scheduled Tribes)
- OBC (Other Backward Class)
- Prefer not to say

Which is the most important religious festival for you?

How important is religion in your life?

- Not at all important
- Not very important
- Somewhat important
- Very Important
- I do not want to answer

Next

Figure D.11: Demographics

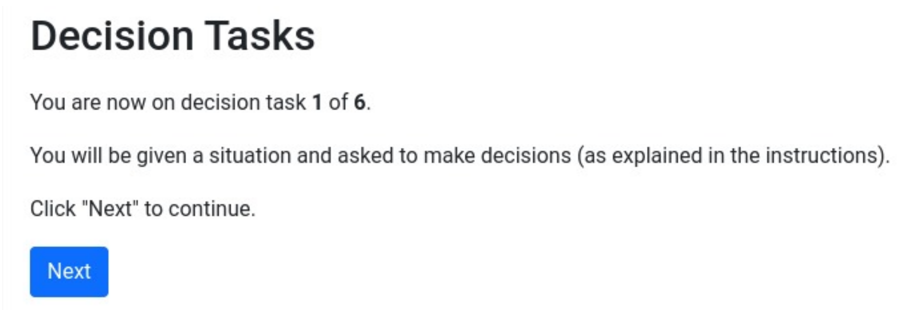

Figure D.12: In-between decision tasks

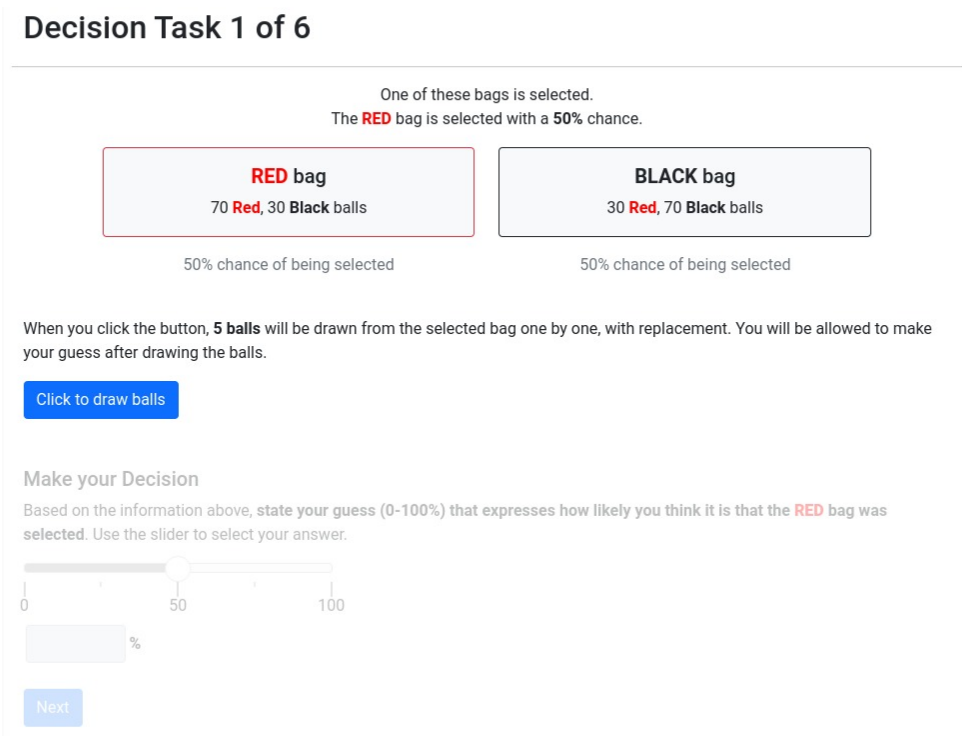

Figure D.13: Decision Tasks - First guess

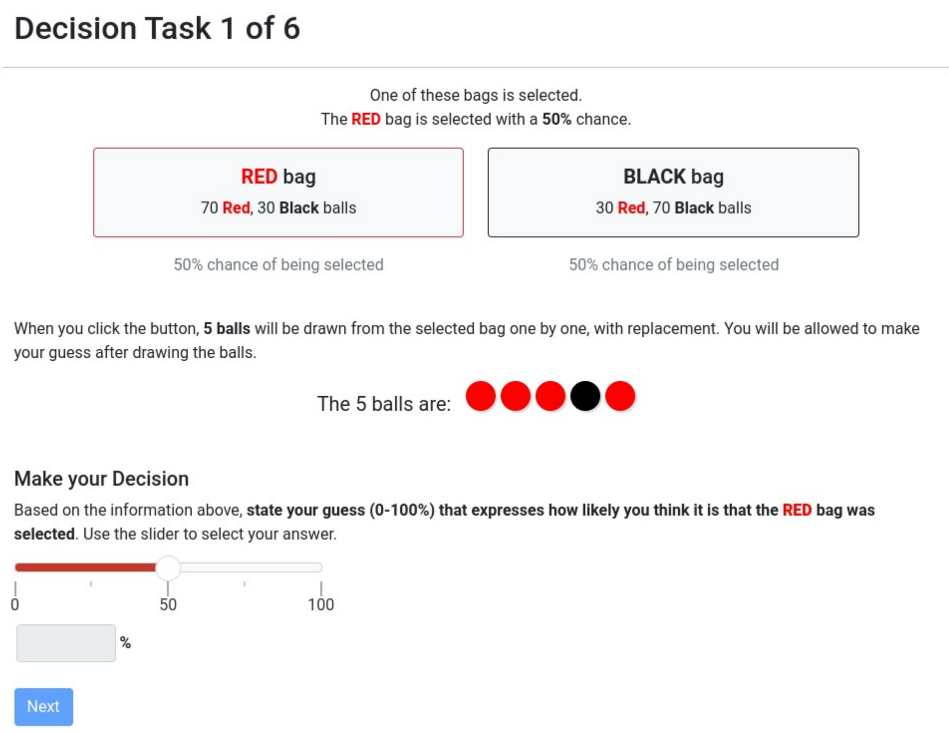

Figure D.14: Decision Tasks - First guess

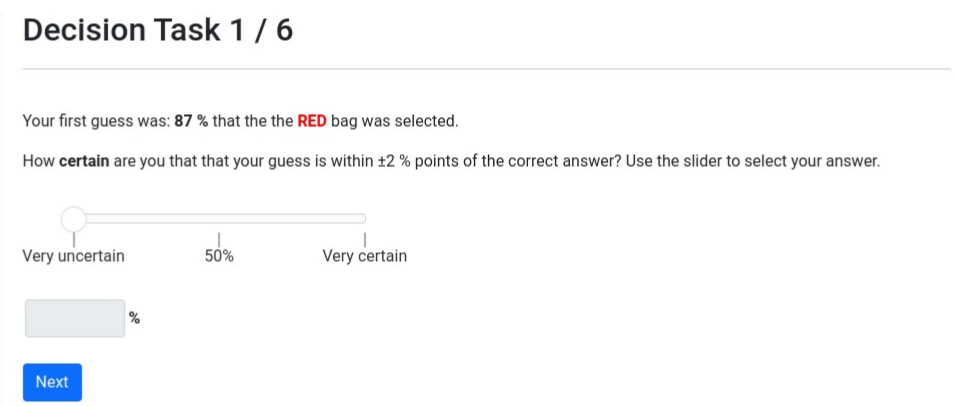

Figure D.15: Decision Tasks - certainty elicitation

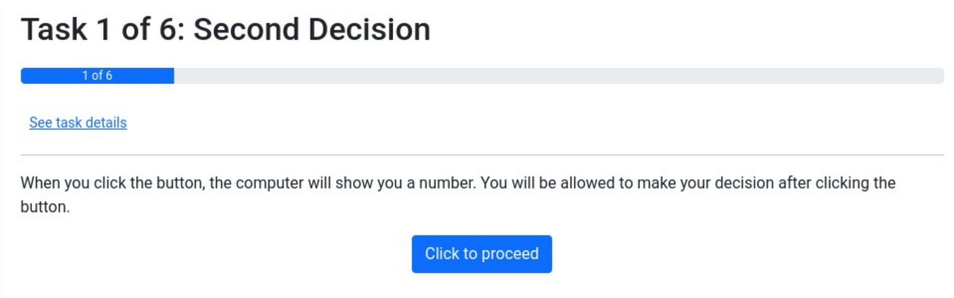

Figure D.16: Decision Tasks - Second decision

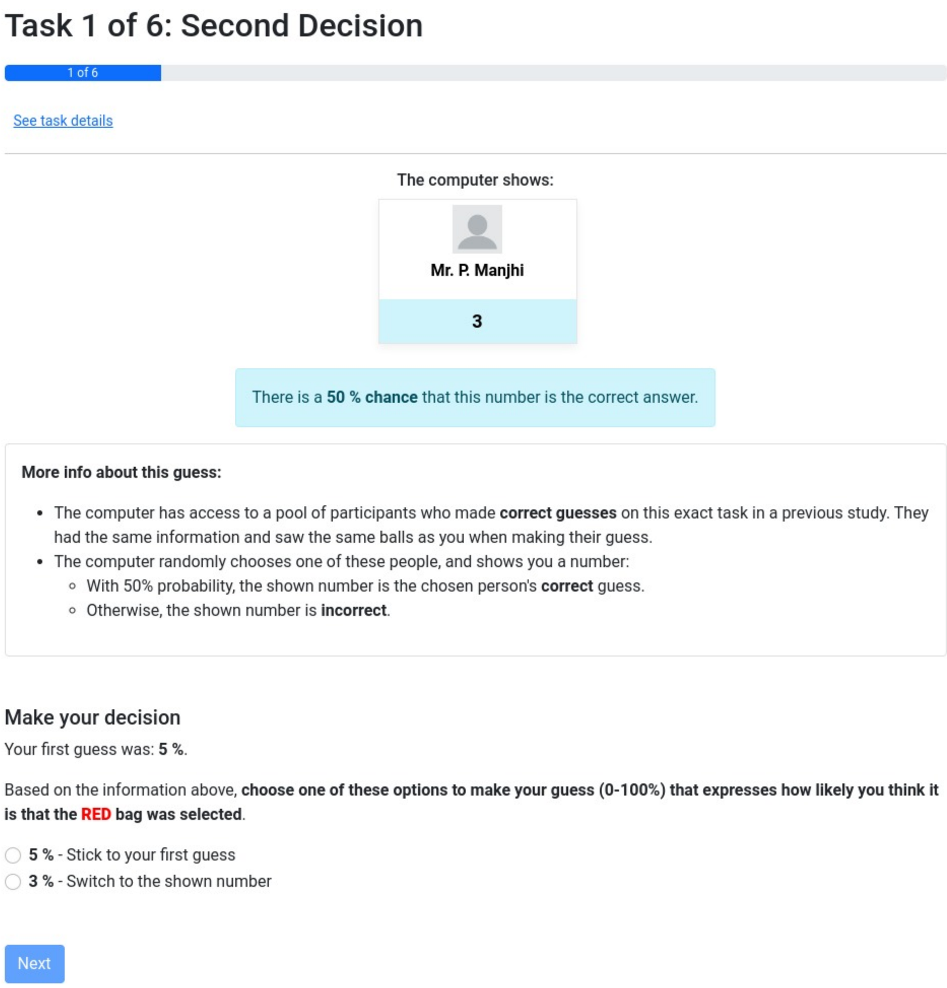

Figure D.17: Decision Tasks - Second decision

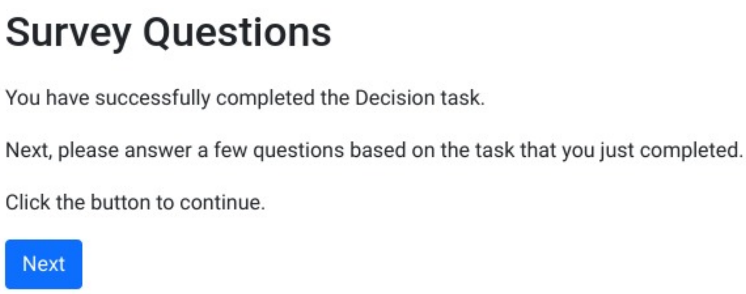

Figure D.18: Survey section

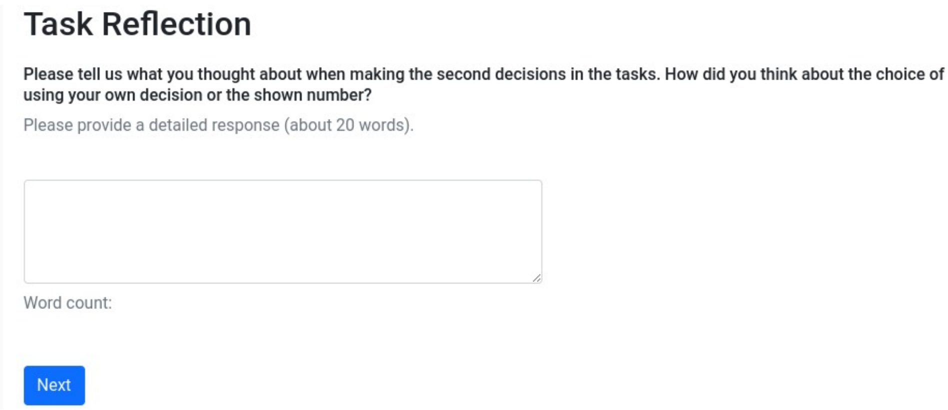

Figure D.19: Reflection question

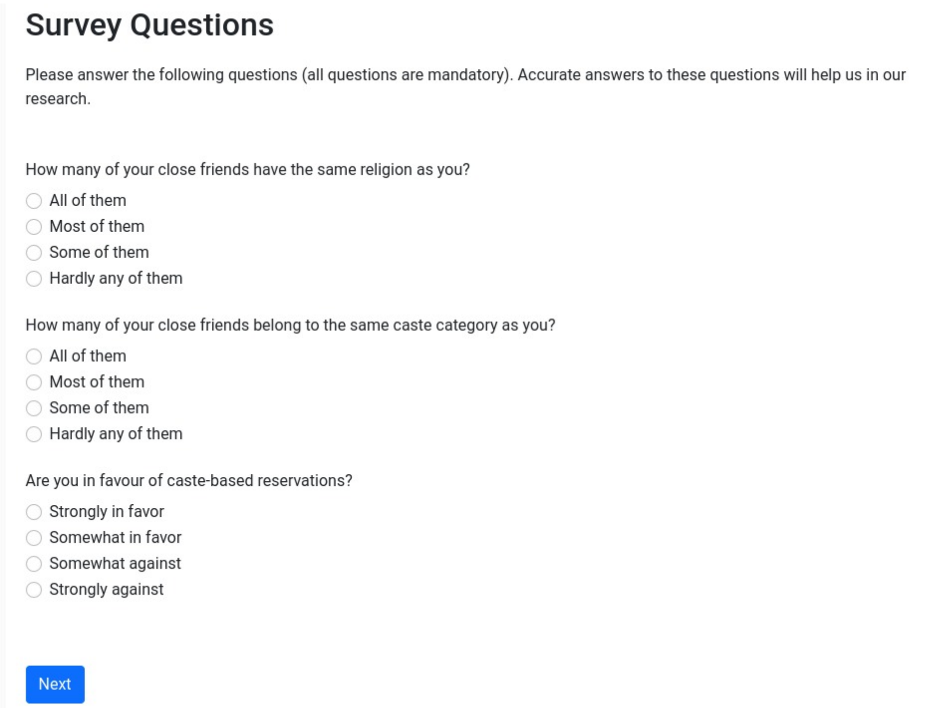

Figure D.20: Survey questions

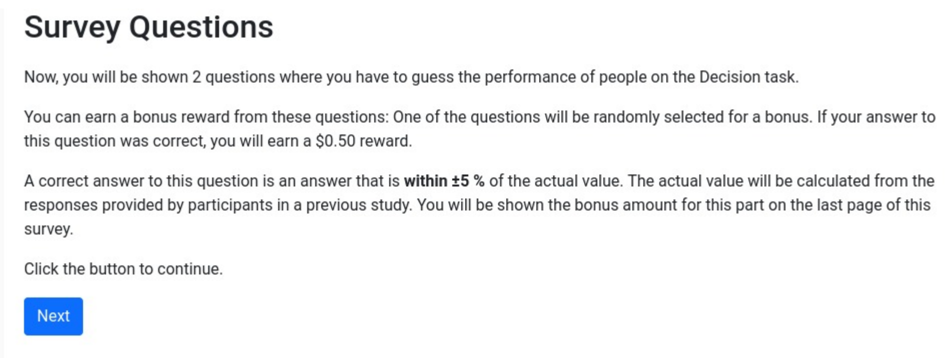

Figure D.21: Belief elicitation

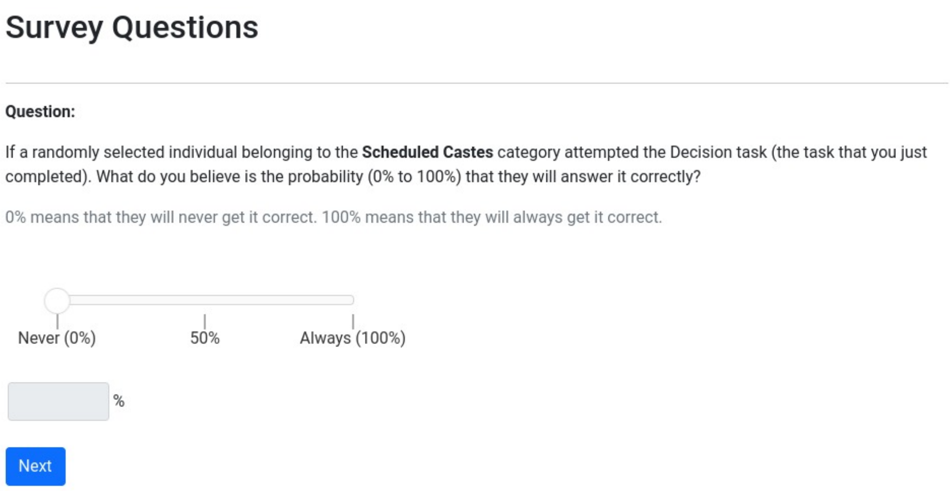

Figure D.22: Belief elicitation

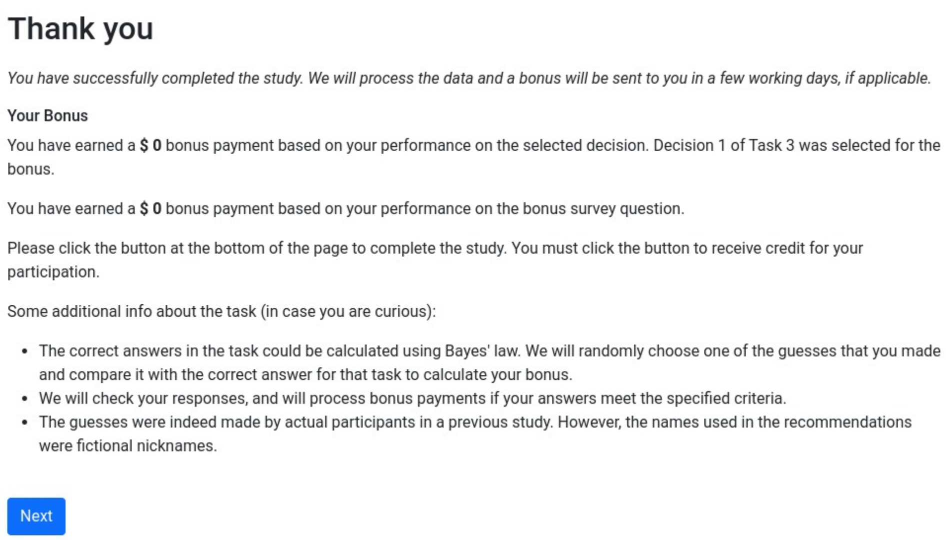

Figure D.23: Thank you and debriefing page

Caste – No Quality

Instructions for the Decision Tasks (2/3)

After making the first decision, you will be shown a number.

How is this number generated?

- The computer has access to a pool of participants who made guesses on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess.
- The computer randomly chooses one of these people and shows you the chosen person's guess.

The Second Decision

You have to make a choice:

1. Stick to your first guess, OR
2. Switch to the shown number.

Next, you will learn about how you can earn bonuses on these tasks.

Next

Figure D.24: Instructions 2/3

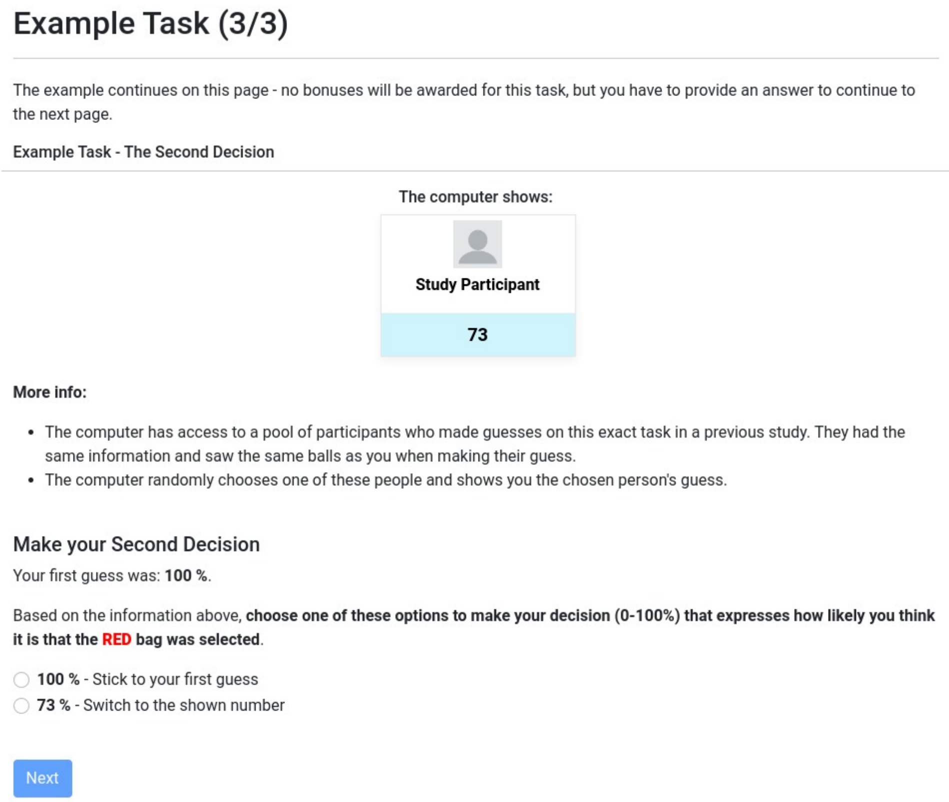

Example Task (3/3)

The example continues on this page - no bonuses will be awarded for this task, but you have to provide an answer to continue to the next page.

Example Task - The Second Decision

The computer shows:

Study Participant

73

More info:

- The computer has access to a pool of participants who made guesses on this exact task in a previous study. They had the same information and saw the same balls as you when making their guess.
- The computer randomly chooses one of these people and shows you the chosen person's guess.

Make your Second Decision

Your first guess was: **100 %**.

Based on the information above, **choose one of these options to make your decision (0-100%) that expresses how likely you think it is that the RED bag was selected**.

- **100 %** - Stick to your first guess
- **73 %** - Switch to the shown number

Next

Figure D.25: Example 3/3

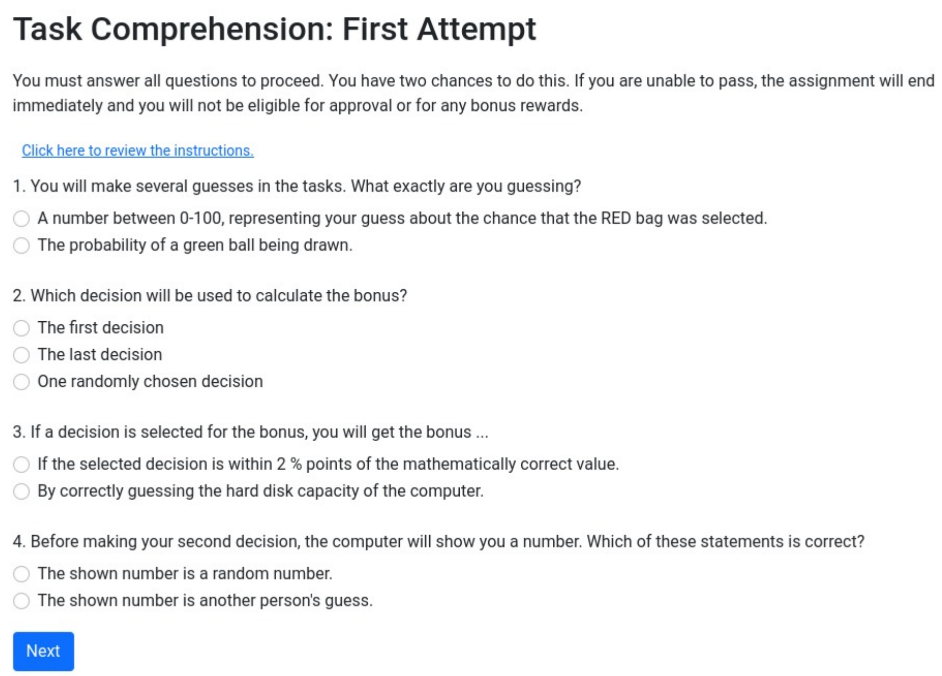

Task Comprehension: First Attempt

You must answer all questions to proceed. You have two chances to do this. If you are unable to pass, the assignment will end immediately and you will not be eligible for approval or for any bonus rewards.

Click here to review the instructions.

1. You will make several guesses in the tasks. What exactly are you guessing?
- A number between 0-100, representing your guess about the chance that the RED bag was selected.
- The probability of a green ball being drawn.

2. Which decision will be used to calculate the bonus?
- The first decision
- The last decision
- One randomly chosen decision

3. If a decision is selected for the bonus, you will get the bonus ...
- If the selected decision is within 2 % points of the mathematically correct value.
- By correctly guessing the hard disk capacity of the computer.

4. Before making your second decision, the computer will show you a number. Which of these statements is correct?
- The shown number is a random number.
- The shown number is another person's guess.

Next

Figure D.26: Comprehension - No Quality

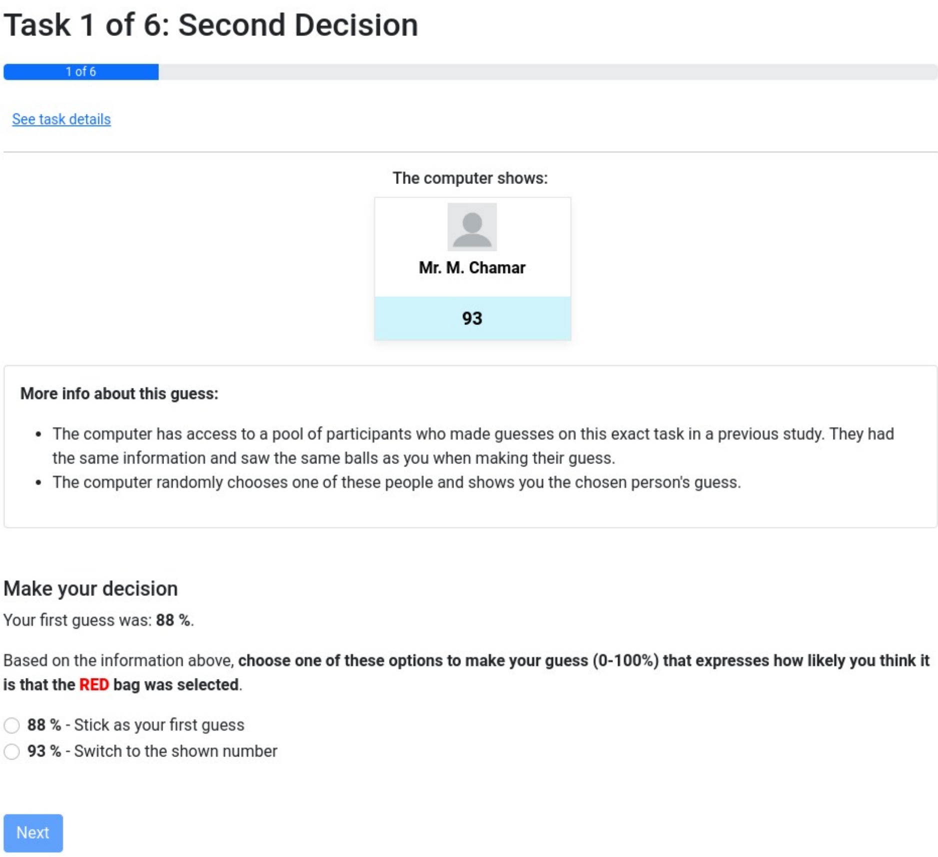

Figure D.27: Decision Tasks in *No Quality* - Second decision

Differences in other experiments

The text used to communicate the quality of the source's estimate was different in experiment *Religion, Minimal, Choice,* and *Computer*. The following screenshots are the instruction screens used in those experiments.

Instructions for the Decision Tasks (4/4)

The example continues on this page - no bonuses will be awarded for this task, but you have to provide an answer to continue to the next page. Below, you will see how to make the second decision.

The second decision

Once you submit your first decision, you will see a page similar to the one below. Here you will see another person's guess on the **same** task. This guess is drawn from a pool of guesses on this task, made by participants in a previous study.

This guess will be randomly chosen by a computer program. There is a chance with which the selected guess will be within ± 2% of the correct answer. You then have to make a choice - your second decision can be:

1. The same as your first guess, OR
2. The same as the guess made by someone else.

Remember: Each decision you make is equally likely to be chosen for the bonus, so choose the value that you think is correct.

To proceed with making your decision, click the button below. You will see a guess made by someone else on the same task.

Click to proceed

Figure D.28: Decision Tasks in *Religion* - Instructions

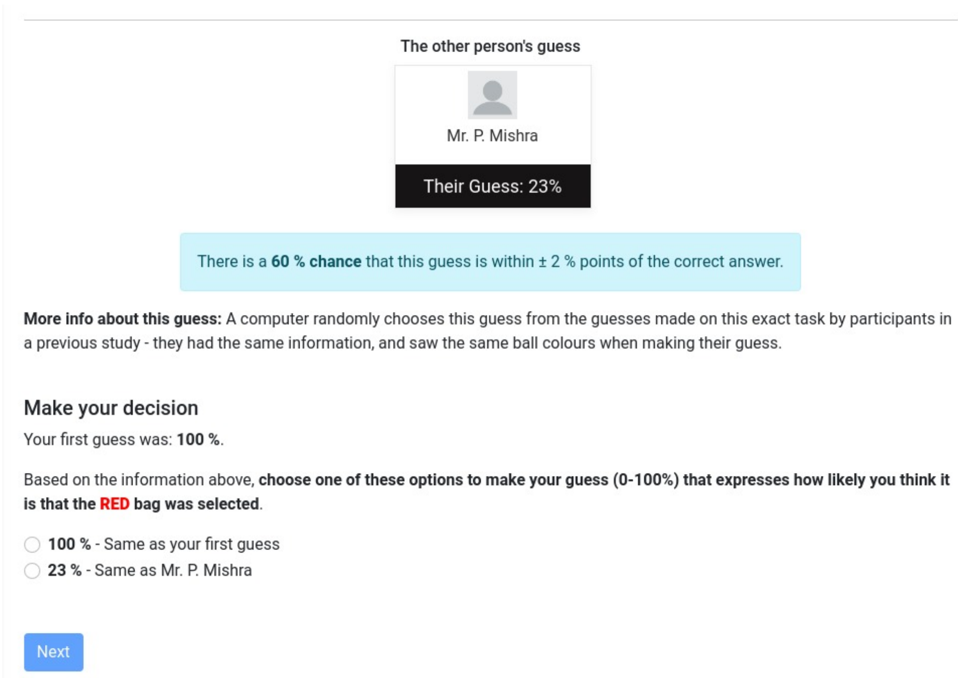

Figure D.29: Decision Tasks in *Religion* - Second guess

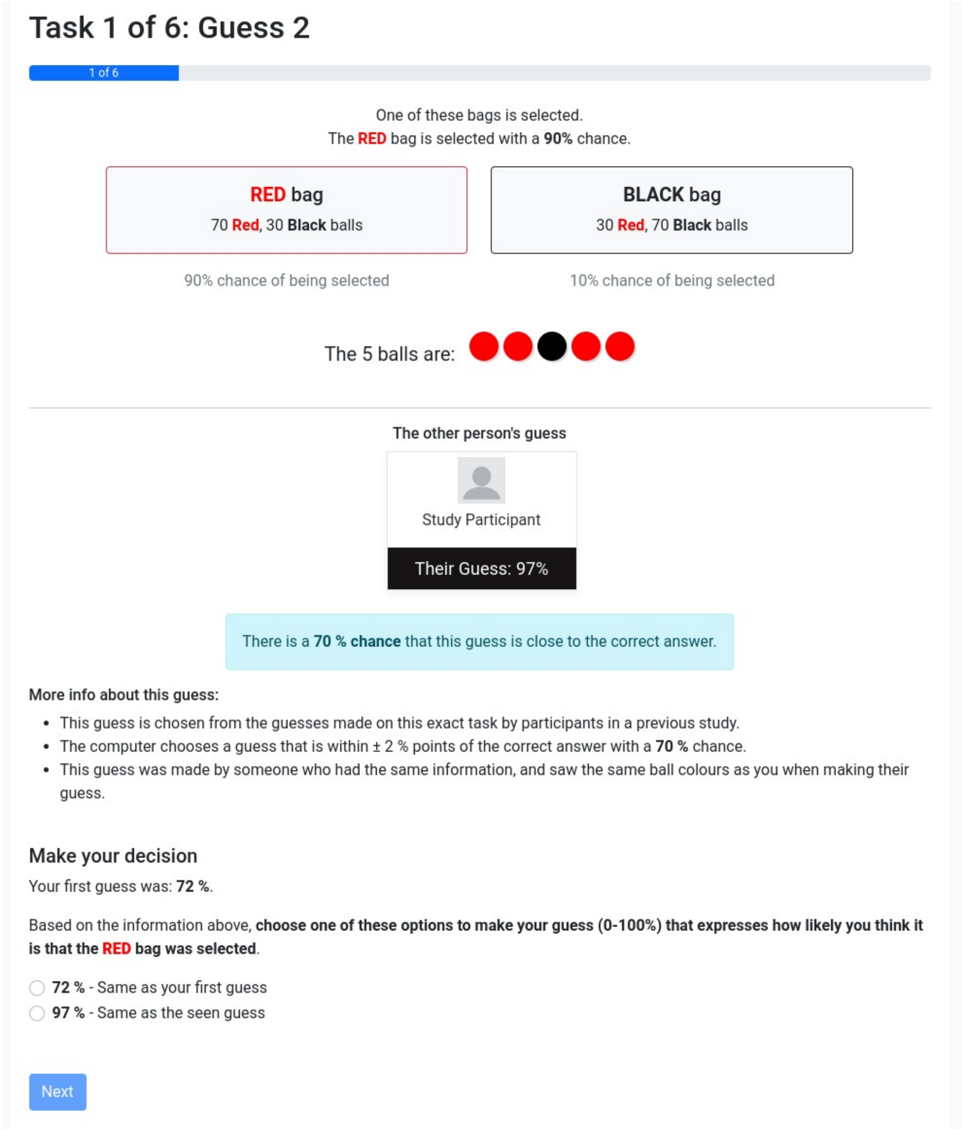

Figure D.30: Decision Tasks in *Choice* - Second guess

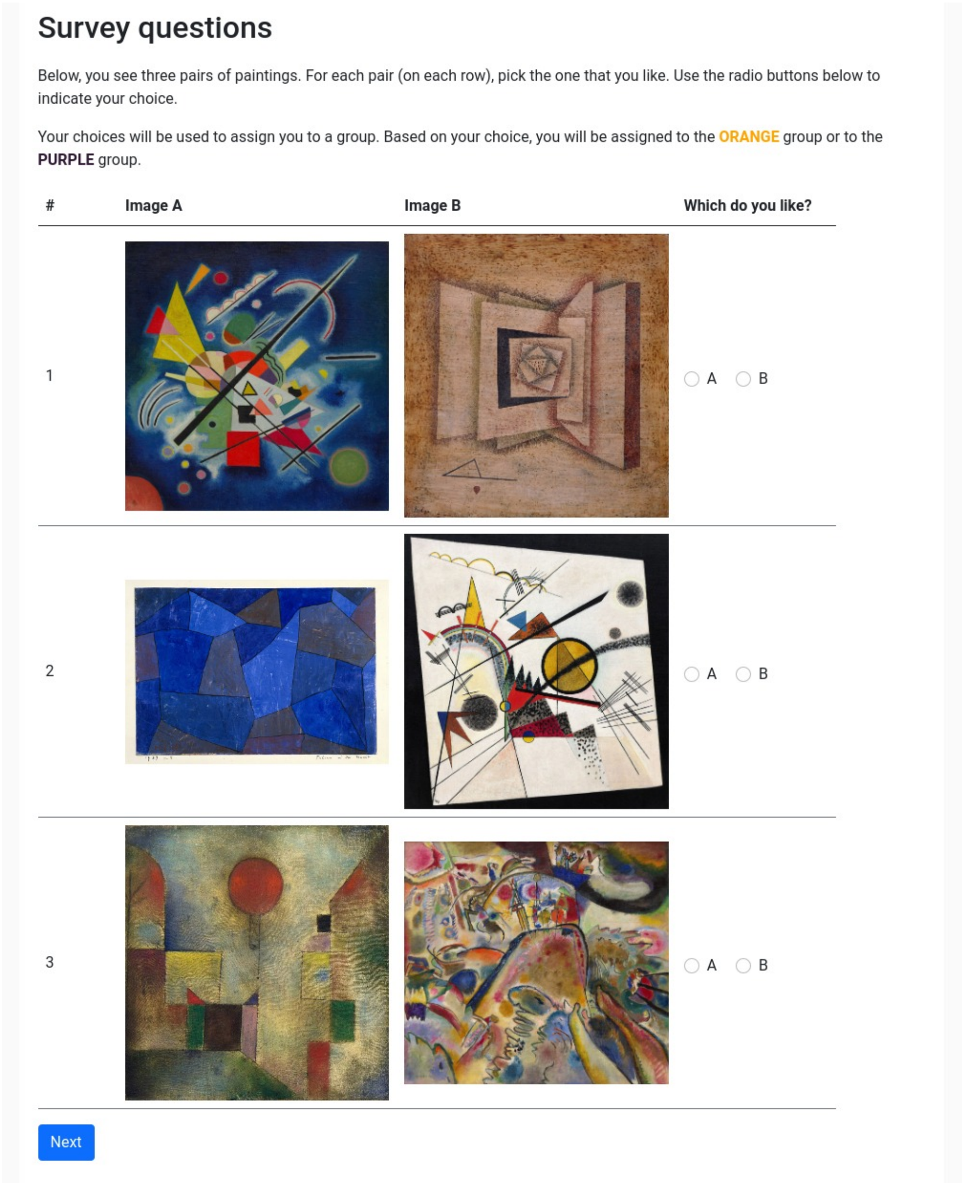

Figure D.31: Decision Tasks in *Minimal* - Klee/Kandinsky task

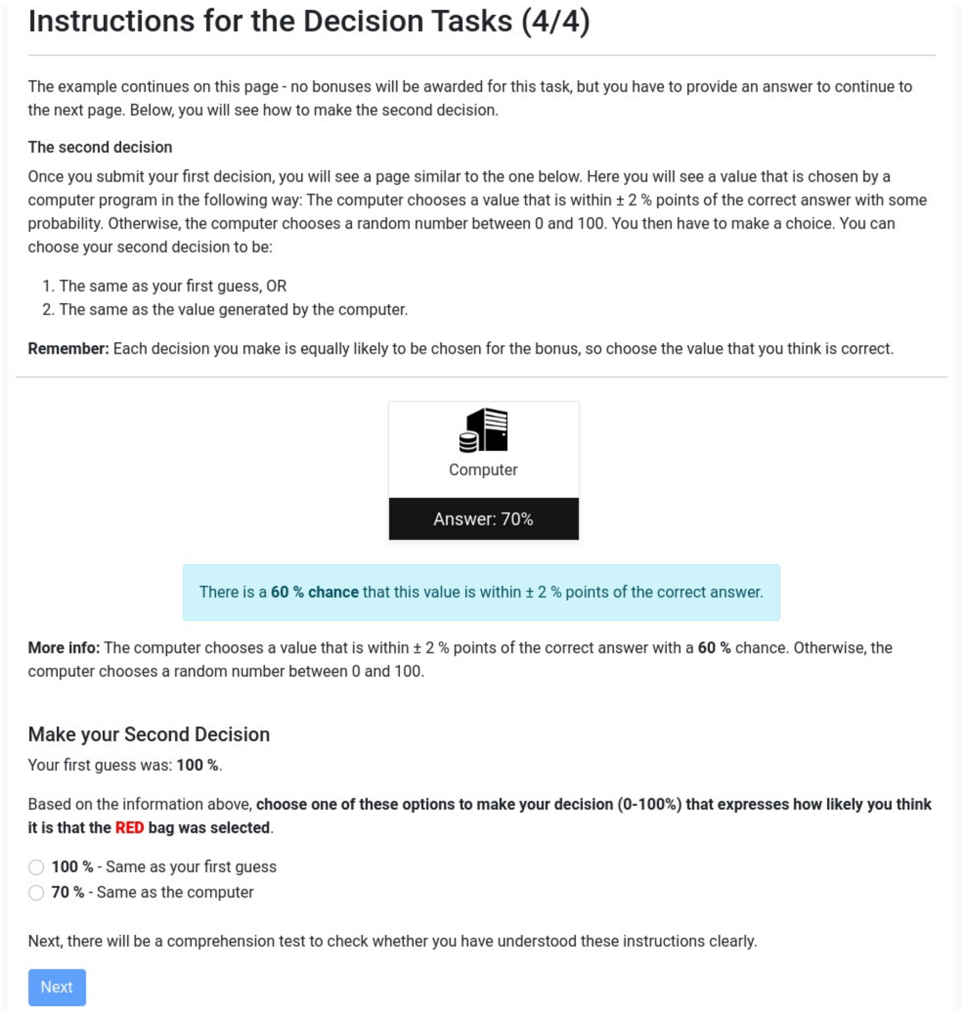

Instructions for the Decision Tasks (4/4)

The example continues on this page - no bonuses will be awarded for this task, but you have to provide an answer to continue to the next page. Below, you will see how to make the second decision.

The second decision

Once you submit your first decision, you will see a page similar to the one below. Here you will see a value that is chosen by a computer program in the following way: The computer chooses a value that is within ± 2 % points of the correct answer with some probability. Otherwise, the computer chooses a random number between 0 and 100. You then have to make a choice. You can choose your second decision to be:

1. The same as your first guess, OR
2. The same as the value generated by the computer.

Remember: Each decision you make is equally likely to be chosen for the bonus, so choose the value that you think is correct.

Computer

Answer: 70%

There is a **60 % chance** that this value is within ± 2 % points of the correct answer.

More info: The computer chooses a value that is within ± 2 % points of the correct answer with a **60 %** chance. Otherwise, the computer chooses a random number between 0 and 100.

Make your Second Decision

Your first guess was: **100 %**.

Based on the information above, **choose one of these options to make your decision (0-100%) that expresses how likely you think it is that the RED bag was selected**.

- **100 %** - Same as your first guess
- **70 %** - Same as the computer

Next, there will be a comprehension test to check whether you have understood these instructions clearly.

Next

Figure D.32: Instructions in *Computer* - Second guess

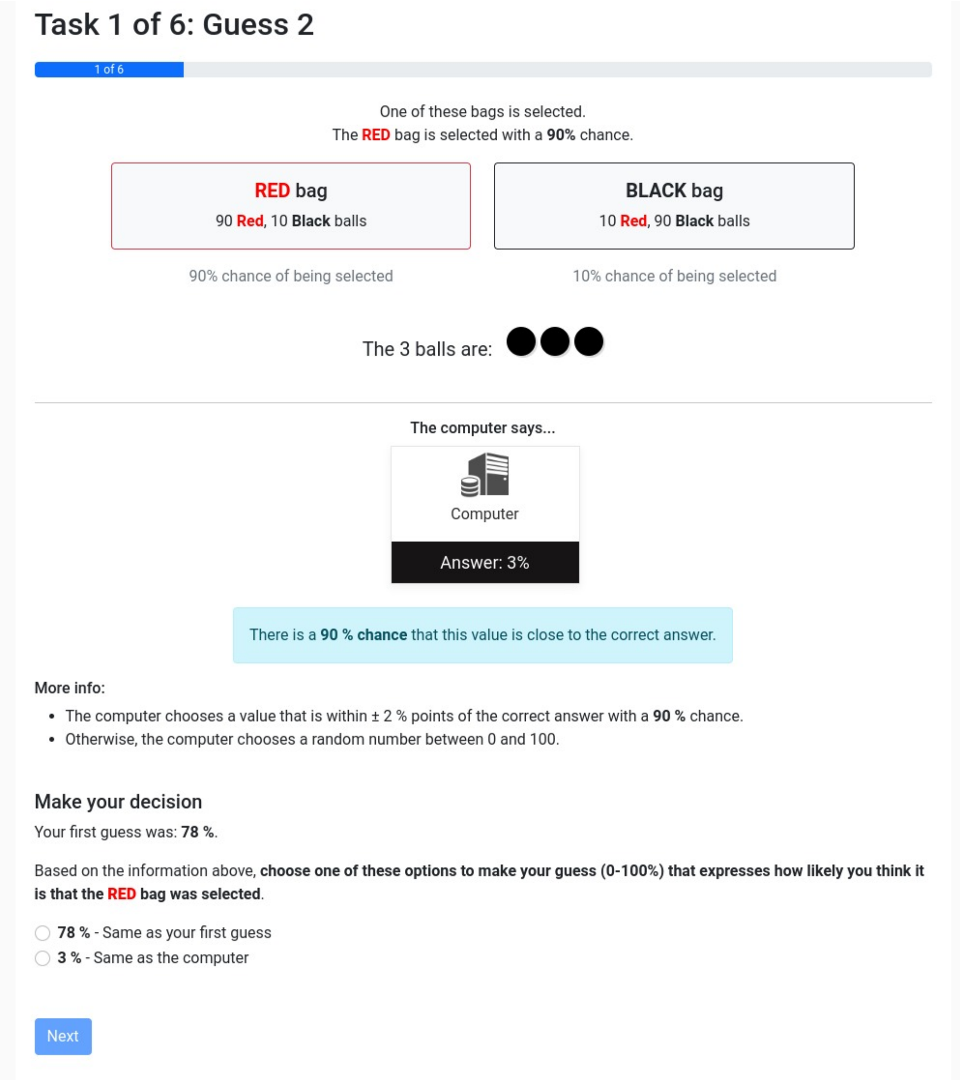

Figure D.33: Decision Tasks in *Computer* - Second guess