

How Group Identification Distorts Beliefs

M. Paula Cacault*

University of Geneva

Manuel Grieder†

ETH Zurich & ZHAW

May 21, 2019

Abstract

This paper investigates how group identification distorts people's beliefs about the ability of their peers in social groups. We find that experimentally manipulated identification with a randomly composed group leads to overconfident beliefs about fellow group members' performance on an intelligence test. This result cannot be explained by individual overconfidence, i.e., participants overconfident in their own skill believing that their group performed better because of them, as this was ruled out by experimental design. Moreover, we find that participants with stronger group identification put more weight on positive signals about their group when updating their beliefs. These in-group biases in beliefs can have important economic consequences when group membership is used to make inference about an individual's characteristics as, for instance, in hiring decisions.

Keywords: social identity, overconfidence, self-image, belief updating, discrimination

JEL codes: C92, J71, D01, C91

*University of Geneva, Geneva School of Economics and Management, Switzerland, email: maria.cacault@unige.ch

†ETH Zurich, Department of Humanities, Social and Political Sciences, Chair of Economics, and Zurich University of Applied Sciences (ZHAW), School of Management and Law, Switzerland, email: manuel.grieder@econ.gess.ethz.ch (corresponding author)

Funding: This work was financially supported by the HEC Lausanne intramural research fund.

Ethics statement: The experiment reported in this paper was approved by the HEC Lausanne ethics committee.

Declarations of interest: none

1 Introduction

This paper investigates whether and how identification with a group leads to overconfident beliefs about the ability of *other* members of that group. Models of ego-utility or self-image protection posit that people systematically inflate ego-relevant beliefs in order to feel good about themselves (e.g., Bénabou & Tirole, 2002; Kőszegi, 2006; Weinberg, 2009). Such motives can explain why individuals are on average overconfident about their personal skill or ability. We hypothesize that a similar effect occurs for social groups because people care about the image of the groups they belong to and identify with (Crocker & Luhtanen, 1990; Aberson et al., 2000; Hewstone et al., 2002). After all, the ego is not created in isolation, and an individual’s identity and self-image are decisively shaped by belonging to social groups (Tajfel & Turner, 1986).¹

Biased beliefs about the abilities of different social groups are problematic whenever people use aggregated group-level information to make inference about the individual characteristics of a particular person. Models of statistical discrimination—for instance Arrow’s (1973) theory of discrimination in the labor market—imply that in situations of incomplete information about an individual’s characteristics (e.g., ability or intelligence), rational decision makers should use knowledge about the social groups that individual belongs to for updating their beliefs.² This is optimal if, on average, group membership is correlated with the relevant characteristic.

To investigate how group identification distorts beliefs about group ability, we implement a laboratory experiment in which we randomly assign participants to groups, manipulate identification with the group via group building tasks that involve outcome dependency and competition between groups, and measure beliefs about fellow group members’ intelligence. The experiment is characterized by i) the use of an incentive compatible belief elicitation mechanism that makes it costly to distort beliefs, ii) the exogenous manipulation of group identification, iii) the exclusion of the self from the reference groups about which beliefs are elicited (in order to exclude that participants overconfident in their own skill simply believe that their group performed better because of them) and iv) the repeated elicitation of beliefs after the provision of noisy information, which allows observing biases in belief updating over time.

Our results show the causal link from identification with a group to overconfidence in the relative ability of that group. We find that participants in treatments with experimentally induced group identification have more favorable prior beliefs about the performance of their randomly assigned group members in an intelligence test than have participants in a control condition. Moreover, we find that group identification affects how participants process information when updating their

¹Akerlof & Kranton (2000) have integrated this insight into an economic model of identity, and there is a growing number of empirical studies that demonstrate the general relevance of social identities for various economic domains such as, e.g., social preferences (Ben-Ner et al., 2009; Chen & Li, 2009), coordination problems (Chen & Chen, 2011), cooperation in social dilemmas (Goette et al., 2006; Chen et al., 2014), risk and time preferences (Benjamin et al., 2010), or decision making in strategic environments (e.g., Le Coq et al., 2015; List et al., 2016; Rong et al., 2016). See also Lane (2016) for a meta-analysis of experimental work on social identity in economics.

²Empirical evidence shows that group-level information is often used to make judgments and inferences about individuals (Allport, 1954; Fiske, 1993). Moreover, people use group-level information even if information on the individual is complete and the group-level information is redundant (Albrecht et al., 2013; Reuben et al., 2014).

33 beliefs, leading people to overweight positive information about the other members of their group
34 and to discount negative information (the effects for the discounting of negative information are
35 only marginally significant, however). Together, these results suggest that people are motivated
36 to maintain a positive image of a group they identify with and adjust their social perceptions and
37 beliefs accordingly.

38 Several recent experimental papers have documented belief distortions about ego-relevant per-
39 sonal attributes. Eil & Rao (2011) and Mobius et al. (2014) find that people display an asymmetric
40 bias when updating their beliefs about intelligence, overweighting positive information and dis-
41 counting negative information.³ Their results mirror the pattern we find in this paper when people
42 update their beliefs about the skills of a social group they identify with. Schwardmann & van der
43 Weele (2017) show that people inflate beliefs about their own ability because it helps them bluff
44 and appear more skilled to others (see also Trivers, 2011). Relatedly, Di Tella et al. (2015) find that
45 people seem to systematically distort their beliefs about the altruism of others in order to justify
46 selfish acts to themselves.

47 Overconfidence in relative group performance has been investigated by Healy & Pate (2007).
48 Their set-up differs importantly from ours as they study group overconfidence in real-existing groups
49 in the lab (using campus fraternities and sororities), as well as at a scrabble tournament in the field
50 (using groups of friends). Healy & Pate find that both students in the lab and scrabble players in
51 the field display significant overconfidence in the relative performance of their group. Whereas the
52 use of real-existing groups has its advantages (Goette et al., 2012), it can pose challenges for a clean
53 identification of true overconfidence, because the experimenter lacks knowledge and control about
54 the information structure and events (prior to the experiment) that led to the elicited beliefs. This
55 is problematic because overconfidence can emerge as the result of rational information processing
56 in environments where feedback is infrequent and signals are ambiguous (Benoît & Dubra, 2011).
57 Controlling the information structure subjects face is thus key in our design for ensuring a clean
58 identification of true overconfidence. Moreover, unlike our study, Healy & Pate do not exclude
59 the self when eliciting confidence about group-level performance, thus preventing their study from
60 disentangling overconfidence in group performance from the well-documented effect of individual
61 overconfidence (see Dunning et al., 1989; Merkle & Weber, 2011; Burks et al., 2013; Benoît et al.,
62 2015, for empirical evidence on individual overconfidence).⁴ Specifically, including a subject's own
63 performance in the comparisons between groups when eliciting confidence about group performance
64 is likely to lead to an overestimation of the extent of overconfidence in group performance as it is
65 confounded by the subject's overconfidence in his or her own performance.

³Several recent papers fail to find such asymmetric updating, however. Whereas the studies by Barron (2016) and Gotthard-Real (2017) find no evidence for asymmetric updating in financial decision making domains that are not necessarily ego-relevant, neither do the studies by Buser et al. (2016), Coutts (2018), and Schwardmann & van der Weele (2017) for presumably ego-relevant cognitive skills.

⁴While investigating different research questions, also Healy & Pate (2011) and Brookins et al. (2014) find evidence that, compared to a rational benchmark, people are on average overconfident about the relative performance of a group they belong to. Both studies include the self in the group comparisons.

66 Psychologists have long studied how motivated reasoning can lead to distorted perceptions of
67 the world (see Kunda, 1990, for a review) and how group membership affects judgments about the
68 members of in- and out-groups (see Hewstone et al., 2002, for a review). Despite some existing
69 evidence for distortions in beliefs about the in-group’s ability (e.g., Bigler et al., 1997), psycholog-
70 ical studies typically rely on non-incentivized self-report measures that do not make it costly for
71 participants to distort beliefs. They do therefore not allow studying a potential bias in belief forma-
72 tion in situations characterized by a trade-off between the instrumental value of beliefs and other
73 (e.g., self-image) motives. Ensuring that belief distortions have negative monetary consequences
74 for experimental participants is further important because previous studies have shown that results
75 from non-incentivized experiments about biases in belief formation and updating do not necessar-
76 ily replicate when participants have monetary incentives for thinking about and revealing their
77 beliefs (e.g., Grether, 1980, 1992).⁵ With regard to individual overconfidence, Hoelzl & Rustichini
78 (2005) find, for instance, that patterns of overconfidence in personal skill differ importantly between
79 incentivized and non-incentivized treatments.

80 Our paper provides new insights on the effects of group identification on beliefs about social
81 groups by studying not only the static effect of group identification on beliefs but also how distor-
82 tions evolve over time and how group identification affects belief updating. Moreover, our design
83 addresses methodological issues present in earlier related studies. First, we form groups randomly
84 and manipulate group identification exogenously. This allows perfect control over the information
85 structure and ensures a clean identification of belief distortions. Second, we exclude the self from
86 the reference group about which we elicit beliefs in order to avoid a confounding effect of partici-
87 pants’ overconfidence in their personal skill. Third, we provide monetary rewards when measuring
88 beliefs and use an incentive compatible elicitation mechanism. The monetary incentives capture
89 the instrumental value of beliefs and should lead to a better measure of beliefs. The following
90 section presents our experimental design in detail. Section 3 presents the results. In section 4 we
91 conclude and discuss the implications of our findings.

92 2 Experimental Design and Procedure

93 We created an experimental set-up that allows capturing the effect of group identification on be-
94 lief distortions both by comparison with a rational benchmark as well as by comparison between
95 experimental conditions. The underlying characteristic about which we elicited beliefs is *other*
96 participants’ performance on a short IQ test. Hence, our dependent variable does not capture
97 beliefs about a participant’s personal performance on the quiz. Moreover, in order to observe belief
98 updating, participants received partial information about group performance on the test, and we
99 elicited beliefs again after each bit of information.

⁵Relatedly, both Gächter & Renner (2010) and Trautmann & van de Kuilen (2014) find that incentivizing belief elicitation leads to belief measures that are more predictive of participants’ behavior, thus indicating that compared to no incentives, incentivized methods provide a more accurate measure of participants’ actual beliefs on which they base their decisions.

Table 1: Characteristics of implemented experimental conditions

	No group framing	Group framing
No group interaction	<i>Control (C)</i>	<i>Group Framing (GF)</i>
(Anonymous) group interaction	-	<i>Group Framing + Interaction (GF+I)</i>

100 In our experimental treatments we manipulated whether a participant perceives these other
 101 participants to belong to the same group as herself and the extent of identification with that group.
 102 We compare these treatments to a control condition in which there were no groups. We describe
 103 our experimental set-up and the manipulations in detail below.

104 2.1 Setup and Experimental Conditions

105 Participants in all experimental conditions were randomly assigned to groups of four, and each
 106 group (the “in-group”) was matched to another group (the “out-group”). We varied to what
 107 extent this group assignment was made salient to participants and whether it was accompanied
 108 by additional group building tasks involving a limited and anonymous form of interaction between
 109 group members.

110 The experimental design consists of three conditions that are summarized in Table 1. In the
 111 *Control (C)* condition, we made no reference to groups whatsoever. The structure of the belief
 112 elicitation tasks, however, was exactly the same as in the other conditions (see stage C in section 2.2
 113 below for details). Participants in the *Group Framing (GF)* condition were made aware of the group
 114 structure, i.e., that there were two groups and that they had been assigned to one of them (we
 115 referred to the groups as “your group” and “other group”). Moreover, the IQ test was also framed
 116 as a group task and introduced outcome dependency within the in-group as well as competition
 117 with the out-group: The participants took the test individually, but they earned a bonus if the total
 118 score of all members of their in-group was higher than the out-group’s score (see the description of
 119 stage B in in section 2.2 below). Finally, in the *Group Framing + Interaction (GF+I)* condition,
 120 designed to induce the highest level of group identification in our study, participants faced the same
 121 group framing as in GF, and additionally engaged in a number of group building tasks in which
 122 they interacted with the other participants in the group. We intended to structure the interaction
 123 in a way that it did not reveal clear information about other participants’ intelligence (see the
 124 description of stage A in section 2.2 and also Appendix C).

125 The main goal of our treatment conditions GF+I and GF was to induce different levels of group
 126 identification. We intended to induce the highest level of group identification in GF+I, a lower
 127 level in GF and expected practically inexistent group identification in the Control condition. The
 128 most important comparison for testing the effects of group identification on belief distortions is thus
 129 the contrast between GF+I and C. The contrast between GF and C indicates whether lower levels
 130 of group identification induced without any interaction between group participants are enough to

Table 2: Timeline in the experiment

	<i>Group Framing + Interaction (GF+I)</i>	<i>Group Framing (GF)</i>	<i>Control (C)</i>
A. Group Framing	✓	✓	✗
B. Group Interaction Manipulation	✓	✗	✗
C. IQ Test	✓	✓	✓
D. Elicitation of Prior Beliefs about Group Performance	✓	✓	✓
E. Provision of Noisy Signals about Group Performance	✓	✓	✓
F. Elicitation of Posterior Beliefs about Group Performance	✓	✓	✓
G. Elicitation of Beliefs about Individual Performance	✓	✓	✓
H. Manipulation Check	✓	✓	✓

Notes: The table lists the stages of the experiment in chronological order, as they were presented to participants. Note that in the GF+I condition, the IQ test (stage C) was chronologically integrated into the group interaction manipulation (stage B). In all conditions, stages E (noisy signals) and F (posterior beliefs) were repeated three times, with a posterior belief elicited after each of the three signals.

131 trigger biases in belief formation. Importantly, our design allows for beliefs in C to differ from
 132 an unbiased rational benchmark simply because of cognitive limitations (e.g., with regard to the
 133 understanding of the belief elicitation mechanism) or general probability distortions. By comparing
 134 beliefs in GF+I and GF with beliefs in C, we control for these other sources of belief distortions.
 135 Finally, the contrast between GF+I and GF allows isolating the effect of our group building tasks
 136 that involved anonymous interaction between participants.

137 2.2 Stages in the Experiment

138 For each stage, participants received detailed written instructions and, when appropriate, answered
 139 control questions to ensure that they had read and understood the instructions. Appendix H
 140 provides a translation of the instructions. Earnings were expressed in Monetary Units (MU) that
 141 were exchanged into Swiss Francs (CHF) at a rate of 5 MU to 1 CHF. In the following, we provide a
 142 detailed description of the different stages of the experiment and the temporal flow. Table 2 provides
 143 an overview. Stages A (group framing) and B (group interaction manipulation) were unique to the
 144 GF respectively the GF+I condition, all other stages were common to all experimental conditions.

145 **A. Group Framing (GF and GF+I Only)**

146 Participants in GF and GF+I were explicitly informed that they had been assigned to groups of
147 four and that these groups would remain constant throughout the experiment.

148 **B. Group Interaction Manipulation (GF+I Only)**

149 Participants in GF+I performed a first task that consisted of designing the flag that would represent
150 the group during the experiment. Participants found envelopes on their desks containing colored
151 paper (each group had a different color: red, blue, green or yellow) and scissors. They individually
152 cut out a shape of their choice and put it back in the envelope. The experimenter collected the
153 envelopes after four minutes and an assistant outside the room pasted each of the four shapes into
154 the corners of a white paper and photographed the resulting flags. In the meantime, participants
155 read instructions for the next part. We uploaded pictures of the flags into the system and in all
156 remaining stages of the experiment, participants could see their in-group’s and out-group’s flags at
157 the top of their screens.

158 Next, participants played an interactive game that consisted of reducing the size of a colored
159 circle displayed on the screen by clicking on it. The color of the circle coincided with the color that
160 represented the out-group. All members of the group could contribute to reducing the size of the
161 circle by clicking as fast as they could. The task was presented as a competition with the other
162 group: after 30 seconds, the group with the smallest circle would win a reward of 10 MUs. Ties
163 were broken randomly. Participants never saw the other group’s circle, nor were they informed of
164 the outcome of the game until the very end of the experiment.⁶

165 After the first clicking game, participants completed the IQ test (see stage C below).

166 The final task was another competitive game that consisted of enlarging the size of a circle
167 by clicking on it. The color of the circle coincided this time with the color that represented the
168 in-group. In order to make the task slightly more interesting, this time the circle was moving
169 around on the screen. All members of the group could enlarge the size of the circle by clicking
170 on it. After 30 seconds, the group with the biggest circle would win a reward of 10 MUs. Again,
171 participants received no feedback about the outcome of this game and they could not see the out-
172 group’s circle. The main goal of the circle tasks was to create group identification via competition
173 and the experience of working together in reducing or increasing the size of the circle.

174 The group interaction tasks in GF+I were inspired by the Robbers Cave experiment (Sherif
175 et al., 1961). In this field study, conducted at a summer camp, researchers randomly assigned
176 boys to two groups and made them engage in activities aimed at enhancing identification with
177 these groups. Competitive games and the creation of flags were important elements of the identity-
178 building activities. For the purposes of our study, we tried to take these elements to the laboratory
179 in a way that sustained the anonymity of participants and that provided no clear cues about the

⁶One may argue that if participants observe the circle reduce “quickly,” they could infer that their teammates are putting high effort in the task and could therefore expect them to put high effort in the IQ test as well. However, the clicking task was easy by design: for any positive effort level, the circles reduced quickly. Moreover, GF serves as a robustness check; participants in that condition stated overconfident beliefs even though they did not play the circle task (see results in section 3).

180 actual intelligence of the other participants in the in- or the out-group (see also Appendix C).
181 We thus designed the group manipulation tasks in GF+I in a way that avoided any meaningful
182 interaction between participants, thus preventing them from inferring their fellow group members'
183 intelligence at this stage (as it would have been likely, e.g., had we used a manipulation involving
184 a group chat like Chen & Li, 2009). Importantly, no communication between participants took
185 place during the group manipulation stage, and participants were forbidden to talk with other
186 participants during the entire duration of the experiment.

187 **C. IQ Test (All Conditions)**

188 All participants completed an IQ test divided into three sections of eight questions each. The
189 questions came from Cattell (1940)'s culture-free test, and they consisted mostly of finding the
190 image that would fit a certain pattern. One point was added per correct answer to the individual
191 score. Participants had 90 seconds per section to answer as many questions as possible.⁷

192 To strengthen group identification in GF+I and GF further, the payoffs of this task were
193 determined at the group level. This reflects a feature of many relevant groups in real life, in
194 which people depend on each other's performance or ability (we discuss this point in more detail
195 in section 2.3 below). Participants in GF+I and GF thus received a reward of 10 MUs if their
196 in-group's score (the sum of the four members' individual scores) was higher than the out-group's.
197 In the Control condition, participants received the reward if their individual score situated them
198 within the top 50% of all participants in the session. In all conditions, ties were broken randomly.
199 Importantly, tests were taken individually and there was no interaction or cooperation during the
200 IQ test stage between participants in any of the experimental conditions. Moreover, participants
201 received no feedback about individual nor group scores until the very end of the experiment. The
202 scores from the IQ test are our measures of individual intelligence and the sum of individual scores
203 are our measures of group intelligence. In the following stages of the experiment, we elicited beliefs
204 about group intelligence.

205 **D. Elicitation of Prior Beliefs about Group Performance (All Conditions)**

206 In all conditions, subjects were asked to compare the performance in the IQ test of two groups
207 of three participants each. For participants in GF+I and GF, we used the previously introduced
208 group framing and asked participants to state the probability that the *other* three members of their
209 in-group scored better than three (randomly selected) members of the out-group.⁸ For participants
210 in C, in contrast, we used a neutral framing and asked them to compare the score of three randomly
211 selected participants (whom we referred to neutrally as "X, Y, and Z") in the session with the score
212 of three other randomly selected participants ("A, B, and C") in the same session. Notice that the
213 comparison participants make in C is equivalent to the one in GF+I and GF, considering that we

⁷The average number of correct answers in the IQ test was 12.4 out of 24, with a minimum of 6, a maximum of 20, and a standard deviation of 2.5. At the level of the groups of three participants whose scores were compared in the belief elicitation stages, the average group-score was 37.2 with a minimum of 26, a maximum of 51, and a standard deviation of 4.7. Ties in the overall group scores occurred in 2 out of 96 cases. There was thus enough variance between individuals and groups on the intelligence measure for meaningful comparisons.

⁸The comparison with three members of the out-group keeps the size of the reference groups constant.

214 formed groups randomly and that no interaction that could reveal anything about the fellow group
215 members’ intelligence took place before the belief elicitation.

216 To elicit beliefs we implemented a mechanism based on reservation probabilities (Karni, 2009).⁹
217 We presented two mutually exclusive options to participants: a “lottery” and a “tournament”. By
218 playing the lottery, participants could earn a reward of 10 MU with probability $\mu\%$. By playing
219 the tournament, participants could earn the same reward if their in-group scored higher than the
220 out-group in the IQ test. Their task was to state a value μ above which they preferred to play the
221 lottery rather than the tournament. A random number $y \in [0, 100]$ was drawn and participants
222 received the reward with probability $y\%$ if $\mu \leq y$, else they received the reward if their in-group
223 scored better than the out-group. They did not get information about realized earnings until the
224 end of the experiment. In Appendix B we provide a more detailed discussion of the implications of
225 the belief elicitation mechanism on participants’ monetary incentives in the experiment.

226 The value μ is our variable of interest and corresponds to a participant’s subjective probability
227 that the in-group scored higher than the out-group. An unbiased prior should be $\mu_0 = 50\%$ for
228 two reasons. First, we formed groups randomly, which implies that the more intelligent individuals
229 have the same probability to be in either group. Second, we excluded the self from the reference
230 group, which implies that individuals cannot extrapolate their knowledge about their individual IQ
231 to the in-group.¹⁰

232 Importantly, to ensure that participants correctly understood the belief elicitation mechanism
233 and the consequences of their decisions, they answered a series of control questions that involved
234 the computation of payoffs of this stage given different decision and case scenarios. Participants
235 could not start until they, and all other participants, had correctly answered all control questions.

236 **E. Noisy Signals about Group Performance (All Conditions)**

237 Directly after eliciting prior beliefs μ_0 , we showed participants three binary signals about actual
238 group performance on the IQ test and we elicited posterior beliefs, μ_1 , μ_2 and μ_3 , after each one of
239 them (see stage F below). The computer constructed signals by randomly selecting three questions
240 from the IQ test (without replacement) and comparing groups’ scores on these three questions. A
241 positive (negative) signal meant that the in-group scored higher (lower) than the out-group on the
242 selected questions. Ties were broken randomly, and we informed participants in the instructions
243 about the tie-breaking procedure. In this design a positive signal for one group implies a negative
244 signal for another group in the same session, and the number of positive signals thus equals the
245 number of negative signals in each session and in all experimental conditions.

246 Notice that the random assignment of participants to groups makes it likely that we observe a
247 large number of ties,¹¹ which means that the signals are not very informative. This is an interesting

⁹This belief-elicitation method is similar in spirit to the multiple price-list approach by Becker et al. (1964) that is frequently used to elicit reservation prices in experiments. Ducharme & Donnell (1973), Allen (1987), and Holt (2007) have proposed variants of the same approach. Schlag et al. (2015) provide a more detailed discussion and a comparison to other belief elicitation mechanisms.

¹⁰Appendix C provides evidence that participants in GF+I could not infer relative performance of their in-group in the IQ test from the tasks they performed in the group manipulation stage.

¹¹Ties occurred in 66 out of 288 score comparisons in the three periods. Table A2 in Appendix A shows the

248 case, however, because also in many real-life settings, e.g., on the labor market, the signals decision
249 makers observe with regard to the underlying qualities and skills of members of certain social groups
250 are often very noisy and hard to interpret. In fact, our set-up mirrors a situation in which there
251 are *no* actual differences in performance between groups on average to be expected (because of
252 the random assignment of individuals to groups). The experiment thus tests whether even in such
253 situations, distorted beliefs about groups may emerge when individuals identify with a group.

254 **F. Elicitation of Posterior Beliefs about Group Performance (All Conditions)**

255 After each signal, we elicited posterior beliefs (μ_1, μ_2, μ_3) using the same mechanism as for the
256 elicitation of prior beliefs in stage D. There was no additional waiting time induced before the signals
257 were presented and posterior beliefs were elicited. The experiment proceeded with the presentation
258 of the first (respectively next) signal as soon as all participants in a session had entered their beliefs.

259 **G. Elicitation of Beliefs about Individual Performance (All Conditions)**

260 In this last stage of the experiment, using the same mechanism as in stage C, we elicited participants’
261 beliefs about the probability that they obtained an individual IQ score that placed them among
262 the top-half of all participants in the session. Appendix D shows that beliefs about individual
263 performance were not significantly correlated with beliefs about group performance.¹²

264 **H. Manipulation Check (All Conditions)**

265 After collecting all incentivized measures on beliefs, in the post-experimental questionnaire we also
266 implemented a one-item self-report manipulation check measuring participants’ group identification.
267 We used the same wording as Chen & Li (2009). In the GF+I and GF conditions, the question was:
268 “On a scale from 1 to 10, please rate how closely attached you felt to your own group throughout
269 the experiment.” For the Control condition, where we did not make reference to groups, but
270 participants reported their beliefs about whether the randomly selected participants “X, Y, and
271 Z” performed better than three other randomly selected participants (“A, B, and C”), we adapted
272 the question accordingly: “On a scale from 1 to 10, please rate how closely attached you felt to the
273 participants X, Y and Z throughout the experiment.”¹³

number of ties per period and condition, with associated p-values from t-tests. The Table shows evidence that ties were balanced across conditions, except in period 3 in which ties in GF+I were twice as frequent as in C.

¹²Participants in GF+I completed three more stages after the elicitation of beliefs: a hiring stage, in which they could “hire” a (random) member of the in- or the out-group (and were paid according to that member’s IQ score); a stage in which we elicited beliefs about the relative performance of the remaining two (“neutral”) groups in the session (intended originally to be a within-subject control condition); and a willingness-to-pay stage where participants could buy information about the relative performance of their in-group and/or their individual relative performance. We abandoned the within design, and with it the additional stages, because of concerns about order effects. Importantly, the additional stages in GF+I came only after stages C and D, in which we elicited the beliefs reported here.

¹³Note that as the manipulation check occurred at the very end of each session, the three additional stages that participants in the GF+I condition completed might be potentially problematic for the comparability of the manipulation check data from GF+I with the other conditions. We therefore only rely on the manipulation check data as additional evidence. However, we believe that if anything the additional stages were likely to weaken the group identification measured in the manipulation check in GF+I, as they increased the time elapsed between the group identification manipulation and the measure.

2.3 Discussion of the Experimental Design

Some elements of our experimental design and manipulations merit a specific emphasis and discussion. First, in the GF and GF+I treatments the participants' payoffs depended on group performance, whereas in the Control condition, they were determined by individual performance. This creates an outcome dependency, which has been previously used as a group identification manipulation (Charness et al., 2007). Outcome dependency is a feature of many relevant social groups. An important example are work teams, where outcomes and thus payoffs typically depend on group performance and can often not be linked to individual performance. In fact, in any (business or other) organization a certain outcome dependency between members exists, as all members of an organization benefit if the organization performs well, even if the individual contributions to organizational performance may differ. Moreover, statistical discrimination can lead to outcome dependency for members of a social group, e.g., with regard to their labor market outcomes. If decision makers use group membership to form—even in an unbiased fashion—beliefs about the skills of an individual, the outcome of one member of the group will be correlated with that of other members of the same group (see, e.g., Coate & Loury, 1993). Thus, outcome dependency may occur in social groups even if the members are not linked by membership to a formal organization.

Second, competitive elements (present mainly in the circle clicking tasks but also in the IQ test) were part of our manipulations designed to increase group identification. We implemented competition as an element of our group manipulation, because the existing literature indicates that competition between groups leads to a higher salience of groups as a source of social identification (e.g., Campbell, 1965; Tajfel, 1982) which increases group identification (e.g., Friedkin & Simpson, 1985; Eifert et al., 2010) and can lead to belief distortions in the form of prejudice (e.g., Allport, 1954; Esses et al., 2005).

Thus, our group identification manipulations in GF+I and GF consist of different elements (in particular, group building through interaction and “working together” in the the flag and circle clicking tasks, and outcome dependency and competition in the circle clicking and IQ tasks) that work together to increase group identification. The goal of our study was not to test the causal effect of each of these tasks individually. Instead, we wanted to test the effects of identification with a group on beliefs about that group. To do so, we needed to have a group manipulation that was strong enough to affect group identification in an anonymous laboratory setting without any real interaction between participants. This is why we opted for a combination of the different elements we used.

2.4 Data Collection

We ran 12 sessions (4 per condition) with 16 participants each, yielding a total of 192 participants (46% women). We recruited participants from the subject pool for behavioral experiments at the University of Lausanne, Switzerland, using ORSEE (Greiner, 2015), and assuring that each subject participated only once. The subject pool includes undergraduates from all disciplines and we did not specify any exclusion restrictions. At the beginning of each session, participants signed a

312 consent form in line with the requirements of the university’s Ethics Committee. Interactions in
313 the laboratory were anonymous and computerized using z-Tree (Fischbacher, 2007).

314 Sessions lasted on average roughly for an hour (with sessions in the C and GF treatments being
315 closer to 55 minutes and sessions in the GF+I treatment being closer to 65 minutes). Participants
316 earned on average 20.10 CHF (including a show-up fee of 8 CHF). Participants received their
317 payments in cash. We ran sessions from May to October 2013 (GF+I and C), and in November
318 2014 (GF).

319 **3 Results**

320 **3.1 Manipulation Check**

321 The answers to the self-report manipulation check question at the end of the experiment confirm
322 that participants identified themselves more with their group in GF+I and GF than in C. The mean
323 identification in GF+I was 5.45 (on a 1-10 scale) whereas it was 3.47 in C, which is significantly
324 lower ($t(126) = 4.13, p < .001$).¹⁴ In GF the average was 4.80, which is also significantly higher
325 than in C ($t(126) = 2.81, p = .006$), but not significantly lower than in GF+I ($t(126) = 1.41,$
326 $p = .161$).¹⁵

327 In Appendix G we use this group identification measure as main explanatory variable for prior
328 and posterior beliefs in an instrumental variable specification. The IV specification has some caveats
329 that we discuss in that Appendix, but it provides a robustness check for the results we report in
330 the remainder of this section.

331 **3.2 Prior Beliefs**

332 On average, participants in C reported a prior belief of 54.92%. In the GF condition, the average
333 reported prior was 60.81%, which is significantly different from C ($t(126) = 2.30, p = .023$). The
334 average reported prior in GF+I was 60.56%, which is also significantly different from C ($t(126) =$
335 $2.33, p = .022$).

336 Figure 1 plots the empirical cumulative distribution functions of prior beliefs (μ_0) by condition.
337 Beliefs in GF+I present a flatter distribution of priors than beliefs in C (two-sample Kolmogorov-
338 Smirnov test for equality of distributions: $p = .018$), which means that the average prior belief
339 is higher in GF+I than in C. In particular, the median in C coincides with the unbiased prior of
340 50%, as the cumulative distribution function crosses 0.5 at a prior of 50%. The median prior in

¹⁴Test statistics and p -values from two-sample t -tests with equal variances. All p -values reported in this paper are for two-tailed tests.

¹⁵Lane’s (2016) meta-analysis finds no significant effect of group building exercises (such as in GF+I) when compared to minimal group manipulations without any group building stage. Our manipulation check data are consistent with this result. Moreover, Zaunbrecher et al. (2017) tested the effects of two separate elements of our GF+I manipulations (the circle tasks and the flag building) on social preferences towards in- and out-groups in isolation, and find no significant difference compared to a minimal group condition. However, as we will see, we did find different belief updating patterns in GF+I compared to GF, which may indicate that the group identification fades out faster in GF than in GF+I.




Figure 1: Empirical cumulative distribution functions of prior beliefs by condition.

The Figure plots the empirical cumulative distribution function of participants' prior beliefs (μ_0) about the probability that the in-group scored higher than the out-group on the IQ test. Two-sample Kolmogorov-Smirnov tests for equality of distributions reject that priors in the *group framing + interaction* and *control* condition are the same ($p = .018$) and that priors in the *group framing* and *control* condition are the same ($p = .047$).

341 GF+I is 60%. Interestingly, the empirical distribution of prior beliefs in GF is also flatter than the
 342 distribution of priors in C ($p = .047$), which suggests that group framing is sufficient to create a
 343 sense of group belonging that favors the distortion of prior beliefs. The distributions of prior beliefs
 344 in GF+I and GF are not significantly different ($p = .677$). Taken together, this is first evidence
 345 that beliefs are distorted in an in-group enhancing way.

346 Table 3 provides regression results that confirm this interpretation. Column 1 reports a simple
 347 specification without control variables: the means of the prior beliefs in GF+I ($p = .021$) and
 348 GF ($p = .022$) are significantly higher than in C.¹⁶ There is no significant difference in average
 349 priors between GF+I and GF ($p = .922$). In column 2 we report estimates of a model that
 350 includes as control variables the individual IQ score (standardized), gender with a dummy for
 351 female participants, age in years and a measure of risk aversion. The risk aversion measure is a
 352 self-reported level of risk aversion collected in the post-experimental questionnaire using the item
 353 suggested by Dohmen et al. (2011). The original variable is integer and ranges from 0 (totally
 354 risk-loving) to 10 (totally risk-averse), but we used the standardized version in the regressions.

¹⁶We do not cluster standard errors at the group level in the regressions reported in this section because all decisions were made individually and at no point participants received any relevant information about other participants' decisions or characteristics. This is also the case in GF+I, as the group manipulation tasks did not reveal any meaningful information about the other members of the in- or out-group. See section C of the Appendix for a more detailed discussion of the appropriate level of clustering. Nevertheless, Appendix F reproduces the main regression analyses but reporting standard errors clustered by group. All our main results are robust to clustering at the group level.

355 Table A1 in Appendix A provides descriptive statistics of these variables for each experimental
 356 condition. The treatment effects estimated are robust to adding these variables, and the control
 357 variables' coefficients are all insignificant.

Table 3: OLS regressions for μ_0

	(1)	(2)
GF+I	5.641** (2.424)	5.700** (2.441)
GF	5.891** (2.560)	5.803** (2.540)
female		2.457 (2.113)
age		0.399 (0.955)
IQ score		0.354 (1.134)
risk aversion		0.888 (0.991)
Constant	54.922*** (1.728)	53.793*** (2.124)
Observations	192	192
R^2	0.036	0.047

Notes: *** p<.01, ** p<.05, * p<.10. Robust SE in parentheses.

GF+I=1 if participant in GF+I condition; GF=1 if participant in GF condition; GF+I=GF=0 if participant in C condition.

Age and IQ score standardized.

Risk aversion standardized based on self-reported 0-10 scale.

358 We argue that prior beliefs in GF+I and GF are higher than in C because participants identify
 359 more strongly with fellow in-group members in these conditions. One might think, however, that a
 360 potential confound in this interpretation is that participants' incentives for performance in the IQ
 361 test are not identical in the control and the other two conditions. Whereas participants in C are
 362 rewarded according to their individual relative performance, participants in the other conditions
 363 are rewarded according to their groups' relative performance. If people have to exert effort in
 364 attaining scores, one might argue that the observed prior beliefs represent beliefs about higher
 365 effort contributed in group contests, rather than overconfidence in the in-group's ability. Whereas
 366 the actual performance in the IQ test does not support this interpretation (average IQ scores are
 367 12.4, 12.2 and 12.6 in C, GF, and GF+I respectively, with no statistically significant differences),
 368 we cannot be sure that participants' beliefs about others' performance correspond to others' actual
 369 performance. However, if people do expect performance to be higher in group contests, they should




Figure 2: Dynamics of beliefs by condition

The Figure plots average beliefs about the probability that the in-group scored higher than the out-group in the IQ test before and after the realisation of the signals. μ_0 is the prior belief and μ_t is the posterior belief reported after the realisation of the t -th signal.

370 expect the same “extra effort” from both the in- and the out-group, which would then not translate
 371 into a prior in-group bias. Because expecting the out-group to provide the same extra effort as the
 372 in-group might require a higher level of reasoning, and because the level of reasoning should be
 373 correlated with a participant’s score in the IQ test, we can try to test for the severity of this problem
 374 by analyzing whether the bias in prior beliefs that we find for the GF and the GF+I treatment
 375 is comparable for high and low IQ participants. Indeed, based on regression analyses (reported in
 376 Table A3 in Appendix A), in which we interact participants’ IQ scores with the treatment dummies,
 377 we find that the bias in prior beliefs does not depend on participants’ scores in the IQ test. This
 378 indicates that the treatment effects are unlikely to be driven by beliefs about other’s effort.

379 3.3 Posterior Beliefs and Information Processing

380 Figure 2 displays the dynamics of average beliefs as participants incorporate the information from
 381 the signals they receive into their posterior beliefs. Remember that by design there is an equal
 382 number of positive and negative signals in all experimental conditions (because a positive signal for
 383 one group implies a negative signal for another group). The Figure shows that the average posterior
 384 belief in GF+I is always above those in the other two conditions. Moreover, it remains above 50%
 385 until the elicitation of the third and last posterior. In C, in contrast, the beliefs fluctuate somewhat
 386 randomly around the average rational benchmark of 50%.¹⁷

387 The beliefs of participants in GF converge almost linearly to the posteriors of control partici-

¹⁷Note that in Figure 2 we do not show standard error bars. The reason is that standard errors necessarily grow bigger over time as participants receive more information, because members of successful groups (who receive positive signals) become more confident and members of unsuccessful groups (who receive negative signals) become less confident with more information.

388 pants. Thus, whereas the group framing is effective in generating a biased prior, the effect of group
389 framing in GF seems to disappear much faster than the effect of the group manipulation tasks in
390 GF+I (see Chen & Li, 2009, for a similar pattern).

391 In the following, we provide an econometric analysis of belief updating in response to “good”
392 and “bad” signals. After the elicitation of prior beliefs, participants received three binary signals
393 about the relative performance of their in-group. Signals were constructed as follows. We randomly
394 drew (without replacement) three questions from the IQ test, and summed the number of correct
395 answers to these three questions for each participant. We then aggregated the score of the three
396 participants in each group and compared the aggregate scores of the two groups. A signal in period
397 t is good ($s_t = G$) if the in-group scored higher in the 3 randomly selected questions, otherwise the
398 signal is bad ($s_t = B$).

399 We estimate the following regression model:

$$\begin{aligned} \tilde{\mu}_{it} = & \alpha_0^B s_{it}^B + \alpha_0^G s_{it}^G + \beta_0 \tilde{\mu}_{i,t-1} + \\ & (\alpha_1^B s_{it}^B + \alpha_1^G s_{it}^G + \beta_1 \tilde{\mu}_{i,t-1}) \times GF + I_i + \\ & (\alpha_2^B s_{it}^B + \alpha_2^G s_{it}^G + \beta_2 \tilde{\mu}_{i,t-1}) \times GF_i + \epsilon_{it} \end{aligned} \quad (1)$$

400 for $t \geq 1$, where $\tilde{\mu}_{it} \equiv \ln(\frac{\mu_{it}}{1-\mu_{it}})$. μ_{it} is participant i 's posterior belief in period t after receiving
401 either a good (s_{it}^G) or a bad signal (s_{it}^B).¹⁸ Given that signals are informative, s^G is always positive
402 (or zero) and s^B is always negative (or zero). Therefore, a bad signal enters negatively in equation
403 (1) and hence a positive coefficient α^B implies a *decrease* of posterior beliefs. Similarly, a good
404 signal enters positively in equation (1) and hence a positive coefficient α^G implies an *increase* of
405 posterior beliefs. We allow all coefficients to differ between conditions according to the indicators
406 $GF + I_i = 1$ if participant i is in the GF+I condition, $GF_i = 1$ if she is in GF, and $GF + I_i = GF_i = 0$
407 if she is in C.

408 The effect of the GF+I treatment on beliefs is given by

$$E[\tilde{\mu}_{it}|GF + I_i = 1] - E[\tilde{\mu}_{it}|GF + I_i = 0, GF_i = 0] = \alpha_1^B s_{it}^B + \alpha_1^G s_{it}^G + \beta_1 \tilde{\mu}_{i,t-1} \quad (2)$$

409 And the effect of the GF treatment on beliefs is given by

$$E[\tilde{\mu}_{it}|GF_i = 1] - E[\tilde{\mu}_{it}|GF + I_i = 0, GF_i = 0] = \alpha_2^B s_{it}^B + \alpha_2^G s_{it}^G + \beta_2 \tilde{\mu}_{i,t-1} \quad (3)$$

410 We are interested in the difference in information processing between conditions, rather than
411 deviations from Bayesian updating. Thus, our main focus is on the significance of the coefficients in
412 (2) and (3).¹⁹ Specifically, the difference in posteriors between the GF+I and the Control condition

¹⁸Note that $s_{it}^G = \ln \frac{p_t}{1-p_t} \mathbf{1}(s_{it} = G)$ and $s_{it}^B = \ln \frac{1-p_t}{p_t} \mathbf{1}(s_{it} = B)$, where p_t is the informativeness of the signal and $\mathbf{1}(s_t = S)$ is an indicator that takes value one if the signal is S . See Appendix E for a detailed derivation of regression equation (1).

¹⁹The presence of individual effects in the errors will render the OLS estimated coefficient β_0 biased and inconsistent. No inference will be made about this coefficient. Nevertheless, our parameters of interest are consistent due to

Table 4: Interpretation of parameters

	Control (C)	Group Framing + Interaction (GF+I)
Perfect Bayesian updating	$\alpha_0^G = \alpha_0^B = 1$	$\alpha_0^G + \alpha_1^G = \alpha_0^B + \alpha_1^B = 1$
Conservative updating	$\alpha_0^G, \alpha_0^B < 1$	$(\alpha_0^G + \alpha_1^G), (\alpha_0^B + \alpha_1^B) < 1$
Asymmetric updating	$\alpha_0^G > \alpha_0^B$	$\alpha_0^G + \alpha_1^G > \alpha_0^B + \alpha_1^B$
No in-group bias in updating		$\alpha_1^G, \alpha_1^B = 0$

Note: The interpretation of parameters in the *Group Framing* condition (GF) is analogous to the reported *Group Framing + Interaction* column by replacing the subscript 1 by 2.

413 after a good signal is equal to α_1^G and the difference in posteriors between the two conditions after
 414 a bad signal is equal to α_1^B . Despite focusing on treatment differences in our analyses, model (1)
 415 also allows us to gain insights into the belief updating process: if $\alpha < 1$ participants are updating
 416 conservatively with respect to the rational Bayesian prescription, whereas if $\alpha^G \neq \alpha^B$ they give
 417 different weight to good and bad signals and update asymmetrically (Mobius et al., 2014). Table
 418 4 provides guidance to interpret the parameters in model (1): if the updating process is perfectly
 419 Bayesian, all parameters α and β should be equal to 1 (see Appendix E for more details).

420 Column 1 of Table 5 presents the OLS estimates of a pooled version of model (1) that does
 421 not separate the effects by condition. Overall, participants are in line with Bayesian updating and
 422 incorporate information correctly: both good and bad signals are considered with equal strength
 423 and the estimated α_0 -coefficients are close to one ($p = .829$ for $\alpha_0^G \neq 1$ and $p = .697$ for $\alpha_0^B \neq 1$ in
 424 the pooled model). When we include control variables, estimates in column 2 change slightly, but
 425 they are still statistically equal to one ($p = .510$ for $\alpha_0^G \neq 1$ and $p = .286$ for $\alpha_0^B \neq 1$).

426 The picture is more interesting when looking at column 3, which reports estimates of the inter-
 427 acted model (1). Participants in GF+I update differently than control participants: participants in
 428 GF+I react significantly more to positive and marginally significantly less to negative information
 429 than control participants do, that is $\alpha_1^G > 0$ ($p = .040$) and $\alpha_1^B < 0$ ($p = .085$). Moreover, we
 430 also find marginally significant evidence of asymmetric updating of beliefs within GF+I by itself:
 431 participants in this condition put more weight on positive signals about the group than they put
 432 on negative signals ($\alpha_0^G + \alpha_1^G > \alpha_0^B + \alpha_1^B$, $p = .088$).²⁰

433 In column 4 we include control variables to the interacted specification. The reaction of partici-
 434 pants in GF+I to good signals is robust to this specification: α_1^G is marginally significant ($p = .050$).
 435 However, they no longer react to bad signals differently from participants in C: α_1^B is not significant
 436 ($p = .174$). There is still statistically significant evidence that participants in GF+I update beliefs

exogenous assignment to conditions and exogeneity of signals.

²⁰In the Control condition, our data suggest a reverse asymmetry ($\alpha_0^G < \alpha_0^B$, $p = .048$ based on column 3 estimates from Table 5). Baumeister et al. (2001) argue that, without the updating process being ego-relevant in some way, a general perceptual principle applies according to which negative information looms larger in the human mind than positive information. However, this asymmetry in the Control condition becomes statistically non-significant when adding control variables ($p = .228$ based on column 4 estimates from Table 5).

437 asymmetrically ($\alpha_0^G + \alpha_1^G > \alpha_0^B + \alpha_1^B$, $p = .042$).

Table 5: OLS regressions for $\tilde{\mu}_{it} \equiv \ln[\mu_{it}/(1 - \mu_{it})]$

	(1)	(2)	(3)	(4)
α_0^G : Good signal	1.047*** (0.217)	1.139*** (0.211)	0.455 (0.330)	0.595* (0.318)
α_0^B : Bad signal	0.927*** (0.188)	0.783*** (0.203)	1.409*** (0.404)	1.148*** (0.403)
β_0 : Lagged belief	0.382*** (0.108)	0.363*** (0.109)	0.372*** (0.112)	0.343*** (0.106)
α_1^G : Good signal \times GF+I			0.918** (0.443)	0.832** (0.421)
α_1^B : Bad signal \times GF+I			-0.827* (0.477)	-0.653 (0.479)
β_1 : Lagged belief			0.021 (0.236)	0.040 (0.233)
α_2^G : Good signal \times GF			0.854 (0.649)	0.827 (0.645)
α_2^B : Bad signal \times GF			-0.593 (0.540)	-0.424 (0.553)
β_2 : Lagged belief			-0.039 (0.239)	-0.012 (0.240)
female		-0.097 (0.104)		-0.108 (0.106)
age		-0.099** (0.046)		-0.079 (0.049)
IQ score		0.003 (0.043)		0.001 (0.043)
risk aversion		0.122*** (0.046)		0.104** (0.044)
Observations	489	489	489	489
Participants	173	173	173	173
R^2	0.212	0.226	0.227	0.237

Notes: *** $p < .01$, ** $p < .05$, * $p < .10$. SE clustered by participant in parentheses. Observations with extreme current or lagged beliefs are dropped from the estimations because $\tilde{\mu}$ is not defined when $\mu = 0$ or $\mu = 1$. age, IQ score and risk aversion standardized based on original scales.

438 When focusing on beliefs in GF, the estimates reported in column (3) of Table 5 suggest that
 439 there is a slight asymmetry in the updating process, as participants in GF tend to react more to
 440 good signals than to bad signals. However, these effects are not statistically significant ($p = .190$
 441 for α_2^G and $p = .274$ for α_2^B) which implies that participants in GF do not update differently than
 442 participants in C. This suggests that group framing leads to overconfidence in priors that does not

443 survive belief updating, whereas the stronger group identification in GF+I renders overconfidence
444 persistent.²¹

445 When adding control variables in column 4, we arrive at similar conclusions regarding updating
446 of participants in the GF condition. The tests on the α_2 coefficients of the GF condition are
447 similar to the ones reported in the previous paragraph ($\alpha_2^G > 0, p = .202$; $\alpha_2^B > 0, p = .444$; and
448 $\alpha_0^G + \alpha_2^G > \alpha_0^B + \alpha_2^B, p = .229$).

449 The fact that participants in GF+I interpret good signals as being more informative about
450 in-group performance than bad signals, may explain the observed persistence of the prior in-group
451 bias. To test this channel, we performed the following thought experiment: given the observed
452 priors, what would be the (counterfactual) posteriors of participants in GF+I if they processed
453 information in the same way as participants in C? To answer this question, we simply set to zero
454 the coefficients with subscript 1 and 2 in model (1) and predict beliefs using the estimates in column
455 2 of Table 5. Figure 3 shows that without any in-group bias in information processing, beliefs of
456 participants in GF+I would rapidly converge to those of participants in C.

457 4 Discussion and Conclusion

458 In this paper we study how identification with a group leads to overconfident beliefs in the relative
459 ability of that group. We randomly assigned participants to an anonymous four-person group, their
460 “in-group” that was matched to an “out-group”, and experimentally manipulated the intensity of
461 group identification. We elicited participants’ beliefs about the probability that their fellow in-group
462 members performed better on an intelligence test than the out-group using an incentive-compatible
463 mechanism that rewards truth-telling and accuracy of beliefs. Crucially, the design excludes the self
464 from the reference group, which prevents individually overconfident participants—who are simply
465 overconfident in their own skill—from believing that their group performed better because of them.
466 We further explored how participants updated their beliefs by revealing partial information about
467 the groups’ relative performance.

468 We find evidence of, i) a “prior in-group bias”, as participants in GF+I and GF are initially
469 more confident about their in-group’s relative performance than are participants in C; and ii) a
470 “dynamic in-group bias”, as participants in GF+I do not process information about the true ability
471 of their fellow in-group members the same way as *control* participants do. Specifically, they put
472 more weight on positive information about the group and (marginally significantly) less weight

²¹On the one hand, the theory of “confirmation bias” predicts that people overweight information that is consistent with their priors and discount information that goes against them. Whereas this could explain the updating pattern observed in GF+I, confirmation bias is unable to explain why the pattern is different in GF, because the priors are similar in both conditions. On the other hand, the outcome dependency present in the GF+I and GF treatments meant that participants in these treatments had some stake in the outcome about which beliefs were elicited, whereas this was not the case in C. “Wishful thinking” may induce them to place a higher probability on the outcome “my in-group scored higher than the out-group” (Mayraz, 2011). This could explain why prior beliefs in GF+I and GF are higher than in C, but it cannot explain why posterior beliefs in GF converge faster than beliefs in GF+I. Moreover, the results by Barron (2016) and Gotthard-Real (2017) indicate that outcome desirability (without ego-relevance) is by itself not strong enough to lead to asymmetric belief updating. Finally, the outcome dependency could also have affected beliefs by triggering reciprocity considerations.

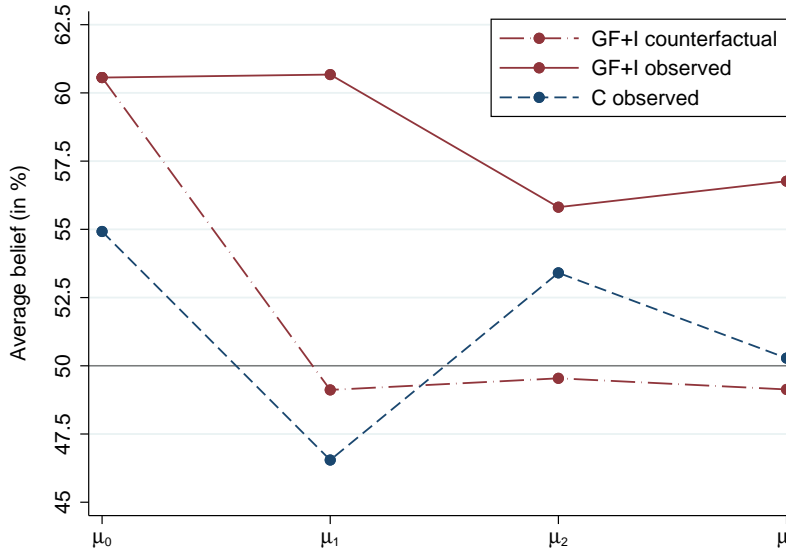


Figure 3: Counterfactual beliefs in GF+I vs. observed beliefs in GF+I and C

The Figure plots average beliefs about the probability that the in-group scored higher than the out-group in the IQ test before and after the realisation of the signals. μ_0 is the prior belief and μ_t is the posterior belief after the realisation of the t -th signal. The dash-dotted line represents the counterfactual beliefs of participants in GF+I. We computed these as the posterior belief that, given their priors, participants in GF+I would have if they updated beliefs like participants in C. The solid line represents observed beliefs in GF+I and the dashed line observed beliefs in C.

473 on negative information. This dynamic bias counteracts the effect of the two forces that should
 474 reduce in-group overconfidence over time: namely, i) learning and ii) the deterioration in group
 475 identification and attachment as decisions take place at later times after the group manipulation
 476 tasks.²² While learning should eventually erase initial differences in prior beliefs, asymmetric
 477 updating reduces the speed at which it does. Moreover, if signals arrive at a low pace or not at all,
 478 e.g., if the counter-factual is not observable as in hiring decisions, the prior in-group bias could be
 479 perpetuated.

480 The effects we show in our experiment are important because biased beliefs about group-level
 481 ability can have critical consequences when statistical discrimination based on group membership
 482 occurs, and people make forecasts about an individual's performance using group membership as
 483 additional information. Economic models of statistical discrimination usually assume that beliefs
 484 about group-level ability are rationally formed to maximize their instrumental value, and that they
 485 are therefore on average correct (e.g., Phelps, 1972; Arrow, 1973; Aigner & Cain, 1977; Coate &
 486 Loury, 1993). By showing that such beliefs may be systematically distorted in an in-group favoring
 487 way, our results shed doubt on this assumption. This means, for instance, that employers who
 488 belong to and identify with a certain social group, may systematically make mistakes in the way
 489 they infer workers' ability based on group information.²³ The discriminatory impact of such biased

²²For supporting evidence see Chen & Li (2009).

²³Consider, for instance, the results by Mobius et al. (2016) who find in a field experiment in China that majority-group employers seem to systematically underestimate the performance of minority group candidates. Similarly, Arrow (1973) discusses that employers' beliefs do not necessarily correspond to real differences between groups. In his

490 beliefs could be further magnified if prior beliefs also affect the willingness to acquire and process
491 individualized information about a candidate (“attention discrimination”, see Bartoš et al., 2016).

492 Our results add a new angle to economic theories of discrimination and contribute to a bet-
493 ter understanding of how and why discrimination occurs. From a policy making perspective this
494 is important, because different mechanisms require different responses when it comes to acting
495 against discrimination (Guryan & Charles, 2013). For instance, if people distort their beliefs about
496 group-level characteristics systematically, policy interventions targeted at changing the fundamen-
497 tals underlying these characteristics, e.g., by trying to increase human capital investments of mi-
498 nority groups, may have only limited impact, and other (potentially simple) measures such as
499 raising awareness of the bias (see, e.g., Pope et al., 2018) may be necessary to ensure the success
500 of such policies. From a theoretical perspective, our findings could provide a building stone for a
501 more nuanced economic model of discrimination, going beyond purely taste-based (Becker, 1957)
502 or statistical accounts. For this purpose, models of ego-utility (Kőszegi, 2006) could be enriched
503 by a component capturing the importance of group-level beliefs to serve as a micro-foundation for
504 modeling why and how people distort beliefs about groups. This could provide interesting new
505 insights for future theoretical and empirical work.

Acknowledgments

We thank Björn Bartling, Adrian Bruhin, Yan Chen, Florian Ederer, Claus Ghesla, Lorenz Goette, Annette Harms, Holger Herz, Michael Kosfeld, Rafael Lalive, Arno Riedl, Dominic Rohner, Luís Santos-Pinto, Jan Schmitz, Vera Shuman, Matthias Sutter, Henrik Zaunbrecher, Christian Zehnder, and Ro’i Zultan, as well as seminar participants at the Universities of Bern and Lausanne, AEW 2013 in Rome, BEER 2013 in Lausanne, SMYE 2014 in Vienna, THEEM 2014 in Stein am Rhein, the Workshop for Experimental and Behavioral Economic Research 2014 in Zurich, ASFEE 2015 in Paris, and the European Winter Meeting of the Econometric Society 2015 in Milan for valuable comments.

words, statistical discrimination occurs based on “perception of reality” that merely reflects employers’ “preconceived ideas” (p. 23) about disparities between groups.

References

- Aberson, C. L., Healy, M., & Romero, V. (2000). Ingroup bias and self-esteem: A meta-analysis. *Personality and Social Psychology Review*, 4(2), 157–173.
- Aigner, D. J. & Cain, G. G. (1977). Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review*, 30(2), 175–187.
- Akerlof, G. A. & Kranton, R. E. (2000). Economics and identity. *Quarterly Journal of Economics*, 115(3), 715–753.
- Albrecht, K., Von Essen, E., Parys, J., & Szech, N. (2013). Updating, self-confidence, and discrimination. *European Economic Review*, 60, 144–169.
- Allen, F. (1987). Discovering personal probabilities when utility functions are unknown. *Management Science*, 33(4), 542–544.
- Allport, G. W. (1954). *The Nature of Prejudice*. Oxford: Addison-Wesley.
- Arrow, K. J. (1973). The theory of discrimination. In O. Ashenfelter & A. Rees (Eds.), *Discrimination in Labor Markets* (pp. 3–33). Princeton University Press.
- Barron, K. (2016). Belief updating: Does the 'good-news, bad-news' asymmetry extend to purely financial domains? *WZB Discussion Paper, SP II 2016–309*.
- Bartoš, V., Bauer, M., Chytilová, J., & Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6), 1437–1475.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323–370.
- Becker, G., DeGroot, M., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9, 226–232.
- Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press.
- Ben-Ner, A., McCall, B. P., Stephane, M., & Wang, H. (2009). Identity and in-group/out-group differentiation in work and giving behaviors: Experimental evidence. *Journal of Economic Behavior & Organization*, 72(1), 153–170.
- Bénabou, R. & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3), 871–915.
- Benjamin, D. J., Choi, J. J., & Strickland, A. J. (2010). Social identity and preferences. *American Economic Review*, 100, 1913–1928.

- Benoît, J.-P. & Dubra, J. (2011). Apparent overconfidence. *Econometrica*, 79(5), 1591–1625.
- Benoît, J.-P., Dubra, J., & Moore, D. (2015). Does the better-than-average effect show that people are overconfident?: Two experiments. *Journal of the European Economic Association*, 13(2).
- Bigler, R., Jones, L., & Lobliner, D. (1997). Social categorization and the formation of intergroup attitudes in children. *Child Development*, 68(3), 530–543.
- Brookins, P., Lucas, A., & Ryvkin, D. (2014). Reducing within-group overconfidence through group identity and between-group confidence judgments. *Journal of Economic Psychology*, 44, 1–12.
- Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2013). Overconfidence and social signalling. *Review of Economic Studies*, 80(3), 949–983.
- Buser, T., Gerhards, L., & Van Der Weele, J. J. (2016). Measuring responsiveness to feedback as a personal trait. *Tinbergen Institute Discussion Paper, No. 16-043/I*.
- Campbell, D. T. (1965). Ethnocentric and other altruistic motives. In *Nebraska Symposium on Motivation*, volume 13, (pp. 283–311).
- Cattell, R. B. (1940). A culture-free intelligence test. *Journal of Educational Psychology*, 31(3), 161–179.
- Charness, G., Rigotti, L., & Rustichini, A. (2007). Individual behavior and group membership. *American Economic Review*, 97(4), 1340.
- Chen, R. & Chen, Y. (2011). The potential of social identity for equilibrium selection. *American Economic Review*, 101(6), 2562–2589.
- Chen, Y. & Li, S. X. (2009). Group identity and social preferences. *American Economic Review*, 99(1), 431–457.
- Chen, Y., Li, S. X., Liu, T. X., & Shih, M. (2014). Which hat to wear? impact of natural identities on coordination and cooperation. *Games and Economic Behavior*, 84, 58–86.
- Coate, S. & Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 83(5), 1220–1240.
- Coutts, A. (2018). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*.
- Crocker, J. & Luhtanen, R. (1990). Collective self-esteem and ingroup bias. *Journal of Personality and Social Psychology*, 58(1), 60–67.
- Di Tella, R., Perez-Truglia, R., Babino, A., & Sigman, M. (2015). Conveniently upset: Avoiding altruism by distorting beliefs about others' altruism. *The American Economic Review*, 105(11), 3416–3442.

- Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3), 522–550.
- Ducharme, W. M. & Donnell, M. L. (1973). Intrasubject comparison of four response modes for “subjective probability” assessment. *Organizational Behavior and Human Performance*, 10(1), 108–117.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57(6), 1082–1090.
- Eifert, B., Miguel, E., & Posner, D. N. (2010). Political competition and ethnic identification in Africa. *American Journal of Political Science*, 54(2), 494–510.
- Eil, D. & Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2), 114–138.
- Esses, V. M., Jackson, L. M., Dovidio, J. F., & Hodson, G. (2005). Instrumental relations among groups: Group competition, conflict, and prejudice. In J. F. Dovidio, P. Glick, & L. A. Rudman (Eds.), *On the nature of prejudice: Fifty years after Allport* (pp. 227–243). Wiley-Blackwell, Malden MA.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Fiske, S. T. (1993). Social cognition and social perception. *Annual Review of Psychology*, 44(1), 155–194.
- Friedkin, N. E. & Simpson, M. J. (1985). Effects of competition on members’ identification with their subunits. *Administrative Science Quarterly*, 30(3), 377–394.
- Gächter, S. & Renner, E. (2010). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics*, 13(3), 364–377.
- Goette, L., Huffman, D., & Meier, S. (2006). The impact of group membership on cooperation and norm enforcement: Evidence using random assignment to real social groups. *American Economic Review (Papers & Proceedings)*, 96(2), 212 – 216.
- Goette, L., Huffman, D., & Meier, S. (2012). The impact of social ties on group interactions: Evidence from minimal groups and randomly assigned real groups. *American Economic Journal: Microeconomics*, 4(1), 101–115.
- Gotthard-Real, A. (2017). Desirability and information processing: An experimental study. *Economics Letters*, 152, 96–99.

- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125.
- Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics*, 95(3), 537–557.
- Grether, D. M. (1992). Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, 17(1), 31–57.
- Guryan, J. & Charles, K. K. (2013). Taste-based or statistical discrimination: The economics of discrimination returns to its roots. *Economic Journal*, 123(572), F417–F432.
- Healy, A. & Pate, J. (2007). Overconfidence, social groups, and gender: Evidence from the lab and field. *Working Paper*, Loyola Marymount University.
- Healy, A. & Pate, J. (2011). Can teams help to close the gender competition gap? *Economic Journal*, 121(555), 1192–1204.
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, 53(1), 575–604.
- Hoelzl, E. & Rustichini, A. (2005). Overconfident: Do you put your money on it? *Economic Journal*, 115(503), 305–318.
- Holt, C. A. (2007). *Markets, games, & strategic behavior*. Pearson Addison Wesley Boston, MA.
- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2), 603–606.
- Kőszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4), 673–707.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480.
- Lane, T. (2016). Discrimination in the laboratory: A meta-analysis of economics experiments. *European Economic Review*, 90, 375–402.
- Le Coq, C., Tremewan, J., & Wagner, A. K. (2015). On the effects of group identity in strategic environments. *European Economic Review*, 76, 239–252.
- List, J. A., Neilson, W. S., & Price, M. K. (2016). The effects of group composition in a strategic environment: Evidence from a field experiment. *European Economic Review*, 90, 67–85.
- Mayraz, G. (2011). Wishful thinking. *CEP Discussion Paper No 1092*.
- Merkle, C. & Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Decision Processes*, 116(2), 262–271.

- Mobius, M., Niederle, M., Niehaus, P., & Rosenblat, T. (2014). Managing self-confidence. *Working Paper*.
- Mobius, M., Rosenblat, T., & Wang, Q. (2016). Ethnic discrimination: Evidence from China. *European Economic Review*, *90*, 165–177.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, *62*, 659–661.
- Pope, D. G., Price, J., & Wolfers, J. (2018). Awareness reduces racial bias. *Management Science*, *64*(11), 4988–4995.
- Reuben, E., Sapienza, P., & Zingales, L. (2014). How stereotypes impair women’s careers in science. *Proceedings of the National Academy of Sciences*, *111*(12), 4403–4408.
- Rong, R., Houser, D., & Dai, A. Y. (2016). Money or friends: Social identity and deception in networks. *European Economic Review*, *90*, 56–66.
- Schlag, K. H., Tremewan, J., & Van der Weele, J. J. (2015). A penny for your thoughts: A survey of methods for eliciting beliefs. *Experimental Economics*, *18*(3), 457–490.
- Schwardmann, P. & van der Weele, J. (2017). Deception and self-deception. *Working Paper*.
- Sherif, M., Harvey, O. J., White, B. J., Hood, W. R., Sherif, C. W., et al. (1961). *Intergroup conflict and cooperation: The Robbers Cave experiment*, volume 10. Norman, OK: University Book Exchange.
- Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology*, *33*(1), 1–39.
- Tajfel, H. & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. Austin (Eds.), *Psychology of Intergroup Relations* (pp. 33–48). Chicago: Nelson-Hall.
- Trautmann, S. T. & van de Kuilen, G. (2014). Belief elicitation: A horse race among truth serums. *The Economic Journal*, *125*(589), 2116–2135.
- Trivers, R. (2011). *Deceit and Self-Deception: Fooling Yourself the Better to Fool Others*. Penguin.
- Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.
- Weinberg, B. (2009). A model of overconfidence. *Pacific Economic Review*, *14*(4), 502–515.
- Zaunbrecher, H. W., Williams, T., & Riedl, A. (2017). Social identity in the lab: A horse race between different methods to induce social identity. *Working Paper*.

Online Appendix

A Additional Descriptive Statistics and Analyses

Table A1: Distribution of control variables by condition

	(1) All		(2) C		(3) GF		(4) GF+I	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Female	0.464	0.500	0.484	0.504	0.484	0.504	0.422	0.498
Age	21.021	2.589	21.578	3.351	20.719	2.250	20.766	1.883
p -value ^a					0.551		0.211	
IQ score	12.411	2.522	12.406	2.505	12.641	2.560	12.188	2.519
p -value ^a					0.99		0.941	
Risk aversion ^b	6.281	1.998	5.953	2.073	6.375	1.786	6.516	2.108
p -value ^a					0.699		0.551	
Observations	192		64		64		64	

a) p -values from Kolmogorov-Smirnov tests of equality of distributions (against Control condition).

b) Self reported measure on a 0-10 scale (higher as risk aversion increases).

Table A2: Number of ties in the signal generation process

Condition	Period1	Period2	Period3
Control	10	7	5
Group Framing	7	8	5
Group Framing + Interaction	6	7	11
t -test GF vs C, p -value	.233	.679	1
t -test GF+I vs C, p -value	.104	1	.014

Table A3: OLS regressions for prior beliefs μ_0 by IQ

	(1)	(2)
GF+I	5.641** (2.424)	5.603** (2.430)
GF	5.891** (2.560)	5.954** (2.589)
GF+I×IQ		-2.073 (2.362)
GF×IQ		-2.419 (3.124)
IQ score		1.684 (1.748)
Constant	54.922*** (1.728)	54.925*** (1.728)
Observations	192	192
R^2	0.036	0.042

Notes: *** p<.01, ** p<.05, * p<.10. Robust SE in parentheses.

GF+I=1 if participant in GF+I condition; GF=1 if participant in GF condition; GF+I=GF=0 if participant in C condition.

IQ score measure standardized

B Cost of Belief Distortion

In this section, we describe in more detail the implications of the belief elicitation mechanism we used for the cost of distorting beliefs (away from the rational benchmark) to participants. Given the design of the experiment, every group has the same probability of scoring higher than the out-group in the IQ test and thus of “winning” the tournament (i.e., $P(W) = 0.5$). Therefore, the expected payoff of the tournament (π^T) is

$$E(\pi^T) = 10P(W) = 5$$

In each time period, the individual participant reports a belief μ_t and the elicitation mechanism selects a random number $r_t \sim U[0, 1]$ ²⁴. If $r \leq \mu$ she plays the tournament and receives the payoff $E(\pi^T)$. If $r > \mu$ she plays the lottery and receives the lottery payoff (π^L), which is in expectations:

$$\begin{aligned} E(\pi^L|r > \mu) &= 10E(r|r > \mu) \\ &= 10[.5(1 + \mu)] \\ &= 5(1 + \mu) \end{aligned}$$

The expected per-period payoff is thus

$$E(\pi|\mu) = E(\pi^T)P(r \leq \mu) + E(\pi^L|r > \mu)P(r > \mu)$$

Since r is uniformly distributed, $P(r \leq \mu) = F(\mu) = \mu$. Therefore

$$\begin{aligned} E(\pi|\mu) &= E(\pi^T)\mu + E(\pi^L|r > \mu)(1 - \mu) \\ &= 5\mu + 5(1 + \mu)(1 - \mu) \\ &= 5(1 + \mu - \mu^2) \end{aligned}$$

It is easy to see that the per-period payoff is maximized at $\mu^* = .5$, which yields $E(\pi|\mu^*) = 6.25$. Figure B1 plots the cost of belief distortion, that is the difference between *biased* expected payoffs $E(\pi|\mu)$ and *unbiased* expected payoffs $E(\pi|\mu^*)$ relative to these latter. The cost function is convex: for small departures around μ^* , the loss is rather small, whereas it increases rapidly with μ . We can see that reporting $\mu = .6$, for instance, represents a loss of 0.8% in expected earnings compared to reporting $\mu = \mu^* = .5$, while reporting $\mu = 1$ entails a 20% loss in expected earnings.

What was the realized cost of distorting beliefs for the participants in our experiment? Table B1 displays the average earnings (in experimental points) per period and condition. It shows that participants in the GF+I treatment earned significantly less overall than the participants in the control treatment. For the GF treatment, the difference is not statistically significant overall.

²⁴For simplicity, we remove the time subscript as the per-period payoffs are independent between periods.

Figure B1: Cost of belief distortion (proportional to unbiased expected payoff)




Table B1: Earnings in the belief elicitation stage (in points)

Condition	Period0	Period1	Period2	Period3	Total
Control	6.88	8.44	7.66	8.59	31.56
Group Framing	7.81	8.44	7.81	7.19	31.25
Group Framing + Interaction	7.03	5.47	7.50	7.03	27.03
<i>p</i> -value GF+I < C	.575	<.001	.419	.016	.025
<i>p</i> -value GF < C	.883	.500	.583	.026	.440

Note: one-sided *p*-values from difference-in-means *t*-tests with equal variances.

C Checking for Private Information

To ensure that participants could make no reliable inference about group intelligence before we measured their prior beliefs, and that decisions are independent (not correlated within groups) in $GF + I$, we check that participants in this condition could not infer their in-group’s intelligence by seeing the in- and the out-group’s flag on their screens or by observing the size of their own group’s circles at the end of the circle tasks in stage A.

C.1 Flags

Figure C1 shows the group flags constructed by the participants in GF+I. The flags right next to each other are the flags of the groups that were in competition with each other in stage A of the experiment. In order to test whether it is possible to infer a group’s intelligence from these flags, we presented the eight pairs of flags to an additional sample of 45 student participants, and asked them, based on the flags only, to guess which of the two groups they believed scored higher on an intelligence test. Participants were not incentivised in this follow-up study. On average, these 45 additional participants guessed correctly 55 percent of the time. This is only marginally significantly different from a chance level of 50 percent ($p = .060$).²⁵ On the level of the individual comparisons, participants guessed significantly correctly for comparisons 2.b, 3.a, and 3.b. They were significantly wrong for comparison 1.b.²⁶

For each group, we compute a variable “flag-score” equal to the proportion of “won the tournament” guesses among the 45 participants in the additional sample. In other words, the flag-score of the group is the probability (assigned by participants in the additional sample) that the group scored higher than the other group in the IQ test.

Figure C2 does not show an evident relationship between the individual prior beliefs of participants in the experiment and the in-group’s flag-score. The regressions in Table C1 further corroborate this statement. Columns 1 and 2 show OLS estimates of parameters in a regression of prior beliefs on the flag-score of the group, with and without control variables. Columns 3 and 4 use as regressor a binary variable that equals 1 when the flag-score is above .5. In none of these specifications there appears to be a significant effect of the flag-score on prior beliefs.

Overall, we interpret this as evidence that it is not possible to reliably infer group intelligence from seeing the flags that the two groups constructed in stage A.

C.2 Circles

Figure C3 plots the in-group’s score in the IQ test (excluding the self) against the final size of the in-group’s circle (at the end of the clicking game), for the two clicking games of stage A. In none

²⁵p-value from OLS regression clustered by individual. A non-clustered binomial probability test yields a p-value of $p = .065$.

²⁶For comparison 1.b 82.2 percent guessed falsely that the green group was more intelligent; for 2.b 71.1 percent guessed correctly that the yellow group was more intelligent; for 3.a 82.2 percent guessed correctly that the red group was more intelligent, and for 3.b 71.1 percent guessed correctly that the green group was more intelligent.

of the two games there seems to be a correlation between the size of the circles and the in-group's score in the IQ test. If any, there is a positive one in the circle-reduce game, though the slope coefficient is clearly not significant ($t = 0.51$, $p > .10$).

The lack of correlation between the outcome of the circle tasks and the groups' scores on the IQ test, plus the fact that intelligence cannot be reliably inferred from observing the flags, allow us to rule out within-group correlation in the elicitation of beliefs as a result of the group manipulation tasks in stage A. Moreover, participants do not get any feedback about the outcomes of the games of stage A: they do not know if they won against the other group in the circle tasks and they do not know how they scored in the IQ test individually nor at the group level. Hence, when we elicit priors, their decisions can be considered as independent. When eliciting posteriors, they are only informed about whether their group scored better or not than the other group at three randomly drawn questions. They do not see the other participants' beliefs, nor their scores, nor their signals. Thus, also for GF+I, we can consider the individual participant as the independent level of observation, and we do therefore not cluster standard errors on the group-level in the regression analyses presented in the results section of the paper.²⁷

Table C1: OLS regressions for prior beliefs μ_0 by flag-score of group

	(1)	(2)	(3)	(4)
flag-score	1.016 (8.475)	3.532 (9.113)		
female		4.213 (3.003)		4.352 (2.722)
age		-0.733 (1.594)		-0.787 (1.539)
IQ score		-0.497 (1.925)		-0.407 (2.039)
risk aversion		1.515 (1.477)		1.565 (1.447)
high flag-score			0.875 (3.812)	1.988 (3.987)
Constant	60.054*** (5.201)	56.725*** (6.189)	60.125*** (3.445)	57.435*** (3.997)
Observations	64	64	64	64
R^2	0.000	0.038	0.001	0.041

Notes: *** $p < .01$, ** $p < .05$, * $p < .10$. SE clustered by group in parentheses. High flag-score = 1 if flag-score $> .5$. age, IQ score and Risk aversion measures standardized.

²⁷For the C and GF conditions clustering is evidently not necessary, as participants did not engage in any of the interactive group tasks of stage A.




Figure C1: Flags constructed by participants in GF+I




Figure C2: Individual prior beliefs in $GF + I$ as a function of the group flag-score




Figure C3: In-group score in the IQ test by size of circles

D In-group Overconfidence vs. Individual Overconfidence

Figure D1 plots the subjective prior belief that the in-group scored higher than the out-group in the IQ test, i.e., μ_0 (elicited in stage C of the experiment) against the measure of individual overconfidence (elicited in stage E), by condition. Albeit positive, the relationship between these two measures is not statistically significant in GF+I ($t = 1.34, p > 0.1$). It is significant in the other conditions ($t = 2.01, p = .049$ in C and $t = 2.12, p = .038$ in GF). The fact that the correlation is significant in the *Control* condition, but non-significant in our main treatment condition (GF+I) means that individual self-confidence is unlikely to be the driver behind the treatment differences we find (specifically that prior beliefs about group performance are higher in GF+I than in C). It does not seem to be the case that overconfidence in one's group and overconfidence in the self are strongly related. The small positive correlations we find in two of our three treatments may be driven by anchoring effects (when entering their choice regarding confidence in the self in stage E, participants may remember their earlier decisions about confidence in the group in stages C and D of the experiment, and use them as an anchor) or also by some individuals consistently misinterpreting the belief elicitation mechanism. By comparing the results from our GF+I condition with C, our design allows taking such potential misunderstandings of the belief elicitation mechanism into account and prevents them from biasing our results.

Another possibility is that information and beliefs about group performance affect beliefs about individual performance (i.e., an effect in the opposite direction than the one discussed above). Specifically, participants could interpret the signals they received about group performance as informative about themselves. Even though with perfectly rational information processing, this should not happen (as the signals contained only information about the performance of the other group members and not about the self), it is possible that identification with the group actually drives such an effect.

To analyze the relationship between individual overconfidence and the signals across the treatment conditions, we computed the number of good signals each participant observed. In Figure D2 below we plot the belief about own performance in the IQ test against the number of good signals observed. There appears to be a positive relationship in the GF and GF+I conditions, but the confidence intervals overlap, suggesting it is not significant.

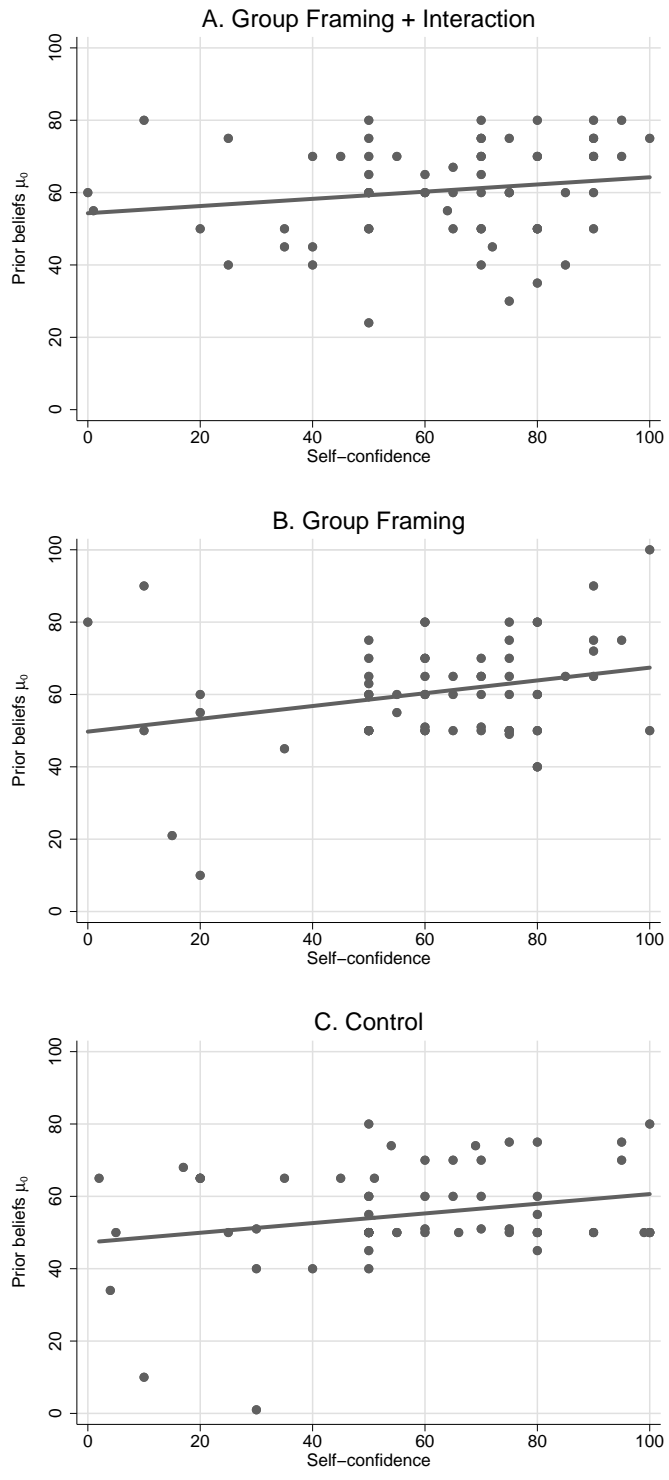


Figure D1: Prior beliefs and self-confidence

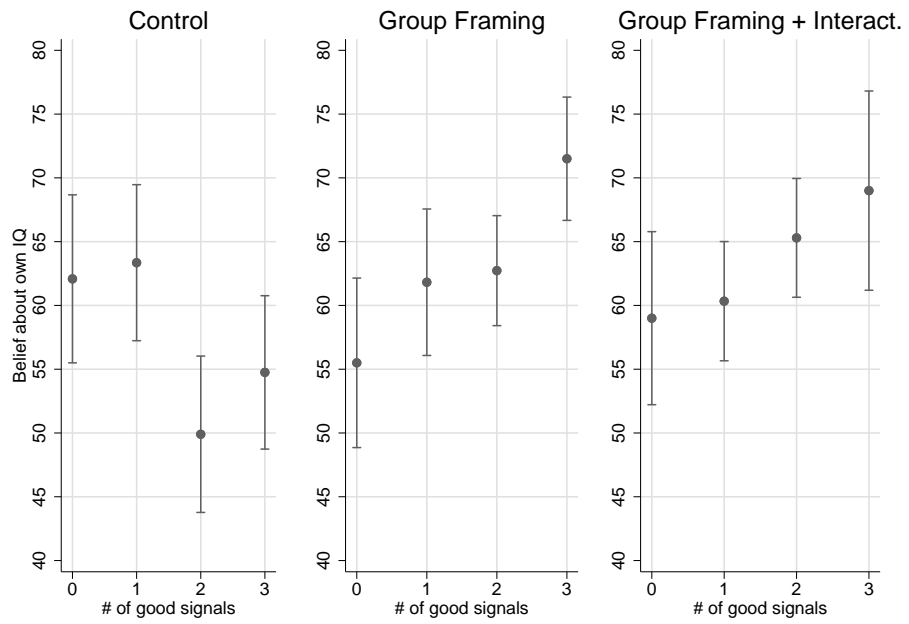


Figure D2: Belief about own intelligence as a function of the number of good signals observed

E Econometric Model of Belief Updating

To identify a potential in-group bias in the way participants process information, we begin by considering the Bayesian posterior belief

$$\mu_t \equiv P(W|s_t, \dots, s_1) = \frac{P(s_t|W, s_{t-1}, \dots, s_1)}{P(s_t|s_{t-1}, \dots, s_1)} \mu_{t-1} \quad (4)$$

where W (for “winning”) is the event “the in-group scored higher than the out-group in the IQ test” and s is the signal received about relative performance of the in-group. Let p_t denote the informativeness of a good signal, i.e. $p_t = P(s_t = G|W, s_{t-1}, \dots, s_1) = 1 - P(s_t = G|L, s_{t-1}, \dots, s_1)$,²⁸ where L (for “losing”) is the event “the in-group scored lower than the out-group in the IQ test.” For good and bad signals, the likelihood ratios of posteriors are

$$\begin{aligned} \frac{\mu_t^G}{1 - \mu_t^G} &= \frac{P(W|s_t = G, s_{t-1}, \dots, s_1)}{P(L|s_t = G, s_{t-1}, \dots, s_1)} = \frac{p_t}{(1 - p_t)} \frac{\mu_{t-1}}{1 - \mu_{t-1}} \\ \frac{\mu_t^B}{1 - \mu_t^B} &= \frac{P(W|s_t = B, s_{t-1}, \dots, s_1)}{P(L|s_t = B, s_{t-1}, \dots, s_1)} = \frac{(1 - p_t)}{p_t} \frac{\mu_{t-1}}{1 - \mu_{t-1}} \end{aligned}$$

Taking logs and combining these two equations we can write the posterior odds as

$$\ln \left(\frac{\mu_t}{1 - \mu_t} \right) = \ln \left(\frac{p_t}{1 - p_t} \right) \mathbb{1}(s_t = G) + \ln \left(\frac{1 - p_t}{p_t} \right) \mathbb{1}(s_t = B) + \ln \left(\frac{\mu_{t-1}}{1 - \mu_{t-1}} \right)$$

with $\mathbb{1}(s_t = S)$ an indicator that takes value one if the signal is S . Allowing all coefficients to differ between conditions according to the indicators $GF + I_i = 1$ if participant i is in GF+I, $GF_i = 1$ if participant i is in GF, and $GF + I_i = GF_i = 0$ if she is in C, we arrive at regression model (1).

Before estimating the parameters in model (1) we computed the empirical informativeness (p_t) of the signals received by participants in the experiment. For the first period, we drew 2024 signals by selecting each possible triplet out of the 24 questions in the IQ test and comparing the scores of each group; p_1 is equal to the proportion of positive signals among the “winning” groups (i.e., groups that scored higher overall than their out-group). For period two, we eliminated the actual questions that were drawn in the first signal and considered the remaining 21 questions in the test; we computed p_2 as the proportion of positive signals among the “winning” groups *conditional* on the realized first signal. Finally, we eliminated the questions drawn in the second signal and drew all possible third signals from the remaining 18 questions and computed p_3 , again *conditional* on the history of signals.²⁹

²⁸Notice that this equality is given by experimental design.

²⁹We do not claim that participants were able to perfectly infer the informativeness of the signal. However, our analysis does not aim at judging the descriptive validity of Bayesian updating (for this see, e.g., Tversky & Kahneman, 1974, and the discussion in Mobius et al., 2014). Rather, our aim is to isolate a potential in-group bias by comparing updating behavior in GF+I with updating in C, and by comparing reactions to positive and negative signals.

F Clustering Standard Errors by Group

This section presents additional results for readers that are not convinced by the argument regarding clustering in Appendix C .

In tables F1 and F2 below we reproduce respectively our main tables 3 and 5, but showing standard errors clustered by group.

Indeed, the results are robust to clustering at the group level for all parameters except α_1^B that loses its marginal significance. This is not problematic as the significance of α_1^G still implies that participants in the GF+I condition update asymmetrically with respect to participants in the Control condition.

Table F1: OLS regressions for μ_0

	(1)	(2)
GF+I	5.641** (2.660)	5.700** (2.731)
GF	5.891** (2.654)	5.803** (2.613)
female		2.457 (2.086)
age		0.399 (1.009)
IQ score		0.354 (1.171)
risk aversion		0.888 (0.876)
Constant	54.922*** (1.898)	53.793*** (2.422)
Observations	192	192
Groups	48	48
R^2	0.036	0.047

Notes: *** p<.01, ** p<.05, * p<.10. SE clustered by group in parentheses.

GF+I=1 if participant in GF+I condition; GF=1 if participant in GF condition; GF+I=GF=0 if participant in C condition. age, IQ score and risk aversion standardized.

Table F2: OLS regressions for $\tilde{\mu}_{it} \equiv \ln[\mu_{it}/(1 - \mu_{it})]$

	(1)	(2)	(3)	(4)
α_0^G : Good signal	1.047*** (0.222)	1.139*** (0.219)	0.455 (0.315)	0.595* (0.302)
α_0^B : Bad signal	0.927*** (0.233)	0.783*** (0.238)	1.409*** (0.463)	1.148*** (0.425)
β_0 : Lagged belief	0.382*** (0.108)	0.363*** (0.109)	0.372*** (0.101)	0.343*** (0.097)
α_1^G : Good signal \times GF+I			0.918** (0.426)	0.832** (0.396)
α_1^B : Bad signal \times GF+I			-0.827 (0.539)	-0.653 (0.506)
β_1 : Lagged belief			0.021 (0.233)	0.040 (0.230)
α_2^G : Good signal \times GF			0.854 (0.662)	0.827 (0.647)
α_2^B : Bad signal \times GF			-0.593 (0.666)	-0.424 (0.651)
β_2 : Lagged belief			-0.039 (0.237)	-0.012 (0.240)
female		-0.097 (0.109)		-0.108 (0.111)
age		-0.099** (0.043)		-0.079 (0.048)
IQ score		0.003 (0.042)		0.001 (0.040)
risk aversion		0.122*** (0.038)		0.104*** (0.036)
Observations	489	489	489	489
Groups	48	48	48	48
R^2	0.212	0.226	0.227	0.237

Notes: *** p<.01, ** p<.05, * p<.10. SE clustered by group in parentheses.
 Observations with extreme current or lagged beliefs are dropped from the estimations because $\tilde{\mu}$ is not defined when $\mu = 0$ or $\mu = 1$.
 age, IQ score and risk aversion standardized.

G Instrumental Variables Estimation

In section 3.1 we discuss the group identification measure we elicited in the post-experimental survey. This measure is likely endogenous in a reduced-form specification. One reason is that it is measured after beliefs, i.e. the dependent variable, and hence may be affected by the way participants updated their beliefs producing reverse causality. Since participants were randomly assigned to conditions with different levels of induced group identification, we can use treatment assignment as an instrument for the group identification measure.

In Table G1, we estimate the model for prior beliefs by two-stage least-squares (2SLS), instrumenting the group identification measure with two dummies for treatment assignment, GF+I and GF, both equal to zero for the Control condition. The first stage regression (with and without further control variables) shows a weak effect of treatment assignment on group identification, as suggested by the F-statistics on the excluded instruments. The second stage shows a strong positive effect of group identification on prior beliefs: a one standard deviation increase in identification leads to a 9 percentage points increase in reported prior beliefs.

Table G1: IV regressions for μ_0

	(1)	(2)
Group identification	8.550** (3.748)	9.239** (4.167)
female		3.326 (2.308)
age		0.970 (1.227)
IQ score		0.745 (1.157)
risk aversion		0.181 (1.257)
Constant	58.766*** (1.113)	57.224*** (1.551)
Observations	192	192
F-stat 1st stage	9.143	7.498

Notes: *** $p < .01$, ** $p < .05$, * $p < .10$. Robust SE in parentheses.

Group identification = self-reported level of identification to the ingroup (1-10 scale) standardized.

age, IQ score and risk aversion standardized based on original scales.

In Table G2, we instrument the three group identification interactions (with good signal, bad signal and lagged beliefs) with the six treatment assignment interactions (two dummies for GF+I and GF, interacted with good signal, bad signal and lagged beliefs), and regress posterior beliefs

Table G2: IV regressions for $\tilde{\mu}_{it} \equiv \ln[\mu_{it}/(1 - \mu_{it})]$

	(1)	(2)
Good signal	0.820*** (0.295)	0.929*** (0.296)
Bad signal	0.753*** (0.231)	0.668** (0.259)
Lagged belief	0.257** (0.125)	0.251* (0.129)
Good signal \times Group identif.	1.596* (0.927)	1.425* (0.863)
Bad signal \times Group identif.	-1.728 (1.111)	-1.573 (1.256)
Lagged belief \times Group identif.	0.003 (0.269)	0.032 (0.270)
female		-0.093 (0.136)
age		-0.063 (0.063)
IQ score		0.062 (0.087)
risk aversion		0.075 (0.060)
Observations	489	489
Participants	173	173
F-stat 1st stage (good signal)	2.31	2.27
F-stat 1st stage (bad signal)	1.58	0.88
F-stat 1st stage (lagged belief)	2.02	2.03

Notes: *** p<.01, ** p<.05, * p<.10. SE clustered by participant in parentheses.

Group identification = self-reported level of identification to the ingroup (1-10 scale) standardized.

Group identification instrumented with condition dummies GF+I and GF.

Observations with extreme current or lagged beliefs are dropped from the estimations because $\tilde{\mu}$ is not defined when $\mu = 0$ or $\mu = 1$. age, IQ score and risk aversion standardized.

on it. The F-statistics from the first stage regressions (with and without further control variables) show evidence that the instruments are very weak. The second stage shows that one standard deviation increase in group-identification leads to stronger reaction to positive information, that is $\alpha_1^G > 0$ ($p = .085$), but no differential reaction to negative information; $\alpha_1^B = 0$ ($p = .120$). However, there is evidence of asymmetric updating of beliefs: a one standard deviation increase in group identification leads participants to put more weight on positive signals about the group than they put on negative signals ($\alpha_0^G + \alpha_1^G > \alpha_0^B + \alpha_1^B$, $p = .022$). When adding control variables, the significance of the α -coefficients vanishes ($p = .101$ and $p = .159$ for good and bad signals respectively), but there is still evidence of asymmetric updating ($p = .022$).

The IV estimation strategy is unreliable for the two reasons mentioned above: the group identification measure suffers from measurement error, and treatment assignment is a weak instrument for it.

H Experimental Instructions

(The following instructions were originally in French.)

Welcome to this experiment! You will make decisions that will affect your earnings. Although we express all earnings in terms of coins, these coins will be exchanged for swiss francs at the end of the experiment using the following exchange rate:

$$5 \text{ coins} = \text{CHF } 1.-$$

Independently from your decisions during the experiment, you will receive a fix amount of 8 CHF for your participation. The final amounts will be rounded to the nearest integer.

It is **strictly forbidden to talk with other participants**. It is important that you respect this rule for the experiment to run smoothly. If you have any questions, please raise your hand to contact the assistants. If you do not follow this rule, we will have to exclude you from the experiment.

H.1 Control condition

The following part was contained in the instructions distributed to participants in the Control condition.

This experiment is divided in three parts (A, B and C). We will now explain what you will do in the first part (part A).

PART A

What is it about? You will do a test of logic based on images. Your objective is to answer correctly as many questions as possible.

What will you do? The test is composed of 3 sections. Each section of the test will be described in detail directly on the screen at the beginning of each section. We will also provide you with examples of the type of questions contained in the section. Each section is composed of 8 questions. There are thus a total of 24 questions in the test. You have 1 minute 30 seconds per section to answer as many questions as possible. For every correct answer, one point will be added to your score. If your score places you among the 50% best participants in the room, you earn a prize of 10 coins. In case of a tie, the computer will determine randomly who earns the prize.

Be aware that certain questions, towards the end of each section, are very difficult and almost nobody can expect to answer them on time. Even if you ignore the right answer, try and guess, because an empty answer is considered as false.

You will be informed of your earnings in this part at the end of the experiment.

Did you understand? At the beginning of each section, you will be able to practice with two trial questions. The answer to these questions will not influence your score.

PART B

What is it about? In this 4-section part, we propose you two games, game A and game B. Your single task is to fix a rule that determines which game you will play.

Game A is a **comparison** game, that consists in comparing the score in the image test of part A (“the test”) obtained by three randomly-selected participants in the room with the score of three other randomly-selected participants in the room. For ease of explanation, we will refer to the first three participants as “persons X,Y,Z” and the other three participants as “persons A,B,C.” Notice that you are not included among these six persons.

Game B is a **lottery** game, whose rules are explained below.

What will you do?

Section 1: The two games are explained in the table below:

Game A: comparison game	Game B: lottery game
If the persons XYZ have obtained MORE points in total in the image test of part A than persons ABC, you earn 10 coins.	A random device in the computer will determine if you earn the 10 coins or if you earn 0 coins.
If the persons XYZ have obtained LESS points in total in the image test of part A than persons ABC, you earn 0 coins.	The probability of earning the 10 coins is X% (see explanation below).
If XYZ and ABC have obtained the same number of points in the test, the computer will toss a coin to determine who has “obtained” more points.	

Game B offers the possibility of earning 10 coins with a probability already determined by the computer (a value between 0 and 100%, that we call here X%). What does this mean? If X is equal to 0%, for exemple, it means that you have no chance of winning in game B. On the contrary, if X is equal to 100%,you will systematically win in game B.

You don’t know what is exactly the probability X of winning the 10 coins in the lottery game B, so how to chose which game you prefer to play? Actually, you will fix the rule that determines which game you will play, by telling us a value between 0 and 100%. The rule is:

- If the probability X determined by the computer is lower or equal to the value that you tell us, you play game A.
- If the probability X determined by the computer is greater than the value that you tell us, you play game B.

- To sum up, the rule is:

$X \leq$ value that you tell us \rightarrow game A

$X >$ value that you tell us \rightarrow game B

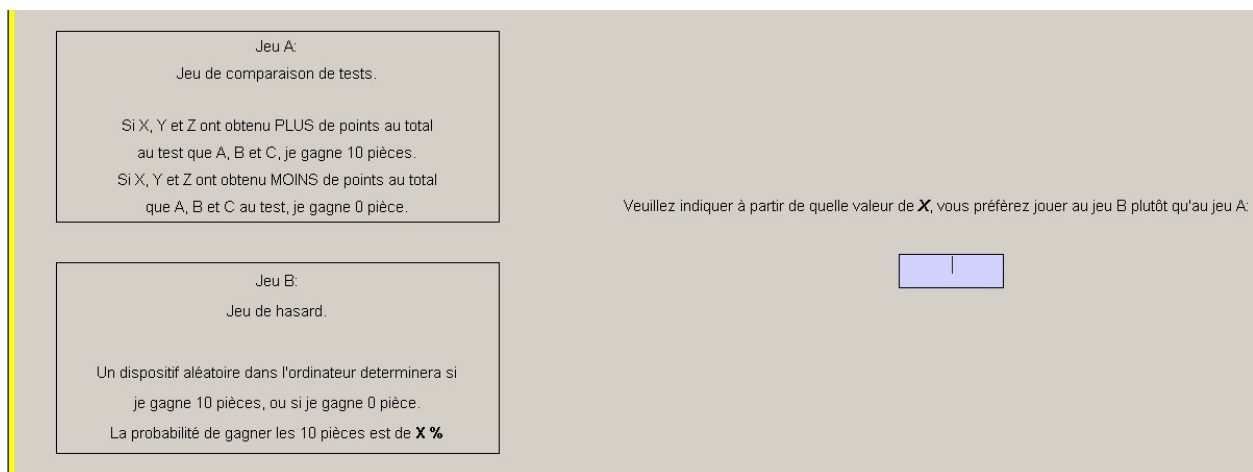
In other words, **you will tell us from which value of the probability X of winning in game B, you prefer to play game B rather than game A.**

Please, look at the examples in the table below to better understand the consequences of your decision.

Probability X (determined by the computer)	Value that you tell us	You play
25%	20%	Game B: you earn 10 coins with a 25% probability
25%	30%	Game A: you earn 10 coins if the score of XYZ in the images test is higher than the score of ABC
75%	30%	Game B: you earn 10 coins with a 75% probability
75%	40%	Game B: you earn 10 coins with a 75% probability
75%	90%	Game A: you earn 10 coins if the score of XYZ in the images test is higher than the score of ABC

These examples show you that, if you want to maximize your chances of earning the 10 coins, you have to think carefully and tell us from which value of the probability X of winning in the lottery game you really prefer to play game B rather than game A.

The following figure shows you the screen where you will tell us this value. After having entered your value click on the button “continue”.



Section 2: This section is almost identical to the previous one. However, in the beginning of this section, the computer will inform you of the outcome of a “match” between ABC and XYZ.

To determine the outcome of the “match,” the computer will randomly draw three out of the 24 questions of the image test of part B and will compute the number of points obtained in total at these three questions by persons ABC and persons XYZ. The three persons that have obtained the most points in these three questions win the “match.” If ABC and XYZ have obtained the same number of points, the computer will toss a coin to determine who wins the “match.” You will see whether XYZ have WON or LOST the match on top of your screen. After having learned the outcome of the match, you will tell us, just like in Section 1, from which value of the probability X , you prefer to play game B rather than game A.

Sections 3 and 4: These sections are identical to Section 2. Just like in Section 2, in the beginning of these sections the computer will inform you of the outcome of a “match” between ABC and XYZ. For every match the computer will consider the answers to three out of the 24 questions of the test. However, the three questions drawn randomly by the computer are different from the ones drawn in the previous matches. The answers to a question are thus never part of more than one match.

Also in sections 3 and 4, after having learned the outcome of the match, you will tell us from which value of the probability X , you prefer to play game B rather than game A.

You will be informed of your earnings in this part at the end of the experiment.

Did you understand? Before starting, we want to make sure that you and everybody else have understood correctly what you will do. For this purpose, please answer the control questions that will appear on your screen. These questions have no influence on your earnings.

PART C

What is it about? In this part, we will propose you two games: game A and game B. Just like in part B, you will fix the rule that determines which game you will play.

Game A is a **ranking** game, where you earn coins if the total number of points that you have obtained in the image test of part A (your individual score) places you among the best 50% of participants in the room in terms of this score.

Game B is a **lottery** game, identical to the one in part B.

What will you do? We will classify the individual scores of each of the 16 participants in this room in two halves: the top 50% (the 8 highest scores) and the bottom 50% (the 8 lowest scores).

In the table below we explain the two games:

Game A: ranking game	Game B: lottery game
If your individual score places you among the top 50% in the room, you earn 10 coins.	A random device in the computer will determine if you earn the 10 coins or if you earn 0 coins.
If your individual score places you among the bottom 50% in the room, you earn 0 coins.	The probability of earning the 10 coins is X% (see explanation below).
If your score is right in the middle, the computer will toss a coin to determine if your score is among the top or among the bottom 50%.	

Just like before, you will tell us from which value of the probability X of winning at the lottery game, you prefer to play game B, rather than game A. Once you have entered your value, click on the button “continue.”

Your earnings will be computed in the same way as in part B. You will be informed of these earnings at the end of the experiment.

Did you understand? If you have finished reading these instructions and you have no questions, click on “start”. If you have questions, raise your hand to call one of the assistants.

H.2 Group Framing condition

The following part was contained in the instructions distributed to participants in the Group Framing condition.

In this study, each participant is part of a group composed of four participants. Your group will always be the same throughout the entire experiment.

This experiment is divided in three parts (A, B and C). We will now explain what you will do in the first part (part A).

PART A

What is it about? Your group is related to another group in the room. You will do a test of logic based on images. The objective of your group is to obtain a score in the test higher than the other group.

What will you do? The test is composed of 3 sections. Each section of the test will be described in detail directly on the screen at the beginning of each section. We will also provide you with examples of the type of questions contained in the section. Each section is composed of 8 questions. There are thus a total of 24 questions in the test. You have 1 minute 30 seconds per section to answer as many questions as possible. For every correct answer, one point will be added to your group's score. The group with the highest score in the test earns a prize of 10 coins for each group member. The members of the losing group do not earn anything. In case of a tie, the computer will determine randomly who earns the prize.

Be aware that certain questions, towards the end of each section, are very difficult and almost nobody can expect to answer them on time. Even if you ignore the right answer, try and guess, because an empty answer is considered as false.

You will be informed of your earnings in this part at the end of the experiment.

Did you understand? At the beginning of each section, you will be able to practice with two trial questions. The answer to these questions will not influence your score.

PART B

What is it about? In this 4-section part, we propose you two games, game A and game B. Your single task is to fix a rule that determines which game you will play.

Game A is a **comparison** game, that consists in comparing the score in the image test of part A ("the test") obtained by the three other members of your group ("your group") with the total score of three members of the other group ("the other group"). The three members of the other group that will be considered for the score comparison have been randomly chosen among the four members of the other group. These three members of the other group will always be the same persons until the end of the experiment.

Game B is a **lottery** game, whose rules are explained below.

What will you do?

Section 1: The two games are explained in the table below:

Game A: comparison game	Game B: lottery game
If your group has obtained MORE points in total in the image test of part A than the other group, you earn 10 coins.	A random device in the computer will determine if you earn the 10 coins or if you earn 0 coins.
If your group has obtained LESS points in total in the image test of part A than the other group, you earn 0 coins.	The probability of earning the 10 coins is X% (see explanation below).
If your group and the other group have obtained the same number of points in the test, the computer will randomly determine which group has “obtained” more points.	

Game B offers the possibility of earning 10 coins with a probability already determined by the computer (a value between 0 and 100%, that we call here X%). What does it mean? If X is equal to 0%, for exemple, it means that you have no chance of winning in game B. On the contrary, if X is equal to 100%, you will systematically win in game B.

You don't know what is exactly the probability X of winning the 10 coins in the lottery game B, so how to chose which game you prefer to play? Actually, you will fix the rule that determines which game you will play, by telling us a value between 0 and 100%. The rule is:

- If the probability X determined by the computer is lower or equal to the value that you tell us, you play game A.
- If the probability X determined by the computer is greater than the value that you tell us, you play game B.
- To sum up, the rule is:

$$X \leq \text{value that you tell us} \rightarrow \text{game A}$$

$$X > \text{value that you tell us} \rightarrow \text{game B}$$

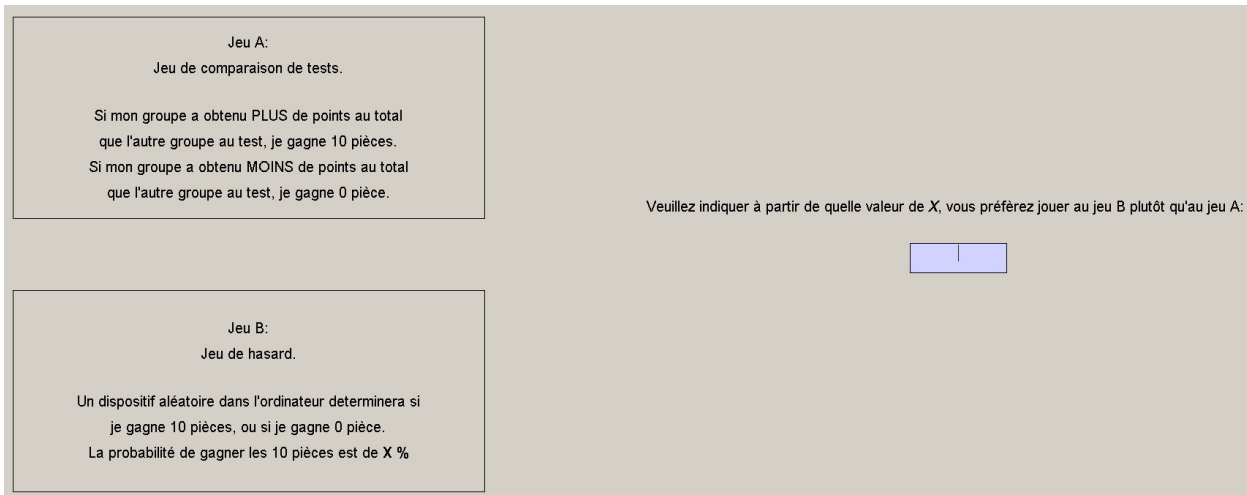
In other words, **you will tell us from which value of the probability X of winning in game B, you prefer to play game B rather than game A.**

Please, look at the examples in the table below to better understand the consequences of your decision.

Probability X (determined by the computer)	Value that you tell us	You play
25%	20%	Game B: you earn 10 coins with a 25% probability
25%	30%	Game A: you earn 10 coins if the score of your group in the images test is higher than the score of the other group
75%	30%	Game B: you earn 10 coins with a 75% probability
75%	40%	Game B: you earn 10 coins with a 75% probability
75%	90%	Game A: you earn 10 coins if the score of your group in the images test is higher than the score of the other group

These examples show you that, if you want to maximize your chances of earning the 10 coins, you have to think carefully and tell us from which value of the probability X of winning in the lottery game you really prefer to play game B rather than game A.

The following figure shows you the screen where you will tell us this value. After having entered your value click on the button “continue”.



Section 2: This section is almost identical to the previous one. However, in the beginning of this section, the computer will inform you of the outcome of a “match” between the two groups.

To determine the outcome of the “match,” the computer will randomly draw three out of the 24 questions of the image test of part B and will compute the number of points obtained in total at these three questions by the three members of your group and by the three members of the other group. The group that has obtained the most points in these three questions wins the “match.” If

the two groups have obtained the same number of points, the computer will toss a coin to determine which group wins the “match.” You will see whether your group has WON or LOST the match on top of your screen. After having learned the outcome of the match, you will tell us, just like in Section 1, from which value of the probability X , you prefer to play game B rather than game A.

Sections 3 and 4: These sections are identical to Section 2. Just like in Section 2, in the beginning of these sections the computer will inform you of the outcome of a “match” between the two groups. For every match the computer will consider the answers to three out of the 24 questions of the test. However, the three questions drawn randomly by the computer are different from the ones drawn in the previous matches. The answers to a question are thus never part of more than one match.

Also in sections 3 and 4, after having learned the outcome of the match, you will tell us from which value of the probability X , you prefer to play game B rather than game A.

You will be informed of your earnings in this part at the end of the experiment.

Did you understand? Before starting, we want to make sure that you and everybody else have understood correctly what you will do. For this purpose, please answer the control questions that will appear on your screen. These questions have no influence on your earnings.

PART C

(This part is identical in all conditions. Refer to subsection G.1, Part C, for instructions)

H.3 Group Framing + Interaction condition

The following part was contained in the instructions distributed to participants in the Group Framing + Interaction condition.

In this study, each participant is part of a group composed of four participants. Your group will always be the same throughout the entire experiment.

This experiment is divided in seven parts (A to G). We will now explain what you will do in the first part (part A).

PART A

What is it about? In this part, the other members of your group and yourself will construct the flag that will represent you throughout the experiment. For this purpose, each of you will be in charge of decorating one of the four corners of the flag using colored paper (each group has a different color). In other words, the flag of your group will be composed of a representative element of each one of you.

What will you do? On your table, in front of you, you will find a pair of scissors and an envelope containing colored paper. Start by cutting out a shape. You are free to cut out any shape you want, but you cannot cut out words, nor can you write on the shape. The shape must be composed of a single element (you cannot cut out two separate shapes).

When you finish, write your place number on the back of the shape and put it inside the envelope.

After 5 minutes, the assistants will collect the envelopes containing the shapes. They will then paste the shapes on the corners of the flag. You will see what the flag of your group looks like in the next part of the experiment.

You can now open the envelope and start working.

PART B

What is it about? From now on, your group is related to another group in the room. You will take part in a series of group games. The objective of your group is to win against the other group in each of these games.

What will you do? Firstly, the flag of your group that you have just created in part A will appear on your screen. After having seen your flag, click on “OK” so that the first game can start.

Game 1: reducing the size of circles

You and the other members of your group will see in the middle of your screen a very big circle. The members of the other group will see a circle of the same size on their screens. The aim of the game is to reduce the size of the circle. To reduce the size of the circle you just have to click on it. Every click of a member of your group reduces the size of the circle. The game is over after 30 seconds. The group that finishes with the smallest circle, wins the game, and each member of the winning group earns a reward of 10 coins. The members of the losing group do not earn anything. In case of a tie, the computer will toss a coin to determine which group earns the reward.

Game 2: image test

The second group game consists of a logic test based on images. The test is composed of 3 sections. Each section of the test will be described in detail directly on the screen at the beginning of each section. We will also provide you with examples of the type of questions contained in the section. Each section is composed of 8 questions. There are thus a total of 24 questions in the test. You have 1 minute 30 seconds per section to answer as many questions as possible. For every correct answer, one point will be added to your group's score. The group with the highest score in the test earns a prize of 10 coins for each group member. The members of the losing group do not earn anything. In case of a tie, the computer will determine randomly who earns the prize.

Be aware that certain questions, towards the end of each section, are very difficult and almost nobody can expect to answer them on time. Even if you ignore the right answer, try and guess, because an empty answer is considered as false.

At the beginning of each section, you will be able to practice with two trial questions. The answer to these questions will not influence your score.

Game 3: enlarging the size of circles

In this game, you and the other members of your group will see in the middle of your screen a small circle. The members of the other group will see a circle of the same size on their screens. The aim of the game is to enlarge the circle. To enlarge the circle you just have to click on it. Every click of a member of your group enlarges the circle. You will notice that the circle will move around the screen. You will have to carefully target the moving circle to be able to enlarge it. The game is over after 30 seconds. The group that finishes with the largest circle, wins the game, and each member of the winning group earns a reward of 10 coins. The members of the losing group do not earn anything. In case of a tie, the computer will toss a coin to determine which group earns the reward.

You will be informed of your earnings in this part at the end of the experiment.

PART C

(This part is identical to part B of the Group Framing condition. Refer to subsection G.2, Part B, for instructions.)

PART F

(This part is identical to part C of the Control and Group Framing conditions. Refer to subsection G.1, Part C, for instructions.)