Self-confidence and Strategic Behavior

Gary Charness

University of California at Santa Barbara

Aldo Rustichini

University of Minnesota

Jeroen van de Ven

University of Amsterdam

ABSTRACT. We suggest that overconfidence (conscious or unconscious) is motivated in part by strategic considerations, and test this experimentally. We find compelling supporting evidence in the behavior of participants who send and respond to others' statements of confidence about how well they have scored on an IQ test. In two-player tournaments where the higher score wins, a player is very likely to choose to compete when he knows that his own stated confidence is higher than the other player's, but rarely when the reverse is true. Consistent with this behavior, stated confidence is inflated by males when deterrence is strategically optimal and is instead deflated (by males and females) when luring (encouraging entry) is strategically optimal. This behavior is consistent with the equilibrium of the corresponding signaling game. Overconfident statements are used in environments that seem familiar, and we present evidence that suggests that this can occur on an unconscious level.

Keywords: Self-confidence, overconfidence, strategic deterrence, unconscious behavior, self-deception, luring, experiment

JEL Classifications: A12, C91, D03, D82

Date: April 27, 2017.

Contact: Gary Charness, charness@econ.ucsb.edu; Aldo Rustichini, arust@econ.umn.edu; Jeroen van de Ven, j.vandeven@uva.nl.

Acknowledgements: We acknowledge helpful comments from three anonymous referees, the Editor (Lise Vesterlund), and from Alex Brown, Juan Carrillo, Peter Cramton, Catherine Eckel, Enrique Fatas, Sanjeev Goyal, Philippe Jehiel, Juanjuan Meng, Markus Möbius, Klaus Schmidt, Lones Smith, Vernon Smith, Joel Sobel, Anton Suvorov, Cheng-Zhong Qin, Daniel Zizzo, and seminar participants at several universities including SITE (Stanford), USC, Toulouse, Peking University, Jinan University, University of Texas A&M, University of Aarhus, Ca'Foscari, Chapman University, University of East Anglia, Cambridge University, Goethe University (Frankfurt), New Economic School (Moscow), University of Maryland, CesIfo, Maastricht University, Tilburg University, Stockholm University, University of Michigan, Pompeu Fabra, Autonoma, University of Southampton, Washington University (St. Louis), and University of Wisconsin.

1. Introduction

Belief about one's abilities is an important ingredient in many decisions, including making career choices, undertaking enterprises, and taking risks. There is considerable evidence that statements people make about their abilities often don't accurately reflect their real abilities. Well-known studies in psychology and economics claim that people are overconfident in their ability (e.g., Svenson, 1981; Dunning et al., 1989).[1] However, the roots of such apparent overconfidence in relative ability and the corresponding benefits that might explain the persistence of the phenomenon are not yet clear. One personal benefit is the consumption value of the belief that one is talented (or "ego utility," Kőszegi, 2006). In this view, people feel better with a favorable self-perception, even at the cost of being overconfident and thus making wrong choices. People also derive benefit from overconfidence in relation to social image concerns, as they desire that others see them as being skilled. Recent studies support this view (Burks et al, 2013; Ewers and Zimmermann, 2014).

We consider here an additional explanation, postulating that overconfidence may also have a strategic foundation. Statements about one's beliefs are often made to affect the decisions of others. This occurs in strategic situations, which are common in social life. For instance, appearing more confident is likely to increase one's chances to be hired or to receive a promotion, and may discourage others from competing for that same position or promotion. It may elicit cooperation by others if they seek talented colleagues to start a joint project. However, in other situations it may pay to appear less skillful than one actually is, as in the case of a pool hustler. In the workplace, this strategy can be employed to elicit help by others. We emphasize that the notion that people may inflate expressed overconfidence for strategic purposes does not exclude other motivations also playing a role, and the different motivations may be complementary to each other.

One main contribution of our paper is to provide a strategic foundation for expressed overconfidence (and under-confidence) by demonstrating that participants inflate (or deflate) their stated confidence when it is strategically beneficial for them to do so. We present experimental data showing that people are responsive to confidence statements by others in strategic decisions, and that

[1]We note that while overconfidence is found in many studies, there is mixed evidence for it (see for instance Clark and Friesen, 2009) and might be overestimated (Heger and Papageorge, 2015). Its prevalence depends on factors such as experience (Weinstein, 1980) or task-difficulty (Kruger, 1999; Hoelzl and Rustichini, 2005).

these confidence statements are adjusted in ways that are consistent with strategic considerations. We also propose a model that can explain these results.

In our experiment, we elicit confidence in one's relative ability on a cognitive task. This is followed by a tournament stage, where scores on that task are compared. In the baseline treatment, all subjects automatically enter the tournament and there are no strategic reasons to distort the reported confidence relative to their true beliefs. In other treatments, we divide each pair of subjects into a sender and a receiver. In those treatments, receivers can opt out of the tournament after observing the sender's reported confidence, giving senders incentives to distort their reported confidence. We implemented two such strategic treatments. In *Deter*, the sender prefers that the receiver does not enter the tournament. We show that in equilibrium, senders over-report to discourage receivers from entering the tournament. In *Lure*, the sender always prefers that the receiver enter the tournament, and senders under-report in equilibrium.

Our main findings are as follows. In the baseline treatment, in which there are no strategic considerations to distort reported confidence, the mean reported confidence about being in the top 50 percent is well above 50 (63.4), consistent with overconfidence, or over-placement (Moore and Healy, 2008).[2] In both strategic treatments, receivers are very responsive to the reported confidence of the senders. Senders exploit this fact. In *Deter*, we find evidence that male (although not female) senders inflate reported confidence. In *Lure*, we find evidence of under-reporting by both male and female senders, again consistent with the equilibrium of the game. Surprisingly, in *Deter* we also find an increase in reported confidence by male *receivers*, even though receivers have no possible strategic advantage from over-reporting; senders automatically enter into the tournament and are not even informed of the receiver's confidence statement. By contrast, receivers in the *Lure* treatment do not adjust their reported confidence in comparison to the baseline treatment.

Finally, we find that females are less likely than males to enter the tournament in *Deter*, despite very similar performance levels. This effect is driven by stated confidence levels in *Deter*, as there is

[2] In an environment of incomplete information, this is not per se conclusive evidence of overconfidence (Benoit and Dubra, 2011; Burks, Carpenter, Goette and Rustichini, 2013).

no significant difference in entry rates when we control for these. In *Lure*, both the stated confidence levels and entry rates for males and females were almost the same (slightly higher for females).

While overconfidence may have large and negative economic consequences on individuals, we find that individuals adjust their reported confidence for strategic reasons to adapt to the situation they are facing, perhaps mitigating the negative impact of overconfidence. Being (or at least appearing) overconfident (or under-confident) can have benefits in strategic interactions, and individuals understand this to a fair degree. That said, our participants do not always adjust their confidence optimally, and the degree to which they adjust their reported confidence is somewhat gender specific.

The paper is structured as follows. In section 2, we provide a literature review. We describe our hypotheses and our experimental design in section 3. We present our results in section 4, and we discuss the motivation of biased confidence in section 5. We conclude in section 6.

2. Background and literature review

Social psychology has long considered the issues of self-esteem, overconfidence, and self-deception. Baumeister (1998) provides an extensive review of the overconfidence phenomenon. Further evidence and discussion on self-esteem can be found in Leary et al. (1995) and Leary (1999), where image concerns lead to a selective demand for information. Berglas and Jones (1978) and Kolditz and Arkin (1982) study how self-handicapping is related to social saliency: Kolditz and Arkin (1982) find that subjects take performance-impoverishing drugs after receiving positive feedback about their past performance when their choice of drugs is visible to the experimenter, but not when this choice is made in private. This suggests that performance/confidence is a social signal.

Rabin and Schrag (1999) provide a model of confirmatory bias, where people misinterpret new information as supporting previously held views. This bias induces overconfidence. An agent may come to believe with near certainty in a false hypothesis, despite receiving an unlimited amount of information. Kőszegi (2006) provides a formal model of overconfidence and ego utility, in which an agent derives internal benefits from positive views about his or her ability. The mechanism in this model is that each person receives an initial signal about own ability and can seek information if desired. Zábojník (2004) shows that endogenous information seeking can result in systematically skewed beliefs about abilities even without ego utility concerns, if seeking for information is less costly for people who expect to have low ability. Compte and Postlewaite (2004) assume that performance depends on a person's confidence and that distorted beliefs can be welfare enhancing. Van den Steen (2004) shows how over-optimism can result if people have different priors.

A number of other papers examine overconfidence experimentally. The focus is typically on establishing overconfidence without considering the strategic value of appearing overconfident or underconfident and the response to confidence statements by others. For instance, Camerer and Lovallo (1999) study entry decisions in a tournament in the context of overconfidence. In contrast to our experiment, participants in their experiment do not observe reported confidence levels of others, and there is no strategic reason to appear under- or overconfident. In some recent experiments, participants receive information about the stated confidence of others (e.g., Vialle et al., 2011; Ewers, 2012), but they do not study if participants use this strategically: confidence levels were elicited before participants were told that their reported confidence would be shared with others.

Burks, Carpenter, Goette and Rustichini (2012), based on data in Burks et al. (2009), investigate whether concerns for self-image contribute to overconfidence and whether confidence judgments are consistent with Bayesian information processing starting from a common prior. They reject both hypotheses. They find that individuals with higher beliefs about their skills are more likely to demand information, rather than less likely. These results clearly reject self-image concerns as a mechanism that yields overconfident judgments, and are consistent with the hypothesis that overconfidence is a form of social signaling. Ewers and Zimmerman (2014) also find evidence of social-image concerns. In their experiment, participants are more likely to make a confident statement if they know that others will observe the statement, compared to a situation in which their statement remains private. However, in neither experimental design is there a strategic environment that can affect confidence. In this paper we introduce the strategic environment explicitly, and study the strategic motivation underlying such signaling.

3. Experimental Design, Model, and Hypotheses

3.1. **Experimental design.** In every treatment, participants were randomly allocated to groups of four individuals. In each group, two players were randomly given the role of senders and the other two the role of receiver (in the instructions we used neutral labels "A" and "B"); each sender was randomly matched with one receiver. All participants received the same 15 questions taken from Raven's APM, a measure of cognitive ability (Raven, 2000). Participants had eight minutes to answer questions, and did not get any feedback on the number of questions they answered correctly.

When taking the test, participants only knew that they would be asked to evaluate their performance later and that every sender would be matched to a receiver with a possibility for the player with the higher rank to earn €10. Upon completion, participants were informed about all the subsequent steps in the experiment. First, one was asked to indicate one's confidence of having a score in the top two of their group, on a probability scale from 0 percent to 100 percent. They received payment for accuracy according to a quadratic scoring rule; for a stated probability p (their report divided by 100), a subject was paid €10 times $1-(1-p)^2$ if he really was in the top 2, and €10 times $1-p^2$ if he was not. We provided assurances that this mechanism favored accurate reporting for this part of the experiment.

Table 1 describes the treatments. In the baseline treatment, no one could see the confidence of another player; each receiver (R) could observe the reported confidence by the paired sender (S) in the other treatments. In all treatments there was a possible competition between the paired S and R, based on their Raven scores. In Treatments 1-3, the player with higher rank received €10 and the other received nothing. Entry by both players was mandatory in *Baseline* and *Social*, but each R faced a strategic decision in *Deter* and *Lure*: After observing S's reported confidence, R chose whether or not to enter a tournament. In the low-outside-option version of *Deter*, R received €3.5 by staying out, while in the high-outside-option version of *Deter*, R received €5.5 for doing so.[3] In these treatments, S preferred that R opted out of the tournament since that would secure €10. In *Lure*, if R chose not to enter, R received €5.5 and S received €10. If R chose to enter and won,

[3]We initially used an outside option of €3.5, but found that 28 of 30 receivers entered the tournament. We then switched to an outside option of €5.5.

then R received €10 and S received €15, while if R entered and S won, then R received 0 and S received €25; thus, S preferred that R enter the tournament. In *Deter* and *Lure* participants must trade off honest reporting against trying to influence the opponent's entry decision.

The description we have just given, including whether or not any player could see the reported confidence of others, or whether player R could choose to compete or not, was common information and provided to all subjects before they reported their confidence. They were also told that they would find out at the end of the session who had the higher rank between the two matched players, but would learn neither their rank in the group of four nor the number of correct answers.[4]

Table 1: *Overview of treatments*

Treatment	Receiver observes Sender's reported confidence?	Payoffs if receiver opts out of tournament (S, R)	Payoffs if receiver enters tournament (S, R)	
			Sender wins	Receiver wins
1: Baseline	No	N/A	(10,0)	(0,10)
2: Social	Yes	N/A	(10,0)	(0,10)
3a. *Deter* (low)	Yes	(10, 3.5)	(10,0)	(0,10)
3b. *Deter* (high)	Yes	(10, 5.5)	(10,0)	(0,10)
4. Lure	Yes	(10, 5.5)	(25,0)	(15,10)

Notes: S stands for Sender, R for Receiver.

3.2. **Experimental Procedures.** Sessions were conducted in Amsterdam with 16 to 28 participants depending on the number of subjects showing up. We ran a total of 22 sessions with a total of 464 subjects.[5] Sessions were run in 2009 and 2010, with the exception of Treatment 4, which was run in 2012. Sessions ended with a questionnaire, and lasted for 40-50 minutes.

Overall, 44 percent of subjects were female. Table 2 shows the gender composition by treatment and role. We tried to get a balanced gender composition, but there are some differences between treatments and roles depending on show-up and random allocation to roles. Most notably, the fraction of females is lower in *Social*. We cannot reject that the different samples are drawn from the same population with respect to the gender composition (Kruskal-Wallis test, p=0.540).

[4]Treatment 1 had some additional components in which we presented updating tasks to participants. We will not describe the results of that part in this paper. Details can be found in Charness, Rustichini, and van de Ven (2015).

[5]We conducted more sessions of Treatment 1 because we were interested in updating behavior after being given noisy feedback about performance. Our basic finding is that people do relatively poorly when the feedback is about their own performance, but do quite well in an isomorphic task that does not involve their own performance. Again, details can be found in Charness, Rustichini, and van de Ven (2015).

Table 2: *Gender composition*

Role	Baseline	Social	Deter	Lure
All	40.97	36.76	48.08	48.96
Senders		38.23	44.87	45.83
Receivers		35.29	51.28	52.08
N	144	68	156	96

Note: Percentage of females.

Participants were undergraduate students (average age 22 years, with the majority studying economics or business, see Appendix C of the Online Supplementary Materials for more details). We also measured risk attitudes using a price list. Participants were presented a choice between a certain amount and a lottery that yielded €10 with 2/3 chance and €0 with 1/3 chance. The certain amount ranged from €4.50 to €7.00 in steps of 0.5. We use the number of safe choices as the degree of risk aversion.[6]

Instructions were displayed on a computer screen and read aloud (see Appendix B of the Online Supplementary Materials for the instructions). Participants were told that their decisions would remain anonymous to the other people present unless explicitly indicated otherwise, and that they would receive their earnings in an envelope from a person in a different room who could only see login numbers and could not match these numbers to names or faces. Payments were presented in points: One point was worth one euro. Subjects were told that the experiment consisted of several parts and that they would be paid for one randomly chosen part. The latter was done to avoid any hedging across the different parts. The average payment was €14 (of which €7 was a show-up fee).

3.3. **Model.** The key tool we use to test our hypotheses is a tournament game where players can send explicit statements on their ability. A simple model may illustrate this game. Here, we only give a brief sketch of the outline of the model. A full characterization is provided in Appendix A of the online Supplementary Materials.

There are two players. One of them (the sender, player 1) sends a message about her ability. The other player (the receiver, player 2) then decides whether to enter a tournament with the

[6]We started collecting risk attitudes only after a while. We measured this among all participants in treatment Deter with the high outside option, in treatment Lure, and in a few sessions of Baseline.

sender. Players can have different types, reflecting different abilities. The type of player i, $\theta^i \in \Theta^i$, is chosen according to some probability distribution, and is private information to the player. To simplify the exposition we assume that the set of types has only two elements, $\Theta^i \equiv \{\theta^i_0, \theta^i_1\}$ with θ_1 of higher ability than θ_0, and that the prior of players is that both types are equally likely.[7]

The sender moves first, and makes a claim about her ability by sending a message $t \in T = \Theta^1$. The message need not be truthful, but sending a false message has a lying cost $c > 0$. After observing the message, the receiver can choose an action from the set $\{In, Out\}$. If the receiver chooses Out, both players receive their outside option O^i. If the receiver chooses In, both players compete in a tournament, and their payoffs are determined by their abilities and are given by:

(1)

	θ^2_0	θ^2_1
θ^1_0	$0, 0$	b, a
θ^1_1	a, b	d, d

We focus on the case for which a player is better off if the opponent is weaker ($b \leq 0$, $d \leq a$) and if she herself is stronger ($a \geq 0, d \geq b$). To avoid trivial cases we assume $a \geq O^2 \geq d$, so that a strong receiver weakly prefers playing the tournament to the outside option if he knows that the sender is a weak type, but prefers the outside option otherwise. We also assume that $d \geq 0$, implying that a weak receiver always weakly prefers to stay out (since $O^2 \geq d \geq 0 \geq b$).

This game reflects situations in which people can strategically manipulate how confident they appear to others. Under the assumptions made, a weak receiver will always opt out of the tournament, but for a strong receiver this choice will depend on his beliefs about the sender's type. The best strategy for the sender depends crucially on her outside option. If her outside option is high, she is better off when she does not have to compete with the receiver in the tournament. The sender can try to achieve this by appearing strong, i.e., over-report, to convince the receiver to opt out. On the other hand, if her outside option is low, she prefers that the receiver compete with her. In this case, the sender can try to achieve this by appearing weak, i.e., under-report. Indeed, both

[7]Of course we have more than two types in our experiment, so that this is a simplification. Nevertheless, we feel that this simple model captures the essential features of the environment.

over- and under-reporting may occur in equilibrium (see Appendix A of the Online Supplementary Materials for details).[8]

In the experiment, we implemented two strategic conditions. In *Deter*, parameters are such that, provided the cost of lying are not too high, senders over-report in equilibrium (claim to be a higher type than they really are). In *Lure*, senders under-report in equilibrium (for sufficiently low cost of lying). In both conditions, if the cost of lying are sufficiently high, truth-telling will result as an equilibrium outcome. While our design accommodates both under and over-reporting, we argue (and provide survey evidence) that the luring environment is relatively rare. In our game, if the payoff of players is higher if the opponent is a stronger type, only equilibria with over-reporting exist for reasonable parameters (i.e., sufficiently low cost of lying). *Mating* games are a prominent example of these games: players compete for mating with strong types. These games are very common in nature so that overconfidence can be expected to be advantageous more frequently than under-confidence. An example at the workplace that has this structure is when an employee wants to convince co-workers that (s)he is talented, so that (s)he will be chosen to collaborate on a joint project.

3.4. **Hypotheses.** We consider that statements of confidence are typically social signals of intentions or private information, and individuals take them into account when they observe the self-evaluations of others. We therefore predict that tournament entry choices will be sensitive to the sender's confidence statements, where higher stated confidence levels will discourage competitors from entering the tournament, and lower stated levels will encourage competitors to enter.

H1: *The likelihood that a receiver enters the tournament decreases in the sender's reported confidence.*

Senders may anticipate this effect and adjust their signals accordingly. So accurate reporting of confidence levels is not an optimal strategy, even with the incentives for accurate reporting.

[8]Kartik (2009) also analyzes a sender-receiver game with lying costs. His setup closely resembles that of Crawford and Sobel (1982), with the addition of lying costs. This transforms the cheap talk model of Crawford and Sobel into a signalling game. Kartik shows that in his setup senders almost always claim to be more confident than they really are. His game has a different setup, so we cannot directly apply his results. For instance, the action set of the receiver is different and we have two-sided instead of one-sided private information. The payoff structure also differs, such that our game is a monotonic signalling game whereas the one by Kartik is not.

Both stated overconfidence and under-confidence can be motivated by strategic considerations, and subjects behave according to the equilibrium predictions. Specifically, senders will increase (decrease) stated confidence when they wish to deter (encourage) tournament entry.

H2: *Compared to the baseline treatment, senders report higher confidence in the Deter treatment and lower confidence in the Lure treatment.*

Since senders are always automatically entered into the tournament, there is no strategic reason for receivers to adjust their reported confidence depending on the environment they are in. Thus, we do not expect their reported confidence levels to differ between treatments.

H3: *For receivers, there are no treatment differences with respect to reported confidence.*

In light of documented gender differences in overconfidence (see e.g., Lundenberg et al., 1994; Barber and Odean, 2001) and the evidence that males and females behave differently in competitive environments (see e.g., Gneezy and Rustichini, 2004; Gneezy et al., 2003; Niederle and Vesterlund, 2007), we also consider the possibility that confidence display differs across genders. We predicted that males are more likely to enter the tournament than females.

H4: *Males will exhibit higher stated confidence levels and are more likely to enter the tournament than females.*

4. Experimental Results

4.1. **Performance and Confidence.** The mean number of correct answers is 8.75 (st. dev. 2.40). The distribution of correct answers (out of 15) is approximately normal, and comparable across genders (8.78 for males and 8.71 for females). The mean reported confidence is 63.75 (st. dev. 21.89). No more than 27 percent report a confidence level below 50 percent in any of our conditions; note however that the rate in the *Lure* treatment (27 percent) was nearly double that in the *Deter* treatment (14 percent). In data pooled over the conditions, 71 percent of the people report a confidence level above 50 percent and only 20 percent report a confidence level below 50 percent; a binomial test finds this asymmetry to be highly significant ($Z = 17.00$, $p < 0.001$).

People in the baseline treatment have no strategic reason to over- or underreport their confidence. Even in that treatment, the mean reported confidence is 63 percent, well above 50 percent ($p < 0.001$, two-sided t-test).[9] The mean reported confidence is 65 percent for males and 61 percent for females, a difference that is not statistically significant ($p = 0.221$, two-sided t-test). The degree of overconfidence varies with ability. Figure 1 shows the mean reported confidence for different performance levels (in bins, for smoothing purposes). The figure also plots the actual probability of being in the top two of the group, based on the actual distribution of performance levels (recall, however, that participants did not get feedback on their performance). It appears that people with low performance are the most overconfident, while the best performers are somewhat underconfident.

FIGURE 1. **Confidence in Baseline, by performance.** "Actual" is the probability of being in the top 2 based on the observed distribution of performance levels.

[INSERT FIGURE 1 HERE]

Figure 2 shows the confidence of senders and receivers in each treatment. We did not expect to find a difference in stated confidence between senders and receivers in the baseline treatment, since their roles do not differ in that treatment, and indeed we do not find any: the confidence of senders is 63 and that of receivers is 64. We therefore pool the observations of senders and receivers in the baseline treatment.

Senders in *Deter* report higher confidence than in *Baseline* but the difference is not significant ($Z = 1.467$, $p = 0.142$). Senders in *Lure* report a significantly lower confidence ($Z = 2.949$, $p = 0.003$). The reported confidence in *Social* is not significantly different from either *Baseline* or *Deter* ($p = 0.509$ and $p = 0.652$, respectively).

FIGURE 2. **Confidence of senders and receivers, by treatment.** Roles pooled in the baseline treatment. Error bars: +/- SE

[9] While this suggests overconfidence, we cannot reject rational Bayesian updating using the Burks et al. (2013) allocation function. This may reflect our having only two intervals, either above or below the median.

[INSERT FIGURE 2 HERE]

Given the considerable degree of difference in confidence levels and entry rates for males and females in, for instance, Niederle and Vesterlund (2007), we naturally gathered data on gender. In fact, we find interesting differences across gender and so we next report treatment effects by gender.

Figure 3 shows the confidence of senders and receivers in each treatment by gender. We again find no differences between senders and receivers in the baseline treatment, and we therefore pool the observations of senders and receivers in the baseline treatment. Male senders in the *Social* treatment report similar confidence levels to the ones reported by senders in the *Baseline* treatment.[10] However, they do report significantly higher confidence in the *Deter* treatments (73 percent, low and high outside option combined, $Z = -2.342$, $p = 0.019$) and lower confidence in the *Lure* treatment (53 percent, $Z = 2.007$, $p = 0.045$).[11] Male *receivers* also report significantly higher confidence in the *Deter* treatments (76 percent, $Z = -2.949$, $p = 0.003$), while their reported confidence in the *Lure* treatment is comparable to that of males in the baseline treatment (65 percent, $Z = -0.437$, $p = 0.662$). Stated confidence by male receivers in *Deter* is also significantly higher than in *Social* ($Z = 2.317$, $p = 0.021$). The reported confidence of females is not statistically different from that in the baseline treatment in any of the other treatments, except that female senders report lower confidence of 51 in the *Lure* treatment ($Z = 1.779$, $p = 0.075$, two-tailed test).[12] Testing for gender differences across treatments, we find that stated confidence is significantly higher for males than for females in *Deter* ($Z = 4.336$, $p < 0.001$) but not in any other treatment (*Baseline*: $Z = 1.595$, $p < 0.111$; *Social*: $Z = 0.550$, $p = 0.582$; *Lure*: $Z = 0.683$, $p = 0.495$).

[10] Note that Ewers and Zimmerman (2014) do find evidence that participants report higher confidence if an audience can observe the reports. Possibly, the effects are stronger in their study because participants could identify other participants, while in our experiment participants remain anonymous to each other. However, they also find no audience effect when that audience receives feedback about the participant's actual performance. That treatment is most similar to our setup, because in our setup the receiver also receives feedback, and learns whether or not the sender had a higher rank than the receiver.

[11] The difference in reported confidence by male senders in *Deter* and *Social* is not significant ($Z = 0.883$, $p = 0.378$).

[12] If we split reported confidence in treatment *Deter* by the low and high outside option, we get: Male senders, low outside option: 72; male senders, high outside option: 74; male receivers, low outside option: 75; male receivers, high outside option: 76; female senders, low outside option: 58; female senders, high outside option: 61; female receivers, low outside option: 66; female receivers, high outside option: 57.

Figure 3. **Confidence of senders and receivers, by gender and treatment.** Roles pooled in the baseline treatment. Error bars: +/- SE

[INSERT FIGURE 3 HERE]

OLS estimates of the determinants of confidence are presented in Table 3; the baseline condition reflects male behavior in the baseline treatment (Treatment 1). Specification (1) shows that the number of correct answers is a strong predictor of confidence, adding about three percentage points for each correct answer; this result is robust over different specifications.

Since subjects were not told their number of correct answers, the effect of correct answers on stated confidence can only be based on one's own perceived relative ability. In addition, we find significantly lower stated confidence in the *Lure* treatment, although the effect of the *Deter* treatment is not quite significant. Specification (2) adds controls for the role of the participant (sender or receiver) and interaction terms for the treatment and role. Being a Sender has no effect by itself, nor is there a significant interaction effect with the *Social* and *Deter* treatments. However, there is a large interaction effect in the *Lure* treatment, indicating that the decrease in stated confidence is entirely due to senders; in fact, the coefficient on *Lure* is now positive, although not significant.

We introduce a dummy for gender and interaction effects for gender and treatment in specification (3). The results are consistent with the picture of the nonparametric tests. We find a negative but insignificant direct gender effect. However, there is a significant treatment effect: Reported confidence increases by almost 10 percentage points in the *Deter* treatment. This effect is only present for males, as the coefficient of the interaction between *Deter* and Female shows a negative coefficient of about the same size as the treatment coefficient. On the other hand, there is no such difference by gender in the *Lure* treatment, indicating that both male and female senders deflate stated confidence. There is no significant difference for sender or with the interaction of either *Deter* or *Lure* and Sender, and none for the three-way interactions. The results are also robust to including a set of other controls, such as age, study, and number of siblings (specification (4)).

Figure 4 shows the reported confidence for different performance level (using bins for smoothing), by gender and treatment. Compared to *Baseline*, Reported confidence tends to be higher in *Deter* (for males) and lower in *Lure* for most performance levels. Note, however, that the number of observations at each point can be very low, after splitting by treatment, gender, and performance.

Table 3: *Determinants of confidence* (scale 0 – 100)

Dependent var.: Confidence	(1)	(2)	(3)	(4)
Number of correct answers	3.50***	3.52***	3.52***	3.32***
	(0.39)	(0.39)	(0.38)	(0.40)
Social	0.29	0.05	0.08	-1.74
	(2.94)	(4.15)	(4.07)	(4.12)
Deter (low and high)	3.54	3.06	9.21**	9.26**
	(2.31)	(3.25)	(4.02)	(4.07)
Lure	-5.55**	1.06	0.57	1.06
	(2.64)	(3.70)	(4.77)	(4.79)
Sender		0.26	0.65	1.02
		(3.31)	(3.27)	(3.26)
Social × Sender		0.47	0.16	1.82
		(5.88)	(5.77)	(5.82)
Deter × Sender		0.96	-0.00	-0.57
		(4.59)	(5.43)	(5.48)
Lure × Sender		-13.21**	-13.21**	-15.18**
		(5.24)	(6.47)	(6.46)
Female			-3.06	-2.57
			(2.76)	(2.78)
Deter × Female			-11.00**	-12.30**
			(5.20)	(5.24)
Deter × Sender × Female			-0.71	-0.58
			(6.26)	(6.31)
Lure × Female			1.96	-0.04
			(6.27)	(6.28)
Lure × Sender × Female			-1.00	0.51
			(7.97)	(8.04)
Familiar with conditional probs.				4.46**
				(1.93)
Constant	33.05***	32.78***	33.76***	31.54***
	(3.75)	(4.16)	(4.22)	(8.97)
Observations	464	464	464	462
R-squared	0.17	0.19	0.23	0.26
Adj. R-squared	0.16	0.18	0.21	0.23

Notes: OLS estimates. Other control variables in model (4) are: familiarity with Raven test, study category, age, number of siblings, birth order, member of sports club, entity theory question. s.e. in parentheses. *** $p < 0.01$, **$p < 0.05$, * $p < 0.1$.

FIGURE 4. **Senders' Confidence by performance bins.** Top panel: Males; bottom panel: Females. Roles pooled in *Baseline*.

[INSERT FIGURE 4 HERE]

Results (Confidence)

(1) *The real performance of participants, measured by the (unknown to the participants) number of correct answers, significantly correlates with reported confidence in the expected directions.*

(2) *Men report a significant 10 percentage points higher confidence in the* Deter *treatment compared to the baseline. There is no significant treatment effect for women.*

(3) *There is a significant treatment effect in the* Lure *treatment, as both male and female senders deflate their stated confidence by about 15 percentage points compared to the baseline.*

(4) *In the* Deter *treatment, male receivers increase confidence as much as male senders do compared to the baseline. Receivers in the the* Lure *treatment do not deflate stated confidence.*

The confidence reports in the *Deter* and *Lure* treatments will be discussed again in the analysis of the strategic behavior of participants.

4.2. **Voluntary tournament entry.** In the *Deter* and *Lure* treatments, player R chooses whether or not to enter a tournament with player S. What determines player R's choice?

Our data show that with the high outside option in the *Deter* treatment and in the *Lure* treatment, player R is much more likely to enter the tournament when own confidence is higher and when the opponent's confidence is lower.[13] Indeed, as is shown in Figure 5, we observe that relative confidence is a very good predictor of entry. In the *Deter* case, 23 of 25 receivers (92.0 percent) enter when their confidence level is at least as large as the paired sender's reported confidence level, while only four of 23 receivers (17.4 percent) enter when their confidence level is lower; a highly significant difference ($Z = 5.21$, $p = 0.000$). The corresponding data for the *Lure* treatment show that 28 of 32 (87.5 percent) choose to enter with higher stated confidence and five of 16 (31.2 percent) choose to enter with lower stated confidence, again a significant difference ($Z = 3.96$, $p = 0.000$). Thus, there is strong potential for senders to influence the receiver's decision.

Our analysis suggests that receivers follow a simple rule to make their entry decision, entering if and only if their own confidence is at least as high as the reported confidence of the sender. This rule correctly classifies 87.5 percent of the receivers' decisions in the *Deter* treatment (high outside

[13] We focus primarily on entry with the high outside option, since 28 of 30 receivers chose entry with the low outside option, so that statistical tests have little power.

option), and 81.25 percent in the *Lure* treatment. In Figure 6 we plot the predicted probability of entering as a function of the difference in confidence (own confidence - opponent's confidence). The predicted probability is based on a probit estimation regressing the entry decision on the difference in confidence and allowing for a discontinuity at zero. In both treatments, we see a sharp increase in the predicted probability of entering around 0, consistent with the simple rule to enter if and only if the receiver's own confidence is at least as high as the reported confidence of the sender.[14] Hence, receivers appear to take the confidence statements at face value instead of deflating them in the *Deter* treatment or inflating them in the *Lure* treatment. We cannot, however, rule out that they somewhat correct for senders' statements; our data is not sufficiently rich to distinguish whether receivers take senders' statements at face value or correct statements by some small amount.

FIGURE 5. **Entry by lower confidence.** Error bars: +/- SE

[INSERT FIGURE 5 HERE]

FIGURE 6. **Predicted probability of entry.** Difference in confidence is own confidence minus the opponent's confidence. Predictions based on a probit model regressing entry on difference in confidence allowing for a discontinuity at zero.

[INSERT FIGURE 6 HERE]

In the *Deter* treatment, we find that males enter twice as frequently as do females, 75.0 percent versus 37.5 percent ($Z = 2.62$, $p = 0.009$, two-tailed test), as is shown in Figure 7. However, this does not reflect a difference in performance: females in the R role in the high-option condition do nearly as well as males on the Raven test (the mean score is 9.12 for males and 8.88 for females; Wilcoxon ranksum test: $Z = 0.28$, $p = 0.779$, two-tailed test). Female receivers state significantly lower confidence levels than do male receivers, 56.63 versus 75.83 ($Z = 3.07$, $p = 0.002$, two-tailed

[14] To be more specific, we estimate: $Prob(Enters|\cdot) = \Phi(\beta_0 + \beta_1 * \Delta Confidence + \beta_2 * \mathbb{1}_{\Delta Confidence \geq 0} + \varepsilon)$, where $\Delta Confidence$ is the difference in confidence. The coefficient β_2 is significantly different from zero in *Lure* but not in *Deter*. Estimates from a linear probability model yields significant coefficients in both treatments.

test), so that the entry rate reflects a higher stated confidence level. This effect is only seen for people who choose to enter the tournament; the average stated confidence level for male entrants is 84.17 versus 69.89 for female entrants, while this comparison is 50.83 versus 48.67 for male non-entrants and female non-entrants, respectively.

The results are quite different in the *Lure* treatment, where the entry rate for males (65.22 percent) is actually *lower* than the entry rate for females (72.00 percent). The average stated confidence level for male entrants is 73.87 versus 69.00 for female entrants, while the average stated confidence level for male non-entrants is 49.62 versus 49.00 for female non-entrants, so we can see that the stated confidence levels for receivers in the *Lure* treatment is largely unaffected by gender. The performance level was 8.82 for males and 8.56 for females, not significantly different (Wilcoxon ranksum test: $Z = 0.20$, $p = 0.843$, two-tailed test).

FIGURE 7. **Entry by gender.** Error bars: +/- SE

[INSERT FIGURE 7 HERE]

Table 4 reports the probit estimates of the decisions to enter the tournament.[15] The first four columns apply to the high outside-option sessions of the *Deter* treatment, and the last four columns apply to the *Lure* treatment. Specification (1) shows that own confidence increases the likelihood of entering the tournament, and the confidence of the opponent decreases it. Each variable substantially affects the probability of entering. Specification (2) used the difference in confidence as an independent variable, showing that a larger difference increases the likelihood of entering. Specification (3) includes a dummy variable that simply compares if own confidence is higher or lower than that of the opponent. Controls in (4) for gender, number of correct answers, and risk aversion have no significant effect. Thus, the lower entry rate by females seems driven by lower confidence, a result that is consistent with past findings (Niederle and Vesterlund, 2007). This also suggests that males are not just reporting higher confidence, but also *feel* more confident. If they were just

[15]Estimates from a Linear Probability Model are qualitatively very similar.

reporting higher confidence without believing it, then, controlling for confidence, males should have been less likely to enter the tournament. The results for the *Lure* treatment are similar, but with smaller magnitudes.

Table 4: *Determinants of entering*

Dependent var: *choice is In*	*Deter* (high outside option)				*Lure*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Number correct answers	-0.041	-0.010	-0.022	-0.019	0.027	0.028	0.019	0.016
	(0.046)	(0.041)	(0.042)	(0.042)	(0.043)	(0.042)	(0.045)	(0.046)
Own confidence	0.048***		0.016**	0.014*	0.012***		0.006	0.008
	(0.014)		(0.007)	(0.008)	(0.004)		(0.005)	(0.005)
Opponent's confidence	-0.031***				-0.004			
	(0.011)				(0.003)			
Own - Opponent's confidence		0.034***				0.007***		
		(0.010)				(0.002)		
Confidence is lower‡			-0.618***	-0.621***			-0.397**	-0.414**
			(0.139)	(0.141)			(0.186)	(0.190)
Female				-0.096				0.040
				(0.203)				(0.149)
Risk aversion†				-0.030				0.044
				(0.080)				(0.044)
Observations	48	48	48	48	48	48	48	48
Pseudo R-squared	0.60	0.52	0.57	0.57	0.26	0.22	0.30	0.32

Notes: Probit estimates, reporting marginal effects. ‡Lower confidence is equal to 1 if receiver's confidence is lower than the paired sender's, and 0 otherwise.†Eight missing observations were replaced by the mean in model (3). St. err. in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Results (Tournament Entry)

(5) *When deciding whether to enter the tournament, participants are more likely to enter when their reported confidence is higher*

(6) *When deciding whether to enter the tournament, participants are sensitive to the confidence reported by the opponent. In particular, participants are much less likely to enter if their own reported confidence is lower than that of the opponent.*

(7) *Females are less likely to enter the competition in the* Deter *treatment, but this effect is mainly due to the difference in confidence. Once we control for confidence, the entry rate*

of women is not significantly lower. There is no significant entry difference in the Lure treatment, regardless of whether or not we control for confidence.

5. Discussion

5.1. **Motivations for reported overconfidence.** What do these results tell us about the origin and motivation of overconfidence? One key motivation for being overconfident is the ego utility that one derives, producing an increase in self-esteem. In our Baseline treatment we observe substantial overconfidence even when it is known that the other player cannot observe one's stated confidence level. This finding suggests that people receive some internal benefit from this inflated belief. We do not find evidence of social image concerns, but this is a mutually-anonymous environment and social observation is quite attenuated.

Our results are consistent with our general hypothesis that views strategic concerns as a source of overconfidence. In fact we see strong evidence that an increase (decrease) in a sender's reported confidence can have deterrent (encouragement) value in terms of inducing the receiver into (or out of) the tournament. We also see some evidence (see Figure 2 and Table 3) that males report higher confidence in the strategic condition than in the baseline treatment.

Let us now consider our hypotheses in light of the data. First, statements of confidence do indeed appear to be social signals of intentions or private information. Receivers take these statements into account when they decide to enter the tournament, in support of H1 (see Result 5). Nevertheless, one potential puzzle is why receivers don't take into account that senders might be inflating (or deflating) confidence levels.[16] To explain this, we appeal to the prevailing view that people have difficulty with higher-order reasoning, as discussed extensively in the literature on k-level thinking (e.g., Crawford and Iriberri, 2007). It is much easier for a sender to realize that a receiver could be swayed by a strategically-biased confidence report than it is for a receiver to realize that the sender realizes that the receiver could be swayed by the confidence report. We do suspect that these results could be different if there were multiple periods of play, particularly if the receiver was given feedback about the sender's actual Raven score.

[16]With the caveat that the data is not sufficiently rich to exclude that receivers discount signals to a small degree.

Second, since we do find that stated confidence levels affect the likelihood that receivers enter the tournament, one would expect to see inflated confidence reports in *Deter* and deflated ones in *Lure*, despite the incentives for accurate reporting (which would support H2). In fact, we see this in all cases, except for females in the *Deter* treatment. Why is there this exception?

While we can only speculate, a 'stylized fact' is that women are much more reluctant to appear (or even feel) overconfident than men. According to Kay and Shipman (2009), "There is a particular crisis for women - a vast confidence gap that separates the sexes. Compared with men, women ... generally underestimate their abilities. This disparity stems from factors ranging from upbringing to biology."[17] Perhaps this reflects a historical belief that men are not interested in women who appear overconfident, so that any such tendencies are reined in. In any case, we suspect that the failure of women to appear overconfident is related to the lack of inflated confidence reports in the *Deter* treatment. We note that there is no such reluctance to appear underconfident, so that women do deflate confidence reports in this environment.

Third, we find that stated overconfidence occurs even when no one else is watching, as in the baseline treatment where competition is absent. In addition, in contrast to H3, we find this for male receivers in *Deter* when there is clearly no cognitive reason to engage in this behavior.

Finally, we find mixed support for H4. In the *Deter* treatment, men do enter at twice the rate of women, although this difference vanishes when controlling for confidence. Women enter the tournament slightly more frequently than men in the *Lure* treatment, but the effect is not significant.

5.2. **Optimality of decisions.** We showed that the entry decisions of the receivers are largely explained by the rule to enter the tournament if and only if his own confidence is at least as high as the sender's reported confidence. If we assume that receivers use this rule, and that senders anticipate this, we can analyze the best response of senders.

FIGURE 8. **Optimal reporting by senders given the behavior of receivers.**

[17]Furthermore, when an article on the dangers of overconfidence appears, it may well be published in a women's journal. See, for example, "When Confidence Can Be a Bad Thing," in Women's Health Magazine online, http://www.womenshealthmag.com/life/overconfidence.

[INSERT FIGURE 8 HERE]

Figure 8 plots the optimal reporting for senders under this assumptions (see Appendix E of the Online Supplementary Materials for details). The thin solid line represents truthful reporting. The thick solid line represents the optimal report in the *Deter* treatment, and the dashed line for the *Lure* treatment. In both treatments it is optimal for senders to deviate substantially from their true belief. For instance, the optimal report for a sender with a confidence of 60 is 80 in the *Deter* treatment, and 20 in the *Lure* treatment. In both cases the optimal deviation from truthful reporting is substantially larger than what the data show. Based on reported confidence levels in the baseline treatment, we should expect that the average confidence is about 20 higher in the *Deter* treatment (we find roughly zero for females and 10 for males), and 45 lower in the *Lure* treatment (we find about 10-15 lower). We should also find no reported confidence below 60 in the *Deter* treatment (because even for type 0 the optimal report is above 60), or above 30 in the *Lure* treatment (because even for type 100 the optimal report is below 30), but we see quite a few examples of this in the data.

We conclude that the behavior of senders goes in the right direction, but not far enough.[18] The fact that they do not exactly match the estimated optimal levels for reports is not so surprising. They must form expectations about several parameters, e.g., those related to the distribution of types, and they only play the game once.

Receivers, in their turn, appear to take the confidence statements more or less at face value instead of deflating them in the *Deter* treatment or inflating them in the *Lure* treatment, as they simply seem to compare their own confidence to the reported confidence by the sender. This behavior is an indication that receivers also do not anticipate a level of over- and under-inflating as high as our estimated optimal reports.

5.3. **Salient Perturbations.** Male receivers in the Deter treatment inflate their confidence levels to about the same degree as the senders (the coefficient for the interaction variable Deter*Sender

[18]The exact magnitude depends on the assumptions we make. In particular, the distribution of types matters. We have also estimated optimal reporting for alternative distributions (assuming a different mean or a uniform instead of normal distribution) but the under-inflating and deflating seems robust to different specifications.

in Table 3 is small and insignificant). The receivers inflated levels of stated confidence cannot deter senders from entry and is known to not even be observed, so this cannot reflect deliberate cognitive planning. One possibility is that male receivers try to justify their intention to enter into the tournament by reporting a high confidence level. If this were the case, we would expect to see different behavior between male receivers in the low and high outside option, as entry rates are very different. This is not supported by the data, which shows that the reported confidence by male receivers is very similar across those two conditions.

One theory that organizes our data rather well relies upon the theoretical literature on bounded rationality. Myerson (1991) proposes that apparent suboptimal behavior can sometimes be understood by assuming that observed behavior is optimal in a related but more familiar environment, which he calls a *salient perturbation*. This notion suggests that behavior by male receivers reflects unconscious motivations generated by the competitive setting, so that people may not be flexible enough to adjust their behavior to their contingent role in the *Deter* environment.

The starting assumption of the notion of salient perturbations is that cognitive limitations prevent people from calculating optimal behavior in each and every situation. Instead, people make the same decisions in different situations that appear similar. This idea is presented in Myerson (1991) and further developed in formal models by Samuelson (2001) and Jehiel (2005); it is also discussed in Gurdal, Miller, and Rustichini (2014). This concept could potentially explain how observed non-optimal behavior can be explained as behavior that would be optimal in a related and familiar environment, but is not optimal in the actual environment. By contrast, unfamiliar environments are more likely to induce deliberate and reflective behavior, inducing people to calculate optimal behavior. To illustrate, people drive automatically on a familiar route but pay more attention to their environment on an unfamiliar road.

The discipline imposed by the theory of salient perturbations is that the environment in which behavior is non-optimal must satisfy three conditions. First, the salient perturbation has to be similar to the situation the individual is really facing (has to be a perturbation). Second (familiarity), the perturbation must be more familiar to the subject than the real situation. The degree of familiarity is measured by the frequency with which one faces a particular situation. Finally (optimality), the observed behavior must be optimal in the salient perturbation of the actual game.

Applied to our context, we hypothesize that the familiar situation is one in which appearing confident is optimal. In this case, male receivers may over-report in Deter, even without a direct benefit for doing so, because the situation looks familiar to them. By contrast, the Lure treatment may be perceived as less familiar, so that as a result a more reflective type of behavior is triggered, in which receivers do not adjust their reported confidence in comparison to the baseline treatment. This interpretation presumes that environments where overconfidence is effective are widespread and familiar, and that environments where under-confidence is effective are relatively rare. The purpose of the Lure treatment was to test this prediction, on the presumption that situations in which it is beneficial to appear under-confident are unfamiliar. In the Social treatment, there is no strategic benefit of appearing confident for either role, so this does not cue a competitive situation where overconfidence is optimal. Receiver behavior is indeed consistent with this prediction.

The presumption that environments where it is beneficial to appear overconfident are more familiar than environments where it is beneficial to appear under-confident seems reasonable and is also supported by additional survey evidence that we collected after analyzing the other data. We asked a new set of participants to rate the familiarity of the two types of situations on a 5-point scale, ranging from very rarely to very frequently (82 participants, 49 percent female, recruited from the same subject pool as for the main experiment).[19] The results are reported in Figure 9. The modal responses of the participants are that situations involving under-confidence happen rarely and that situations involving overconfidence are quite frequent. The hypothesis that the distributions are equal is rejected (Kolmogorov-Smirnov test, $D = 0.346$, $p < 0.001$). Thus, the results provide clear evidence that appearing overconfident is more familiar to our participants. Our survey evidence suggests that each of these is indeed the case in the subjects experience. Thus, in our experiment, the three conditions for a salient perturbation are satisfied: our treatments use similar games that differ in only one respect, and overconfidence is optimal in the familiar environment.

[19]These sessions were run in 2014. See Appendix D of the Online Supplementary Materials for more details.

Figure 9. **Perceived frequency of situations where appearing underconfident (left) or overconfident (right) can be effective.** The category "(very) rare" pools the answers "rare" and "very rare," and the category "(very) frequent" pools the answers "frequent" and "very frequent."

[INSERT FIGURE 9 HERE]

It is also worth noting that the over-reporting of confidence may not be conscious. In fact, the observation that male receivers over-report in Deter indicates that it is not. This suggests that subjects seem to believe their own inflated reports (at least at a conscious level), so this in fact appears to reflect actual overconfidence. This can be advantageous to senders: Trivers (1985) suggests that it is much easier for such a signal to be convincing if the sender believes it. To be clear, we are not proposing that a strategic imperative is the only consideration that generates overconfidence. Indeed, ego utility may well play a part in the behavior and we take the point about social-image concerns having a role, as shown in Burks et al. (2013) and Ewers and Zimmermann (2014). Neither do we propose that the concept of salient perturbations is the only possible explanation for the observed patterns in the data. Some of the findings are not readily explained by salient perturbations, in particular the fact that females under-report in Lure but do not over-report in Deter. Nevertheless, we demonstrate that stated overconfidence (and underconfidence, where appropriate) reflects strategic considerations and that salient perturbations do a rather good job of organizing our data. And again, interestingly, it appears that some of the observed behavior is determined at a non-conscious level.

6. Conclusion

Our experiments examined the determinants of self-confidence, and the degree to which it reflects strategic concerns about social image. Our main conclusion is that levels of stated confidence are likely to be influenced by strategic interest, perhaps unconsciously processed. We see evidence that people will inflate or deflate statements of confidence levels when doing so is strategically beneficial.

Our novel strategic environment (in which another party observes the stated confidence level of another and then chooses whether or not to enter a tournament with this other person) allows a direct test of the strategic-interest hypothesis. First, the social signal is perceived and has consequences: subjects in our experiment do respond to statements about confidence made by others, taking that information into account when choosing whether or not to enter. In the *Deter* treatment, male (but not female) participants on average report significantly higher confidence levels than in the non-strategic treatments. Inflated confidence serves as an effective deterrent. Interestingly, males (but not females) do so in both roles, even when deterrence is impossible; this suggests processing on an unconscious level. We suggest that inflating confidence when not strategically beneficial is consistent with the notion of salient perturbations. In familiar situations overconfidence often has strategic value so that we may also observe it in non-strategic environments that are similar to the familiar situation.

In the less-familiar *Lure* environment, we observe deflated confidence for both men and women in the role of senders, which serves to encourage entry. We argue that conscious cognition is present in this less-familiar environment, and indeed receivers do not deflate their own reports. Strategic deterrence and luring are consistent with the equilibrium we characterize; the degree to which one engages in costly strategic distortion depends on the values of the parameters in the game.

When inflated reported confidence is strategic, it is natural to find gender differences in our participants' behavior, given the evidence of other gender differences such as with respect to financial risk preferences (Charness and Gneezy 2010, 2012), competition (Gneezy, Niederle, and Rustichini 2003), and even shame (Ludwig and Thoma 2012). But since luring is a much less familiar environment and strategic distortion is presumably driven by cognitive ability (which is the same for men and women on the Raven test), we see men and women engaging equally in this behavior.

Men choose to enter a tournament much more frequently than women in the *Deter* treatment, although we find no difference when controlling for confidence. There is no gender difference in entry rates or stated confidence in the *Lure* treatment.

There are a number of directions for future research. A prominent one concerns the degree in which individuals are aware of the strategic implications of their signaling. Is some of the observed behavior truly unconscious? To what extent is self-deception present? We are only able to measure statements of confidence. Do those people who make overconfident statements (or underconfident statements in the lure treatment) actually believe these statements? At least in the cases of the baseline treatment or the receiver role in the other treatments, there is no cognitive reason to misrepresent beliefs, so one might claim that people believe their reports. We suspect that senders in the *Deter* treatment also believe their statements, in the same manner as the receivers. On the other hand, we speculate that senders in the *Lure* treatment, where we have argued that cognitive resources are engaged, do not believe their own statements of confidence. One approach to address this more directly would be to elicit two reports from each sender, with the difference that only one is sent to the receiver.[20] If people use their reported strategically and consciously, the report that is sent should vary with the treatment whereas the other report should not. A potential concern could be that any belief gap might solely result from the dual elicitation.

References

Barber, B. M., and T. Odean, T. (2001), "Boys will be boys: Gender, overconfidence, and common stock investment, " *Quarterly Journal of Economics* 116(1), 261-292.

Baumeister, R. F. (1998), "The Self," in D. Gilbert, S. Fiske and G. Lindzey (Eds.), *The Handbook of Social Psychology*, Boston: McGraw-Hill.

Benoit, Jean-Pierre and Juan Dubra (2011), "Apparent Overconfidence," *Econometrica*, 79(5) 1591-1625.

Berglas, S. and E. Jones (1978), Drug choice as a self-handicapping strategy in response to non-contingent success, *Journal of Personality and Social Psychology* 36, 405-417.

Burks, S., J. Carpenter, L. Goette, and A. Rustichini (2009), "Cognitive abilities explain economic preferences, strategic behavior and job performance," *Proceedings of the National Academy of Sciences* 106, 19, 7745-7750.

Burks, S., J. Carpenter, L. Goette, and A. Rustichini (2013), "Overconfidence and Social Signaling," *The Review of Economic Studies*, 80(3), 949-983.

Camerer, C. and D. Lovallo (1999), "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review* 89, 306-318.

[20] We thank a referee for suggesting this idea.

Charness, G. and U. Gneezy (2010), "Portfolio Choices and Risk Attitudes," *Economic Inquiry* 48, 133-146.

Charness, G. and U. Gneezy (2012), "Strong Evidence for Gender Differences in Risk-taking," *Journal of Economic Behavior and Organization* 83, 50-58.

Charness, G., A. Rustichini, and J. van de Ven (2015), "Self-confidence and Strategic Behavior," *Mimeo.*

Cho, I.K. and D. Kreps (1987), "Signaling Games and Stable Equilibria," *Quarterly Journal of Econonomics* 102, 179-221.

Compte, O. and A. Postlewaite (2004), "Confidence-Enhanced Performance," *American Economic Review* 94(5), 1536-1557.

Crawford, V. and N. Iriberri (2007), "Level-k Auctions: Can Boundedly Rational Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions," *Econometrica* 75, 1721-1770.

Dichgans, J., and T. Brandt (1978), "Visual-vestibular interaction: effects on self-motion perception and postural control," In: *Handbook of Sensory Physiology*, edited by R. Held, H. Leibowitz, and H-L. Teuber. Berlin: Springer, 755-804.

Dunning, D., J.A. Meyerowitz, and A.D. Holzberg (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.

Ewers, M. (2012), "Information and Competition Entry," *IZA Discussion Paper* 6411.

Ewers, M. and F. Zimmerman (2014), "Image and Misreporting," *Mimeo.*

Gilboa, I. and D. Schmeidler (1995), "Case-based Decision Theory," *Quarterly Journal of Economics* 110(3), 605-639

Gneezy, U., M. Niederle, and A. Rustichini (2003), "Performance in Competitive Environments: Gender Differences," *Quarterly Journal of Economics* 118(3), 1049-1074.

Gneezy, U. and A. Rustichini (2004), "Gender and Competition at a Young Age," *American Economic Review* 94(2), 377-381.

Heger, S. and N. Papageorge (2015), “We Should Totally Open a Restaurant: How Optimism and Overconfidence Affect Beliefs,” *Mimeo.*

Hoelzl, E. and A. Rustichini (2005), “Overconfident: Do you put your money on it?,” *Economic Journal* 115(503), 305-318.

Jehiel, P. (2005), “Analogy-based expectation equilibrium,” *Journal of Economic theory* 123(2), 81-104.

Kartik, N. (2009), “Strategic Communication with Lying costs,” *Review of Economic Studies* 76, 1359-1395.

Kolditz, T. A., and R. M. Arkin (1982), “An impression management interpretation of the self-handicapping strategy,” *Journal of Personality and Social Psychology* 43, 492-502.

Koszegi, B. (2006), “Ego Utility, Overconfidence, and Task Choice,” *Journal of the European Economic Association* 4(4), 673-707.

Kruger, J. (1999), “Lake Wobegon Be Gone! The Below-Average Effect and the Egocentric Nature of Comparative Ability Judgements,” *Journal of Personality and Social Psychology* 77(2), 221-232.

Leary, M. (1999), “The social and psychological importance of self-esteem,” in R. M.Kowalski and M. R. Leary, eds., *The social psychology of emotional and behavioral problems: interfaces of social and clinical psychology*, 197-222.

Leary, M. R. Tambor, S. Terdal, and D. Downs (1995), “Self-esteem as an interpersonal monitor,” *Journal of Personality and Social Psychology* 68, 518-530.

Liberman, V., S. Samuels, and L. Ross (2004), “The name of the game: predictive power of reputations versus situational labels in determining prisoner’s dilemma game moves,” *Personality and social psychology bulletin* 30, 1175-85.

Ludwig, S. and C. Thoma (2012). “Do Women Have More Shame than Men: An Experiment,” unpublished manuscript.

Lundeberg, M., P. Fox, and J. Puncochar (1994). “Highly confident but wrong: gender differences and similarities in confidence judgments, ” *Journal of educational psychology*, 86(1), 114-121.

Moore, D. and D. Cain (2007), "Overconfidence & Underconfidence: When and Why People Underestimate (and Overestimate) the Competition," *Organizational Behavior and Human Decision Processes* 103(2), 197-213.

Moore, D. and P. J. Healy (2008), "The Trouble with Overconfidence," *Psychological Review* 115(2), 502–517.

Myerson, R. (1991), "*Game Theory: Analysis of Conflict*," Harvard University Press: Cambridge, MA.

Niederle, M., and L. Vesterlund (2007), "Do Women Shy Away from Competition? Do Men Compete Too Much?," *Quarterly Journal of Economics* 122(3), 1067-1101.

Rabin, M. and J. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics* 114, 37-82.

Raven, J.C. (2000), *Ravens Advanced Progressive Matrices* (APM), Pearson.

Samuelson, L. (2001), "Analogies, Adaptation, and Anomalies," *Journal of Economic Theory* 97(2), 320-366.

Shipman, C. and K. Kay (2009), "Womenomics,"Harper Collins: New York.

Svenson, O. (1981), "Are we all less risky and more skillful than our fellow drivers?," *Acta Psychologica* 47, 143-148.

Trivers R. (1985), *Social Evolution*, Benjamin Cummings. Menlo Park, CA.

Van den Steen, E. (2004), "Rational Overoptimism (and Other Biases)," *American Economic Review* 94(4), 1141-1151.

Vialle, I., L. Santos-Pinto, and J.L. Rullière (2011), "Self-Confidence and Teamwork: An Experimental Test, "Mimeo.

Weinstein, N. D. (1980), "Unrealistic optimism about future life events," *Journal of Personality and Social Psychology* 39(5), 806-820.

Wolosin, R. J., S. Sherman, and A. Till (1973), "Effects of Cooperation and Competition on Responsibility Attribution After Success and Failure," *Journal of Experimental Social Psychology* 9, 220-23

Zábojník, J. (2004), "A Model of Rational Bias in Self-Assessments," *Economic Theory* 23(2), 259-282.

Supplementary Materials (for online publication only)

Appendix A. Model

In this appendix we characterize the equilibrium set of the signaling game described in section 3.3. We characterize the set of equilibria, and prove the following properties of equilibrium behavior in the games used in our experiment: (i) In the *Deter* treatment, where the sender has a relatively high outside option, he will over-report to appear strong and deter the receiver from entering the tournament, (ii) In the *Lure* treatment, where the sender has a relatively low outside option, he will underreport to appear weak and encourage the receiver to enter the tournament.

We start with repeating the payoff structure and introducing some notation and definitions. We then analyze the equilibrium set when the payoffs are symmetric, as in the *Deter* treatment. After establishing in section A.2 the conditions under which any type of equilibrium may occur (pooling, separating, partial separating), we summarize in section A.3 the entire characterization of the equilibrium set (see Theorem A.9). In section A.4 we compute the predictions of the model for the payoff used in the experiment. We first (section A.4) use the parameters adopted in the deterrence treatment of our experiment to show that over-reporting is part of the equilibrium in that treatment. We then (section A.4) do the same for the *Lure* treatment: we extend the analysis to asymmetric payoffs and then show that underreporting is equilibrium behavior with the parameters of the *Lure* treatment. The final section A.5 presents the intuitive reason for the existence of the deterrence and lure equilibrium. A reader who does not want to follow the computational detail can get a good idea of the argument from the two sections A.3 and A.5.

A.1. **Preliminaries.** Recall that θ^i_j denotes player $i \in \{1,2\}$ of type $j \in \{0,1\}$. The sender is indexed as player 1 and the receiver as player 2. A weak player is indexed as 0, a strong player as 1. We assume that the payoffs of players are symmetric as in the following payoff table:

(2)

	θ^2_0	θ^2_1
θ^1_0	$0,0$	b,a
θ^1_1	a,b	d,d

This is the case for the *Deter* treatment. We will return later to the *Lure* treatment which has asymmetric payoffs. We consider:

$$a \geq O^2 \geq d \geq 0 \geq b, \tag{3}$$

where O^2 is player 2's outside option. This implies that Out strictly dominates In for a weak player 2. Strategies of player 1 are functions from type to probability on signals: $\sigma^1(\theta^1;\cdot) \in \Delta(T)$. The cost of lying from not reporting truthfully (i.e., when $t_j \neq \theta_j$) are given by c. Strategies of player 2 are functions from type and signal of player 1 to probability on actions: $\sigma^2(\theta^2,t;\cdot) \in \Delta(\{In, Out\}))$. To lighten notation we call in the following $\sigma(\theta^2_1,t_0;In) = r, \sigma(\theta^2_1,t_1;In) = s$, where t_0 and t_1 are the low and high message respectively.

We now make precise what we mean by over- and underreporting.

Definition A.1. ***(Over- and underreporting)*** *We call an equilibrium in our game underreporting if:*

$$\sigma^1(\theta^1_0;t_0) = 1; \sigma^1(\theta^1_1;t_0) \equiv \tau \in (0,1). \tag{4}$$

that is if the low type only reports a low type, and the high type reports a low type with positive probability. We call an equilibrium over-reporting equilibrium if:

$$\sigma^1(\theta_1^1; t_1) = 1; \sigma^1(\theta_0^1; t_1) \equiv \sigma \in (0, 1). \tag{5}$$

A.2. **Types of equilibria.**

A.2.1. *Monotonic equilibria.* We first examine monotonic equilibria, i.e., those where the function $\theta^1 \to \sigma^1$ is increasing (higher types give higher signal). We then show that non-monotonic equilibria do not exist. In our simple model an equilibrium is monotonic if:

$$\sigma^1(\theta_1^1; t_1) \geq \sigma^1(\theta_0^1; t_1) \tag{6}$$

and we say it is strictly monotonic if the inequality (6) is strict.

The equilibrium set is easily characterized if we take into account the following. Take the $(\sigma, \tau) \in [0, 1]^2$ pairs describing as in (4) and (5) the strategy of player 1.

Lemma A.2. *If (3) holds, then for generic payoffs the only monotonic equilibria are either the fully revealing truthful, or the two pooling (at the low and high type respectively) or the under or over-reporting equilibria.*

Proof. Suppose that an equilibrium exists with $(\sigma, \tau) \in (0, 1)^2$. This implies that the two signals are indifferent for both types of player 1. This is equivalent to

$$(s - r)(b - O^1) = (r - s)(d - O^1) = 2c$$

which can only hold if $s \neq r$ and which in turn implies

$$O^1 = (b + d)/2. \tag{7}$$

So except for non-generic cases in which the equality (7) holds, there is no equilibrium where $(\sigma, \tau) \in (0, 1)^2$, so equilibria are on the boundary of the unit square. Monotonicity requires $1 - \tau \geq \sigma$, which excludes the strategies with $\sigma + \tau > 1$ (such as the "reverse fully revealing" equilibrium $(\sigma, \tau) = (1, 1)$). So we have either one of the three residual corners of the square, or a point in the two sides $\{0\} \times (0, 1)$ (under-reporting equilibria) and $(0, 1) \times \{0\}$ (over-reporting equilibria). □

A.2.2. *Fully pooling and fully separating equilibria.*

Lemma A.3. *If (3) holds a fully revealing equilibrium exists if and only if:*

$$2c \geq \max\{O^1 - b, d - O^1\} \tag{8}$$

Proof. Suppose that player 1 is following the truthful strategy. At the best response of player 2, he chooses Out when he is low type for any signal, and when of high type chooses In if and only if the signal is t_0, because:

$$m(\theta_0^1 | t_0) = 1, m(\theta_0^1 | t_1) = 0,$$

where $m(\cdot, t)$ denotes the posterior belief upon observing message t. So the best response of 2 to the truthful strategy of player 1 has $r = 1, s = 0$. To determine the set of parameters for which the fully revealing strategy of player 1 is part of an equilibrium we determine now when this strategy is a best response to $r = 1, s = 0$. With this strategy of 2, type θ_0^1 prefers t_0 to t_1 if and only if $2c \geq O^1 - b$; and type θ_1^1 prefers t_1 to t_0 if and only if $2c \geq d - O^1$. These conditions are equivalent to (8). □

As intuitively clear this equilibrium exists for all costs large enough. For the high signal pooling we have (if we ignore the case $(a + d)/2 = O^2$):

Lemma A.4. *If (3) holds a pooling equilibrium at the high signal exists if and only if either:*

$$(a+d)/2 < O^2 \textit{ and } O^1 - b \geq \max\{2c, d - O^1\} \tag{9}$$

or:

$$(a+d)/2 > O^2 \textit{ and } b - O^1 \geq \max\{2c, O^1 - d\} \tag{10}$$

Proof. Take a pooling equilibrium at the high signal. The best response of player 2 is determined by r and s given the posterior belief at t. Note that $m(\theta_0^1|t_1) = 1/2$, and we can take $m(\theta_0^1|t_0) \in [0,1]$ since the event t_0 has probability zero in this equilibrium. At the best response given these posteriors,

$$r \in [0,1] \text{ and } s \in sign(\frac{1}{2} - \frac{O^2 - d}{a - d}), \tag{11}$$

(where the *sign* correspondence is 1, 0 at positive and negative values, and the unit interval at 0.)

Consider now the best response of player 1 to such pairs (r, s). The condition that t_1 is preferred to t_0 by both types of player 1 is equivalent to

$$(s - r)(b - O^1) \geq 2c \geq (r - s)(d - O^1). \tag{12}$$

Thus, a pooling equilibrium at high type exists if and only if there is a pair (r, s) that satisfies both (11) and (12). We consider the two cases.

(1) If $(a + d)/2 < O^2$ then $s = 0$ and (12) is now equivalent to $r(O^1 - b) \geq 2c \geq r(d - O^1)$ which is equivalent to (9);
(2) If $(a + d)/2 > O^2$ then $s = 1$ and (12) is now equivalent to $(1 - r)(b - O^1) \geq 2c \geq (1 - r)(O^1 - d)$ which is equivalent to (10).

□

For the low signal pooling we have:

Lemma A.5. *If (3) holds a pooling equilibrium at the low signal exists if and only if either:*

$$O^2 \leq (a+d)/2 \textit{ and } d - O^1 \geq \max\{2c, O^1 - b\} \tag{13}$$

or:

$$O^2 \geq (a+d)/2 \textit{ and } O^1 - d \geq \max\{2c, b - O^1\} \tag{14}$$

Proof. With r, s denoting as usual the probability that the high type Player 2 chooses In at t_0 and t_1 respectively, an equilibrium pooling on the low type exists if and only if with $r \in sign((a+d)/2 - O^2)$, $s \in [0,1]$ the inequality

$$(r - s)(d - O^1) \geq 2c \geq (s - r)(b - O^1)$$

holds. These hold if and only if (13) or (14) holds. □

In the following we can then focus on the under and over reporting equilibria; remember that we have excluded by definition the fully pooling or fully separating equilibria from this set.

A.2.3. *Under-reporting equilibria.*

Lemma A.6. *An under-reporting equilibrium exists if and only if the inequalities in 3 and*

$$d - O^1 \geq \max\{2c, O^1 - b\} \tag{15}$$

for player 1 and

(16) $$O^2 \geq (a+d)/2$$

for player 2 hold.

Proof. We already know that at all equilibria, player 2 chooses Out at θ_0^2 irrespective of the signal. We check whether an equilibrium exists with $\tau \in (0,1]$. At θ_1^2 with such a strategy of player 1, player 2 chooses In at t in a best response if and only if the posterior $m(\theta_0^1|t) \geq \frac{O^2-d}{a-d}$. This is never the case if the signal is t_1, and it holds at t_0 when $\frac{1}{1+\tau} \geq \frac{O^2-d}{a-d}$. So the strategy of player 2 has only one indeterminate value $r \equiv \sigma(\theta_1^2, t_0; In)$ and we know that $r \in sign(\frac{1}{1+\tau} - \frac{O^2-d}{a-d})$.

Our last step is to check when the best response of player 1 to such strategy has the under-reporting form, with $\tau \in (0,1]$. Since as we have seen player 2 exits at t_1, choosing t_1 gives $O^1 - c$ to the type θ_0^1 and gives O^1 to the type θ_1^1. For a given r, t_0 is better than t_1 for type θ_0^1 if $O^1 - c < (1-(r/2))O^1 + (r/2)b$, and t^1 is indifferent to t_0 for type θ_1^1 if $O^1 = (1-(r/2))O^1 + (r/2)d - c$. These two conditions are satisfied if for some $r \in (0,1)$:

$$r(d - O^1) = 2c > r(O^1 - b)$$

which is equivalent to the additional condition (15). $\square$

The equilibrium is based on the fact that player 1 may be willing to pay the cost of signaling t_0 to lure player 2 in the tournament to get the payoff d; the gain is $(r/2)(d - O^1)$ and is equal to the cost c. Player 2 at θ_1^2 may choose In at the low signal t_0 because may get the high payoff a from θ_0^1 or the lower payoff d from θ_1^1, but overall this is the same as the Out payoff O^2.

A.2.4. *Over-reporting equilibria.* A similar analysis yields:

Lemma A.7. *An over-reporting equilibrium exists if and only if the inequalities in 3 and*

(17) $$O^1 - b \geq \max\{2c, d - O^1\}$$

for player 1 and

(18) $$O^2 \leq (a+d)/2$$

for player 2 hold.

Proof. In this case at equilibrium $\sigma^1(\theta_1^1; t_1) = 1$, and $\sigma^1(\theta_0^1; t_1) \equiv \sigma$. The posterior beliefs are $m(\theta_0^1|t_0) = 1$ and $m(\theta_0^1|t_1) = \frac{\sigma}{1+\sigma}$. Player 2 weakly prefers In to Out if and only if $\frac{\sigma}{1+\sigma} \geq \frac{O^2-d}{a-d}$, which is equivalent to $\sigma = \frac{O^2-d}{a-O^2} = 1/\tau$. His strategy is to choose Out at θ_0^2 and $\sigma^2(\theta_1^2, t_0; In) = 1; \sigma^2(\theta_1^2, t_1; In) = s$. The variables determining the equilibrium are s and σ; s is constrained to:

$$s \in sign(\frac{\sigma}{1+\sigma} - \frac{O^2-d}{a-O^2}),$$

and σ is determined by the best response of player 1. He prefers t_1 to t_0 at θ_1^1 and is indifferent between t_1 and t_0 at θ_0^1 if for some $s \in (0,1)$, $2c > (1-s)(d-O^1)$ and $2c = (1-s)(O^1 - b)$. Together these conditions are equivalent to (17) above. $\square$

To complete the full characterization of the equilibrium set, we examine non-monotonic equilibria.

Lemma A.8. *Non-monotonic equilibria do not exist.*

Proof. An equilibrium is non-monotonic only if $\sigma(\theta_1^1; t_1) < \sigma(\theta_0^1; t_1)$, i.e., if $1 - \tau < \sigma$. It is easy to see that the fully revealing non-monotonic equilibrium (both players 1 are always dishonest) does not exist. In that case, $\tau = \sigma = 1$. In that case, a strong player 2 chooses Out after t_0 and In after t_1. It is easy to verify that player 1 does not deviate at θ_0^1 if $O^1 \leq b - 2c$, and does not deviate at θ_1^1 if $O^1 \geq d + 2c$. This requires $d + 2c \leq b - 2c$ which is incompatible with $b \leq d$ for any $c > 0$.

Consider next the case with $\tau \in (0, 1)$ and $\sigma = 1$. A strong player 2 chooses Out following t_0 as this can now only come from a strong player 1. If player 1 is indifferent between t_0 and t_1 after θ_1^1, and prefers t_1 after θ_0^1, we must have:

$$s(b - O^1) \geq 2c = s(O^1 - d),$$

but these can never be simultaneously satisfied for $d \geq b$ and $c > 0$. Intuitively, indifference by a strong player 1 requires that he is willing to incur costs c to encourage a strong player 2 to choose Out. This requires that he values the outside option more than competing with a strong player 2, so $O^1 > d$. But then a weak player 1 must surely also like player 2 to choose Out, since for him the payoff from competing with a strong player 2 is even worse: $b < d$, and would then deviate to t_0.

Consider next the case with $\tau \in 1$ and $\sigma \in (0, 1)$. A strong player 2 chooses In following t_1 as this can now only come from a weak player 1. If player 1 is indifferent after θ_0^1 and prefers t_0 after θ_1^1, we must have:

$$r(O^1 - d) \geq 2c = r(b - O^1),$$

but this can never be satisfied for $d \geq b$ and $c > 0$. Intuitively, if even a weak player 1 is willing to pay a cost to encourage a strong player 2 to choose In, then a strong player 1 also prefers the strong player 2 to choose In and would deviate to t_1. □

A.3. **Summary.** We can summarize the previous sections characterizing the equilibria. A verbal description of the equilibrium set may be helpful. When the cost of lying is high compared to the other payoffs, the only equilibrium is the fully revealing. With smaller lying cost, the equilibrium correspondence separates into two branches, under and over reporting respectively, depending on whether $(d + b)/2 \geq O^1$ (in which case we have the under-reporting branch, see equation (20)) or the opposite holds. The behavior in these two branches is similar. In the under-reporting, when $O^2 \geq (a + d)/2$ we have a truly under-reporting equilibrium, where the high type reports the low type with positive probability. When O^2 is smaller, the equilibrium becomes a pooling equilibrium. The behavior of the over-reporting branch is similar. Figure 10 describes the type of equilibrium for different values of the parameters.

FIGURE 10. **Equilibrium set and parameter values.** The figure describes the type of equilibrium for values of the outside option O^1 and the lying cost c. The left panel reports the case where $O^2 < (a + d)/2$; the right panel reports the case where $O^2 > (a + d)/2$.

[INSERT FIGURE 10 HERE]

The following theorem gives a complete characterization of the equilibrium.

Theorem A.9. *For generic payoffs, in the interesting case $a \geq O^2 \geq d \geq 0 \geq b$,*

(a) ***(Fully revealing truthful branch)***: *There is a fully revealing equilibrium if*

(19) $$2c \geq \max\{O^1 - b, d - O^1\}$$

(b) ***(Under-reporting branch)***: *There is an under-reporting equilibrium if and only if:*

(20) $$d - O^1 \geq \max\{2c, O^1 - b\}$$

(21) $$O^2 \geq (a+d)/2$$

with strategies $\sigma^1(\theta_0^1; t_0) = 1, \sigma^1(\theta_1^1; t_0) = \tau = \frac{a-O^2}{a-d}$ *and*

$$\sigma^2(\theta_1^2, t_0; In) = r = \frac{2c}{d - O^1}, \sigma^2(\theta_1^2, t_1; In) = s = 0.$$

There is a pooling equilibrium at the low signal if (i) condition (20) holds and $O^2 \leq (a+d)/2$, *sustained by* $\sigma^2(\theta_1^2, t_0; In) = 1, \sigma^2(\theta_1^2, t_1; In) = 0$, *or (ii) condition (21) holds and* $O^1 - d \geq \max\{2c, b - O^1\}$ *sustained by* $\sigma^2(\theta_1^2, t_0; In) = 0, \sigma^2(\theta_1^2, t_1; In) = 1$.

(c) ***(Over-reporting branch)***: *There is an over-reporting equilibrium if and only if:*

(22) $$O^1 - b \geq \max\{2c, d - O^1\}$$

and

(23) $$O^2 \leq (a+d)/2$$

with strategies $\sigma^1(\theta_0^1, t_1) = \sigma = \frac{O^2-d}{a-d}, \sigma^1(\theta_1^1, t_0) = 0$ *and*

$$\sigma^2(\theta_1^2, t_0; In) = r = 1; \sigma^2(\theta_1^2, t_1; In) = s = 1 - \frac{2c}{O^1 - b}.$$

There is a pooling equilibrium at the high signal if (i) condition (22) holds and $O^2 \geq (a+d)/2$, *sustained by* $\sigma^2(\theta_1^2, t_0; In) = 1, \sigma^2(\theta_1^2, t_1; In) = 0$, *or (ii) condition (23) holds and* $b - O^1 \geq \max\{2c, O^1 - d\}$ *sustained by* $\sigma^2(\theta_1^2, t_0; In) = 0, \sigma^2(\theta_1^2, t_1; In) = 1$.

Remark A.10. *Note that some of the pooling equilibria are counterintuitive: they are sustained by the belief that a strong player 2 chooses Out after the low signal but In after the high signal. While these types of pooling equilibria are Perfect Bayesian Equilibria, it is easy to show that they do not survive equilibrium refinements such as "D1" (Cho and Kreps, 1987).*

A.4. **The experimental design.** We finally consider a game which is closer to the one used in our experimental sessions.

The strategic Deter *treatment.* In the strategic *Deter* treatment payoffs were symmetric, so the analysis of above applies. Specifically, if we subtract 5 from all payoffs, and make use of the fact that ties are broken by random assignment of the *win* outcome, the payoffs were:

(24)

	θ_0^2	θ_1^2
θ_0^1	$0, 0$	$-5, 5$
θ_1^1	$5, -5$	$0, 0$

and outside options $O^1 = 5$ for player 1 and $O^2 = 0.5$ for player 2 in the treatment with the high outside option. This is the case $O^1 - b = 10 \geq \max\{2c, d - O^1\} = 2c$ because $d - O^1 = -5$, and

$(a+d)/2 = 2.5 > O^2 = 0.5$, hence (provided costs c are sufficiently low) we are in the over reporting branch (see theorem A.9).[21]

The Lure *treatment.* In the *Lure* treatment the payoffs from entering the tournament were not symmetric, and we represent them as follows:

$$\begin{array}{c|cc} & \theta^2_0 & \theta^2_1 \\ \hline \theta^1_0 & e, d & g, f \\ \theta^1_1 & a, b & e, d \end{array} \tag{25}$$

where $a = 25, b = 0, d = 5, e = 20, f = 10, g = 15$, and $O^1 = 10, O^2 = 5.5$ (again using the fact that ties are broken by random assignment of the *win* outcome). Note that

$$\text{for all } \theta, O^1 < v^1(\theta); f > O^2 > d > b \tag{26}$$

We look for equilibria of the under-reporting form, as in lemma (A.6).

Lemma A.11. *An under-reporting equilibrium of the game with tournament payoffs (25) exists if (26) and*

$$e - O^1 \geq 2c \tag{27}$$

hold.

The proof is a simple computation.

A.5. ***Deter*****rence and luring equilibria.** The intuition for the strategic deterrence equilibrium is clear, and has been presented in the section A.3. For the lure equilibrium, note that player 2 of the low type prefers *Out* to *In* irrespective of what outcome he is expecting in the tournament. The high type will choose *In* if he gives enough weight to the event that he is facing a low type. At a monotonic under-reporting equilibrium the high signal t_1 reveals that the type is high, so player 1 of low type will not incur the cost of lying when by doing so he could only tempt player 2 to choose *Out*. The high type player 1 may be made indifferent between telling the truth and thus forcing player 2 out, or luring him by stating the low type, paying the cost, and getting with enough probability to reap the benefit of a match with a high type player 2. He can be made indifferent between these two options when the extra gain from luring the receiver (the quantity $r(e - O^1)$) can be made at least equal to the cost $2c$, for some r the probability that the high type player 2 plays *In* after a low signal. This is what condition (27) insures.

[21]In the *Deter* treatment with the low outside option for player 2 the corresponding payoff was $O^2 = -1.5$. Note that in this case $a \geq d \geq 0 \geq O^2 \geq b$ so that condition (3) is not satisfied. It is easy to show, however, that the results are qualitatively similar in this case. With the payoffs in the experiment, we would be in the case with pooling at the high signal, a limit case of overreporting.

Appendix B. Instructions

The comments in square brackets are meant to illustrate instructions to the reader and were not part of the instructions.

General instructions

Introduction Welcome to our experiment. You will receive €7 for showing up, regardless of the results. The instructions are simple. If you follow them carefully, you can earn a substantial amount of money in addition to your show up fee. Throughout the stages we will ask you to answer questions. At each stage, you will receive more detailed instructions.

You will be part of a group of 4 persons. You don't know who the other persons are, and you will remain anonymous to them. All your choices and the amount you will earn will remain confidential and anonymous, except if explicitly indicated otherwise. You will receive your earnings in an envelope. The person that puts the money in the envelopes can only see the login number that has randomly been assigned to you, and cannot match any names, student numbers, or faces with the login numbers and the decisions made.

Payments There are several items in the experiment for which you can earn points. At the end of the experiment, one item is randomly chosen and your points for that item are paid in addition to the show-up fee (1 point is worth € 1). One of the participants is randomly chosen to be an assistant during the experiment. There is a random component in the experiment. The task of the assisting person will be to throw a dice which will determine the outcome.

No deception Remember, we have a strict no deception policy in this lab.

Questions Please remain seated and raise your hand if you have any questions, and wait for the experimenter. Please remain silent throughout the experiment.

Part 1.
In the first stage, all group members receive the same 15 questions. You will see a matrix with one missing segment at the bottom right. Your task is to identify the segment that would logically fit at the position of the missing segment, by choosing from the suggested answers. You can make your choice by clicking the corresponding number on the right of your screen. [A screen shot with an example question was provided.]

You can go back and forth between the questions. There is a time limit of 8 minutes. The time remaining is indicated on your screen.

After the time limit, we will rank all 4 people in your group depending on the number of questions answered correctly. The person with the highest score will get rank 1, and the person with the lowest score will get rank 4. In case of ties, the computer will randomly determine who gets the higher rank. After this, you will get some questions regarding how well you think you did.

We then randomly divide the group in 2 players A and 2 players B. Every player A will be matched against a player B. If your rank is higher than the player with which you are matched, you can receive 10 points.

Part 2.
All four group members have now finished with the questions, and we have determined the rank of every person.

We now ask you to indicate how likely you think that you are among the top 2 of your group. You can indicate this on a scale from 0 to 100%. Indicating 0% means that you are sure you are not among the best 2 of your group, while indicating 100% means that you are sure you are among the top 2 of your group. Similarly, 50% indicates that you think it is equally likely that you are among the best 2 of your group, or that you are not among the best 2 of your group.

We will pay you for the accuracy of your estimate. You earn more points for this item if your estimate is more accurate. The formula that is used to calculate the amount of money you earn is chosen in such a way that your expected earnings are highest when you report to us what you really believe. Reporting any value that differs from what you believe decreases your expected score for this item. If you are interested, you can find some detailed examples of this to see how this works. [An explanation with examples was available to participants, see below.]

The role of player A and player B We matched you with one other randomly chosen person from your group. You are either Player A or B, and this is randomly determined.

[*baseline*] None of the players can see the other player's estimate of being in the top 2.

[*social*] Player A will not see the estimate by player B that he or she is among the best two in the group, but player B will see the estimate by Player A that he or she is among the best two of the group.

[*baseline and social*] Later on in the experiment, we will compare the rank of player A with the rank of player B, and for that item the player with the highest rank receives 10 points, the other nothing. Both of you will see who has the highest rank, and this ends the stage.

[*strategic deterrence and lure*] Player A will not see the estimate by player B that he or she is among the best two in the group, but player B will see the estimate by Player A that he or she is among the best two of the group.

Later on in the experiment, after player B has observed the estimate of player A, player B will choose between two options: IN or OUT.

If player B chooses OUT, then for that item player B receives 3.5 [5.5] points and player A automatically receives 10 points. Both players will see who has the highest rank, and this ends the stage.

[*deterrence*] If player B chooses IN, we will compare the rank of player A with the rank of player B, and for that item the player with the highest rank receives 10 points, the other nothing. Both of you will see who has the highest rank, and this ends the stage.

[*lure*] If player B chooses IN, we will compare the rank of player A with the rank of player B. For this item, Player A receives 25 points if (s)he is the highest ranked player, and 15 points if (s)he is not the highest ranked player. Player B receives 10 points if (s)he is the highest ranked player, and 0 points if (s)he is not the highest ranked player. Both of you will see who has the highest rank, and this ends the stage.

You can see what role you have on the top left of your screen (see the example below). [Participants could see their role on the next screen.]

Determination of your score What follows is a brief explanation about the determination of your score, showing that it is in your interest to report truthfully what you believe in order to maximize your expected earnings.

The score is determined as follows. You start with 10 points. We subtract points depending on how close your reported belief is to the outcome. The outcome is set to 1 if you are in the top 2, and to 0 if you are not.

For instance, if you report 70% (.7), and you are in the top 2 (outcome is 1), you are .3 away from the outcome, while if you are not in the top 2 (outcome is 0), you are .7 away from the outcome.

The difference with the outcome is squared and multiplied by 10, and then subtracted from the 10 points that you start with. Thus in the example with 70%: if you are in the top 2, this gives you $10 - 10(.3)^2 = 9.1$. If you are not in the top 2, this gives $10 - 10(.7)^2 = 5.1$. You would weight these two scores by your belief about the likelihood of each occurring.

Larger differences between your reports and the outcome decrease your score proportionally more than small differences. To minimize the expected difference, and maximize your expected score, you should report what you believe.

The following examples illustrate that your expected score is highest when you report your true beliefs. *All numbers used are for illustrations only and are no indication for the decisions for you to take.*

Example 1
You believe 50% and report 50%. As a simple example: if you believe there is a 50% chance you are in the top 2, and you report 50%, then there is always a difference of .5 with the outcome, and since this is squared we always subtract 10 times $(0.5)^2$ points from your score, i.e. 2.5 points. Your expected score is 7.5.

You believe 50% but you report 100%. If you report 100%, then in one case there is no difference (if you are in the top 2) and no points are subtracted. But in the other case the difference is 1 (if you are not in the top 2), and then we subtract 10 times $(1)^2$ from your score. If you believe the likelihood of being in the top 2 is 50%, you expect this to happen in 50% of the cases, so the amount subtracted would be $10(0.5) = 5$. This gives you an expected score of 5, which is lower than if you report your belief of 50%.

Example 2
You believe 70% and report 70%. As another example, suppose that you think there is a 70%

likelihood that you are among the best 2. If you report 70%, your score will be either 9.1 (if you are in the top 2) or 5.1 (if you are not in the top 2). You believe that with 70% chance your score will be 9.1, and with 30% your score will be 5.1. So your expected score is $0.7(9.1) + 0.3(5.1) = 7.9$.

You believe 70% and report 100%. Now suppose that, instead of reporting this belief of 70%, you report another number. For instance, you report 100% (1). This means that if you are in the top 2, the outcome is as predicted, and you get $10 - 10(0)^2 = 10$ points. If you're not in the top 2, you are 1 away from the outcome, and your score will be $10 - 10(1)^2 = 0$. Since you actually expect to be in the top 2 with 70% chance, your expected score is 7. This is lower than if you would have reported 70%.

You believe 70% and report 20%. The same is true if you report a number below your belief, for instance 20% (.2). If you are in the top 2, your score would be $10 - 10(0.8)^2 = 3.6$ points. If you're not in the top 2, your score will be $10 - 10(0.2)^2 = 9.6$. Since you actually expect to be in the top 2 with 70% chance, your expected score is $0.7(3.6) + 0.3(9.6) = 5.4$, again lower than if you would have reported 70%.

The table below shows the expected scores for some more possible beliefs you may have and reports you give. As you can see, expected scores are highest when the reported belief is equal to the true belief (the cells on the diagonal that are highlighted in green).

[INSERT FIGURE TABLE INSTRUCTIONS HERE]

[*baseline*] **Part 3.**
Based on your true ranking in the group, we will send you a report. The report will say if you are among the two best of your group, or if you are not among the two best of your group.

However, *sometimes the report will be incorrect.* The way this works is as follows.

If you *are not* among the top two of your group, then the report will always be correct and inform you that you are not among the best two of your group.

If you *are* among the top two of your group, the report is mistaken in half of the cases. That is, in half of the cases, the report correctly informs you that you are among the top two of your group. In the other half of the cases, the report is incorrect and says you were not among the top two of your group, even if you were.

Whether or not the report you receive is correct when you are among the top two of your group, depends on the outcome of a dice throw by the assistant. You will not see the outcome, but if the assistant throws 1, 2, or 3, you will receive a correct report when you are among the top two. If the assistant throws 4, 5, or 6, you will receive an incorrect report when you are among the best two of your group. (For some groups, the incorrect report is sent after different values of the dice, but in any case the report is incorrect in half of the cases when you are among the best 2.)

After you see the report, we will ask you if you think the report is more likely to be correct or incorrect.

You earn 10 points if you are right.

[*baseline*] **Part 4**
In this part, we ask you some questions about the scenario below. The first part is always the same, but some additional information is given in the question, so please read it carefully. For this part, we randomly choose a question and this is treated as a single item.

Scenario Consider two machines placed in two sides of a large production hall, left side = L and right side = R. The two machines produce rings, good ones and bad ones. Each ring that comes from the left machine, L, has a 50% chance of being a good ring and a 50% chance of being a bad ring. Each ring that comes from the right machine, R, is good. Both machines produce 100 rings every day.

The mechanic visits the production hall every day, and randomly examines one of the machines by taking one ring. On some days, he takes a ring from the left machine, and the other days he takes a ring from the right machine. Suppose the ring he takes is *good*.

We will ask you if it is more likely that the mechanic went to the left or right machine.

Example question 'On 50% of the days, the mechanic takes a ring from the left machine, and the other 50% of the days from the right machine. Of the rings that come from the left machine, on average half are good and half are bad. Each ring that comes from the right machine is good.

Imagine the ring he takes is good. Is it more likely to come from the left or right machine?'

You will get 3 questions like this one. We vary the percentage of days that the mechanic goes to the left or right machine, but everything else remains the same.

You earn 10 points if you are right.

Part 5.

[*baseline and social*] In this part, you are informed if player A or B has the highest rank.

[*strategic deterrence and lure*] Player A will not see the estimate by player B that he or she is among the best two in the group.

Player B will see the estimate by Player A that he or she is among the best two of the group, and then gets the choice between two options: IN or OUT.

[We repeated the instructions of Part 2 in which the payoffs were given for this item].

Appendix C. Background Characteristics

Table 3: *Summary of background characteristics*

	Mean	Std. Error	Min.	Max.
Background characteristics				
Age	21.96	0.14	17	49
Number of siblings	1.48	0.05	0	7
Gender (fraction females)	0.44			
Member of sports club	0.49			
Took Raven test before	0.54			
Familiar with condition probs.	0.61			
Study category				
Economics/Business/Finance	0.58			
Social Sciences and Law	0.15			
Physics, Math, Computer science	0.07			
Other study or not student	0.20			
N	464			

Appendix D. Survey

We recruited 82 participants from the same database that we used for the main experiment. None of the participants participated in any of the previous sessions. There were 5 sessions with between 12 and 22 participants. 40 out of 82 participants were female, and the mean age was 22. Participants received a flat fee of 10, and each session took about 30 minutes.

Participants received the instructions on their screen. The survey consisted of eight questions in total (see below). In the first four questions, we asked participants to list situations in which appearing over or underconfident could be effective. The purpose of these questions was to make them realize how easy/hard it was to think of any such situations. Our main questions of interest were questions five and six, in which we asked them how frequently they find themselves in situations where appearing under or overconfident can be effective.

We reversed the order of questions between participants: half of them received the questions about underconfidence first, the other half received the questions about overconfidence first.

They could answer questions 5-8 on a five point scale (ranging from very rarely to very frequently) and we gave them the option to answer the question with "don't know." The distribution of answers is reported in Table C1. The full instructions and questions are provided below.

Table C1: Distribution of responses (percentages)

Question	1	2	3	4	5	Don't know
Appearing overconfident is effective (Q5)	4.88	14.63	24.39	50.00	4.88	1.22
Appearing underconfident is effective (Q6)	21.95	31.71	24.39	17.07	3.66	1.22
Others appear overconfident (Q7)	1.22	14.63	23.17	45.12	9.76	6.10
Others appear underconfident (Q8)	12.20	30.49	21.95	23.17	7.32	4.88

Notes: Answers are on a scale from 1 (very rarely) to 5 (very frequently).

Instructions and questions Thank you for your participation in this short experiment. It takes about 20-30 minutes to complete, and you will earn 10 for your participation. Talking is not permitted and we ask you to make sure that your cell phone is turned off completely. Please raise your hand if you have any questions. The experiment consists of a short survey. You are not matched to any other participant and your earnings do not depend on the answers that you (or anyone else) give.

Please answer the questions to the best of your ability. Your answers will be treated confidentially.

The questions that follow are about situations in which you interact with one or more other persons.

[version for participants with even numbers]

Question 1) Can you briefly describe situations (real or hypothetical) in which you think it might be useful to appear to others more confident about your ability to succeed in an activity or at a task than you really are? You can list as many different situations as you like, up to a maximum of ten situations. You can, for instance, think of situations in the domains of sports, school, work, or social life, amongst others.

Question 2) For each of the situations you just described under Question 1, have you found it to be effective to appear more confident than you were? Or if the situation you described was hypothetical, do you think it would be effective?

Question 3) Can you briefly describe situations (real or hypothetical) in which you think it might be useful to appear to others less confident about your ability to succeed in an activity or at a task than you really are? You can list as many different situations as you like, up to a maximum of ten situations. You can, for instance, think of situations in the domains of sports, school, work, or social life, amongst others.

Question 4) For each of the situations you just described under Question 3, have you found it to be effective to appear less confident than you were? Or if the situation you described was hypothetical, do you think it would be effective? We will now ask you some questions about how frequently you find yourself in different types of situations. You can indicate your answer on a 5-point scale, ranging from 1 (very rarely) to 5 (very frequently). There is also an option not to give an answer in case you do not know the answer.

Question 5) How frequently do you find yourself in situations where it could be effective to appear to others more confident about your ability to succeed in an activity or at a task than you really are?

Question 6) How frequently do you find yourself in situations where it could be effective to appear to others less confident about your ability to succeed in an activity or at a task than you really are?

Question 7) How frequently do you find yourself in situations where someone else appeared to you to be more confident about his or her ability to succeed in an activity or at a task than (s)he should have been?

Question 8) How frequently do you find yourself in situations where someone else appeared to you to be less confident about his or her ability to succeed in an activity or at a task than (s)he should have been? What is your age? What is your gender (male or female)? What is your field of study?

Thank you for your participation in this experiment. Please remain seated until your table number is called. If your table number is called, please bring the card with your table number with you and you will receive your payment.

Appendix E. Optimality of Decisions

In this Appendix we describe the methodology underlying Figure 8. We assume that players have types, $\theta \in [0, 100]$, that are drawn from a continuous distribution function with density $f(\theta)$. We index players by $i = S, R$ (sender and receiver respectively). Players choose a message $t^i \in T = [0, 100]$, so that the message space is the same as the type space. The message of a player is his reported confidence, and the type is his true belief about his confidence. In our experiment a receiver has no incentives to report a confidence that differs from his type, so we assume $t^r = \theta^r$. After observing t^s, receivers choose an action in the set $\{Out, In\}$. The assumed strategy of the receiver is then to play In if and only if $t^r \geq t^s$.

Let O^S be the sender's outside option payoff if the receiver chooses Out. If the receiver chooses In, the sender's payoff is v^h if he wins and v^l if he loses. The probability that the sender wins is $\Pi(\theta^s, \theta^r)$. We will specify precise functional forms of Π. The expected payoff for the sender of the tournament is then given by:

$$\int_0^{t^s} O^S dF(\theta^r) + \int_{t^s}^{100} (\Pi(\theta^s, \theta^r)(v^h - v^l) + v^l) dF(\theta^r). \tag{28}$$

The reason for reporting an inaccurate confidence level is to change the probability that the receiver chooses In. The optimal reported confidence for a risk-neutral sender is determined by:

$$f(t^s)[O^S - \Pi(\theta^s, t^s)(v^h - v^l) - v^l] = c(t^s - \theta^s). \tag{29}$$

The RHS reflects the fact that players have the incentive provided by the quadratic scoring rule to report truthfully, creating costs when their reported confidence differs from their true belief (where $c = 2/10,000$). In the *Deter* treatment, the term in brackets on the LHS is positive so that over-reporting is optimal ($O^S = v^h = 10, v^l = 0$), while in the *Lure* treatment this term is negative so that underreporting is optimal ($O^S = 10, v^h = 25, v^l = 15$).

We specify

$$\Pi(\theta^s, \theta^r) = 1/(1 + e^{-\delta(\theta^s - \theta^r)}), \tag{30}$$

where $\delta = .021$ is estimated from the data of the baseline treatment in which there are no incentives to over-report. For $F(\theta)$, we assume that players believe that types are normally distributed (truncated at 0 and 100) with mean 50 and standard deviation 21. The value of the standard deviation is estimated from the data and we take a mean of 50 to reflect that players do not believe that other players are on average overconfident.