

Amelio, Andrea

Working Paper

Cognitive uncertainty and overconfidence

ECONtribute Discussion Paper, No. 173

Provided in Cooperation with:

Reinhard Selten Institute (RSI), University of Bonn and University of Cologne

Suggested Citation: Amelio, Andrea (2022) : Cognitive uncertainty and overconfidence, ECONtribute Discussion Paper, No. 173, University of Bonn and University of Cologne, Reinhard Selten Institute (RSI), Bonn and Cologne

This Version is available at:

<https://hdl.handle.net/10419/262335>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

ECONtribute
Discussion Paper No. 173

Cognitive Uncertainty and Overconfidence

Andrea Amelio

June 2022

www.econtribute.de



Cognitive Uncertainty and Overconfidence^{*}

Andrea Amelio[†]

June 20, 2022

Abstract

Overconfidence is one of the most ubiquitous cognitive bias. There is copious evidence of overconfidence being relevant in a diverse set of economic domains. In this paper, we relate the recent concept of cognitive uncertainty with overconfidence. Cognitive uncertainty represents a decision maker’s uncertainty about her action optimality. We present a simple model of overconfidence based on the concept of cognitive uncertainty. The model relates the concepts theoretically and generates testable predictions. We propose an experimental paradigm to cleanly identify such theoretical relationships. In particular, we focus on overplacement and we find that, as predicted, cognitive uncertainty is inversely related to overplacement. Exogenously manipulating cognitive uncertainty through compound choices, we are able to show a causal relationship with overplacement. Evidence on these relationships allows to link overplacement with other behavioral anomalies explained through cognitive uncertainty.

Keywords— *Cognitive Uncertainty, Overconfidence, Overplacement, Cognitive Noise, Experiments*

JEL Codes: D91, C91, D83

^{*}I am grateful to Dominik Damast, Armin Falk, Ximeng Fang, Lorenz Götte, Thomas Graeber, Luca Henkel, Martin Kornejew, Robin Musolff and Rafael Suchy for fruitful discussion and helpful comments. A special thanks to Florian Zimmermann for his supervision and guidance. Thanks to BGSE (Bonn Graduate School of Economics) and briq (Institute on Behavior & Inequality) for financial support. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC2126/1-390838866.

[†]University of Bonn: andrea.amelio@briq-institute.org

1 Introduction

Uncertainty undeniably plays a central role in the economics literature, as it permeates every aspect of economic decision-making, such as stock market investments, innovation decisions, consumption choices, and many more. However, most of the economics literature focuses on uncertainty stemming from the environment, that is *external uncertainty*. An example of decision making under this kind of uncertainty may be booking a holiday to Paris: Linda would like the weather to be sunny while she visits the city, so she considers factors that impact the chance of any given day to be rainy when picking the dates for the trip; hence, the uncertainty is generated by factors external to the decision-maker.

On the other hand, a recent and growing branch of literature started focusing on the uncertainty that is not originated by environmental conditions but from the cognitive processes involved in undertaking a decision. Woodford (2020) provides a review of key ideas from psychophysics, with a focus on economic applications of what is defined as *imprecision*. Khaw, Li, and Woodford (2017) and Gabaix (2019) both propose a theoretical framework where some form of cognitive noise is generated when a decision-maker undertakes any decision. The idea is that this noise is not rooted in an imperfectly observable environment but possibly caused by the complexity of the problem in the process of elaborating the inputs and providing an answer. Building on previous works, Enke and Graeber (2021) (EG henceforth) define the concept of *cognitive uncertainty* (CU henceforth) as an agent's uncertainty about own's action optimality: the decision-maker is aware of the existence of the cognitive noise, which impairs her ability to take the optimal decision, and is hence uncertain whether the action she picked is the optimal one. In other words, agents are aware that they may commit mistakes and they hold doubts about having done the right choice. Very importantly, EG employ this concept to unify several boundedly rational behaviors documented in the literature and provide experimental evidence of the role of cognitive uncertainty in moderating such behaviors. This paper aims to extend this process, establishing a link between cognitive uncertainty and overconfidence,¹ with a focus on overplacement,² both theoretically and empirically.

The relevance of overconfidence in economic decision-making is well established. Notable examples are Malmendier and Tate's (2005) paper, in which the authors show how overconfidence induces CEOs to undertake sub-optimal investment decisions, or Barber and Odean

¹Although it is possible to argue that both internal and external uncertainty play a role in giving rise to overconfidence, we focus on how internal uncertainty, through CU, contributes to the phenomenon.

²As Moore and Healy (2008) argue, what is commonly defined as overconfidence comprises different constructs, which is wiser to treat separately. In [Section 2.2](#) this point is laid out more extensively.

(2001), who show how overconfidence (related to gender) may lead to excess trading on the stock market, with a negative impact on returns. In the words of Kahneman (2011), overconfidence is “[...] *the most significant of the cognitive biases*”. A series of papers in economics, part of the literature on ego-based utility and motivated beliefs,³ investigates the structure and the causes of overconfidence. This literature identifies the cause of overconfidence in the fact that positive self-assessments increase agents’ utilities. Nonetheless, explanations of overconfidence based on motivated reasoning leave out some unsolved puzzles and relevant questions. For example, it is not clear how overconfidence is related to other cognitive biases, how overconfidence can persist over time in the presence of feedback⁴ or how overconfidence emerges in not ego-relevant contexts. This suggests that the mechanism behind overconfidence in economic decision-making is still unclear. We aim to, partially, shed light on these aspects, making use of the concept of CU, focusing on overplacement. Overplacement and CU are inversely related. For example, an individual who is highly uncertain about the optimality of her own action will tend to place herself relatively lower, with respect to a less cognitively uncertain individual. Crucially, a form of internal uncertainty is conceptually necessary to rationalize overconfidence-related phenomena: to make self-assessment mistakes, an agent must be uncertain about the optimality of her choice. We use the concept of cognitive uncertainty to justify and formalize this idea.

This paper brings about two key contributions. First, we show how overconfidence, and more extensively under/overplacement, is generated in a cognitive uncertainty framework. This is intended as a step to merge the overlapping parts of the economic literature on imprecision and overconfidence. Second, the model delivers a set of predictions about the impact of CU on two different overplacement measures. We test these predictions experimentally. Our results show that CU and overplacement are negatively related. Also, we manipulate CU experimentally using compound choices, showing the existence of a causal link between CU and overplacement. As a third, minor, contribution, we document a relationship between placement measures and the shape of probability weighting, which, to our knowledge, has not been explored before. The model presented in this paper, jointly with EG’s results, can account for this preliminary evidence.

We conduct an experiment built on the “balls-and-urns” workhorse paradigm (see Benjamin, 2019) to collect evidence of a causal relationship between overplacement and CU. Participants are introduced two fictitious urns and told that one of the two has been picked with some probability. Each urn contains a different number of blue and red balls and, after observing a draw of one or two balls, participants state their probability guess about each

³See Bénabou (2015) and Bénabou and Tirole (2016) for reviews.

⁴On this latter, also the role of memory has been studied, both theoretically (e.g. Bénabou and Tirole (2002)) and empirically (e.g. Huffman, Raymond, and Shvets (2021) or Zimmermann (2020)).

urn. Before formulating the guess we elicit an absolute placement measure, asking them a guess about their rank on a scale from 1 to 100. After the guess, they observe another participant’s answer to the same problem and provide a relative placement guess (probability of having performed better than the other participant). To identify a causal role of CU we follow EG, introducing ambiguity in half of the tasks, presenting the diagnosticity parameter of the problem (number of blue balls in each urn) as a random variable. We interpret this as an exogenous manipulation of CU.

We have two main findings. First, placement and overplacement decrease in CU, that is more cognitive uncertain participants tend to place themselves lower and are less likely to wrongly place themselves higher, relative to other participants. This finding is robust across different measures of placement and overplacement. Second, more cognitively uncertain participants are more likely to change their answers and to a greater extent. These findings are consistent with a formal model of overplacement built on EG’s model of CU. In the model, agents are not sure about the optimality of their actions and the level of uncertainty about their actions’ optimality, that is CU, regulates the extent to which they are under/overconfident.

Besides contributing to the literature on overconfidence and imprecision, this paper further contributes to the economic literature on observational learning, through the structure of our experimental paradigm. Weizsäcker’s (2010) metastudy on social learning shows how individuals fail to effectively learn from others when this would imply to contradict their own initial choice, even if it would be optimal to do so. This evidence can be seen as a form of underreaction to new signals, which is prevalent in social learning experimental contexts. Several other works document this and describe it as a form of overconfidence (e.g. Nöth and Weber (2003); Celen and Kariv (2004); Goeree, Rogers, Palfrey, and McKelvey (2007); and De Filippis, Guarino, Jehiel, and Kitagawa (2017)). On the other hand, the psychology literature offers several instances of *underconfidence* in diverse tasks (Burson, Larrick, and Klayman (2006); Kruger and Dunning (2009); Krueger and Oakes Mueller (2002); Moore and Small (2007)), with the mechanism regulating the presence of over or underconfidence not being clear. Cognitive uncertainty may provide this regulating mechanism, along with a theoretical foundation for that.

The remainder of the paper is organized as follows. [Section 2](#), establishes the theoretical link between CU and overconfidence frameworks, showing to what extent they are equivalent and which insights can the cognitive uncertainty perspective provide. The key predictions of the model are tested in an experimental setting with financially incentivized decisions. The experimental design is described in [Section 3](#) and the analyses and results in [Section 4](#). [Section 5](#) concludes.

2 Theoretical Framework

In this section, we first briefly introduce EG framework of CU. Afterwards, we show the link between CU and overplacement in a simple formal setting, which allows to formulate testable empirical hypotheses.

2.1 Cognitive Uncertainty

The model developed in this section builds on EG illustration of choice under CU, where the decision-maker behaves *as if* she was facing a signal extraction problem, with the noise being internally generated.

Consider an agent with quadratic utility function:

$$u(a, x) = -\frac{1}{2}(a - Bx)^2. \quad (1)$$

Clearly, the optimal action would then be $a^* = Bx$. However, the agent is affected by cognitive noise and behaves *as if* the state variable x was not observed deterministically, but only through a noisy signal $s = x + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. The noise term ε is the noise *perceived* by the agent and may also not correspond to the *true* cognitive noise, denoted by $\tilde{\varepsilon}$. Assuming the agent holds a prior $x \sim \mathcal{N}(x_0, \sigma_x^2)$ about the state, the optimal action would be $a^*(s) = B\lambda s + B(1 - \lambda)x_0$, with the agent's uncertainty about her own action optimality reflected by

$$a^*(x | s) \sim \mathcal{N}(B\lambda s + B(1 - \lambda)x_0, B^2(1 - \lambda)\sigma_x^2), \quad (2)$$

with $\lambda = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2}$. Hence, cognitive uncertainty (σ_{CU}) is defined as the standard deviation of the above probability distribution: the agent's uncertainty about her optimal action. Note that the normality assumption is imposed for simplicity, but the general idea would hold also for different prior and cognitive noise distributions: the agent's internal uncertainty induces a distribution on the space of possible answers, with a certain degree of dispersion that is her cognitive uncertainty.

The theoretical contribution of this paper, strictly related to the empirical investigation, is twofold: linking CU and overplacement and representing an observational learning process in this framework. In the [Appendix](#), we also briefly present MH's benchmark model of overconfidence, showing how a CU-based model of overconfidence can generate equivalent predictions and arguing how a CU-based model may provide additional insights.

2.2 Cognitive Uncertainty and Overconfidence

A series of works⁵ highlights the distinction among three different concepts of overconfidence: overestimation, overprecision, and overplacement. In this literature, it is also stressed how, even though often confused in the vernacular, these phenomena are distinct in their causes and in the conditions under which they manifest. In line with this branch of literature, we present a model that stresses formally the differences in the constructs and their causes. More specifically, we present a model that focuses on nesting overplacement within the CU framework. Note that this does not mean that the model cannot reproduce established results concerning overprecision⁶ or overestimation: as shown in the [Appendix](#), MH’s results can be reproduced within this framework, under some assumptions. In what follows, we work out a link between the concepts of CU and overplacement, which is also the main object of the empirical investigation.

2.2.1 Cognitive Uncertainty and Overplacement

Overplacement is defined as an excessive belief in being better than others. In MH’s paper, the phenomenon is studied in a very specific informational setting: the agent has her own performance revealed and has to assess whether it is higher or lower than an “average” agent. Formally, this means that, having observed her performance, she updates her belief about the mean of the performance distribution, being able to assess the relative goodness of her performance. However, thinking about single tasks, instead of aggregated performance, allows us to study the problem from a different perspective.

Consider an agent, part of a measure one set of *cognitive uncertain* agents, having provided her best answer to a given task. Given her prior x and her signal s , she will hold some belief about her action optimality, following (2). We assume that it is common knowledge that all

⁵Originating in MH. See also Moore and Schatz (2017).

⁶There is a clear relation between the concept of CU and overprecision: the first is a necessary but not sufficient condition for the other to emerge. An individual may be affected by cognitive noise and be aware of that, being uncertain about the optimality of her choice, but will not necessarily exhibit overprecision. The latter may emerge only if the *perceived* cognitive noise is lower than the *actual* one. If an agent’s perceived cognitive noise is less disperse than his actual noise, he will overestimate his performance and the precision of his answer, that is, underestimate his CU. Cognitive uncertainty may hence constitute a building block for a formal model of overprecision. Additionally, it is interesting how this perspective provides a rationale for the phenomenon of overprecision to emerge. Agents are affected by some form of cognitive noise, of which they are aware, but their perception does not necessarily correspond to the actual process generating the noise.

agents have identical preferences, described by a simplified version of (1):

$$u(a, x) = -\frac{1}{2}(x - a)^2, \quad (3)$$

which implies that the optimal action corresponds to the state itself $a^* = x$. This change does not affect the interpretation of the model in any way but simplifies the notation. Also, any order-preserving transformation would not change the results. Moreover, we assume that agents being cognitive uncertain is common knowledge.

If an agent, say i , can observe the action undertaken by another agent, call him j , then the expectation about the placement can be defined as the probability that j is worse off, from i 's perspective:

Definition 1. *Given preferences defined by (1) and some belief distribution on the space of action, the relative placement of agent i , with action a_i^* , with respect to agent j , with action a_j^* , is:*

$$Placement_i(a_i^*, a_j^*) = P(u(a_j^*, x) < u(a_i^*, x)).$$

Agent i holds some beliefs about the potential optimal actions, with a_i^* being the mode (and the mean, under normality) of such distribution. Given that agent i observes another agent's action, she will be able to assess, according to her own beliefs, the probability that agent j performed better than she did. This expression can be interpreted as a *continuous answer* to the question "Did you perform better than agent j ?". This definition is a building block to construct i 's overall ranking measure.

Definition 2. *Let $G_i(a)$ be some CdF representing i 's beliefs about other agents' actions. Then agent's i expected ranking is:*

$$Rank_i(a_i^*) = \mathbb{E}_{G_i}[Placement(a_i^*, a_j^*)]$$

Given the definition above, agent i would need to know the distribution of answers provided by other agents, or to hold some belief about that, to be able to form expectations about other agents' actions and hence about her overall placement.

This assumption is not unrealistic in many applied frameworks. Two examples are a firm setting prices and being able to observe prices set by other firms on a similar product, or a financial market investor observing other agents' decisions of buying or selling certain assets. Moreover, for the results to hold, the distribution does not have to be correct, mirroring the actual distribution of agents' actions. It is just necessary that, when assessing her own ranking, i holds some beliefs about other agents' actions.

Given the preferences described by (3), it follows that $a^* = x$, that is a^* and x may be used interchangeably. In this setting it is possible to state the following, all of which is proven in the [Appendix](#):

Proposition 1. Consider a measure one set of cognitive uncertain agents, with preferences defined by (3), and some agent i , with beliefs $a^* \sim \mathcal{N}(a_i^*, \sigma_{CU}^2)$. Then, for any CdF $G_i(\cdot)$, describing agent's i beliefs about other agents' actions such that $\text{Rank}_i(\cdot)$ is well defined, it holds that:

$$i) \text{ Placement}_i(a_i^*, a_j^*) = \begin{cases} 1 - F_{a_i^*}(\frac{a_i^* + a_j^*}{2}) & \text{if } a_j^* < a_i^* \\ F_{a_i^*}(\frac{a_i^* + a_j^*}{2}) & \text{if } a_j^* \geq a_i^* \end{cases},$$

ii) $\text{Placement}_i(a_i^*)$ is decreasing in cognitive uncertainty for all a_j^* ,

iii) $\text{Rank}_i(a_i^*)$ is decreasing in cognitive uncertainty,

with $F_{a_i^*}(\cdot)$ being the CdF representing i 's beliefs about the optimal action.

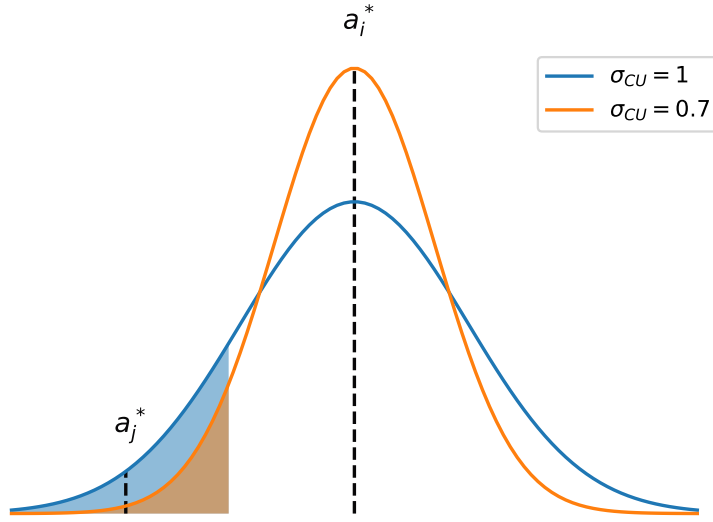


Figure 1: Distribution of beliefs about the optimal action a^* for different levels of CU. The areas represent $1 - \text{Placement}(a_i^*, a_j^*)$ for both beliefs.

The first point of the proposition characterizes i 's relative placement, under this set of assumptions. As shown in the proof, this characterization is an immediate consequence of the distributional assumption and of the quadratic preferences, which formalize an intuitive basic structure: an agents performs better than another if her answer is closer (in a classic Euclidean sense) to the optimal action a^* . Then, the first point of the proposition represents the probability of this event happening, given i 's beliefs about the optimal action a^* . Figure 1 shows this same intuition graphically. The shaded areas represent $(1 - \text{Placement}_i(a_i^*, a_j^*))$, for the case of $a_j^* < a_i^*$, for two different levels of σ_{CU} . The blue area, representing the

higher CU case, is larger, implying that the agent with the higher CU places herself relatively lower.

The second point of the proposition states that the placement of an agent decreases in her cognitive uncertainty and, consequentially, also the overall expected ranking (third point). As uncertainty increases, probability mass is shifted away from a_i^* , the mean of the distribution, towards the tails. Hence, the agent will deem values far from her chosen action more likely to be optimal, decreasing her expected rank.

This result establishes a direct link between cognitive uncertainty and overplacement. An interesting aspect of this result is that it does not depend on beliefs about other agents' actions, as the impact of CU on overplacement does not vary with different specifications of $G_i(\cdot)$. Also, the result in the second point of the Proposition relies on the fact that $Placement_i(\cdot, \cdot)$ is also increasing in σ_{CU} , meaning that the same logic applies for a framework where the agent observes another agent's action.

2.2.2 Persistence

It is possible to model the agent to hold invariant beliefs or to allow for her to update after observing a_j^* . In the first case $F_{a_i^*}(\cdot)$ would be the same CDF prior to observing a_j^* . In the other case, for agent i to be able to update her beliefs, she would have to formulate an assumption about agent's j cognitive uncertainty, denoted by σ_{-i} .

The updated belief about the optimal action would then be:

$$a^* \sim \mathcal{N}\left(\frac{\sigma_{-i}^2}{\sigma_{-i}^2 + \sigma_{CU}^2} a_i^* + \frac{\sigma_{CU}^2}{\sigma_{-i}^2 + \sigma_{CU}^2} a_j^*, \frac{\sigma_{CU}^2 \sigma_{-i}^2}{\sigma_{-i}^2 + \sigma_{CU}^2}\right). \quad (4)$$

In both cases (static or dynamic beliefs) the results from Proposition 1 hold. However, the expression in (4) can be employed to analyze a potential source of overplacement persistence. A long-standing puzzle in the literature is the persistence of overconfidence over time, even in the presence of repeated feedback.⁷ A prominent explanation for this phenomenon has been *motivated beliefs*. In these models, positive self-assessments enter positively inside agents' utilities, under some constraints or costs as to prevent generating infinitely inflated beliefs.⁸ The general idea is that an individual biases his beliefs upwards, as he enjoys holding a positive view of himself, even if this generates (costly) sub-optimal behavior. As compelling as this narrative is, there are arguably frameworks where it may not fit. The motivated beliefs narrative is based on the fact that individuals value a good performance

⁷See, among others, Huffman, Raymond, and Shvets (2021) and Zimmermann (2020).

⁸This literature stems from Bénabou and Tirole's (2002) seminal paper proposing an economic theory of prosocial behavior. See also Köszegi (2006) for a theoretical formulation of utility theory including ego-relevant features. See Bénabou (2015) and Bénabou and Tirole (2016) for literature reviews.

in the task, which may not be the case for neutral tasks or for tasks that people regret participating in. In a series of experiments, Logg, Haran, and Moore (2018) find stronger evidence for a cognitive-based explanation for overconfidence, with the role of motivation being related to vague measures and tasks. Expression 4 suggests an alternative, though not exclusive, way by which persistent overplacement may arise: keeping other factors constant, an agent with a higher assessment of σ_{-i} will hold more conservative beliefs towards her initial guess, resulting in a higher persistence of overplacement. In other words, an agent who *underestimates* others excessively, would be able to keep, over time, excessively high beliefs about her own action optimality.

Figure 2 represents this idea graphically. Agents starting with the same (incorrect) prior, that is with the same belief about the optimal action and the same level of cognitive uncertainty, observing the same action a_j^* , will have different learning paths, for different levels of σ_{-i} : the agent with a larger assessment of σ_{-i} will hold a higher and more persistent belief about his placement over time. Similarly, Figure 3 shows a one-period cross-section of the process shown in Figure 2: holding prior fixed, the agent with the highest assessment of σ_{-i} will hold a posterior such that his placement is higher. The fact that the agent with a larger σ_{-i} has a larger CU after updating⁹, is more than compensated by the fact that the new optimal action is closer to the observed action a_j^* . The figure compares placement functions for two different levels of σ_{-i} , showing the location of the midpoint $\frac{a_i^* + a_j^*}{2}$ for both. This intuition is formalized with the following (proven in Appendix B.2):

Proposition 2. *Consider two agents, i_H and i_L , with identical priors regarding the optimal action $a^* \sim \mathcal{N}(a_i^*, \sigma_{CU}^2)$, but with $\sigma_{-i,H} > \sigma_{-i,L}$. Let a_j^* be some action by agent j observed by both, and let $a_{i_K}^*$, for $K \in \{L, H\}$ be the posterior mean, after having observed a_j^* . Then, $\text{Placement}_{i_H}(a_{i_H}^*, a_j^*) > \text{Placement}_{i_L}(a_{i_L}^*, a_j^*)$.*

Hence, an agent with a higher assessment of σ_{-i} , will in general be more subject to overplacement. The following empirical hypotheses follow from the set of results collected throughout this section:

Hypotheses

1. *Placement decreases in σ_{CU} and increases in σ_{-i} .*
2. *Ranking decreases in σ_{CU} .*
3. *Reaction to others' actions/information increases in σ_{CU} and decreases in σ_{-i} .*

⁹Note that the variance of both agents, after updating their beliefs, is $\frac{\sigma_{CU}\sigma_{-i}}{\sigma_{CU} + \sigma_{-i}}$, given their assessment of σ_{-i} .

In what follows we describe the experimental design, define empirical measures for theoretical quantities, and finally present our analysis strategy and results.

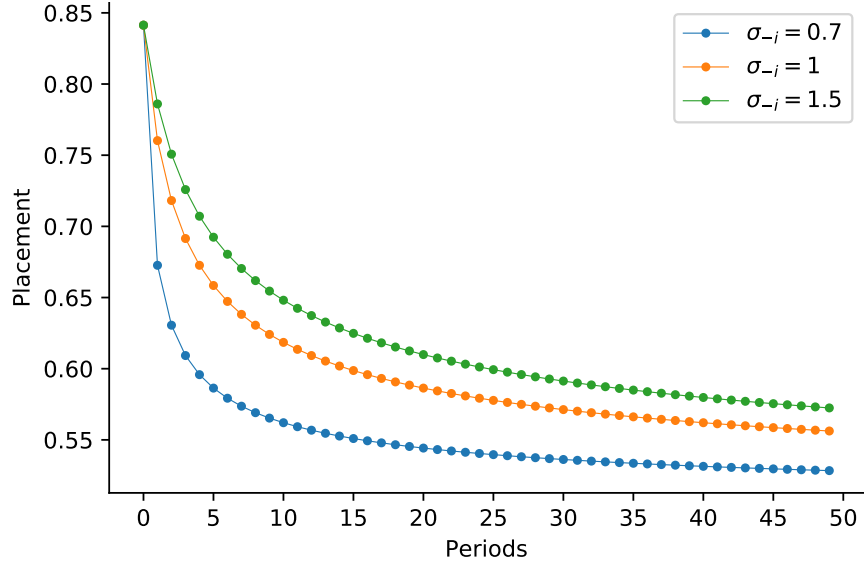


Figure 2: Placement dynamics for different levels of σ_{-i} . All agents start from an identical belief about a^* and observe the same action a_j^* and differ only in the level of σ_{-i} .

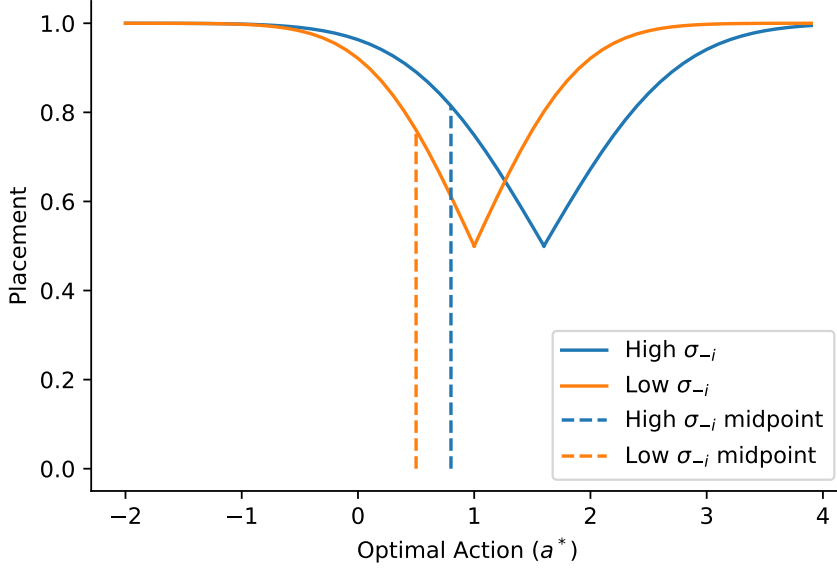


Figure 3: Placement functions: one-period belief updating for different levels of σ_{-i} . Midpoints are the optimal action according to each agent in that period.

3 Experimental Design

The main goal of the experimental design is to test the relationship between cognitive uncertainty and overplacement. More specifically, we aim to test the [hypotheses](#) formulated in the previous section.

The experiment is organized into two main blocks: a set of belief updating tasks and a final survey. Each belief updating task is constructed following Benjamin’s (2019) review, similarly to EG. Participants undergo a classic “balls-and-urns” task, in which they are presented with two hypothetical urns, each containing blue and red balls, in different proportions. After observing a draw of one or more balls, their goal is to provide their best guess of the probability of the draw coming from one of the two urns. Participants are also endowed with a prior probability of either one of the two urns being picked before observing the draw. They are informed that the computer draws a card from a 100 cards deck. Each card is labeled as either one of the two urns. Based on the drawn card, the computer performs the second draw of one or two balls from the selected urn. Moreover, participants are informed about the proportions of the cards in the deck.¹⁰ More formally, participants are

¹⁰For further details about the exact experimental instructions one may refer to [Appendix C.1](#).

provided with the number of "A cards" in the deck, $\mathcal{A} = 100 \cdot P(A)$, as well as the number of blue balls in urn A, $\mathcal{B} = 100 \cdot P(\text{blue} \mid A)$, with the number of blue balls in urn B always set as its complement, that is $100 \cdot P(\text{blue} \mid B) = 100 - \mathcal{B}$. The parameter space, which is unknown to participants, is $\mathcal{A} \in \{30, 50, 70\}$ and $\mathcal{B} \in \{70, 90\}$. Finally, the possible signals are $s \in \{\text{blue}, \text{red}, \text{blue} - \text{blue}, \text{red} - \text{red}, \text{blue} - \text{red}, \text{red} - \text{blue}\}$, with the probability of s being a single ball draw set to 50%.¹¹

Figure 4 represents the task timeline graphically. The green boxes represent the financially incentivized decisions¹². Participants are presented with a balls and urns task with a given parameters specification and formulate their guess. They are also asked to provide a guess about their overall ranking in the task. Afterward, they observe an answer to an identical task provided by another participant and may change their previous guess. Additionally, they are asked to assess their placement relative to that participant and the level of cognitive uncertainty of the other participant when providing the observed answer (σ_{-i}). The key outcomes of interest are the placement measures and the participant deviation from the initial answer, if any such deviation occurs. The steps are repeated for different specifications of the belief elicitation task. The same participant goes through several sessions of the task, each with a different parameters specification. Also, as explained in more detail in the next subsection, half of the sessions would have \mathcal{B} , the diagnosticity parameter, expressed as a random variable. These choices are referred to as *compound* choices. Participants undergo each of the possible 6 parameter specifications. For *compound* choices the parameters are intended in expectations.

Clearly, there is a significant intersection with EG in terms of experimental structure. The key differences are represented by rank and placement elicitation and by the additional steps after CU elicitation, namely: a subject is shown another subject's answer, elicitation about the other subject's CU, and the answer adjustment step. Combining the belief updating task as carried out in EG with these additional steps, represents the novel contribution of this paper from the experimental point of view.

¹¹For more details about the structure and the implementation of the random draws see [Appendix D](#).

¹²For details on how financial incentives are implemented see [Appendix C.4](#)

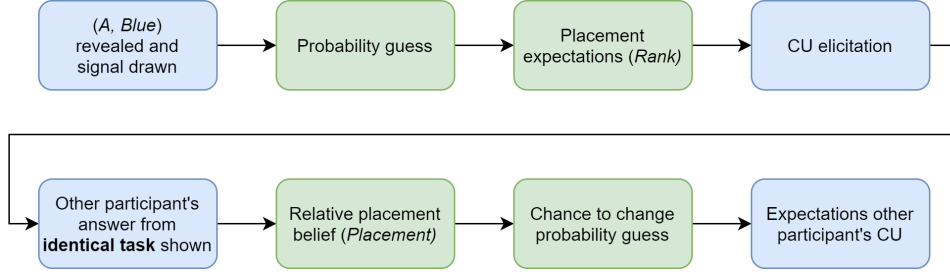


Figure 4: Experimental task timeline.

3.1 Compound Choices

Besides establishing that CU and placement measures are correlated as expected, we are interested in establishing a causal relationship. To do so, we introduce compound choices, as done in EG.

Participants undergo 6 sessions, as described in the previous subsection and shown in Figure 4. In the last 3 sessions, the diagnosticity parameter of the belief updating problem (\mathcal{B}) is presented as a random variable instead of a number. As EG show, introducing compound choices this way is arguably equivalent to manipulate CU, hence helping to establish a causal link between CU and overplacement. To this respect, it should be noted that, to our knowledge, only one work studying the impact of ambiguity on overconfidence is present in the literature (Brenner, Izhakian, and Sade, 2015). However, the evidence presented in the paper is based on a different concept and manipulation of ambiguity. For this reason, in interpreting the results, we assume that variations in placement measures and answers adjustment, in a compound choice framework, would be channeled through the exogenous variation in cognitive uncertainty. To cleanly identify the effect of the manipulation, participants are shown another participant’s answer to an equivalent (reduced) problem without compound parameters. It is stressed that the correct answer for the differently formulated problem is the same. The aim is to keep all other factors constant, compared to the previous condition, including the subject’s beliefs about other subject’s CU: knowing that only her problem is posed in a compound way, the participant should have relatively, but not absolutely, more trust in the observed answer. This condition is implemented intervening on points 1 and 5 of the experiment timeline. A different example is provided in point 1, comparing the new compound task with the previous task. In point 5 subjects are provided with an answer from an equivalent, non-compound task. A preliminary study has been run to collect a sufficiently large pool of answers for the belief updating task. This study excluded the learning component of the belief updating task since the aim was only to gather answers to be used in the next phase of the study. The answers shown in phase 5 of the belief updating task are randomly drawn from the pool of answers gathered in the

preliminary study, conditioning on parameters specification and signals realization.

3.2 CU Elicitation

A key measure in the experiment is the one for cognitive uncertainty. In this, we follow closely EG. Figures 5 and 6 show screenshots from an example task. The only way our elicitation differs from Enke and Graeber operationalization of CU is that we set the uncertainty to grow moving the slider from left to right.

Referring to their work, this operationalization of participants uncertainty has a simplicity advantage over confidence intervals elicitation. This is because participants do not have to understand the concept of confidence intervals¹³ and think about probability in answering the question about cognitive uncertainty. Similarly, eliciting full probability distributions over (range of) outcomes is more complex and require the subjects to have a certain degree of understanding of probability theory.

How certain are you that the optimal guess for **bag A probability** is **exactly 20%**?

You are **sure** that the **optimal guess** is **please click the slider**.

Very Certain

Very Uncertain



Figure 5: Example of Cognitive Uncertainty Elicitation Pre Click

¹³Confidence interval being a proper measure for eliciting participants perceived precision has been widely debated in the literature. Enke and Graeber (2021) also conducted a calibration experiment using confidence intervals, noting how changing the confidence level does not impact much intervals wideness. This was already noted by Alpert and Raiffa (1982), who started this literature. For an extensive review of this problem see Logg, Haran, and Moore (2018).

How certain are you that the optimal guess for **bag A probability** is **exactly 20%**?

You are **sure** that the **optimal guess** is **between 11% and 29%**.

Very Certain Very Uncertain



Figure 6: Example of Cognitive Uncertainty Elicitation After Click

3.3 Rank and Placement Elicitation

As illustrated in Figure 4 rank is elicited following participants probability guess. More specifically, participants are asked to provide their guess about their overall ranking in that specific task, right after providing their probability guess. Figure 7 provides an example.

The computer **drew 1 blue ball**.

Please write down **your guess** for the **probability** (between 0 and 100) of **each bag being chosen**.

Probability of bag A :	<input type="text" value="0"/>
Probability of bag B :	<input type="text" value="0"/>
Total	<input type="text" value="0"/>

Please provide you **guess about your placement** in this **specific task** (between 1 and 100):

Figure 7: Example of Rank Elicitation

For what concerns placement, participants were first shown the answer of another participant in an identical task and then asked how likely it was that they performed better than that participant. Figure 8 provides an screenshot from the experiment.

The other participant answered as follows:

Probability of bag A: 10%

Probability of bag B: 90%

Did you perform better than the other participant?

Please write your guess of the probability that you did (100 means that you are sure to have performed better and 0 means that you are sure that you have not).

Figure 8: Example of Placement Elicitation

3.4 Logistics

The experiment’s participants were recruited using Amazon Mechanical Turk platform (MTurk). Attention checks were put into place to ensure data quality.

We ran a preliminary data collection, in which a total of 176 participants were recruited. Of those, 71 were screened-out, either because they answered incorrectly at least one of the comprehension questions¹⁴ or because they failed the attention check put within the tasks. Hence, a total of 105 participants were kept. The answers of these participants have been used as a pool to draw from in the actual study. The attention check was a guessing task framed in a way such that either urn A or urn B was correct with probability 1. If participants did not answer correctly to that task they were screened-out.¹⁵ The preliminary study took approximately 19 minutes to complete on average. Participants who successfully completed it received, on average, 4.97 USD.

A total of 422 participants were recruited to take part in the study, with 198 being screened

¹⁴The comprehension questions used in the preliminary study are almost the same to the ones used for the non-compound part of the main study, reported in Appendix C.2. The preliminary study contained an additional question ensuring that participants understood what the probability of the sure event is.

¹⁵For more details on the attention check and on structural differences between preliminary and the main study see Appendix C.5.

out for failing to answer correctly comprehension questions, leaving a sample of $N = 224$ participants. Participants were paid 0.5 USD for accepting the task on MTurk and an additional 4.5 USD upon completion. Additionally, they could earn up to a 3 USD bonus, which was determined as previously described.

The study took an average of 23 minutes to complete. Participants who completed the study received, on average, 6.62 USD.

4 Analysis and Results

In this section, we illustrate our data analysis strategy and results. Each of the main results corresponds to one of the previously formulated hypotheses. Moreover, we run additional analyses on two different measures of overplacement and report preliminary evidence on how CU may mediate the relationship between overplacement and probability weighting.

4.1 Hypothesis 1: Placement

Figure 9 provides an overview of the distribution of *Placement*, respectively for a high and low level of CU. The groups are determined by taking the median level of CU in the whole sample as a threshold. The figure provides preliminary evidence in line with hypothesis 1: comparing the two distributions, the *Low CU* group exhibits more mass on the right end of the domain, suggesting that participants with low levels of CU tended to place themselves higher compared to the participants in the *High CU* group.

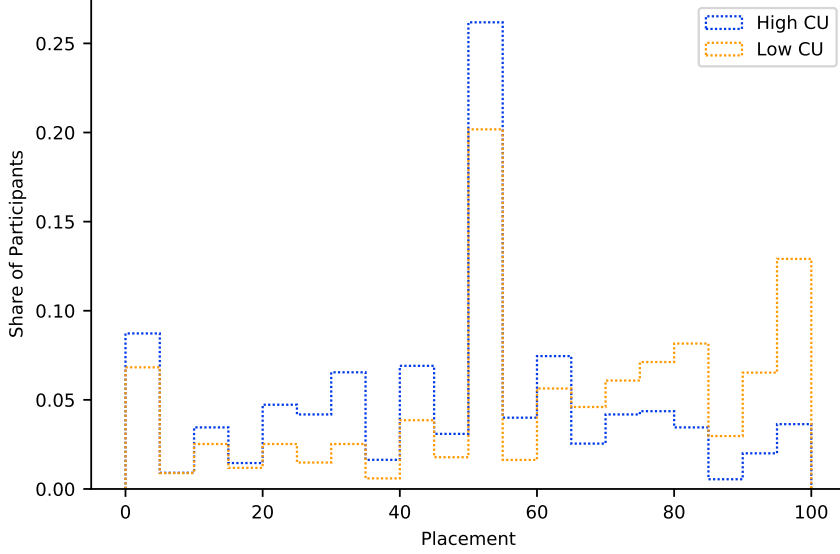


Figure 9: Placement distribution for High/Low CU groups. Groups are determined using median CU as a threshold.

To perform a more rigorous analysis, we estimate the following equation:

$$Placement_i = \alpha + \beta_1 CU_i + \beta_2 I_{compound} + \sum_{k \in K} \beta_k X_{k,i} + \varepsilon_i, \quad (5)$$

with $I_{compound}$ being an indicator for compound choices and K the set of control variables (e.g. survey variables, session fixed effects). The main coefficients of interest are β_1 and β_2 . The first can be interpreted as the estimated average marginal effect of CU on placement. β_2 is interpreted as the effect of manipulating CU through compound choices.¹⁶

Table 1 provides coefficients from linear estimates of elicited placement. Columns (1) and

¹⁶For this interpretation to be valid, it must hold that: (i) CU is significantly higher for compound choices and (ii) any effect of compound choices on placement level is due to variation in CU or σ_{-i} . The second point is argued in detail in the previous section. The core idea is that, to our knowledge, there is no theory relating compound choices to overplacement (or overconfidence in general). Concerning the first point, we find that compound choices increase CU by 17% on average. Figures E.1 and E.2, reported in the Appendix, present this finding graphically. Figure E.1 shows how the distribution of CU changes between compound and baseline choices. Mass is shifted towards higher levels of CU for compound choices, although this change is not sharp. Figure E.2 represents a t-test at the 95% confidence level, comparing the average normalized CU for compound and non-compound choices. Based on this evidence, we conclude that compound choices represent an effective manipulation of CU and that β_2 may be interpreted as suggested.

(2) report the results of regressing placement only on CU and the manipulation dummy,¹⁷ respectively. Column (3) provides estimates of β_1 and β_2 , estimated together, without additional control variables. Columns (4) and (5) separately add sessions fixed effects and demographic controls.¹⁸ Finally, column (6) estimates full equation 5, including also the absolute distance between the participant’s first guess and the shown answer from another participant ($|a_i^* - a_j^*|$).

This analysis shows that CU has a significant effect on elicited placement, in line with hypothesis 1. The compound manipulation allows us to interpret at least part of this effect causally. In addition, the variation in *1 if compound choice* coefficient from column (2) to column (3) suggests exactly that part of the effect of the manipulation is explained by the variation in CU: when adding CU to the specification the coefficient of the compound choices dummy decreases in absolute value. Interestingly, $|a_i^* - a_j^*|$ coefficient is positive and significant: when a participant observes an answer from another participant that is more distant from her initial answer, she will be more likely to place herself higher. This result is consistent with the model proposed in Section 2, in which $Placement_i(\cdot)$ is increasing in $|a_i^* - a_j^*|$.¹⁹ Hypothesis 1 also conjectures an effect of σ_{-i} on placement. To test this, we add σ_{-i} to equation 5. The results of the estimation are reported in Table 2. Columns (1)-(6) of Table 2 perfectly correspond to Table 1 columns, with σ_{-i} added to each specification. The estimated effect of σ_{-i} is positive, as hypothesized, and significant, for each of the 6 specifications. Two additional aspects are worth noting. First, comparing column (2) from Table 1 and 2, it is possible to observe that, as for CU, introducing σ_{-i} in the estimation model decreases the compound choices coefficient, suggesting that part of the estimated effect of the dummy is to be attributed to σ_{-i} . Second, in column (3) of Table 2, the dummy coefficient decreases drastically in absolute value and the model R^2 is doubled (compared to the same column in the previous table). These elements suggest that both CU and σ_{-i} are relevant in assessing placement and that they should be considered jointly, as doing so sharply increases the model explanatory power. Moreover, this suggests that the effect of the compound choices manipulation is channeled through both CU and σ_{-i} .

¹⁷This variable assumed the value of 1 if the observation corresponds to a compound choice.

¹⁸These comprise age and participant’s education level.

¹⁹Note that the number of observations for column (6) is 1218, instead of 1224. This is to be attributed to a technical problem by which it was not possible to keep track of which of the other participant’s answers (a_j^*) was shown to that participant. Hence, in all estimations including $|a_i^* - a_j^*|$, the 6 observations from that participant are dropped.

<i>Dependent variable: placement</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
CU	-1.023*** (0.189)		-0.954*** (0.191)	-0.952*** (0.191)	-0.944*** (0.186)	-0.915*** (0.176)
1 if compound choice		-10.593*** (1.472)	-9.099*** (1.446)	-9.102*** (1.447)	-9.114*** (1.444)	-9.487*** (1.411)
$ a_i^* - a_j^* $						0.381*** (0.054)
Session FE				0.424 (0.702)		0.453 (0.689)
Demographic Controls	✗	✗	✗	✗	✓	✓
Observations	1,224	1,224	1,224	1,224	1,224	1,218
R^2	0.070	0.037	0.097	0.097	0.100	0.154

Notes. OLS estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's placement level, that is the elicited probability of performing better than another subject whose action is observed. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1: Effect of CU on placement

<i>Dependent variable: placement</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
CU	-1.513*** (0.158)		-1.437*** (0.160)	-1.435*** (0.160)	-1.424*** (0.157)	-1.333*** (0.157)
σ_{-i}	1.323*** (0.143)	0.692*** (0.158)	1.241*** (0.144)	1.242*** (0.144)	1.234*** (0.146)	1.045*** (0.143)
1 if compound choice		-9.274*** (1.427)	-5.976*** (1.377)	-5.979*** (1.377)	-6.011*** (1.376)	-6.691*** (1.355)
$ a_i^* - a_j^* $						0.266*** (0.052)
Session FE				0.507 (0.670)		0.535 (0.669)
Demographic Controls	✗	✗	✗	✗	✓	✓
Observations	1,224	1,224	1,224	1,224	1,224	1,218
R^2	0.174	0.069	0.185	0.185	0.187	0.211

Notes. OLS estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's placement level, that is the elicited probability of performing better than another subject whose action is observed. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Effect of CU and σ_{-i} on placement

4.2 Hypothesis 2: Rank

Similarly to what we showed for placement, Figure 10 depicts how rank distribution differs for *High/Low* CU levels. In this case, the cut between the two distributions is less sharp, but the *Low* CU group exhibits more mass on the left,²⁰ as hypothesis 2 would imply.

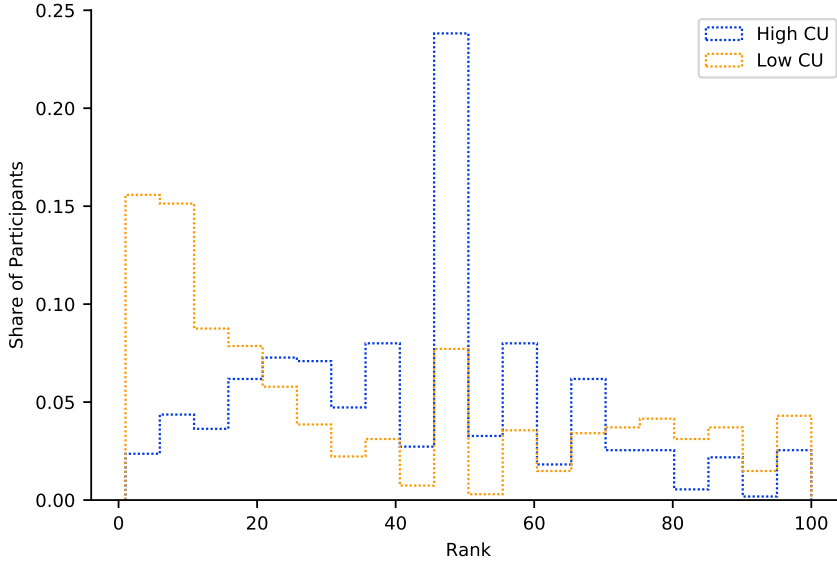


Figure 10: Rank distribution for high/low CU. Groups are determined using median CU as a threshold.

We perform an econometric analysis, estimating an equation equivalent to equation 5, with the exception of the dependent variable being rank, instead of placement, that is the expected placement without observing other participants answers. The results of the estimation procedure are reported in Table 3. Both the estimated effects of CU and of compound choices are positive and significant, for all specifications. Similar to what we note for placement, it is possible to see that the compound choice dummy coefficient decreases comparing columns (2) and (3). This reinforces the interpretation of compound choice serving as a manipulation for CU.

²⁰Note that the way rank is operationalized implies that a higher value for the variable is interpreted as a lower probability of being better-off. For example, a participant that assumes to be ranked 10th expects to be better off than a participant who assumes to be ranked 15th.

<i>Dependent variable: rank</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
CU	0.734*** (0.219)		0.713*** (0.221)	0.713*** (0.221)	0.754*** (0.219)	0.754*** (0.219)
1 if compound choice		3.892*** (1.075)	2.775*** (1.030)	2.775*** (1.031)	2.711*** (1.036)	2.711*** (1.036)
Session FE				0.010 (0.573)		0.024 (0.576)
Demographic Controls	✗	✗	✗	✗	✓	✓
Observations	1,224	1,224	1,224	1,224	1,224	1,224
R^2	0.036	0.005	0.039	0.039	0.052	0.052

Notes. OLS estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's rank level, that is the elicited expected ranking in the current task, from 1 (first) to 100 (last). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effect of CU on rank

4.3 Hypothesis 3: Answer Adjustment

The third hypothesis is not directly related to overplacement measures, but to how CU and σ_{-i} are related to participants' reactions after observing another participant's answer. Although not the main goal of the paper, this represents a further way to test the proposed model.

We take as a dependent variable the absolute difference between the probability guess provided in the first part of the task and the guess provided after observing the other participant's answer. Importantly, only 243 observations out of 1224 have non-zero answer adjustments. This raises concerns of OLS estimates being driven by the observations in which no adjustment took place. For this reason, we employed two approaches in testing hypothesis 3: OLS and probit estimation.

We first analyze the effect of CU and σ_{-i} on answer adjustment estimating the following using OLS:

$$ans_adj = \alpha + \beta_1 CU_i + \beta_2 I_{compound} + \beta_3 \sigma_{-i} + \sum_{k \in K}^K \beta_k X_{k,i} + \varepsilon_i. \quad (6)$$

The results of the estimation are reported in Table 4. Afterward, we estimated a probit model using the same variables of equation 6. Table 5 reports the estimates of this exercise. Overall, OLS estimates are in line with our hypothesis: answer adjustments increase with CU and decrease with σ_{-i} , on average. In the full specification, the magnitude of the estimated effect of σ_{-i} is approximately 25% larger than that of CU. This is the opposite for the case of placement, in which the estimated effect of CU is approximately 22% larger. Concerning probit estimates, it is interesting to see how CU is highly significant only in the full specification of the model, unlike σ_{-i} , which is always significant. This indicates that variation in CU impacts the estimated probability of adjusting the answer less than σ_{-i} . Hence, when not considering the magnitude of the adjustment, as in the OLS case, but only the probability of the adjustment taking place, CU has less impact. This may be interpreted as follows: once a participant decides to adjust her answer, her level of cognitive uncertainty matters to determine how much she will deviate from her initial answer. However, CU is less impactful concerning the decision of changing the answer or not.

<i>Dependent variable: answer adjustment</i>				
	(1)	(2)	(3)	(4)
CU	0.187*** (0.036)		0.155*** (0.037)	0.195*** (0.035)
σ_{-i}	-0.187*** (0.040)		-0.153*** (0.039)	-0.243*** (0.048)
1 if compound choice		3.042*** (0.447)	2.508*** (0.425)	2.043*** (0.360)
$ a_i^* - a_j^* $				0.137*** (0.032)
Session FE				-0.193 (0.221)
Demographic Controls	X	X	X	✓
Observations	1,224	1,224	1,224	1,218
R^2	0.040	0.040	0.066	0.155

Notes. OLS estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's answer adjustment, that is the absolute difference between the first and the second choice in the probability guessing task. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Effect of CU and σ_{-i} on answer adjustment

<i>Dependent variable: answer adjustment probability</i>				
	(1)	(2)	(3)	(4)
CU	0.010*		0.008	0.038***
	(0.005)		(0.006)	(0.007)
σ_{-i}	-0.089***		-0.090***	-0.062***
	(0.007)		(0.007)	(0.008)
1 if compound choice		-0.513***	0.084	0.559***
		(0.053)	(0.073)	(0.089)
$ a_i^* - a_j^* $				0.008***
				(0.003)
Session FE				-0.212***
				(0.047)
Demographic Controls	X	X	X	✓
Observations	1,224	1,224	1,224	1,218

Notes. Probit estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's answer adjustment probability, that is an indicator for the subject having changed answer after observing another participant's answer. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Effect of CU and σ_{-i} on answer adjustment probability

4.4 Overplacement

In the previous sections, we studied how the empirical hypotheses descending from the model met with our experimental data. The hypotheses concerned how *placement* (and *ranking*) varied with CU and σ_{-i} , but not how those impacted *overplacement*. This is because, in our theoretical framework, to state if and how much overplacement takes place, it is necessary to formulate additional assumptions about the structure of the cognitive noise.²¹ However, given that an increase in CU (σ_{-i}) decreases (increases) placement, it will impact overplacement both on the extensive margin (whether a participant overplaces herself or not) and on the intensive margin (the extent to which a participant overplaces herself).

We run two additional analyses to test this hypothesis, that is assessing the effect of CU and σ_{-i} on *overplacement*. The two analyses correspond to two different measures we propose, corresponding to the extensive and intensive margin of overplacement. Both measures are constructed using the placement decision of participants after observing the other participant's answer. The first measure is a dichotomic variable, taking the value of 1 if the participant overplaced herself and 0 otherwise. A participant i , with answer a_i^* and observed answer a_j^* , overplaced herself if her placement decision was above 50% and $|a^* - a_i^*| > |a^* - a_j^*|$. In other words, a participant overplaced herself if she stated that it was more likely to have performed better than the other participant when she did not. Table 6 reports the results of running a probit regression on this measure of overplacement, which can be interpreted as overplacement probability. The four specifications are the same as the ones in the previous sections. Both CU and compound choices have a highly significant effect on overplacement probability. On the other hand, σ_{-i} seems to have either no effect or a quite small one. This suggests that beliefs in other participant's cognitive uncertainty play no role in determining whether someone will overplace herself or not.

²¹More specifically, it would be necessary to assume how the variance of the cognitive noise is distributed among agents.

<i>Dependent variable: overplacement probability</i>				
	(1)	(2)	(3)	(4)
CU	-0.063*** (0.006)		-0.051*** (0.006)	-0.031*** (0.007)
σ_{-i}	0.004 (0.005)		0.005 (0.005)	0.011* (0.006)
1 if compound choice		-0.832*** (0.058)	-0.367*** (0.072)	-0.211** (0.082)
$ a_i^* - a_j^* $				0.015*** (0.002)
Session FE				-0.082* (0.043)
Demographic Controls	✗	✗	✗	✓
Observations	1,224	1,224	1,224	1,218

Notes. Probit estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's overplacement probability, that is an indicator for the subject assessing her probability of performing better than the other subject higher than 0.5 and having performed worse. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Effect of CU and σ_{-i} on overplacement probability

It is important to note, that this first measure does not take into account two relevant factors. First, individuals who exhibit underplacement are codified in the same way as the ones who correctly place themselves, i.e. with a 0. Secondly, this measure does not consider the magnitude of overplacement (or underplacement) for those who do overplace (underplace) themselves. To make up for these limitations, we build a second measure of overplacement. The measure is built in two steps. We first codify underplacement in our variable. Similarly to overplacement, a participant underplaced herself if she performed better than the other participant, but thought she did not. This is first codified with a -1, opposed to a 1 for overplacement. To address the second concern, we weigh all observations that exhibit underplacement or overplacement by their distance from 50%. This way we differentiate participants by overplacement (underplacement) level. To clarify, consider two participants who overplaced themselves: if one answered 90% and the other 60%, the first would be "overplacing herself more" in our measure. Table 7 reports the results of regressing this measure of overplacement on our variables of interest. CU, σ_{-i} and compound choice dummy are significant in all specifications and their sign in line with the model. Hence, the evidence suggests that cognitive uncertainty plays a role in regulating both the probability of overplacement and the extent of such overplacement.

<i>Dependent variable: weighted overplacement</i>				
	(1)	(2)	(3)	(4)
CU	-0.595*** (0.116)		-0.574*** (0.118)	-0.459*** (0.113)
σ_{-i}	0.615*** (0.113)		0.592*** (0.115)	0.378*** (0.106)
1 if compound choice		-3.681*** (1.241)	-1.653 (1.265)	-2.596** (1.253)
$ a_i^* - a_j^* $				0.291*** (0.043)
Session FE				0.381 (0.602)
Demographic Controls	X	X	X	✓
Observations	1,224	1,224	1,224	1,218
R^2	0.055	0.008	0.056	0.113

Notes. OLS estimates, robust standard errors are clustered at the subject level. The dependent variable is a subject's weighted overplacement. For details on how the measure is constructed see [Section 4.4](#). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of CU and σ_{-i} on weighted overplacement

4.5 Relation to Posterior Compression

In their work, EG document how CU can account empirically for a series of previously unrelated and well-established empirical regularities. One of such regularities is an inverse S-shaped relationship between Bayesian posteriors and posterior beliefs reported by participants in classic "balls-and-urns" tasks. In line with the evidence of the rest of the paper, EG show how participants with high levels of CU exhibit a more pronounced inverse S-shaped pattern. In other words, participants with lower levels of CU report priors that are closer to the Bayesian benchmark.

In our paper, we postulate and investigate empirically a relationship between CU and overplacement, using different measures. If a variation in CU impacts both overplacement and reported posterior compression towards a 50-50 mental default, we would expect to observe a relationship also between overplacement and CU. We explore this relationship in Figure 11, which reports the relationship between Bayesian posterior and stated posterior, separately for participants with a "High Rank" and a "Low Rank". Each marker in the figure represents the average probability guess by participants for a given Bayesian posterior and a given rank level. The mean in the "High Rank" ("Low Rank") group for each Bayesian posterior is computed considering participants who rank themselves in the bottom (top) half of the distribution.²² Figure 11 shows that, on average, "High Rank" participants exhibit a more pronounced S-shaped pattern, while "Low Rank" participants have on average posteriors that are flatter towards the 45-degree line, representing the Bayesian benchmark. This difference is consistent with our findings concerning the relationship between rank and CU. Participants with higher CU also tend to rank themselves higher (that is rank themselves worse) and hence we observe this relation between rank and posterior compression. The idea is that CU regulates both phenomena, which are in turn correlated. To our knowledge, we are the first to document this kind of relationship, and, relatedly, we are not aware of any other theory that can account for this relation. However, it is important to stress that this evidence is extremely preliminary and potentially not robust, as our experimental paradigm was not designed to identify it. Indeed, running the same type of graphical analysis using placement as a threshold to generate groups (Figure E.3 is reported in the Appendix) does not suggest that any relationship between placement and compression. Hence, we believe that documenting this relationship between rank and compression of reported posterior corroborates the rest of our findings and their relation with EG's findings, but further investigation is required to develop a better understanding of a potential relationship between overplacement and probability weighting.

²²This is because participants who believe to have a top-half performance, would provide a small number, as the best rank is 1.

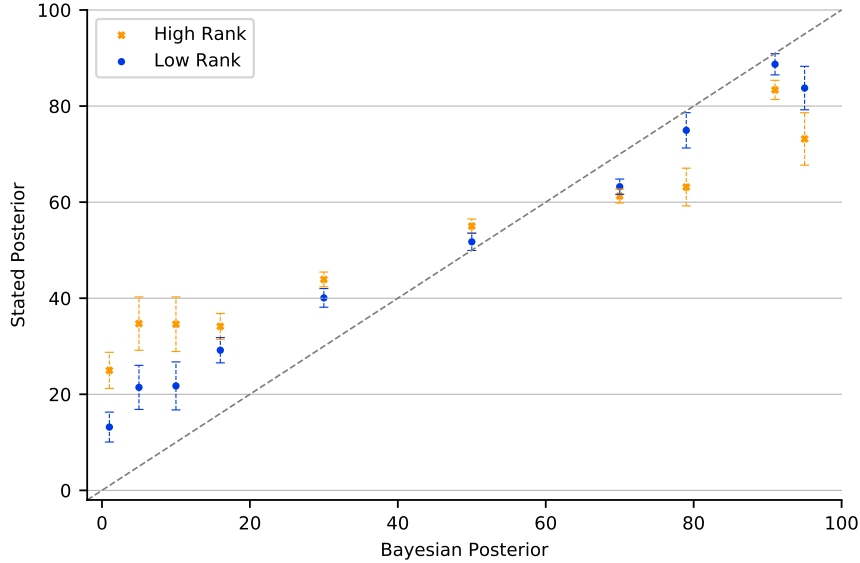


Figure 11: Average reported posteriors for high/low ranks. Groups are determined using median rank as a threshold. The error bars represent the standard error of the means. Bayesian posteriors are rounded to the nearest integer. We show only buckets that contains at least 20 observations.

5 Conclusion

In this paper we propose a model of overconfidence based on the concept of cognitive uncertainty: in the process of solving complex tasks, agents are uncertain about the optimality of their choice, uncertainty caused by a noisy cognitive process. More specifically, we focus on the phenomenon of overplacement. We derive an inverse relationship with cognitive uncertainty and show how persistent overplacement may arise in this framework. Finally, we show how this relationship holds empirically, through an online experiment, based on belief updating tasks. The evidence obtained through the experiment suggests that an increase in cognitive uncertainty induces participants to place themselves lower, relative to other participants, and to react more strongly to information inferred by observing other participants' choices. These results, besides confirming our hypotheses, imply that overplacement, and overconfidence in general, may be related to other behavioral biases through cognitive uncertainty. We present preliminary evidence of this idea, documenting a relationship between our placement measures and the shape of probability weighting. We observe that participants who rank themselves lower exhibit a more compressed probability weighting function. Investigating if and how cognitive uncertainty modulates the relationship between

overconfidence and other behavioral anomalies, such as probability weighting, is left to future work.

On top of overplacement, we also briefly discussed a new perspective to approach the concept of overprecision, but with a quite shallow contribution. To provide deeper insights, it would be of particular interest to test empirically the relationship between actual cognitive noise, cognitive uncertainty, and overprecision. Hence, we believe an operationalization of actual cognitive noise would represent a relevant step in this investigation. The scope of the empirical investigation may be broadened, including tasks more traditionally used in the overconfidence literature. This would strengthen the link with this literature and allow to formulate more general claims about the validity of the theory. Finally, extending the theory to feature discrete action spaces and hence non-Gaussian beliefs may provide interesting insights, especially when considering discrete applications.

References

- Alpert, Marc, and Howard Raiffa.** 1982. "A progress report on the training of probability assessors." In *Daniel Kahneman, Paul Slovic and Amos Tversky (eds.), Judgment Under Uncertainty: Heuristics and Biases*. 294–305. [15]
- Barber, Brad, and Terrance Odean.** 2001. "Boys Will Be Boys: Gender, Overconfidence, And Common Stock Investment." *Quarterly Journal of Economics* 116 (02): 261–92. [2]
- Bénabou, Roland.** 2015. "The Economics of Motivated Beliefs." *Revue d'économie politique* 125 (09): 665–85. [3, 9]
- Bénabou, Roland, and Jean Tirole.** 2002. "Self-Confidence And Personal Motivation." *Quarterly Journal of Economics* 117 (02): 871–915. [3, 9]
- Bénabou, Roland, and Jean Tirole.** 2016. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives* 30 (08): 141–64. [3, 9]
- Benjamin, Daniel.** 2019. "Errors in Probabilistic Reasoning and Judgmental Biases." *Handbook of Behavioral Economics*, [3, 12]
- Brenner, Menachem, Yehuda Izhakian, and Orly Sade.** 2015. "Ambiguity and Overconfidence." *Working Paper*, [14]
- Burson, Katherine, Richard Larrick, and Joshua Klayman.** 2006. "Skilled or Unskilled, but Still Unaware of It: How Perceptions of Difficulty Drive Miscalibration in Relative Comparisons." *Journal of personality and social psychology* 90 (02): 60–77. [4]
- Celen, Boğaçhan, and Shachar Kariv.** 2004. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, [4]
- De Filippis, Roberta, Antonio Guarino, Philippe Jehiel, and Toru Kitagawa.** 2017. "Updating ambiguous beliefs in a social learning experiment." *CeMMAP working papers*, [4]
- Enke, Benjamin, and Thomas Graeber.** 2021. "Cognitive Uncertainty." *Working Paper*, [2, 15, 44]
- Gabaix, Xavier.** 2019. "Behavioural inattention." in *Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., Handbook of Behavioral Economics-Foundations and Applications* 2, 261–343. [2]
- Glaisher, James Whitbread Lee.** 1871. "On a class of definite integrals." *London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 42: 421–36. [41]
- Goeree, Jacob, Brian Rogers, Thomas Palfrey, and Richard D. McKelvey.** 2007. "Self-Correcting Information Cascades." *Review of Economic Studies* 74 (07): 733–62. [4]
- Hossain, Tanjim, and Ryo Okui.** 2013. "The Binarized Scoring Rule." *Review of Economic Studies* 80 (3): 984–1001. [51]
- Huffman, David, Collin Raymond, and Julia Shvets.** 2021. "Persistent Overconfidence and Biased Memory: Evidence from Managers." *Working Paper*, [3, 9]
- Kahneman, Daniel.** 2011. *Thinking Fast and Slow*. New York: Farrar, Straus, and Giroux. [3]

- Khaw, Mel Win, Ziang Li, and Michael Woodford.** 2017. "Risk Aversion as a Perceptual Bias." Working Paper 23294. National Bureau of Economic Research. [2]
- Köszegi, Botond.** 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association* 4 (06): 673–707. [9]
- Krueger, Joachim, and Ross Oakes Mueller.** 2002. "Unskilled, unaware, or both? The better-than-average heuristic and statistical regression predict errors in estimates of own performance." *Journal of Personality Social Psychology* 82 (01): 180–88. [4]
- Kruger, J., and David Dunning.** 2009. "Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments." *Journal of Personality and Social Psychology* 1 (01): 30–46. [4]
- Logg, Jennifer, Uriel Haran, and Don Moore.** 2018. "Is overconfidence a motivated bias? Experimental evidence." *Journal of Experimental Psychology: General* 147 (10): 1445–65. [10, 15]
- Malmendier, Ulrike, and Geoffrey Tate.** 2005. "CEO Overconfidence and Corporate Investment." *Journal of Finance* 60 (12): 2661–700. [2]
- Moore, Don, and Derek Schatz.** 2017. "The three faces of overconfidence." *Social and Personality Psychology Compass* 11 (08): [6]
- Moore, Don, and Deborah Small.** 2007. "Error and Bias in Comparative Judgment: On Being Both Better and Worse Than We Think We Are." *Journal of personality and social psychology* 92 (07): 972–89. [4]
- Moore, Don A., and Paul J. Healy.** 2008. "The Trouble With Overconfidence." *Psychological Review* 115 (2): 502–17. [2, 38, 40]
- Nöth, Markus, and Martin Weber.** 2003. "Information Aggregation with Random Ordering: Cascades and Overconfidence." *Economic Journal* 113 (01): 166–89. [4]
- Weizsäcker, Georg.** 2010. "Do We Follow Others When We Should? A Simple Test of Rational Expectations." *American Economic Review* 100 (5): 2340–60. [4]
- Woodford, Michael.** 2020. "Modeling Imprecision in Perception, Valuation, and Choice." *Annual Review of Economics* 12 (08): 579–601. [2]
- Zimmermann, Florian.** 2020. "The Dynamics of Motivated Beliefs." *American Economic Review* 110 (2): 337–61. [3, 9]

Appendix A Moore and Healy Model

A.1 Overconfidence

In their seminal paper MH propose a classification, according to which overconfidence can be split in three sub-phenomena: *overestimation*, *overplacement* and *overprecision*. The first can be defined as an upward bias in assessing one’s own performances (downward for underestimation), the second as an inflated belief about one’s relative performance, and the last as an excessive belief in the fact that one knows the truth.

They also develop a model of overconfidence able to provide a foundation for some important puzzles in the literature, such as the *hard-easy* effect in overestimation and the inconsistency between overestimation and overplacement. We here provide a sketch of the model and how it accounts for said puzzles.

The problem of assessing one’s performance is modeled as a signal extraction problem: agent i assumes her performance (in a test) is a realization of a random variable $x_i = \mu + \gamma_i$, with μ representing the average performance and γ_i some (not necessarily) zero-mean noise term. Hence, x_i distribution represents the agent’s prior about her performance.

Once the agent undertakes the test, she receives a signal, a "gut feeling", $s_i = x_i + \rho_i$ about how she performed in the said test. Once again, it is assumed for ρ_i to have zero mean, but no specific distributional assumption is necessary for the main intuition to hold: when an agent is assessing her performance under this information structure, her updated belief will be a weighted average of the signal s_i and μ .

This first element may account for the *hard-easy* effect: an *easy* test ($s_i > \mu$) will induce an updated belief $\mathbb{E}[x_i \mid s_i] \in [\mu, s_i]$, mechanically generating underestimation. The opposite would hold for hard tests, mechanically generating overestimation.

Before proceeding with the MH model, a remark is due. The fact that the authors model agents’ performance assessment as a signal extraction problem, implicitly assumes the existence of a source of uncertainty, from which the noise comes, with two points in common with the CU model sketched so far. First, this source is, at least partially, *internally generated*: the agent is still uncertain about her performance also after taking the test, with the γ term representing the (cognitive) noise. Assuming that the mapping from correct answers to performance is not where uncertainty is generated, then the source must be internal. Second, the agent is aware of the existence of the noise: instead of taking the signal s_i at face value, she updates her belief. If the agent was not aware of the existence of her (cognitive) noise, she would have no reason to act in that sense. Hence, it is already possible to see how this benchmark model shares, also implicitly, some key assumptions with the CU model.

MH also show how this model may account for the negative relation between overestimation and overplacement. For this argument to hold, however, the agent must have a different

information set. It is now assumed that the agent observes x_i and is required to evaluate how she performed compared to another random test-taker. In other words, the agent will compute $\mathbb{E}[\mu \mid x_i]$, the updated expectations about μ , given that she observed the realization of her performance x_i . For an argument similar to the one for overestimation, an easy test will generate overplacement and a hard test underplacement. For example, an easy test, that is an higher than expected performance ($x_i > \mu$), will mechanically induce an expected average performance $\mathbb{E}[\mu_i \mid x_i] \in [\mu_i, x_i]$, generating overplacement.

It is worth stressing that MH theory of overconfidence can relate overprecision to the other two sub-phenomena²³, but stays silent as to what may be the mechanism behind it emerging. In the authors' words: *"As to the question of when we should expect overprecision, our theory has little to say."* This represents a first direction in which modeling overconfidence within a CU framework may be an advantage, as it is clarified later.

A.2 Relating Models

From a mathematical perspective, the two models considered in this Section are very similar in that both are (Gaussian) signal extraction problems. However, what distinguishes them is the object of the inference. In one case, the MH model, the agent is trying to assess her performance in a set of tasks (e.g. a test). In the other, the agent inference concerns the optimality of her own action. Clearly, the domains are closely related but do not directly overlap.

To cleanly relate the two domains, a mapping from the action space to the performance space is needed. The equivalence of the two models will then depend on the properties of this mapping and the distributional assumptions. Indeed, the equivalence of the two models is expressed in terms of the resulting distribution on the performance space.

Let A be the set of feasible *actions* and P the set of possible resulting *performances* or *outcomes*. In principle, both sets are unrestricted and may be dense subsets of the real numbers as well as natural numbers. In a previous paragraph illustrating the basic structure of CU model, for example, $A = (-\infty, +\infty)$. Similarly in Moore and Healy model, $P = (-\infty, +\infty)$, even though a discrete space would have probably better suited the test framework of their example. For coherence and simplicity, we will also assume that the beliefs about actions or performances can be represented by normally distributed random variables.

A *performance function* $p : A^n \rightarrow P$ is a mapping from the n -ary Cartesian product of the action space to the performance space, where $n < \infty$ is the number of decisions that the

²³Precision may be generally thought of as the noise variance, meaning that as precision goes up (variance goes down) the agent will trust her signal more, increasingly neglecting her prior.

agent has to undertake. Hence, the framework consists of an agent facing n tasks with the same action space A , attempting to maximize a quadratic utility function as in equation (1), where the utility-maximizing action is (possibly) different for each task. The agent's aggregate preferences (over the Cartesian product of all actions) can be represented by a utility function that is the sum of the utilities, or by any order-preserving transformation. We formulate the following assumption that does not impact the interpretation of the model, but that buys tractability.

Assumption 1. (*linearity*) Assume that the performance mapping $p : A^n \rightarrow P$ is linear.

In the following result, this is a key assumption, to preserve normality when aggregating the beliefs about actions to beliefs about performance. Linearity does not change the key interpretation of belief aggregation of the performance function and also presents it in a fairly intuitive perspective: when assessing beliefs about a phenomenon, an individual "sums" his beliefs about the single sub-components of it.

Before stating the main result of the Section, I define the auxiliary concept of *consistency*:

Definition 3. Consider a set of priors $\{x_0, x_1, \dots, x_n\}$, signals (or gut feelings) $\{s_0, s_1, \dots, s_n\}$, all with same support X . Consider a mapping $p : X^n \rightarrow X$ and denote any posterior distribution $x_i \mid s_i$ with z_i . Then $\{x_i, s_i\}_{i=1}^n$ are said to be consistent with $\{x_0, s_0\}$ under $p(\cdot)$ if the following hold:

- i) $p(x_1, \dots, x_n) = x_0$ (priors consistency)
- ii) $p(s_1, \dots, s_n) = s_0$ (signals consistency)
- iii) $p(z_1, \dots, z_n) = z_0$ (posterior consistency)

The idea behind consistency is that all the elements in the belief updating process are related through the performance mapping. In principle, it is possible to obtain a given posterior distribution with infinitely many signals and priors. Consistency restrict the focus on the set of priors, signals and posteriors there are related through the performance mapping. Having defined consistency we formulate two additional assumptions:

Assumption 2. (*dimensionality or solvability*) The dimension of the Cartesian product of the action space is $n \geq 3$.

The intuition behind this assumption is that the number of tasks must be large enough as to be able to satisfy all consistency requirements that are defined above and ensure the existence of a solution for the system of equations that it induces.

Assumption 3. (*sufficient noise*) $\sigma_{x_0}^2$ or $\sigma_{s_0}^2$ are "large enough".

This assumption is made more clear in the proof. Essentially, in solving the polynomial system of equations induced by this setting, a lower bound condition arises on some variances, and hence a lower bound must be imposed on either $\sigma_{x_0}^2$ or $\sigma_{s_0}^2$, to ensure the existence of positive solutions for the set of $\{\sigma_{y_i}\}_{i=1}^n$ with $y \in \{x, s\}$. A way to interpret this assumption is that, as the beliefs about the actions are linearly combined into the belief about the

overall performance, the lower the noise of the performance the smaller is the set of beliefs about the actions that may have generated it. If the noise is too small the set is reduced to the empty set.

The following Proposition, which we prove in the [next section](#), characterizes the kind of performance function for which it is possible to represent the overconfidence model à la Moore and Healy (2008) with a cognitive uncertainty model.

Proposition 3. *Consider the case of normally distributed performance prior x_0 , gut feeling s_0 and posterior z_0 . Fix any (linear) performance mapping $p(\cdot) : A^n \subseteq \mathbb{R}^n \rightarrow P \subseteq \mathbb{R}$. Then, there exist infinitely many sets triplets $\{x_t, s_t, z_t\}_{t=1}^n$ of consistent independent priors, signals and induced posteriors.*

This result implies that for any given linear performance function and any belief about performance generated in a framework à la MH, it is possible to find a set of beliefs about single tasks that induce the same belief about performance. In other words, under the linearity restriction on the performance function, it is always possible to specify a CU model that induces the same beliefs on performance. The main implication is that, under the stated assumptions, all predictions generated under the MH model of overconfidence can be generated in a CU framework.

Appendix B Proofs

B.1 Proof of Proposition 1

Given how preferences are defined by (3), (2) may be rewritten as:

$$a^*(x | s) \sim \mathcal{N}(a_i^*, \sigma_{CU,i}^2), \quad (\text{B.1})$$

with $a_i^* = \lambda_i s_i + (1 - \lambda_i) x_{0,i}$, that is a_i^* is the optimal action for agent i and $\lambda_i = \frac{\sigma_{x_i}^2}{\sigma_{x_i}^2 + \sigma_{\varepsilon_i}^2}$ is i 's shrinkage factor.

Now, fix any a_j^* , $j \neq i$, and note that, under preferences described by (3) it holds that

$$u_j(a_j^*, x) \leq u_i(a_i^*, x) \iff |a_j^* - x| \geq |a_i^* - x|,$$

with $a^* = x$.

This, in turn, implies that

$$\begin{aligned} \text{Placement}_i(a_i^*, a_j^*) &= P(u_j(a_j^*, x) \leq u_i(a_i^*, x)) = P(|a_j^* - a^*| \geq |a_i^* - a^*|) = \\ &= \begin{cases} P(a^* > \frac{a_j^* + a_i^*}{2}) & \text{if } a_j^* < a_i^* \\ P(a^* \leq \frac{a_j^* + a_i^*}{2}) & \text{if } a_j^* \geq a_i^* \end{cases} = \begin{cases} 1 - F_{a^*}(\frac{a_i^* + a_j^*}{2}) & \text{if } a_j^* < a_i^* \\ F_{a^*}(\frac{a_i^* + a_j^*}{2}) & \text{if } a_j^* \geq a_i^* \end{cases}, \end{aligned}$$

proving the first point of the proposition.

Note that $F_{a^*}(\cdot)$ is the CDF of the random variable representing agent's i beliefs about the optimal action to undertake. In the remainder of the proof, I assume that agent i does not update her own beliefs after observing a_j^* , that is the beliefs are distributed as per (9). Even without this assumption the structure and the conclusion of the proof would remain unchanged.

For the second point of the proposition there are two cases. I will consider the case of $a_i^* \leq a_j^*$, the other case being specular.

From point one of the Proposition²⁴:

$$\begin{aligned} Placement_i(a_i^*, a_j^*) &= F_{a^*}\left(\frac{a_i^* + a_j^*}{2}\right) = \left[1/2 \operatorname{erf}\left(\frac{x - a_i^*}{\sqrt{2}\sigma_{CU,i}}\right)\right]_{-\infty}^{\frac{a_i^* + a_j^*}{2}} = \\ &= 1/2 + 1/2 \operatorname{erf}\left(\frac{a_j^* - a_i^*}{2^{3/2}\sigma_{CU,i}}\right), \end{aligned}$$

with $\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$, and the last equality following from the fact that $\operatorname{erf}(-\infty) = -1$.

Now, note that

1. $\operatorname{erf}(z)$ is monotonically increasing,
2. $\operatorname{erf}(0) = 0$.

$$\begin{aligned} \text{As } a_i^* < a_j^* &\xrightarrow{(1)} \operatorname{erf}\left(\frac{a_j^* - a_i^*}{2^{3/2}\sigma_{CU,i}}\right) > 0 \xrightarrow{(2)} \frac{\delta \operatorname{erf}\left(\frac{a_j^* - a_i^*}{2^{3/2}\sigma_{CU,i}}\right)}{\delta \sigma_{CU,i}} < 0 \\ \implies \frac{\delta Placement_i(a_i^*, a_j^*)}{\delta \sigma_{CU,i}} &< 0. \end{aligned}$$

Finally, note that, as $Placement_i(\cdot)$ is strictly decreasing in σ_{CU} , then, for any $G_i(\cdot)$ such that $\int Placement_i(a_i^*, z) dG_i(z)$ exists, then also such integral will be strictly decreasing in σ_{CU} . Hence, $Rank_i(a_i^*) = \mathbb{E}_{G_i}[Placement(a_i^*, a_j^*)] = \int Placement_i(a_i^*, z) dG_i(z)$ is strictly decreasing in σ_{CU} . \square

B.2 Proof of Proposition 2

After observing a_j^* , agent i beliefs can be described by equation 4. For notational convenience, let the expectations and the variance of the updated beliefs be $a^* = \alpha a_i^* + (1 - \alpha) a_j^*$, with $\alpha = \frac{\sigma_{-i}^2}{\sigma_{-i}^2 + \sigma_{CU}^2}$, and ς^2 respectively.

As shown in Proposition 2,

$$Placement_i(a_i^*, a_j^*) = \begin{cases} 1 - F_{a_i^*}\left(\frac{a_i^* + a_j^*}{2}\right) & \text{if } a_j^* < a_i^* \\ F_{a_i^*}\left(\frac{a_i^* + a_j^*}{2}\right) & \text{if } a_j^* \geq a_i^*, \end{cases}$$

²⁴For the original treatment and derivation of the error function, see Glaisher (1871).

which implies that, for the statement to be true, it must hold that

$$\begin{cases} \frac{\delta F_{a^*}(\frac{a^*+a_j^*}{2})}{\delta \sigma_{-i}} < 0 & \text{if } a_j^* < a^* \\ \frac{\delta F_{a^*}(\frac{a^*+a_j^*}{2})}{\delta \sigma_{-i}} > 0 & \text{if } a_j^* > a^*. \end{cases}$$

Now, as in the previous Proposition proof, recall that for a normal distribution with mean μ and variance σ^2 it holds that

$$F_{\mu,\sigma} = 1/2(1 + \text{erf}(\frac{x-\mu}{\sqrt{2}\sigma})).$$

Substituting $\mu = a^*$ and $\sigma^2 = \varsigma^2$, leads to, after some simplification:

$$F_{a^*,\varsigma} = 1/2 \left[1 + \text{erf} \left(\frac{\alpha^{1/2}(a_j^* - a_i^*)}{2^{3/2}\sigma_{CU}} \right) \right].$$

As the error function is monotonically increasing, the sign of the CdF derivative with respect to σ_{-i} is the same as the sign of $\text{erf}(\cdot)$ argument.

Differentiating $\frac{\alpha^{1/2}(a_j^* - a_i^*)}{2^{3/2}\sigma_{CU}}$ with respect to σ_{-i} leads to:

$$\left[(\sigma_{CU}^2 + \sigma_{-i}^2)^{1/2} - \frac{\sigma_{-i}^2}{(\sigma_{CU}^2 + \sigma_{-i}^2)^{1/2}} \right] \frac{(a_j^* - a_i^*)}{2^{3/2}\sigma_{CU}},$$

which is strictly negative for $a_j^* < a_i^*$ and strictly positive for $a_j^* > a_i^*$ (the placement function is non-differentiable at $a_j^* = a_i^*$). \square

B.3 Proof of Proposition 3

First, note that, under linearity of the performance function²⁵ and normality, the consistency constraints can be rewritten as:

$$\left\{ \begin{array}{l} \sum_{i=1}^n \alpha_i \mu_{x_i} = \mu_{x_0} \\ \sum_{i=1}^n \alpha_i^2 \sigma_{x_i}^2 = \sigma_{x_0}^2 \end{array} \right\} \text{priors consistency}$$

$$\left\{ \begin{array}{l} \sum_{i=1}^n \alpha_i \mu_{s_i} = \mu_{s_0} \\ \sum_{i=1}^n \alpha_i^2 \sigma_{s_i}^2 = \sigma_{s_0}^2 \end{array} \right\} \text{signal consistency}$$

$$\left\{ \begin{array}{l} \sum_{i=1}^n \alpha_i \mu_{z_i} = \mu_{z_0} \\ \sum_{i=1}^n \alpha_i^2 \sigma_{z_i}^2 = \sigma_{z_0}^2 \end{array} \right\} \text{posterior consistency}$$

Moreover, since z_i is the Bayesian posterior of an agent holding x_i as a prior and observing the signal s_i , it holds for all i that:

$$\mu_{z_i} = \frac{\sigma_{s_i}^2}{\sigma_{s_i}^2 + \sigma_{x_i}^2} \mu_{x_i} + \frac{\sigma_{x_i}^2}{\sigma_{s_i}^2 + \sigma_{x_i}^2} \mu_{s_i},$$

$$\sigma_{z_i}^2 = \frac{\sigma_{s_i}^2 \sigma_{x_i}^2}{\sigma_{s_i}^2 + \sigma_{x_i}^2}.$$

²⁵That is $p(x_1, \dots, x_n) = \sum_{i=1}^n \alpha_i x_i$ for some $\alpha_1, \dots, \alpha_n$.

Hence, the posterior consistency conditions can be rewritten as:

$$\begin{cases} \sum_{i=1}^n \frac{\alpha_i}{\sigma_{x_i}^2 + \sigma_{s_i}^2} (\sigma_{s_i}^2 \mu_{x_i} + \sigma_{x_i}^2 \mu_{s_i}) = C, \\ \sum_{i=1}^n \frac{\alpha_i^2 \sigma_{x_i}^2 \sigma_{s_i}^2}{\sigma_{x_i}^2 + \sigma_{s_i}^2} = D, \end{cases}$$

with $C = \frac{\sigma_{s_0}^2}{\sigma_{s_0}^2 + \sigma_{x_0}^2} \mu_{x_0} + \frac{\sigma_{x_0}^2}{\sigma_{s_0}^2 + \sigma_{x_0}^2} \mu_{s_0}$ and $D = \sigma_{x_0}^2 = \frac{\sigma_{s_0}^2 \sigma_{x_0}^2}{\sigma_{s_0}^2 + \sigma_{x_0}^2}$.

Hence, proving the statement is equivalent to prove the existence of a set $\{\mu_{x_i}, \sigma_{s_i}^2, \mu_{s_i}, \sigma_{s_i}^2\}_{i=1}^n$ such that all the consistency constraints hold. In other words, with $n \geq 3$, the aim is to prove the existence of a solution for an underdetermined system of polynomial equations, with the additional constraints that $\sigma_{x_i}^2, \sigma_{s_i}^2 > 0$ for all i .

Without loss of generality solve the first four constraints with respect to $i = 1$, which leads to:

$$\mu_{y_1} = \mu_{y_0} - \sum_{i=2}^n \alpha_i \mu_{y_i}, \quad (\text{B.2})$$

$$\sigma_{y_1}^2 = \sigma_{y_0}^2 - \sum_{i=2}^n \alpha_i^2 \sigma_{y_i}^2, \quad (\text{B.3})$$

for $y \in \{x, s\}$. Clearly, these four conditions have infinitely many solutions. Substituting into the first posterior consistency condition and isolating (without loss of generality) μ_{x_2} , after some algebra, leads to:

$$\begin{aligned} \mu_{x_2} = & \left[\frac{(\sigma_{x_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{x_i}^2)(\mu_{x_0} - \sum_{i=3}^n \alpha_i \mu_{x_i})}{(\sigma_{s_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{s_i}^2) + (\sigma_{x_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{x_i}^2)} \right. \\ & \left. + \frac{(\sigma_{s_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{s_i}^2)(\mu_{s_0} - \sum_{i=2}^n \alpha_i \mu_{s_i})}{(\sigma_{s_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{s_i}^2) + (\sigma_{x_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{x_i}^2)} - C \right] \eta, \end{aligned}$$

$$\text{with } \eta = \left(\frac{\alpha_2(\sigma_{x_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{x_i}^2)}{(\sigma_{s_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{s_i}^2) + (\sigma_{x_0} - \sum_{i=2}^n \alpha_i^2 \sigma_{x_i}^2)} - \frac{\alpha_2 \sigma_{s_2}^2}{\sigma_{s_2}^2 + \sigma_{x_2}^2} \right)^{-1}.$$

Up to now the only constraint is for the variances to be small enough in order for (11) to be positive, for which there would be infinitely many solutions.

Solving the last posterior constraint as a function of one of the variances (w.l.o.g. $\sigma_{x_2}^2$) leads, after quite some tedious algebra, to a quadratic equation in $\sigma_{x_2}^2$:

$$a\sigma_{x_2}^4 + b\sigma_{x_2}^2 + c = 0, \text{ with the coefficients being}$$

$$\begin{aligned} a &= \alpha_2^2 \sigma_{s_1}^2 (\sigma_{s_2}^2 - \alpha_2^2 \beta \sigma_{s_1}^2 - 1), \\ b &= (\sigma_{x_0}^2 - \sum_{i=3}^n \alpha_i^2 \sigma_{x_i}^2) \sigma_{s_1}^2 + \alpha_2^2 [\sigma_{s_2}^2 \gamma - \sigma_{s_1}^2 \sigma_{s_2}^2 + \beta \gamma \sigma_{s_1}^2 + \beta \gamma], \\ c &= \gamma (\sigma_{s_2}^2 - \beta (\sigma_{x_0}^2 - \sum_{i=3}^n \alpha_i^2 \sigma_{x_i}^2)), \end{aligned}$$

with $\beta = D - \sum_{i=3}^n \alpha_i^2 \frac{\sigma_{x_i}^2 \sigma_{s_i}^2}{\sigma_{x_i}^2 + \sigma_{s_i}^2}$ and $\gamma = (\sigma_{x_0}^2 - \sum_{i=3}^n \alpha_i^2 \sigma_{x_i}^2)(\sigma_{s_0}^2 - \sum_{i=2}^n \alpha_i^2 \sigma_{s_i}^2)$. Imposing $\sigma_{s_2}^2 > 1 + \alpha_2^2 \beta \sigma_{s_1}^2$ and $\sigma_{s_2}^2 < \beta (\sigma_{x_0}^2 - \sum_{i=3}^n \alpha_i^2 \sigma_{x_i}^2)$ a positive solution exists.

Hence, the algorithm to find a solution in this case would be:

$$\begin{aligned} & (\sigma_{x_3}^2, \dots, \sigma_{x_n}^2), (\sigma_{s_2}^2, \dots, \sigma_{s_n}^2) \rightarrow \sigma_{s_2}^2 \rightarrow \sigma_{x_2}^2 \rightarrow \\ & \rightarrow (\mu_{x_2}, \dots, \mu_{x_n}), (\mu_{s_2}, \dots, \mu_{s_n}) \rightarrow \mu_{x_1}, \mu_{s_1}. \square \end{aligned}$$

Appendix C Online Experiment Implementation

C.1 Problem Description

The figures below show several screenshots from the experiment. Figure C.1 shows the introductory instructions for participants, while Figure C.2 a more detailed description of the belief updating task (the graphical illustration of the task is taken from Enke and Graeber (2021)). Figures C.3 and C.4 show how the structure of the experiment is illustrated to participants.

Instructions

Please take your time to read the instructions carefully. Your understanding of the instructions will be tested later.

The study is split into two parts.

In this first part, you will be asked to complete **3 similar tasks**.

These tasks can be divided into parts where you **receive relevant information** and parts where you **make decisions**. Importantly, **some** of these decisions may **determine** your **bonus payment**. However, which decision will determine your bonus payment will be picked randomly.

For this reason, you should always **provide your best guess**, to increase the chances that you will receive the bonus payment.

Figure C.1: Experiment Induction 1

Task Description

There are **two bags**, bag A and bag B. Each of the bags contains **100 balls**, which may be blue or red. The number of blue and red balls in each bag will vary in each task, but you will be given information on **how many blue and red balls each bag contains**.

Moreover, there is a deck of **100 cards**. Each card may have "A" or "B" written on it.

The number of "A" cards and "B" cards in the deck will vary in each task, but you will be given information on **how many "A" and "B" cards are present in the deck**.

The computer will randomly select a bag, drawing a card from the deck. You will **not observe which bag was selected** this way. Then, **one or more balls** will be drawn from the selected bag. After observing the balls drawn, you will have to provide a **probabilistic guess** on which bag was selected.

The picture below illustrates the process graphically.

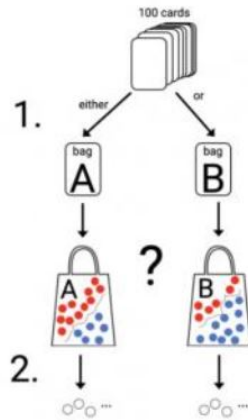


Figure C.2: Experiment Induction 2

Timeline

- 1) You are told how many “A” and “B” cards the deck contains and how many blue and red balls bags A and B contain.
- 2) The computer randomly draws a card from the deck. Each card may be drawn with the same probability.
- 3) If the card is an “A card”, the computer draws one or more balls from bag A and if the card is a “B card” the computer draws one or more balls from bag B. The balls are drawn **with replacement**, meaning that every time a ball of any color is drawn it gets replaced in the bag. Hence, the probability of a red or blue ball being drawn depends only on the bag it is drawn from, but not on previous draws.
- 4) You are **shown** how many blue and red balls were drawn, but not from which bag those were drawn.
- 5) Your task is then to provide a probability between 0% and 100% that the balls were drawn from bag A. Keep in mind that the probability of the sure event is 100% and the probability of the impossible event is 0%. Note that the probability of the balls being drawn from bag B corresponds to 100 minus the probability that the balls were drawn from A (**this decision may determine your bonus payment**).
- 6) After you provide your answer, you will be asked to provide your expectations on your relative performance in this task. For example, you may expect to be in the best 5% or in the worst 10% of the participants, where the performance is measured on the base of how close your answer is to the correct one. To answer, imagine you are competing with other 99 participants that have the same information as you have and try to guess your placement from 1 (first) to 100 (last) (**this decision may determine your bonus payment**).

Figure C.3: Experiment Timeline 1

7) You will be asked **how certain** you are about your answer. This point will be explained thoroughly in the next page.

8) You will be shown an **answer from another participant to the exact same task**. This means that the other participant **observed the same ball(s) draw**, with the deck having the same number of **A and B cards** and the A and B bags having the same number of blue and red balls.

9) You will be asked about **how certain you think** the other participant was about the answer you were shown in the previous point.

10) You will be asked **your guess of the probability of your answer being better than the observed answer**. An answer is **better if it is closer to the correct answer**. (this decision may determine your bonus payment).

11) You are given the chance to change your answer after observing the other participant's answer. You may also leave the answer unchanged (this decision may determine your bonus payment).

Please note:

- In each new task the **number of red and blue balls in the bags and the drawn card change**: please think about each of the tasks **independently**.

- As will be clarified in what follows, there is **no real correct answer** to the decision in points 7 and 9. This is why the bonus payment will not depend on that, **so try to answer truthfully**.

Figure C.4: Experiment Timeline 2

C.2 Comprehension Check

After receiving the instructions and the task description, the participants undergo four comprehension questions to assess their understanding of the information provided to them. If one of the question is answered incorrectly the participant is redirected to an exit screen and may not proceed with the rest of the experiment.

Below two screenshots of the questions are provided.

Comprehension questions

The following question test your understanding of the instructions read so far.

Please note that **falling to answer any of the following correctly**, will result in **not being able to proceed further with the study** and **not earning any bonus**.

Which of the following statement is correct?

The number of "A cards" and "B cards" in the deck vary in each task

The probability of bag A being picked in each task is always the same

Which of the following statement is correct?

When asked about the certainty of my guess there is a correct answer

There is no point in giving a strategic answer when asked about the certainty of my guess

Figure C.5: Comprehension Questions 1

What is the best way to maximize the chances of receiving the bonus?

Providing any guess, the bonus is received randomly

Providing my best guess for the asked probabilities

What is the best placement among the following?

Top 5%

Top 15%

Top 30%

If you answered the questions correctly, pressing the **right arrow** you will first proceed to an example task, where you will be shown the procedure without being required to provide any answer, and after to the actual task.

Figure C.6: Comprehension Questions 2

C.3 Belief Updating Tasks

Figures C.7 and C.8 show the screens of the belief updating task respectively before and after the participants observe the ball(s) draw.

First, the participants have the chance to observe the parameters of the problem and then can trigger the computer draw pressing the right arrow. Afterwards the signal is drawn and they are shown the screen in C.8 and given the chance to answer. Note that they still have access to the problem parameters scrolling the page upwards.

This task consists in **guessing the probability** of Bags A and B being chosen by the computer.

Number of "**bag A**" cards: 30

Number of "**bag B**" cards: 70

Bag A contains 70 blue balls and 30 red balls.

Bag B contains 30 blue balls and 70 red balls.

Next:

1. The computer **randomly selects one bag** by drawing a card from the deck.
2. The computer **draws 1 ball(s)** from the secretly selected bag. Click the arrow for the computer to perform the draw:



Figure C.7: Belief Updating Task Pre Signal

The computer **drew 1 blue ball**.

Please write down **your guess** for the **probability** (between 0 and 100) of **each bag being chosen**.

Probability of **bag A**:

Probability of **bag B**:

Total

Please provide you **guess about your placement** in this **specific task** (between 1 and 100):



Figure C.8: Belief Updating Task Post Signal

C.4 Scoring Rule

Subjects are informed that a subset of the tasks will be randomly picked to determine their final earnings, with earnings for each of the chosen tasks being determined according to a quadratic, incentive-compatible scoring rule (Hossain and Okui, 2013).

The rule is implemented in a slightly different way, depending on the bonus-relevant task that is randomly drawn, being it a placement task or a belief updating task. For both scenarios, the computer draws a random number $n \in \{1, \dots, 2500\}$, where the probability assigned to each draw is the same. Afterwards, depending on the task, the bonus is assigned if the following is true:

$$\begin{cases} P(A)^2 > n, & \text{if } A \\ P(A)^2 \leq n, & \text{if } B, \end{cases} \quad (\text{C.1})$$

$$\begin{cases} Placement^2 > n, & \text{if } Placement > 50 \\ Placement^2 \leq n, & \text{if } Placement \leq 50, \end{cases} \quad (\text{C.2})$$

with $P(A)$ being the probability the participant assigned to bag A and A (B) being true if the bag actually selected by the computer was A (B).

Note that the CU elicitation would not be incentivized in either case, as illustrated in Figure 4. For more details about how the scoring rule is explained to participants, see Figure C.9.

Bonus payment

*You may receive a bonus of 1.5\$ for this part. Whether or not you receive the bonus depends on the choices you make in the **bonus relevant decision parts**, indicated in the previous screen. One of the 4 relevant decisions will be randomly picked by the computer, with equal probability. The “goodness” of that decision will determine whether you will receive the prize or not, according to the following rules.*

*If the decision picked by the computer is **one of the placement decisions** (points 5 and 10 in the timeline), the computer will randomly draw a number n between 0 and 2500 (each number having the same probability of being drawn).*

*For the decision in point 5, if your actual placement is in the top half of participants, or you will receive the bonus if the square of your stated placement is **lower than n** . On the other hand, if your actual placement is in the lower half of participants, you will receive the bonus if your stated placement is **larger than n** .*

*Similarly, for the decision in point 10, if you performed better than the other participant, you will receive the bonus if the square of your stated probability is **larger than n** and, in the other case, if your stated probability is **lower than n** .*

*If the decision picked by the computer is **one of the probability guesses** (points 6 and 9 in the timeline), the computer will randomly draw a number n between 0 and 2500 (each number having the same probability of being drawn).*

*If in the task the drawn bag was A, if the square of the probability (in percent) you assigned to A is **larger than n** , you receive the bonus. However, if B was drawn, you receive the bonus if the square of the probability assigned to A is **lower than n** .*

All these rules mean that you have the incentive to provide your best guess of your placement or the probabilities, since your best guess improves the chances of receiving the bonus.

Figure C.9: Scoring Rule Description to Participants

C.5 Preliminary Study

Some differences characterize the preliminary study, compared to the main design. First, in the preliminary study, each task stops after CU elicitation. Also, all choices are non-compound choices. Moreover, each participant must complete 7, instead of 6, tasks to complete the probability guessing section of the study. The additional task is an attention check.

Participants are informed, within the instructions, that the tasks may contain attention checks and that failing an attention check would result in being discarded from the study. The attention check is identical for each participant and consists of a guessing task with a particular parameter specification and signal. An urn contains 99 blue balls and 1 red ball and B urn contains 99 red balls and 1 blue ball. The participant is informed that the 2 balls are drawn without replacement. The signal is always of 2 blue balls, implying that the

probability of A being the correct urn is 1. If a participant failed to answer this correctly, the observation was excluded from the sample. Figure C.10 shows what the participants saw when undergoing the attention check.

Concerning attention check, a final difference of the preliminary study is that the experimental instructions explain to participants the difference between the draw with or without replacement. This is not necessary for the main study, as all draws are performed with replacement.

This task consists in **guessing the probability** of Bags A and B being chosen by the computer.

The draw is performed **without replacement**.

Number of **"bag A"** cards: **50**

Number of **"bag B"** cards: **50**

Bag A contains **99 blue balls** and **1 red balls**.

Bag B contains **1 blue balls** and **99 red balls**.

Next:

1. The computer **randomly selects one bag** by drawing a card from the deck.
2. The computer **draws 2 ball(s)** from the secretly selected bag. Click the arrow for the computer to perform the draw:

The computer **drew 2 blue balls**.

Figure C.10: Preliminary Study Attention Check

Appendix D Code Implementation

Figures D.1 and D.2 show how the balls draws are implemented for the experiment, using JavaScript.

The "urn_draw" and "signal_draw" variables are both draws from a standard uniform. The other three variables are already determined outside the script. Afterward, based on the uniform draws and on the given parameters, the urn draw (A or B) is determined, which in turn establishes how to use the parameters and finally which will be the ball(s) draw shown to the participant. For example, in the first lines of the set of *if* conditions, if the uniform draw is below the parameter "p_a" (an event that has probability $P(A)$) then the urn is set to A. Then, depending on "N" and on "p_a", the other uniform draw ("signal_draw") determines the signal.

```

var urn_draw = Math.random();
var signal_draw = Math.random();
var p_a = parseInt("${lm://Field/1}"/100;
var b_balls_a = parseInt("${lm://Field/2}"/100;
var N_draw = parseInt("${e://Filed/N}");

if (urn_draw <= p_a)
{
  var drawn_urn = "A";
  if (N_draw == 1)
  {
    if (signal_draw <= b_balls_a)
    {
      var message = "1 blue ball";
    }
    if (signal_draw > b_balls_a)
    {
      var message = "1 red ball";
    }
  }

}
if (N_draw == 2)
{
  if (signal_draw <= (b_balls_a)*(b_balls_a))
  {
    var message = "2 blue balls";
  }
  if (signal_draw > (1 - (1 - b_balls_a)*(1 - b_balls_a)))
  {
    var message = "2 red balls";
  }
  if (signal_draw > (b_balls_a)*(b_balls_a) && signal_draw <= (1 - (1 - b_balls_a)*(1 - b_balls_a)))
  {
    var message = "1 blue ball and 1 red ball";
  }
}
}

```

Figure D.1: Code Implementation of Randomness 1

```

if (urn_draw > p_a)
{
  var drawn_urn = "B";
  if (N_draw == 1)
  {
    if (signal_draw > b_balls_a)
    {
      var message = "1 blue ball";
    }
    if (signal_draw <= b_balls_a)
    {
      var message = "1 red ball";
    }
  }

  if (N_draw == 2)
  {
    if (signal_draw <= (b_balls_a)*(b_balls_a))
    {
      var message = "2 red balls";
    }
    if (signal_draw > (1 - (1 - b_balls_a)*(1 - b_balls_a)))
    {
      var message = "2 blue balls";
    }
    if (signal_draw > (b_balls_a)*(b_balls_a) && signal_draw <= (1 - (1 - b_balls_a)*(1 - b_balls_a)))
    {
      var message = "1 blue ball and 1 red ball";
    }
  }
}
}

```

Figure D.2: Code Implementation of Randomness 2

Appendix E Results: Additional Details

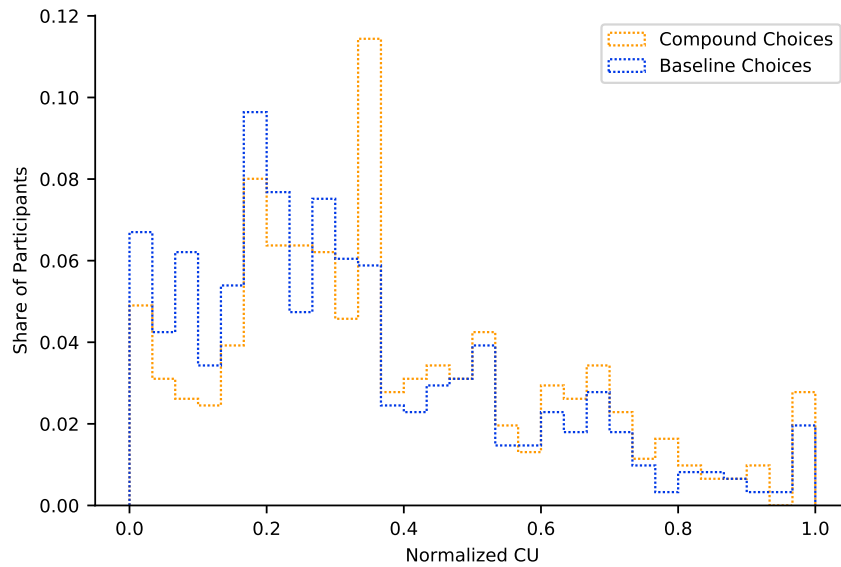


Figure E.1: CU distribution for compound/baseline choices

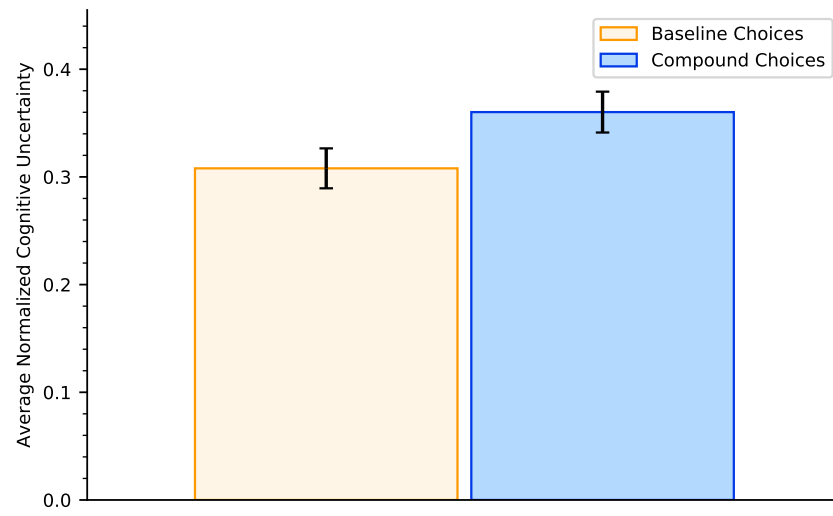


Figure E.2: T-test on CU means for compound/baseline choices

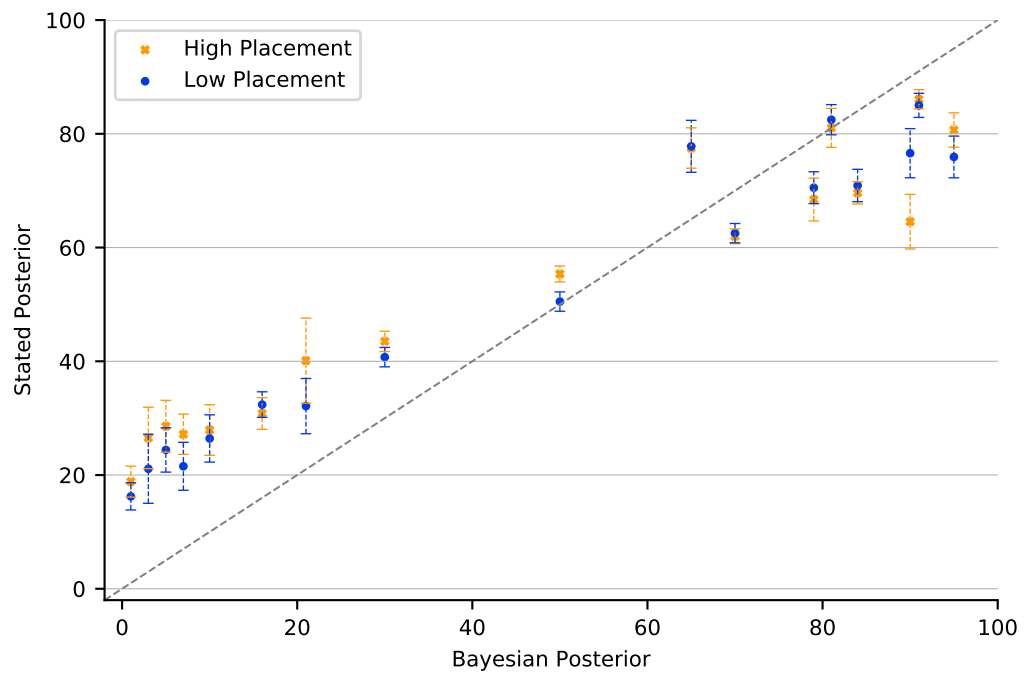


Figure E.3: Average reported posteriors for high/low placements. Groups are determined using median placement as a threshold. The error bars represent the standard error of the means. Bayesian posteriors are rounded to the nearest integer. We show only buckets that contains at least 20 observations.