Working Paper

2025/24/DSC

Individual Differences in Overconfidence and their Psychological Bases

M. Asher Lawson
INSEAD, asher.lawson@insead.edu
Corresponding author

Richard P. Larrick
Duke University, larrick@duke.edu

Jack B. Soll
Duke University, jsoll@duke.edu

This version of the manuscript was submitted February 9, 2025

Judgment and decision making research has typically not explicitly considered whether there are individual differences in overconfidence. In cases where correlations among individual measures are reported, the strength of associations tends to be modest. Despite this, we show that when analyzing measures of overconfidence from an individual differences perspective, there is substantial evidence for trait-level overconfidence—which we call 'core overconfidence'—that is stable over both short and long-time periods. Further, overconfidence has primarily been studied in one of three forms: overestimation, overplacement, or overprecision. We find that measures across these forms are consistently related when each form is measured reliably. These results are shown to be robust over a broad array of design and analytical choices. We also find evidence that differences in both self-enhancement and reflective cognitive processes appear to contribute to this stability: a range of individual difference measures (e.g., narcissism, actively open-minded thinking) significantly correlate with core overconfidence, though these relationships are generally explained distinctly by either a relationship with confidence or accuracy. In additional studies, we show that individual differences in overconfidence are robust when minimizing the correlations in accuracy across tasks, demonstrating that they are not specific to narrow self-views in a domain. Overall, by studying the same respondents' behaviors across domains, forms, and time periods, we find strong evidence for stable, individual differences in overconfidence that are interpretable in terms of common motivated and cognitive processes.

Keywords: Overconfidence; Calibration; Individual Differences.
JEL Classification: D90; D91

We thank Sandra Matz, Hemant Kakkar, Steven Spiller, and attendees of IESE's Managerial Decision Sciences departmental seminar, John's Hopkins's Marketing departmental seminar, and INSEAD's School Wide seminar series for comments on this project. Parts of this work were presented at the 2023 meeting of the European Association for Decision Making and in poster format at the 2023 meeting of the Society for Judgment and Decision Making. This manuscript has been available on SSRN since September 2023. This project was based on the doctoral dissertation of Asher Lawson (Lawson, 2022). We thank Jessica Paek for her assistance in data collection for Study 1a (December 2019).

Author Contributions Statement

M. Asher Lawson: Conceptualization, Methodology, Formal Analysis, Investigation, Data Curation, Writing – Original Draft, Writing – Review & Editing, Visualization
Richard P. Larrick: Conceptualization, Methodology, Investigation, Writing – Review & Editing, Visualization, Funding acquisition
Jack B. Soll: Conceptualization, Methodology, Formal Analysis, Investigation, Writing – Review & Editing, Visualization, Funding acquisition

Data availability

All materials, data, and code associated with this project are available here: https://tinyurl.com/2vsjpn63.

Introduction

Overconfidence is a widely studied phenomenon across academic disciplines ranging from psychology (Oskamp, 1965), to economics (Galasso & Simcoe, 2011), to strategy (Camerer & Lovallo, 1999), accounting (Malmendier & Tate, 2005a), and law (Goodman-Delahunty et al., 2010). To date, different literatures have made different assumptions and reached different conclusions regarding whether an individual who is more overconfident on one task is likely to be more overconfident on another. In this research, we attempt to provide a definitive answer as to whether overconfidence can be regarded as a stable trait such that some individuals are systematically more overconfident than others.

Overconfidence research emerged in the fields of decision analysis and judgment and decision making (JDM) in the 1970s (Fischhoff et al., 1977; Raiffa, 1968). Within JDM, much of the focus has been on the average degree to which people are overconfident in a task and how this varies across tasks that study different *forms* of overconfidence: overestimation, overplacement, and overprecision (Moore & Healy, 2008). The JDM tradition has largely not explored overconfidence as an individual difference. When it has, some have concluded that there is little evidence for an individual difference (Moore & Dev, 2018), whereas others have found consistency in overconfidence within participants in a single time period (Klayman et al., 1999), and that individual measures of overconfidence correlate with narcissism (e.g., Ames & Kammrath, 2004).

In related work, research on overclaiming —which appears to be psychologically similar to overconfidence—has shown that self-enhancement of knowledge is a stable individual difference (Paulhus et al., 2003). There is also a body of related work on metacognition, one aspect of which is the ability to track one's own mental performance (e.g., memory) with feelings of confidence or knowing (Koriat, 1993; Nelson & Narens, 1990). Recent work in this tradition has identified differences in confidence formation that affect this monitoring or metacognitive ability, and researchers have linked these differences to brain structure (Fleming, 2024). Whereas work on metacognitive ability has focused on the monitoring of mental performance, our focus here is on the estimation of the level of performance (we comment further on this distinction in the General Discussion). Outside of psychology, research in economics, finance, and accounting—which we collectively refer to as the economics tradition—typically treats measures of overconfidence as a stable property of individuals. For example, measures of executive overconfidence are used to predict firm decisions in different time periods (Ben-David et al., 2007; Malmendier & Tate, 2005a).

There is, moreover, debate regarding the underlying psychological processes that produce overconfidence. On the cognitive side, overextrapolation from limited information (Erev et al., 1994; Klayman et al., 2006), overreliance on internal cues relative to external cues (Kahneman & Lovallo, 1993), and confirmatory thinking (Koriat et al., 1980; Soll & Klayman, 2004) have all been proposed as mechanisms. In a related test, Logg et al. (2018) found that motivation to possess a target trait inflated overconfidence, though only when rating vague criteria. Combined with overconfidence's relation with narcissism, this points to an additional, self-enhancement related aspect of overconfident beliefs. Despite these proposed explanations, there is no consensus regarding the psychological bases of overconfident beliefs.

In the present research, we conduct a systematic investigation of measures of overconfidence across domains, forms, and time to contribute to this discourse, finding evidence for individual differences in overconfidence across measures that are interpretable in terms of common psychological processes.

Measures of Overconfidence in Psychological Research

Overconfidence can be defined as the fact that people "tend to express confidence in their judgments that exceeds the accuracy of those judgments" (Klayman et al., 1999, p. 217).

Yet given that there are many different expressions of confidence and definitions of accuracy, this scope is quite broad. The most commonly used taxonomy (Moore & Healy, 2008) distinguishes between three forms—overestimation, overplacement, and overprecision. These forms differ both in their elicitation (i.e., how they are measured) and their definition (i.e., what they purport to measure). A substantial amount of research in psychology concerning overconfidence has been conducted using one or more of these three forms.

The Three Main Forms of Overconfidence

Overestimation is described as overestimating "one's actual ability, performance, level of control, or chance of success" (Moore & Healy, 2008, p.502). The classical form of overestimation (e.g., Erev et al., 1994) is captured by differences between people's subjective confidence that an outcome will occur (e.g., that they will answer a question correctly) and the actual outcome (e.g., whether they did answer the question correctly). For example, if someone indicated for each of five questions that they had an 80% chance of being correct, but only 3/5 (60%) of their answers were in fact correct, they would be overconfident by 20 percentage points. Others later proposed that confidence could be elicited by asking participants at the end of a quiz how many questions they answered correctly, and that this could be compared with the true performance on the quiz (Gigerenzer et al., 1991).

Overplacement refers to people's evaluations of their abilities and performance *relative* to others. This form of overconfidence, which is sometimes studied under the name of the 'better-than-average-effect', occurs when people believe themselves to be higher up in an outcome distribution than they are. For example, if someone tells you they are a better driver than 90% of all motorcar drivers but are in fact only a better driver than 70% of the population, this is an overplacement of 20 percentage points of their driving ability (Larrick et al., 2007).

Finally, *overprecision* is a form of overconfidence where people assign too much probability to the outcomes that they view most likely—thus underestimating the uncertainty in what could happen and exaggerating the precision of their own knowledge (Soll et al., 2023). This occurs if people provide overly narrow confidence intervals. Consider the following question: *What do you think is the population of Poland*? In tests of overprecision, participants might provide, for example, a low and high estimate such that they are 80% confident that the true answer lies inside the interval. Typically, subjective 80% confidence intervals contain the true values less than 80% of the time, which indicates overprecision. Having defined the forms of overconfidence, we next turn to describing extant evidence regarding how people behave in tasks of each form.

Levels of Overconfidence Across Forms

It is well established that the different forms of overconfidence can differ substantially in the magnitude of overconfidence observed. For example, studies of overprecision show that 90% subjective confidence intervals contain the truth only about 50% of the time—an impressive gap of 40 percentage points between confidence and accuracy (Klayman et al., 1999; Winman et al., 2004). In contrast, overestimation of performance in forced choice questions between two alternatives is typically on the order of 5-10 percentage points (Klayman et al., 1999), and it is not uncommon to find tasks in which people are underconfident (Moore, 2020). It has also been highlighted that task difficulty—a situational factor—can interact with the form of overconfidence in counterintuitive ways. For instance, more difficult tasks are simultaneously associated with greater overestimation and reduced overplacement (Larrick et al., 2007). These countervailing effects are in part due to statistical processes such as sampling error (Klayman et al., 1999), and in part due to the inferences that people draw based on feelings of difficulty (Burson et al., 2006; Moore & Small, 2007).

The different degrees of over and under-confidence observed for different elicitations have been offered as evidence against individual differences in overconfidence. Moore and Dev (2018) propose that situational factors swamp the effects of individual differences, which they assert may not exist at all. We propose that task variation and individual differences are not contradictory. Mean differences across tasks can occur even if people's rank orderings within those tasks are stable. To test this proposal, one needs to examine whether a person's overconfidence in one task correlates with their overconfidence in a separate task. In the next section, we define the different types of correlations among measures of overconfidence that have been reported in the literature and summarize the evidence found using each type.

Evidence for Individual Differences in Overconfidence

There are four main types of correlations that are reported in the literature: split sample, across domain, across form, and test-retest. We define each of these next before summarizing the associated evidence. To foreshadow the findings across tests, there are many cases of positive correlations among measures of overconfidence, but they are typically of modest strength, with many around $r = .3$.

Split Sample

For this type of correlation, questions from a single, overarching test are randomly divided into multiple subsets (typically two) of questions. If there are individual differences, overconfidence in each of the two subsets should be highly correlated. It generally is, with correlations in the order of $r = .6$ to .9 (Klayman et al., 1999; West & Stanovich, 1997).

For example, consider Experiment 1 from Klayman et al. (1999): Participants answered 120 binary questions, either answering 40 questions in each of 3 domains or 10 questions in each of 12 domains. These 120 questions were provided to each participant in a different random order, before being split into two subsets of 60 questions based on systematically sampling every other question that each participant answered. As a result, the subsets were different for each participant. Overconfidence was then computed on each of the two subsets, subtracting the percentage of questions that participants answered correctly from participants' average confidence, indicated on a scale from 50-100%. The correlation between these two measures of overconfidence was $r = .66$. In this example, the measures of overconfidence were computed across multiple domains but within a single form, but one could also compute split-sample correlations in a single domain or between forms.

Across Domain

In some studies, participants answer questions in different domains of knowledge (e.g., history, sports, geography, etc.). When there are enough questions in each domain, one can test for cross-domain stability. In empirical studies in which the form of overconfidence is held constant and there are measures from multiple domains, these correlations are often in the range of $r = .3$ to .4 (Glaser et al., 2013; Glaser & Weber, 2007; Hilton et al., 2011; Klayman et al., 1999; Pallier et al., 2002). For example, Hilton et al. (2011) found that overprecision in two sets of trivia questions was correlated at $r = .37$.

Across Form

Some studies have reported correlations across different forms of overconfidence. Using a split-sample methodology such that each form was computed on a different set of 10 questions, Larrick et al. (2007) found that overplacement and overestimation were correlated with one another on average at around $r = .25$. Studies in the economics tradition have often failed to find an association between measures of overconfidence across forms, but, notably, they generally compute overplacement as the difference between a respondent's estimated performance and average performance in the sample rather than comparing their estimated performance to their individual performance, so overconfidence is incompletely measured (e.g., Deaves et al., 2009; Glaser & Weber, 2007; Hilton et al., 2011; Merkle, 2017; Oberlechner & Osler, 2012).

Test-retest

If there are individual differences in overconfidence, one immediate implication is that overconfidence in the *same or similar tests* should be stable across time. Such correlations have been estimated for tests administered from days to weeks apart (e.g., Binnendyk & Pennycook, 2024; Stankov & Crawford, 1996). For example, Binnendyk and Pennycook (2024) find test-retest reliabilities of r = .35 to .42 for three individual measures of overconfidence over a 15-day period. In the present research, we will inspect this reliability over a period of years and across forms, further building on Boutros et al. (2020), who found evidence for stability over a 12-year period using data from a survey of Chief Financial Officers (CFOs) using a measure of overprecision.

Aggregation and Empirical Approach

Overall, across extant work, there are many documented associations between measures of overconfidence, but given that most relationships were modest in strength (e.g., r = .25; Larrick et al., 2007) and the risk of publication bias, one might question whether such relationships should be considered as evidence of individual differences in overconfidence. In a notable exception, Klayman et al. (1999) report much stronger correlations among measures of overconfidence (e.g., r = .66). What can explain this difference? Critically, Klayman et al. (1999) used many more items to evaluate accuracy than most other tests—for example, their split-sample analysis compared subsets containing *60 questions* each, as opposed to say, five—meaning that the measurement was likely to be more reliable, thus allowing stronger relationships to emerge. In the present research, we build on this approach to yield further insights regarding the nature of overconfidence.

Generally, researchers that correlate different instantiations of overconfidence compare, for example, measures from two short tests constructed from 10 questions each. Reported correlations are typically between $r = 0$ to .5. Despite using multiple questions, each test only provides a single measure of overconfidence (e.g., estimated number correct less actual score, estimated percentile less actual percentile, etc.) and hence serves as one item when measuring trait-level overconfidence. We propose that to create a reliable measure of overconfidence researchers should average over multiple such items. In our studies we test whether sets of multiple items measuring overconfidence exceed the classic threshold of .7 for Cronbach's alpha (Cronbach, 1951). This would constitute strong evidence against the null model that there is no trait, only noise (Cortina, 1993). To put things in context, if different measures of overconfidence correlated at an average of $r = .2$ with each other, a Cronbach's alpha of .71 could be achieved with a test containing ten items (Cronbach, 1951). This is a length of test that one typically might see for measures of others individual differences (e.g., the HEXACO-60 personality measure (Ashton & Lee, 2009) uses 10 items per facet and the Need for Closure scale uses 7-10 items per facet (Webster & Kruglanski, 1994)).

We call the average across multiple measures of overconfidence—standardized first to account for variation in the extent of overconfidence across tasks—'core overconfidence' to reflect that it constitutes a measure of the common element across tasks. A reliable measure of overconfidence such as this is fundamental for understanding the relationship between overconfidence and other variables. For example, past research has frequently observed a null effect when testing the relationship between a single measure of overconfidence and another psychological trait. But it is unclear whether this lack of relationship represents the true absence of a relationship or whether it reflects the unreliability in the measurement of overconfidence. In the present research, we resolve this ambiguity with our core overconfidence measure.

In sum, as discussed, a number of studies have found evidence that is potentially consistent with overconfidence being an individual difference using a variety of methodologies. Yet the finding is hardly universal. A plethora of null effects (e.g., Moore & Healy, 2008) combined with generally weak correlations suggests that observed individual differences could in fact be a mirage, especially if even weaker correlations go unreported due to the file-drawer problem. In contrast, we believe that the small-scale of many of these ‘null effect’ studies and the lack of sufficient aggregation is at the heart of the field’s mixed findings. We predict that aggregating multiple tests of overconfidence to make a measure we call core overconfidence will reveal a reliable individual difference that is stable over time. In addition, we predict that different forms of overconfidence (overestimation, overplacement, and overprecision) when measured reliably will be associated with each other. Although the forms have been shown to display different *levels* of overconfidence, we propose that *across individuals* they will nonetheless move together. These proposals form our first set of hypotheses:

H1: (a) There will be individual differences in overconfidence; (b) core overconfidence will be stable over time; and (c) the different forms of overconfidence will be related to one another.

Psychological Bases of Overconfidence

If a reliable measure of individual differences in overconfidence can be created, it facilitates exploring a new set of questions using a new methodological approach. In particular, what are the psychological factors that are correlated with *trait-level* overconfidence? Several lines of past research have aimed to test the relationship between overconfidence and individual difference measures and found few significant results (e.g., Blais et al., 2005; Larrick et al., 2007; Moore & Healy, 2008; Moore & Swift, 2010), which has led to some concluding that there may be no consistent personality correlates of overconfidence (Moore & Dev, 2018). These studies, however, have typically relied on single tests of overconfidence to measure these relationships. By using a measure of core overconfidence, the present research can provide a more definitive test of the degree to which psychological factors are associated with overconfidence. Moreover, the present research can test whether different psychological factors relate to different elements of overconfidence, as some psychological variables may be more strongly related to confidence and others more strongly related to accuracy.

Theoretically, the literature highlights two broad categories of difference that may play a role in overconfidence: respondents' self-views and cognitive processes. For example, participants with more positive self-views might tend to indicate higher levels of confidence regardless of their degree of accuracy, which would contribute to individual differences in overconfidence. An aspect of self-views that is particularly relevant to overconfidence is the tendency to favor evidence that maintains positive self-beliefs (Sedikides & Gregg, 2008). Grandiose narcissists—who are defined by self-enhancement, immodesty, and self-promotion (Miller et al., 2017)—for example, may be especially prone to this tendency. By ignoring reasons why their judgment may be wrong, they can continue to have excessive faith in their expertise. There is some evidence in support of this idea: For example, grandiose narcissism has been associated with both overestimation and overplacement (Ames & Kammrath, 2004; Campbell et al., 2004; Macenczak et al., 2016; O'Reilly & Hall, 2021). A different subfactor of narcissism, leadership/authority, has also been associated with overestimation (Macenczak et al., 2016), and several studies have found relationships between extraversion (or related constructs) and overestimation (Anderson & Brion, 2012; Pallier et al., 2002; Schaefer et al., 2004). Yet as mentioned, there are many cases of null results too (e.g., Larrick et al., 2007). Here, we extend this work by including a broad range of individual differences associated with more dispositionally positive self-views and testing the nuances of their relationships with overconfidence; it could be the case that differences in respondents' self-views primarily tend to influence overconfidence via inflating judgments of confidence.

Another pathway by which individual differences might influence overconfidence is through cognitive processes that lead to a broader consideration of uncertainties or to increasing respondents' accuracy. For example, it has been proposed that confirmatory thinking plays a role in both overestimation and overprecision (Koriat et al., 1980; Soll & Klayman, 2004). The conjecture is that people tend to favor evidence that supports an initial answer, neglecting to reflect and consider reasons why their first answer might be wrong. In related tests, Haran et al. (2013) found that individuals who rate higher on actively open-minded thinking exhibit less overestimation and overplacement, Moreover, if reflectiveness drives superior performance across a host of tasks, it might be expected to reduce overconfidence. On this point, Stanovich and West (1998) reported that having a higher SAT score was associated with less overestimation. We will examine evidence for both these potential pathways.

Overall, we examine the role of both respondents' self-views and measures of cognitive style and ability. Our investigation includes not only narcissism and cognitive reflection, but also a host of other measures. Although past findings have been mixed, we believe that there are good theoretical reasons to expect relationships, such as the role of explicit statements of confidence and the role of confirmatory thinking in producing overconfidence. We propose, therefore, that creating our reliable measure of core overconfidence will reveal clear relationships:

H2. Core overconfidence will be (a) positively related to personality measures that capture more positive self-views; (b) negatively related to individual differences that capture more reflective cognitive processes.

Overview of Studies

We test our hypotheses in four studies and one supplemental study containing a total of eight waves of data collection, following best practices identified in the overconfidence literature to mitigate statistical artifacts (Olsson, 2014): We randomly sample questions from domains (Gigerenzer et al., 1991; Klayman et al., 1999), measure each form of overconfidence on a separate set of questions (Klayman et al., 1999), and use multiple domains for generality. The data, materials, code, and preregistrations for all studies are available at the OSF site https://tinyurl.com/2vsjpn63. All studies except Study 1a were pre-registered; for further details of our pre-registrations, refer to the *Supplementary Information* (SI).

Studies 1a and 1b

In our first two studies, we collected evidence regarding the relationships between overconfidence measures across domains, forms of overconfidence, and time periods. We constructed 12 tasks that measured overconfidence, covering each of the three forms across four domains (geography, history, science, and literature). All participants answered these 12 measures in two time periods, and we tested the relationships among them both within a cross-section and over time. For Study 1a, these two time periods were two-and-a-half years apart; for Study 1b they were one-week apart with a higher retention rate (89% for Study 1b versus 43% for Study 1a).

Method

Study 1a Participants

We recruited 200 participants from Amazon's Mechanical Turk (MTurk), who were paid $4 for completing the study (M_{age} = 37.0, 58.5% men). There were 187 complete datasets and some that were missing a subset of the individual difference items. We later invited all 200 of the participants of Study 1a to complete the survey again for $7, two-and-a-half years after first completing it; 85 participants returned for this second wave (M_{age} = 41.8, 57.6% men).

Study 1b Participants

497 MTurk participants completed the first session in exchange for $3 ($M_{age}$ = 40.7, 54.1% men). One week after this first data collection ended, we invited participants to a second online session for $5, which was available for 5 full days. In total, 448 participants returned (M_{age} = 41.3, 53.6% men). Finally, three months later we added a third session of data collection, for which 344 participants returned (M_{age} = 42.8, 53.5% men), receiving $2.50.

Materials

We elicited overconfidence using the three most prevalent methods: with 2-answer-questions (2AQ) to measure overestimation (OE), calculating the difference between perceived and actual percentile of performance to measure overplacement (OPL), and eliciting interval estimates to measure overprecision (OPR). In Study 1a, we also included measures of overclaiming to facilitate a further validation of our measures of overconfidence (Paulhus et al., 2002). Overclaiming is the tendency to overstate one's familiarity with concepts or exemplars within a domain of knowledge. It is measured by seeing if, when asked about a subject such as classical music, people report being familiar not only with Chopin and Mozart (real composers), but also Brandinky and Chartreuse (foils). In constructing our measures of overconfidence and overclaiming, each required selecting questions. To minimize dependency on any one domain, we included questions from four domains of knowledge: geography, history, literature, and science.

An early statistical critique of overconfidence research was that selective, unrepresentative element pairings increase the likelihood of someone exhibiting overconfidence (Gigerenzer et al., 1991). To minimize this problem, we therefore selected the questions as follows, we: i) generated a list of 50 elements in the relevant domain, ii) drew the elements required for each form of overconfidence and overclaiming without replacement, and iii) for OE, drew the comparison pairs at random. Each OE measure required 20 elements to make its 10 questions, the OPL and OPR measures required 10, and the overclaiming measures required 8 in addition to 2 foils that we devised. In total, participants responded to 120 overconfidence domain questions, and the questions were the same for both Study 1a and Study 1b.

Study 1b's materials differed from Study 1a in three ways. First, the overprecision questions now requested a median estimate in addition to lower and upper estimates. We used the median estimate to measure the accuracy of each subject's best guess on a question by calculating the mean absolute deviation (MAD) between this estimate and the true value. Second, following a proposal by Gigerenzer et al. (1991), we added a new question to the end of each overestimation block, asking respondents, "How many of the 10 [domain] questions do you think you got right?" Finally, we modified the wording of our elicitation of the confidence estimate in the overplacement paradigm to use language that reflected our ranking of people's accuracy by MAD rather than number of questions answered correctly.

Study 1a Procedure

In each wave, participants had to pass a comprehension check and provide informed consent. The first four study blocks contained the overestimation form of overconfidence. For each domain, participants answered 10 questions and indicated their chance of being correct using a radio button scale from 50 to 100% in increments of 5%. The next four blocks contained the overplacement form of overconfidence. Participants answered 10 questions and then estimated their percentile in the performance distribution for each domain (see the SI for question wording). Next, participants responded to the four overclaiming blocks. This task asked participants to indicate their familiarity with 10 objects on a Likert scale, of which 8 were real and 2 were fabricated. In the third and final measure of overconfidence, participants provided confidence intervals to measure overprecision. For each quantity, participants were asked to provide a lower and upper estimate (e.g., "*a lower estimate such that you are 90% sure that Norway does not have fewer people than that*", and "*an upper estimate such that you are 90% sure that Norway does not have more people than that*").

In sum, participants responded to four measures of three different forms of overconfidence (OE, OPL, OPR) for a total of 12 measures, plus an additional set of overclaiming tasks. Participants then completed a series of individual difference scales and provided demographic information (age, gender, education, race). The included scales were intellectual humility (Leary et al., 2017), need for closure (decisiveness facet and closed mindedness facet; Kruglanski et al., 1993), the narcissistic personality inventory (Ames et al., 2006), optimism (Scheier et al., 1994), self-efficacy (Chen et al., 2001), desirability of control (Burger & Cooper, 1979), and the dominance facet of extraversion (Wiggins et al., 1988). The first wave of data was collected December 13, 2019. In the second wave of data collection (May 12-18, 2022), participants responded to an identical survey.

Study 1b Procedure

In the first experimental session, respondents were asked to indicate if they were willing to complete both sessions and answered a comprehension check. Participants then responded to the 12 measures of overconfidence (and the questions for the alternate measure of overestimation). The order of the forms was randomized, as was the order of domains within each form, and questions within domain, except for overestimation. After completing these measures, respondents answered some short items that we were piloting (see Lawson et al., 2024 for additional details) and provided demographic information. We also included an item to test whether participants had metacognitive awareness of their own overconfidence: "*How overconfident are you in how you make decisions?*" (0-10), which we comment on in the General Discussion. One week after the first experimental session was completed, the respondents were invited to complete a second session, where they completed the same survey again with the addition of the narcissistic personality inventory (Ames et al., 2006), which facilitates an additional test of the stability of overconfidence over time, as described in a footnote. The narcissism measure came after the overconfidence measures and before the pilot items. The randomized orders of the overconfidence tasks differed across the waves.

The first two sessions of Study 1b were collected May 9-10, 2022 and May 17-21, 2022. The same participants were invited back to complete a third survey three months later (August 17-24, 2022). Here, we counterbalanced the order of administering a finalized three question short measure of core overconfidence (the "overconfidence test" (OCT)) with a string of individual differences measures, including those from Study 1a, as well as four new measures; the cognitive reflection test (Frederick, 2005), the Berlin numeracy test (Cokely et al., 2012), a measure of actively open-minded Thinking (AOT; Baron, 2019), and a short measure of the Dark Triad (Jones & Paulhus, 2014). We summarize the data from the OCT in this manuscript in the General Discussion, but for further details, please see Lawson et al. (2024).

Coding of Overconfidence and Overclaiming Measures

For the overestimation (OE) questions in Study 1a, we calculated how many questions each respondent got right in a domain, and their average confidence across their responses. We then subtracted the percentage of correct responses from the average confidence to compute OE. For overplacement (OPL), we averaged the absolute deviation of each of the respondent's 10 answers in a domain from the true values to get an overall measure of their accuracy (mean absolute deviation; MAD). We then ranked the participants within each domain based on this measure to compute their quantile in the performance distribution and subtracted this quantile from the participant's estimate of their performance quantile to calculate a measure of OPL. For overprecision (OPR), we counted how many of a respondent's confidence intervals contained the true value and compared this to the number that should have contained the true value—which was 8, as we elicited 80% confidence intervals.[1]

Though Paulhus et al. (2003) analyzed overclaiming using signal detection methods, we opted for a single measure. It involves two primary components: the average familiarity with the fabricated elements (F) and the average familiarity with the real elements (R). The formula we used to compute the measure is F(7 – (R – F)). A higher score on this measure can indicate two things: i) a smaller disparity in familiarity between real and fake elements, and ii) a greater tendency to claim familiarity with fake elements. In the SI, we analyze these measures using signal detection theory (Paulhus et al., 2003).

In Study 1b, we included the same overconfidence measures as in Study 1a, but further measured OE using an alternate approach. In this approach, we subtracted the number of questions a respondent answered correctly in a domain (e.g., 6 of 10) from their estimate of how many they answered correctly (e.g., 8 of 10). This example would constitute an OE of 20%. We adopted this measure for our primary analyses in Study 1b following Moore and Healy's (2008) critique that the traditional measure of overestimation is confounded with overprecision; we replicated our analyses with the traditional form in the SI.

---

[1] For the OPL and OPR geography responses in Studies 1a-b, some participants provided responses (e.g., 8,000,000) that were converted into a response on the requested scale of millions (i.e., 8).

In addition, for Study 1b, we constructed measures of participants' confidence and accuracy so that we could analyze them separately. For OE, these were simply the estimated number of questions answered correctly, and the percentage answered correctly; for OPL it was respondents' estimated percentile in the performance distribution and their actual percentile (based on mean absolute deviation); for OPR, it was slightly more complicated. In our OPR questions, we estimated the confidence communicated by the respondents' answers by looking at the mean expected absolute deviation (MEAD) associated with their provided interval estimates (Soll & Klayman, 2004). To do so, we used the three quantiles provided by respondents—the 10^{th}, 50^{th}, and 90^{th} percentile of their subjective probability density function—to estimate the parameters of a beta distribution using the R package 'SHELF' (Oakley, 2021). Using these parameters, we simulated the expected absolute deviation of the observed outcome from the respondent's median estimate. We averaged these across the 10 questions in each domain to get the MEAD. For further details of these analyses, refer to the SI. We computed the MAD of the answers by computing the average absolute deviation between the respondents' median estimates and the true values in each domain. We did not replicate this analysis for Study 1a because we did not collect median estimates in that design. These measures were used later in our analysis of the role of confidence and accuracy in determining the stability of measures of overconfidence.

Individual Difference Measures

In both studies, participants provided responses to a series of individual difference measures. For measures used in both studies (1a, 1b), the alphas were as follows: intellectual humility (.84, .87), need for closure decisiveness (.84, .87), need for closure closed mindedness (.76, .73), the dominance facet of extraversion (.90, .92), and the Narcissistic Personality Inventory (NPI; .87, .85). Notably, the NPI was administered in wave 2 of Study 1b rather than wave 3. Several measures were unique to Study 1a: optimism (.91), self-efficacy (.93), and desire for control (.80). In Study 2b, the short measure of the Dark Triad yielded alphas of .86, .85, and .82 for Machiavellianism, narcissism, and psychopathy, respectively. In all cases, the reliabilities were in the acceptable range.

Results

Levels of Overconfidence

We start by analyzing the levels of overconfidence observed across measures to replicate existing findings. In Study 1a, confidence exceeded accuracy by 10.4 percentage points (pp) across all measures, but this average level of overconfidence concealed significant variation across forms. Whereas there was some overestimation (OE: 5.6 pp) and substantial overprecision (OPR: 35.4 pp), participants were on average underconfident for the overplacement items (OPL: -9.8 pp). In Study 1b—where we focus on the first wave of data to present the cross-sectional results—we observed a similar pattern, with 4.0, 35.6, and -14.0 percentage point gaps indicating overestimation, overprecision, and underplacement, respectively. In Study 1b we found that with the alternate measure of OE, participants underestimated their performance by 21.4 percentage points. Thus, the two measures of OE yielded very different average levels of overconfidence, as others have found (Gigerenzer et al., 1991; Kahneman & Tversky, 1996). Overall, our results are in line with a wealth of extant research (e.g., Klayman et al., 1999) finding that different forms yield different average levels of overconfidence. We next turn to a different question, testing the relationships between different measures of overconfidence.

Individual Differences in Overconfidence

We begin by reporting the pairwise correlations between the individual items (i.e., overconfidence on the individual 10-item tests), as previous reporting has typically looked at this statistic. We used the alternative measure of OE for Study 1b; for the analyses replicated with the traditional measure, see the SI. The average interitem correlation was $r = .37$ (range:

.14 to .78) in Study 1a and *r* = .31 (.09 to .72) in Study 1b. All pairwise correlations were significant at the .05 level, with the median $p < .001$ in both studies. We next calculated Cronbach's alpha for a scale that combined the 12 measures to determine internal reliability—we found alphas of .88 and .84 in Studies 1a and 1b, respectively, which supports H1a. Given this reliability, and as planned, we averaged the 12 measures to compute *core overconfidence*.

To verify that core overconfidence constitutes an individual difference, we measured it in two time periods in both Study 1a (2.5-years apart) and Study 1b (1-week apart). Test-retest reliability was $r = .78$ (95% CI [.67, .85), $p < .001$) and $r = .85$ (95% CI [.82, .87), $p < .001$) in Studies 1a and 1b, respectively. We plot the relationships between core overconfidence at time 1 and time 2 for Studies 1a and 1b in Figure 1a and 1b. These data show that core overconfidence is stable even over long time periods, supporting H1b.[2]

Figure 1a
Core Overconfidence in Time 1 and 2, Study 1a

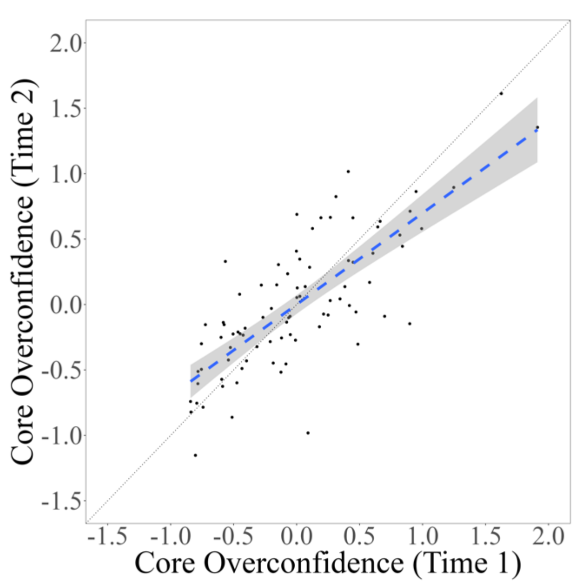

Figure 1b
Core Overconfidence in Time 1 and 2, Study 1b

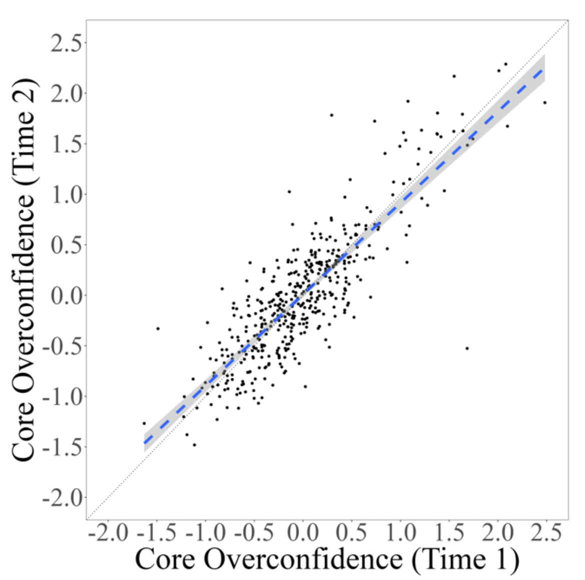

[2] We also provide an additional validation: In Study 1b, we included the NPI in the second but not the first wave of data collection to circumvent the concern that respondents answering the same questions in stage two may recall the contemporaneous situational factor of how they felt during stage one. There was a positive significant correlation between this time 2 measure and core overconfidence at time 1 ($r = .52$, 95% CI [.45, .58], $p < .001$), in support of there being individual differences in overconfidence.

Notably, core overconfidence contains measures of accuracy using specific questions that are intended to proxy for general knowledge in the given domains. As a result, our test-retest measures include the same specific questions across both periods of time. If knowledge in specific questions remains stable, this could inflate the observed test-retest reliability of core overconfidence. To combat this, we tested the correlation between measures of core overconfidence in different time periods that did not contain any of the same questions. At time 1, we sampled 6 measures from the 12, such that there were 2 measures of each form of overconfidence. We then correlated the average of these measures with the average of the other 6 measures administered at time 2. We repeated this exercise for all 216 possible ways of dividing the 12 measures in this fashion. In Study 1a, the correlations between these two 6-item measures of core overconfidence varied from $r = .50$ to .75 ($M = .65$); in Study 1b they varied from $r = .68$ to .76 ($M = .72$). Even with a smaller set of non-overlapping items, core overconfidence persisted over time, in further support of H1b.

Relationships Between Measures of Overconfidence Across Forms

In a test of H1c, we consider the relationships between measures of overconfidence across form. We first averaged across the four domain measures in each form before correlating these averages (Table 1). The different forms of overconfidence were positively related (all *p*s < .001), in support of H1c.

Table 1

The Strength of Correlations Among Forms of Overconfidence

Study 1a correlations

	OE (1)	OPL (1)
OPL (1)	.54** [.43, .63]	
OPR (1)	.59** [.49, .68]	.36** [.23, .48]

Study 1b correlations

	OE (1)	OPL (1)
OPL (1)	.55** [.49, .61]	
OPR (1)	.22** [.14, .30]	.35** [.27, .43]

Note. OE = overestimation, OPL = overplacement, OPR = overprecision. Each entry indicates the correlation between the average of four constituent domain measures, each of which were standardized before being averaged. The number in brackets indicates that all measures were from the first wave of data collection in each study. * $p < .05$; ** $p < .01$. Study 1b used the alternate 'outside view' measure of overestimation. With the traditional measures, the first column correlations were $r = .48$ and .40 (ps < .001).

Given the longitudinal nature of our data, we next turn to studying the relationships between measures of overconfidence both within-form across time and between-form across time (Tables 2a & 2b). The leading diagonals provide evidence in support of within-form stability for each of the three forms: in all cases, the correlations were statistically significant at $p < .001$. This builds on extant research showing the intertemporal stability of measures of OE (e.g., Stankov & Crawford, 1996) and extends the evidence to OPL and OPR. When considering overconfidence measured in different time periods and different forms, there is some variation in the strength of associations. Notably, in Study 1a's samples, OPR in time 1 was not significantly related to OPL in time 2, yet in Study 1b's larger sample with higher retention, this association was significant. Across all relationships tested, we also see many significant relationships across form across time.

Table 2

The Strength of Correlations Among Forms of Overconfidence Across Time

Study 1a correlations

	OE (1)	OPL (1)	OPR (1)
OE (2)	.65** [.51, .76]	.34** [.14, .52]	.43** [.24, .59]
OPL (2)	.15 [-.06, .35]	.58** [.42, .71]	-.03 [-.24, .18]
OPR (2)	.43** [.24, .59]	.26* [.05, .45]	.77** [.66, .84]

Study 1b correlations

	OE (1)	OPL (1)	OPR (1)
OE (2)	.72** [.67, .76]	.52** [.45, .59]	.20** [.11, .29]
OPL (2)	.48** [.40, .54]	.76** [.72, .80]	.37** [.29, .45]
OPR (2)	.19** [.10, .28]	.36** [.27, .44]	.81** [.78, .84]

Note. OE = overestimation, OPL = overplacement, OPR = overprecision. Each entry indicates the correlation between the average of four constituent domain measures, each of which were standardized before being averaged. The number in brackets indicates which wave of data collection the data came from. * $p < .05$; ** $p < .01$. Study 1b used the alternate 'outside view' measure of overestimation.

Having considered the relationships between measures of overconfidence, finding evidence for individual differences as well as for relationships across forms, we next visualize these effects using the data from the first two waves of Study 1b in Figure 2. In Figure 2 we split the participants into deciles based on their core overconfidence score in the first wave of data collection. We then computed the average overconfidence for the 10 groups in each of the three forms in the second wave of data collection. Three things are captured in the resulting visualization: i) the stability of measures of overconfidence over time, such that time 1 core overconfidence is predictive of measures at time 2, ii) the presence of situational effects across forms of overconfidence, as captured by the distances between the lines—for example, OE and OPR at time 2 differed by around 50 percentage points for most values of core overconfidence at time 1, and iii) the importance of individual differences—despite the difference in levels, each line showed steep slopes over the range, varying by 40-50 percentage points.

Figure 2

Average Overestimation, Overplacement, and Overprecision in Time 2 Based on Core Overconfidence Decile in Time 1 (Study 1b)

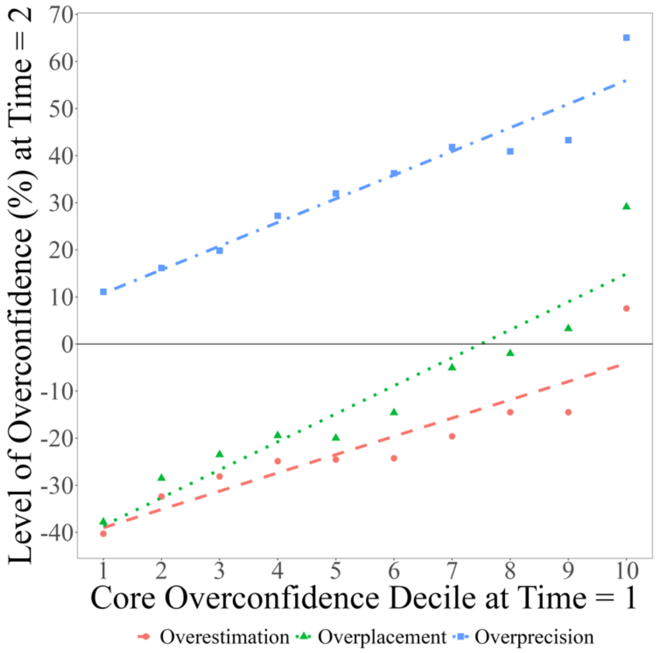

Factor Structure of Overconfidence Across Forms and Domains

Thus far, we have found that core overconfidence is internally reliable, stable over time, and that the different forms of overconfidence are related to one another. In this next section, we will consider the extent to which a person's tendency toward core overconfidence can be predicted by personality traits and cognitive styles. Before doing that, it is important to ask whether core overconfidence can be thought of as a monolithic trait, or whether there are facets that move separately as a function of situational factors such as form and domain. For example, Table 2 shows stronger relationships between measures within-form than between-forms, which could indicate that there are multiple facets.

To answer this question, we turn to considering the McDonald's omega values (McDonald, 1999; Zinbarg et al., 2005) for our 12 measures. There are two such values—omega hierarchical (ω_h), the proportion of variance associated with a single, general factor, and omega total (ω_t), the proportion of variance associated with all factors that were shared by more than one measure (Revelle & Zinbarg, 2009). The difference between the two omega values indicates the variance associated with specific factors. We report here the asymptotic version of omega hierarchical that is adjusted to account for the finite length of test items (Revelle, 2024).

In Study 1a, the omega hierarchical value was $\omega_h = .80$, indicating that 80% of the variance in core overconfidence is attributable to a single, general factor. The ω_t was .91, indicating that an additional 11% of variance was accounted for by specific factors. In Study 1b, a general factor explained 71% of the variance, with an additional 9% captured by specific factors. This analysis indicates that a single, general factor explains the majority of variance in measures of overconfidence; yet the existence of specific factors hints at differences as well that could explain some inconsistencies in the literature. Importantly, the factor structure only uncovers statistical relationships across measures, rather than the psychological factors associated with overconfidence. We turn to these next.

Self-views, Cognitive Style, and Overconfidence

In our next set of analyses, we transition into unpacking *why* core overconfidence appears to be so stable. One possibility is that it depends on respondents' behaviors in producing numeric estimates. Perhaps individuals who are systematically more overconfident than others when reporting numbers are no longer relatively more overconfident on non-numerical tasks. To verify this, we tested the relationship of core overconfidence with overclaiming in Study 1a, which was measured using Likert scales. Here, we found a positive relationship ($r = .70$, 95% CI [.62, .77], $p < .001$) between core overconfidence and the average of the four overclaiming measures, with correlations with individual measures ranging from $r = .55$ to .72. Given that this was measured using simple Likert scale responses, this shows that core overconfidence is predictive of judgments in other contexts that do not involve complex numeric responses. We next turn to testing the psychological trait correlates of core overconfidence to understand its psychological bases.

We tested the relationship between core overconfidence and other psychological traits in both Study 1a and Study 1b. Here, we focus on the results of the larger, pre-registered Study 1b. Study 1a's analyses revealed several non-significant relationships with core overconfidence (e.g., optimism) and some significant relationships that did not replicate in Study 1b (e.g., intellectual humility); for full details of those results, refer to the SI. We further present results not just on the relationships of other traits with core overconfidence, but also with specific forms of overconfidence (e.g., OE) and the components of overconfidence: confidence and accuracy. As discussed earlier, this analysis required constructing separate measures of the two components. This was done in the same way as for core overconfidence—we first standardized each measure of confidence and accuracy separately for the domains and forms, and then averaged across all 24 measures to obtain the "core" average for each participant.[3]

We present the significant relationships between individual differences and overconfidence in Table 3, further disaggregating the relationships by form and component. Inspection of the table reveals that the cognitive process measures—CRT (cognitive reflection), BNT (numeracy), and AOT (actively open-minded thinking)—were related to *lower* core overconfidence, but this pattern was largely attributable to their positive association with accuracy. In contrast, self-view related variables such as dominance and the Dark Triad (which includes narcissism) were related to higher core overconfidence and this pattern was largely attributable to their relationship with higher confidence. Psychopathy was

[3] The components' interpretations are reversed for overprecision (e.g., lower MADs correspond to greater accuracy). Thus, before aggregating, we took the log of the MEAD and MAD and multiplied each of them by minus one. A higher score on the resulting confidence variable indicated greater narrowing of one's confidence intervals (i.e., higher confidence), and greater accuracy (i.e., less absolute deviation between the median estimate provided and the correct answer).

an exception—it both increased confidence *and* reduced accuracy. Table 3 further reveals that these effects appear to be moderated by the form of overconfidence. For example, the cognitive process measures had the greatest impact on measures of overprecision. This was not due to providing wider confidence intervals, which would drive a negative association with core confidence. Rather, individuals high in CRT, for instance, appeared to provide median estimates that were better centered on the truth (i.e., they were just more accurate).

Table 3

Correlations Between Other Psychological Traits and Forms and Components of Overconfidence

Variable	Core OC	OE	OPL	OPR	Core accuracy	Core confidence
CRT	-.14* [-.24, -.03]	.06 [-.05, .16]	-.04 [-.15, .07]	-.30** [-.40, -.20]	.26** [.15, .36]	.03 [-.08, .14]
BNT	-.12* [-.22, -.01]	.11 [-.00, .21]	-.01 [-.12, .10]	-.33** [-.42, -.23]	.31** [.21, .41]	.10 [-.01, .21]
AOT	-.21** [-.31, -.11]	-.09 [-.20, .02]	-.13* [-.24, -.02]	-.23** [-.33, -.13]	.24** [.13, .34]	-.06 [-.17, .05]
Dominance	.33** [.23, .42]	.37** [.27, .46]	.21** [.10, .31]	.16** [.05, .26]	-.01 [-.12, .10]	.35** [.25, .44]
Machiavellianism	.24** [.13, .34]	.32** [.22, .42]	.19** [.08, .29]	.02 [-.08, .13]	-.05 [-.16, .06]	.20** [.09, .30]
Narcissism	.39** [.30, .48]	.41** [.32, .50]	.29** [.19, .39]	.17** [.07, .28]	-.09 [-.20, .02]	.33** [.23, .43]
Psychopathy	.36** [.26, .45]	.37** [.27, .46]	.24** [.13, .34]	.18** [.08, .29]	-.15** [-.26, -.05]	.24** [.14, .34]

Note. Core OC = core overconfidence, OE = overestimation, OPL = overplacement, OPR = overprecision. Each entry indicates the correlation of a psychological trait with the average of all relevant measures across two time periods, each of which were standardized before being averaged. CRT = cognitive reflection test, BNT = Berlin numeracy test, AOT = actively open-minded thinking. * $p < .05$; ** $p < .01$. For further details of the insignificant relationships between traits and core overconfidence, refer to the SI.

Discussion

In Studies 1a and 1b, we found evidence for individual differences in overconfidence that extend across forms of overconfidence, are stable over time, and can be explained by differences in both respondents' self-views and cognitive processes. Yet despite using four separate domains (geography, history, literature, and science), it remains possible that the observed individual differences are limited to a specific aspect of respondents' self-views—in particular, their self-views concerning trivia-type questions—and do not extend to their *chronic* self-views across a wide range of tasks (Ehrlinger & Dunning, 2003). One piece of evidence that these domains might capture a narrow slice of respondents' self-views is the correlations in accuracy across the domain measures. To examine this, we first computed Z-scores for each participant's accuracy, doing this separately for each domain-form combination. We then averaged across the three forms to compute each participant's overall accuracy in a domain. Given that there were four domains, there are six unique pairs. For each of these, we calculated the correlations in accuracy, which provides a measure of the degree to which accuracy in the two domains is related. In Study 1b, there was an average correlation of $r = .51$ across domains. In our next two studies, we address the possibility that individual differences in overconfidence are limited to narrow self-views directly, testing for the presence of individual differences in overconfidence when participants' accuracy across domains is less stable.

Study 2a and 2b

Across Studies 1a-b we found robust evidence for relationships between different measures of overconfidence across domains, forms, and time periods. However, the domains we adopted, while distinct, all contained trivia-type questions on a variety of academic topics. This resulted in highly correlated accuracy scores across domains (average $r = .51$). The similarity of these domains could increase correlations between measures of overconfidence in two ways: i) directly, given that overconfidence is a difference score between confidence and accuracy, and ii) indirectly, given that people may have relatively more stable confidence within domains in which they tend to be accurate or inaccurate.

To explore whether associations among measures of overconfidence remain when accuracy and confidence are less correlated, in Study 2 we intentionally sampled domains to depress the correlations between domain-level measures of both confidence and accuracy to focus on respondents' chronic self-views (Ehrlinger & Dunning, 2003). Specifically, we ran a pilot study where we measured self-rated competency in ten domains and selected the three domains with the weakest associations. This was designed as an edge case to minimize the likelihood that measures of overconfidence would be related. In short, we expect the relationship between measures of overconfidence to be weaker in Study 2; the question is by how much. In extremely disparate tasks, perhaps one would not observe *any* relationships among measures of overconfidence.

Method

Participants

In Study 2a, we recruited 398 participants from Amazon's Mechanical Turk in exchange for $5 ($M_{age}$ = 41.2, 56.3% male). In Study 2b, we recruited 395 participants from Prolific Academic in exchange for $5 ($M_{age}$ = 41.5, 49.6% male).

Materials

As mentioned, we empirically selected three domains where individuals' self-rated competency was minimized across them (see the SI for details). With the three domains that resulted from this process (estimating average commute times between U.S. cities, maximum temperatures in U.S. cities, and movie release dates for the top 1,000 rated movies on the website IMDB), we constructed comprehensive lists (e.g., average commute times in the 146 most populated U.S. metro areas) and randomly drew the elements needed for our questions. We used the same methods to measure overplacement and overprecision as in Study 1b but only used the alternate measure for overestimation—asking people to estimate how many of the questions they answered correctly.

Procedure

Participants first answered three short-form overconfidence items in a randomized order—the overconfidence test, or OCT (Lawson et al., 2024). Next, participants answered the nine measures of overconfidence – 3 forms by 3 domains. Participants received the overestimation, overplacement, and overprecision tasks in a randomly selected order. Domains were randomly ordered within each task, and questions randomly presented within domain, except for in overestimation. Finally, participants provided demographic information.

Results

Study 2a and 2b had identical designs and thus identical analyses: all that differed was the survey platforms. As a result, we summarize their results together. First, note that across both samples we replicated the pattern of *levels* of overconfidence from Study 1b. We observed underestimation of -5.2 and -9.8 percentage points in 2a and 2b, respectively, using the alternate measure. Once again, on average participants underplaced themselves (by -9.5 and -13.9 pp) and were susceptible to overprecision (24.5 and 23.2 pp).

We next examine whether Study 2's domain selection procedure was successful. Using the same approach to quantifying across-domain accuracy relationships described in the Discussion of Study 1b, average across-domain accuracy correlations were r = .30 and .14 in Studies 2a and 2b, respectively. Given the goal of attenuating the relationship, this compares favorably to an average correlation of r = .51 in Study 1b. We also computed the equivalent measures for confidence (Study 1b: r = .67, Study 2a: r = .61, Study 2b: r = .54). Notably, selecting disparate domains had a larger effect on accuracy than on confidence, pointing to respondents possessing stable, chronic self-views that are relatively insensitive to actual performance, as others have proposed (e.g., Ehrlinger & Dunning, 2003).

As expected, we saw correspondingly weaker relationships among the individual item measures than in Studies 1a-b. Correlations in Study 2a ranged between -.10 and .51, with a mean r of .26. The majority, 31 out of 36, of these correlations were both positive and significant with $p < .05$. For Study 2b, the average r was .15, ranging between -.11 and .43, with 22/36 positive and statistically significant. Overall, though there were still many significant relationships between measures of overconfidence, they were somewhat weakened.

We next tested whether there was evidence for individual differences in core overconfidence even with such disparate domains. In Study 2a, the nine measures of overconfidence had a Cronbach's alpha of .76; in Study 2b, alpha was .62. The omega values indicated that 44% of variance was explained by a general factor and 40% by specific factors in Study 2a; the corresponding numbers were 37% and 35% for Study 2a. In sum, in each case, over half of explained variance was associated with a single, general factor. These results support the presence of individual differences, even when the cross-domain relationships were deliberately weakened by construction. (Notably, though the alpha in Study 2b fell below the commonly used standard of .7, this threshold could be attained by increasing the number of overconfidence items from 9 to 14: with an average interitem correlation of $r = .15$, an alpha of .7 would require 14 items.)

Discussion

In sum, we found that the relationships between accuracy and confidence across measures are a moderating factor that determine the strength of individual differences in core overconfidence, but that there was consistent evidence for individual differences in core overconfidence even when these associations were intentionally weakened by selecting disparate stimuli. In the General Discussion next, we further expand on the takeaways from these results.

General Discussion

In the present research, we aimed to reconcile divergent perspectives on whether there are individual differences in core overconfidence. In four primary studies with seven waves of data collection, we examined the stability of overconfidence across time, domains of knowledge, and different forms of overconfidence. We consistently found that though the average levels of overconfidence vary across tasks, the same individuals tend to be relatively more or less overconfident. This provides support for an assumption of stability often made in other fields (e.g., in accounting, Malmendier & Tate, 2005a)—provided that core overconfidence is measured reliably through the use of multiple items.

We found support for five hypotheses. Namely, we found that multiple measures of overconfidence exhibited strong reliability, showing support for individual differences in core overconfidence (H1a), as well as finding that core overconfidence was stable over short and long periods of time (H1b), and that there were significant relations between different forms of overconfidence (H1c). Further, we found that when measured reliably, core overconfidence was consistently related to individual differences both in personality variables related to the positivity of respondents' self-views (H2a)—as well as self-concepts more broadly (e.g., Machiavellianism)—and differences in reflective cognitive processes (H2b).

The Nature and Strength of Individual Differences in Overconfidence

Our central theoretical contribution is that there is a general factor that underpins different measurements of overconfident beliefs and that this contains a dispositional component that is highly stable over time. In the primary tests of our hypotheses in Studies 1a-b, we found omega hierarchical values of .71-.80, indicating that the vast majority of variance in standardized measures of overconfidence—including different forms of overconfidence and different domains of knowledge—was associated with a single, general factor. These data are consistent with other findings, such as that there are individual differences in measures of overestimation (Binnendyk & Pennycook, 2024; Jonsson &

Allwood, 2003; Stankov & Crawford, 1996), expanding on earlier work (Lawson, 2022) to show that these individual differences in core overconfidence emerge across measures of the three most commonly used forms (overestimation, overplacement, and overprecision). We further augmented these data by conducting two multi-wave experiments that enabled us to demonstrate the significant reliability of core overconfidence over time (e.g., with a test-retest reliability of $r = .78$ over a 2.5-year delay between experimental sessions). Taken together, these data clearly demonstrate a dispositional component to overconfident beliefs. As a result, a person who exhibits more overconfidence today on one task is also likely to be relatively more overconfident on different tasks today and to exhibit relatively more overconfidence in the future.

In efforts to reconcile our results with the literature, such as the claim that the three forms of overconfidence are empirically distinct (Moore & Healy, 2008), we highlight several additional analyses that inform our perspective. First, in all studies, additional, specific factors associated with the adopted measures of overconfidence also explained variance. In light of these results, we propose viewing the three forms of overconfidence as *facets* of an overarching construct with a general factor, rather than as *the same construct* or *independent constructs*. Second, a natural question to ask is whether the stability we observed in Studies 1a-b occurred because the selected domains were similar to each other (i.e., all trivia-type questions), and whether one would expect to observe similar effects with more disparate tasks (e.g., a math test versus performance at shooting free throws)? In designing Studies 2a-b we intentionally reduced the stability of confidence and accuracy across measures and found that relatively less variance was explained by a single, general factor—though there was still substantial evidence of individual differences in core overconfidence. We build on this further here.

Consider the question of domain similarity more generally: Given that overconfidence is the difference between confidence and accuracy, stability in either of these components, for example through generally inflated confidence or depressed accuracy, is likely to lead to individual differences in overconfidence. As a result, individual differences in overconfidence will be stronger when accuracy and confidence are more correlated across tasks. Yet we observed less variability in confidence than accuracy across tasks, even when the tasks were selected to be disparate, suggesting that the expected stability of accuracy across tasks in the ecological environment is likely the most critical determinant of the strength of individual differences in overconfidence. We argue that moderate to strong correlations across measures of accuracy are likely to be present for many applications, which could lead to correlations between measures of overconfidence looking more like Studies 1a-b than Studies 2a-b. The reason is that there are differences in cognitive variables (e.g., Frederick, 2005; Meyer et al., 2024; Spearman, 1961) that are likely to drive performance across tasks. This will not be true in all cases, but for example, if you observe a colleague being overconfident in one work task, their performance in that task is likely correlated with their performance in other tasks in the workplace that you might want to predict their overconfidence in when deciding whether to follow their advice, and so individual differences are in overconfidence are likely to be strong. This has important implications for expectations in organizational settings, as well as managers' choices regarding task allocation, and managerial decision making.

In addition to our four primary studies, we conducted a Supplementary Study S1 (N = 992) to examine further possible moderators of the strength of individual differences in core overconfidence (refer to the SI for further details). In particular, we tested the role of incentives, as well as different participant recruitment procedures and response exclusion criteria, among other factors. Across the 24 pre-registered possible specifications, we found that the Cronbach's alpha of the 12 measures of overconfidence always remained above the threshold of .7 (median α = .79, range = [.71, .88]). Surprisingly, incentives did not have any substantial overall effect, though limiting the recruited participants to those without an established track record of successful task completions on the platform or excluding any participants who indicated different types of responses (e.g., three equal quantiles when providing an estimate for a quantity) somewhat weakened the observed stability. Notably, such exclusion criteria also attenuated the reliability of established individual differences such as the Dark Triad (Jones & Paulhus, 2014), and such recruitment restrictions led to less representative samples relative to the U.S. population. That said, it is likely that the stability of core overconfidence will vary across populations as well as measures, depending on the consistency of the behavior of the particular population, which is a noteworthy additional moderator to consider.

Finally, we note that in addition to these moderators of the actual strength of correlations among single item measures of overconfidence—associations in accuracy and confidence driven by tasks and populations—differences in analytical strategy appear to drive many divergences in conclusions in the literature. If many pairwise correlations are in the range of r = .2 to r = .3, researchers only collecting a few measures and not considering aggregating them to form an index may not view such correlations as evidence for individual differences in overconfidence, when in fact they could be. As previously mentioned, such correlations are frequently observed among items of established individual difference (e.g., the HEXACO-60 personality measure (Ashton & Lee, 2009)). Moreover, with relatively weak positive associations between measures, the use of small samples will sometimes yield correlations close to zero or even negative, perhaps dissuading researchers from conducting further inquiry. We hope that our proposed empirical framework—considering individual measures of overconfidence as items that need to be aggregated to study individual differences—will stimulate a re-interpretation of many of these results across the field.

Psychological Explanations of Individual Differences in Overconfidence

In another theoretical contribution, our more reliable measurement of core overconfidence also had implications for our results regarding core overconfidence's relationship with other individual differences. In contrast to earlier suggestions that overconfidence is unrelated to stable personality measures (Larrick et al., 2007), we find that core overconfidence is strongly related with both personality variables that capture self-views and measures of cognitive style but appears to be more consistently related with measures of self-views. Moreover, we found that participants' responses to the question, "*How overconfident are you in how you make decisions?*" correlated with their core overconfidence at $r = .45$, which could challenge the view that overconfidence is an error of awareness, though this measure was asked after completing the overconfidence measures and so may constitute post-hoc reflection rather than respondents' contemporaneous awareness. Put together, this evidence contributes to the debate regarding whether overconfidence is a cognitive bias or a motivated bias. Whereas others have concluded that the role of motivation in overconfidence may be exaggerated when examined as an experimentally-induced state (Logg et al., 2018), these data suggest that self-view variables that display chronic self-enhancement play a key role in determining overconfident beliefs. In the next section, we examine one pathway by which psychological tendencies to self-enhance might impact the calibration of people's confidence to their accuracy.

The Accuracy-Confidence Relation and Metacognition

As well as overconfidence—the extent to which confidence exceeds accuracy—related measures are of interest to researchers, such as the association between accuracy and confidence. In other words, are people who are generally more accurate also likely to be more confident (e.g., Kruger & Dunning, 1999)? In such tests, several researchers have found a curvilinear relationship between accuracy and confidence (Burson et al., 2006; Jansen et al., 2021; Sanchez & Dunning, 2018), which we replicate in our data in Figure 3. In the figure, each point represents the average confidence and accuracy for one participant in Study 1b, with averages computed across the 24 constituent measures described in the methods. Here, a model that includes both linear and quadratic terms for accuracy achieved a superior fit to one solely estimating a linear relationship, F (1, 325) = 24.3, $p < .001$ (Panel A).

Figure 3

The Relationship Between Core Accuracy and Core Confidence, Split Based on Quartile in the Dark Triad

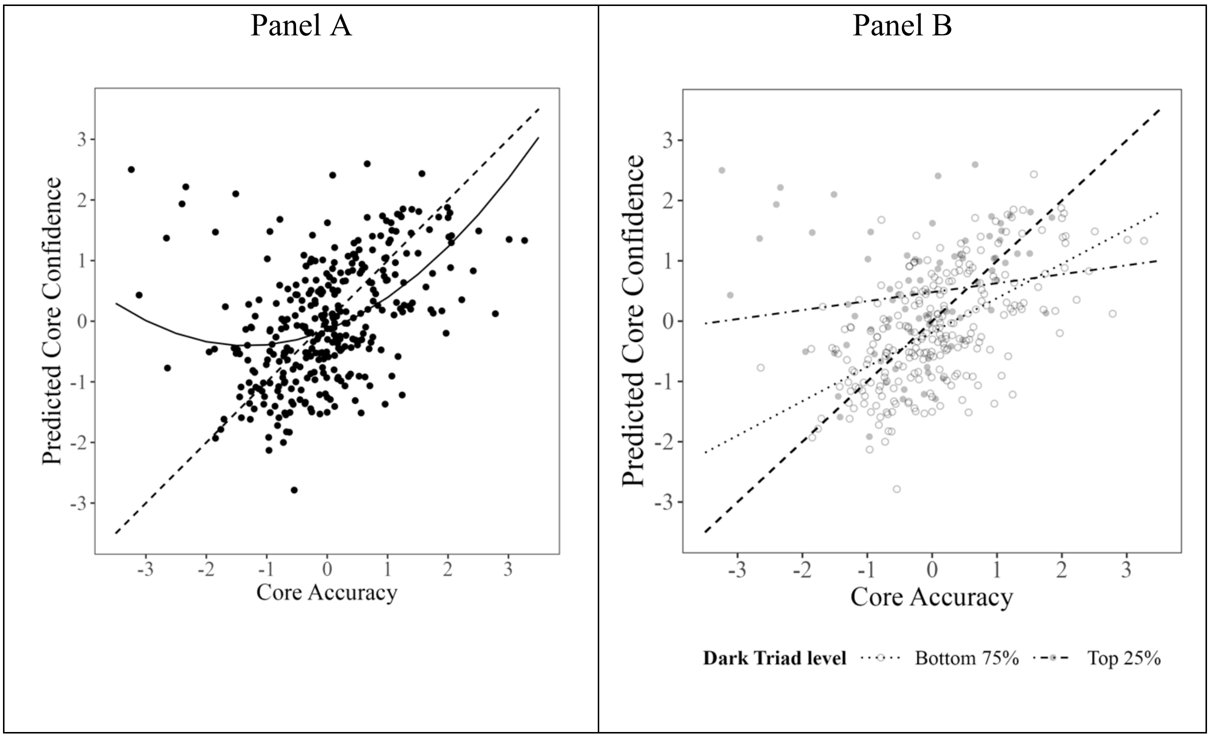

Note. The points indicate the raw data of accuracy-confidence pairings in Study 1b's 3-wave data (N = 328). The dashed line is the identity line. The black curved line plots regression predictions from a quadratic model with no term for Dark Triad score, and the dot-dash line and dotted lines plot regression predictions from a model interacting accuracy with a binary indicator for whether a participant was in the top quartile of Dark Triad scores or not. In the second panel, the filled circles indicate the high Dark Triad respondents and the empty circles the low Dark Triad respondents.

One possible explanation for the quadratic shape in Panel A (Figure 3) is that less knowledgeable or competent individuals are less attuned to how much they know (i.e., the classic Dunning-Kruger effect (Kruger & Dunning, 1999)). However, an alternative explanation is heterogeneity among the participants. In an exploratory analysis, we included a dummy variable for whether an individual was in the top quartile of the average of the three Dark Triad measures or not, and an interaction term between the dummy and accuracy. The result is depicted in Panel B of Figure 3. Here we found a significant interaction effect ($b = -.421$, $SE = .103$, 95% CI [-.624, -.218], $p < .001$) such that the original non-linearity appears to be driven by two types of participants—those who score highly on the Dark Triad whose relation between accuracy and confidence is flatter and those who do not, who exhibit a steeper slope. This result opens the door to future research into how other individual differences—such as the Dark Triad—might impact the calibration of confidence to accuracy.

As well as studying these questions at an aggregate level, a wealth of research considers how well individuals are able to map their confidence to their accuracy across items, such as the ability to assign different levels of confidence to correct and incorrect answers—termed 'resolution' (Murphy, 1973; Yates, 1982). For instance, one way to be well-resolved would be to assign a higher average level of confidence to correct answers than incorrect answers, with minimal variance around these averages (Ronis and Yates (1987) refer to these features of the distribution of confidence judgments as *slope* and *scatter*, respectively). Resolution is sometimes referred to as metacognitive ability, and researchers have been especially interested in isolating metacognitive ability from knowledge (e.g., Grabman & Dodson, 2024), developing specialized measures of *metacognitive efficiency* to do so using a signal detection framework (e.g., Fleming & Lau, 2014). Interestingly, research in this paradigm has found stable individual differences in metacognitive efficiency (Ais et al., 2016). Future research might endeavor to integrate the fields of overconfidence and metacognition to benefit from their synthesis, though many overconfidence measures (e.g., overplacement, overprecision) are not currently amenable to the methods used to study metacognition due to lacking item-level measures of confidence, and thus further methodological advances would first be required.

Implications for Constructs Related to Overconfidence

We further consider the implications of our results for literature concerning adjacent concepts. For example, overclaiming has been connected to both overestimation (Bensch et al., 2019) and overplacement (Luo et al., 2016). Several researchers have considered overclaiming as another instantiation of overconfidence (Anderson & Brion, 2012; Bensch et al., 2019). Our data show that the two are closely related, though further work is required to verify if the *psychological bases* are similar. We would argue that overclaiming is primarily about recognition and not confidence. Someone can perceive familiarity with an item without having confidence that their perception is veridical.

In the economics tradition, overconfidence has been found to be related to a host of important outcomes such as having lower portfolio diversification, higher merger and acquisition activity, and affecting the raising of capital (Ben-David et al., 2007, 2013; Brown & Sarma, 2007; Malmendier et al., 2011; Malmendier & Tate, 2004, 2005a, 2005b). Much of this evidence is indirect because the measures are confounded with risk preferences (Malmendier & Tate, 2008), or involve decisions made in teams (Ben-David et al., 2007). Thus, there may be other mechanisms at work here aside from individual differences in overconfidence, but, critically, the present research supports the assumption of individual differences present in this research, facilitating future research in this tradition.

Measuring Core Overconfidence Parsimoniously

As mentioned, in Study 1b and Studies 2a-b (as well as in Study S1), we included three-items that constitute what we call the "overconfidence test" (OCT), a short test that aims to explain substantial variance in core overconfidence. The items from this test are featured in Appendix A; each requires a response from 0-100%, and the three items are averaged to make an overall score. We found promising evidence for the feasibility of this approach—across studies, its Cronbach's alpha ranged from .37 to .68, and its correlation with core overconfidence from r = .41 to .59, including tests with and without incentives, across time periods, and across two different survey platforms. By developing this short-form measure, we hope that trait-level overconfidence can be measured parsimoniously in a host of research designs. For further details of these analyses, refer to Lawson et al. (2024).

Limitations

There are some notable caveats about our results that both limit our conclusions and offer avenues of future investigation. For example, when we deliberately weakened the correlations in accuracy across domain measures in Studies 2a-b, though the majority of correlations between overconfidence measures remained positive, some dropped below the threshold for statistical significance, and some were negative. As mentioned above, stability in accuracy and stability in people's self-views are both important contributors to individual differences in core overconfidence. Future work can further map the relationships among domains to identify the boundary conditions where average correlations shrink enough that the alpha of core overconfidence falls below .7, which we found once, in Study 2b.

It is also important to highlight several nuances in the interpretations of our analyses of the relationship between core overconfidence and other individual difference measures. First, we presented analyses showing that the relationship between some of the psychological traits and overconfidence depend on the form of overconfidence (i.e., overestimation, overplacement, or overprecision). The exceptions, however, were dominance, narcissism, and psychopathy, which correlated with all forms, and arguably provide the most robust psychological insight about the source of individual differences in core overconfidence. For other psychological traits, however, the pattern is mixed and points to the need for caution in assuming a link between the trait and the dispositional component of overconfidence. Moreover, we note that the relationships between individual differences and core overconfidence may be non-linear. Although we focused on linear correlations, we found some evidence for a small cluster of participants who score highly in all three of the Dark Triad measures (Jones & Paulhus, 2014) who score particularly highly on all measures of overconfidence. Understanding the predictive validity of individual difference measures in different *ranges* of their values is an important area for future investigation.

Another potential limitation is that that popular measures of overconfidence can be confounded with ability due to i) accuracy being an imperfect measure of ability, and ii) predictions regarding performance (i.e., confidence) rationally including prior beliefs about one's own ability (Spiller, 2024). This could potentially explain some of our results. For instance, the stability of core overconfidence may be inflated because the same test was used on both administrations, though we verified its stability was robust when using disjoint sets of questions, so this concern is largely mitigated. Regardless, given that our tests cannot be perfectly correlated with ability (despite being long and comprehensive), error in measuring ability (which is likely stable) would lead to core overconfidence being correlated over time, which could partly explain some of the effect.

One important question is whether we would expect to see such high reliability of measures of overconfidence in high stakes, incentivized situations or in different, more diverse samples. Our data are primarily from Amazon's Mechanical Turk (MTurk) samples. Given that MTurk samples can sometimes suffer from sampling and data quality issues, it is possible that the correlations among measures of overconfidence will be weaker when participants pay closer attention (Chandler et al., 2019; Stewart et al., 2017). That said, we tested the role of incentives in Study S1 and found that they did not have a substantial impact.

We also note that our experiments were conducted on U.S. based participants—and so may not generalize to different populations (Henrich et al., 2010). Future work should aim to verify how the stability of overconfident beliefs varies across nationalities and cultures, as well as in population that differ in all manner of attributes.

Conclusion

In this research, we aimed to reconcile divergent perspectives in the literature, where some areas consider overconfidence an individual difference, and others do not. To this end, we asked the same participants many measures of overconfidence, varying domains, forms, and time periods. What became clear is that though overconfidence varies substantially across tasks, it is the same people who tend to be relatively more or less overconfident across tasks. In subsequent studies we investigated the *nature* of this stability, inspecting both its psychological roots via its relationship with individual differences, as well as highlighting the roles of accuracy and confidence in its observed stability. We conclude by highlighting some implications for overconfidence research stimulated by the following insight: There are stable individual differences in overconfidence.

References

Ais, J., Zylberberg, A., Barttfeld, P., & Sigman, M. (2016). Individual consistency in the accuracy and distribution of confidence judgments. *Cognition*, *146*, 377–386. https://doi.org/10.1016/j.cognition.2015.10.006

Ames, D. R., & Kammrath, L. K. (2004). Mind-Reading and Metacognition: Narcissism, not Actual Competence, Predicts Self-Estimated Ability. *Journal of Nonverbal Behavior*, *28*(3), 187–209. https://doi.org/10.1023/B:JONB.0000039649.20015.0e

Ames, D. R., Rose, P., & Anderson, C. P. (2006). The NPI-16 as a short measure of narcissism. *Journal of Research in Personality*, *40*(4), 440–450. https://doi.org/10.1016/j.jrp.2005.03.002

Anderson, C., & Brion, S. (2012). *Overconfidence and the Attainment of Status in Groups*. 61.

Ashton, M. C., & Lee, K. (2009). The HEXACO–60: A Short Measure of the Major Dimensions of Personality. *Journal of Personality Assessment*, *91*(4), 340–345. https://doi.org/10.1080/00223890902935878

Baron, J. (2019). Actively open-minded thinking in politics. *Cognition*, *188*, 8–18. https://doi.org/10.1016/j.cognition.2018.10.004

Ben-David, I., Graham, J., & Harvey, C. (2007). *Managerial Overconfidence and Corporate Policies*. National Bureau of Economic Research. https://doi.org/10.3386/w13711

Ben-David, I., Graham, J. R., & Harvey, C. R. (2013). Managerial Miscalibration*. *The Quarterly Journal of Economics*, *128*(4), 1547–1584. https://doi.org/10.1093/qje/qjt023

Bensch, D., Paulhus, D. L., Stankov, L., & Ziegler, M. (2019). Teasing Apart Overclaiming, Overconfidence, and Socially Desirable Responding. *Assessment*, *26*(3), 351–363. https://doi.org/10.1177/1073191117700268

Binnendyk, J., & Pennycook, G. (2024). Individual differences in overconfidence: A new measurement approach. *Judgment and Decision Making*, *19*. https://doi.org/10.1017/jdm.2024.22

Blais, A.-R., Thompson, M. M., & Baranski, J. V. (2005). Individual differences in decision processing and confidence judgments in comparative judgment tasks: The role of cognitive styles. *Personality and Individual Differences*, *38*(7), 1701–1713. https://doi.org/10.1016/j.paid.2004.11.004

Boutros, M., Ben-David, I., Graham, J. R., Harvey, C. R., & Payne, J. W. (2020). *The Persistence of Miscalibration* (Working Paper 28010). National Bureau of Economic Research. https://doi.org/10.3386/w28010

Brown, R., & Sarma, N. (2007). CEO overconfidence, CEO dominance and corporate acquisitions. *Journal of Economics and Business*, *59*(5), 358–379. https://doi.org/10.1016/j.jeconbus.2007.04.002

Burger, J. M., & Cooper, H. M. (1979). The desirability of control. *Motivation and Emotion*, *3*(4), 381–393. https://doi.org/10.1007/BF00994052

Burson, K. A., Larrick, R. P., & Klayman, J. (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*(1), 60–77. https://doi.org/10.1037/0022-3514.90.1.60

Camerer, C., & Lovallo, D. (1999). Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review*, *89*(1), 306–318. https://doi.org/10.1257/aer.89.1.306

Campbell, W. K., Goodie, A. S., & Foster, J. D. (2004). Narcissism, confidence, and risk attitude. *Journal of Behavioral Decision Making*, *17*(4), 297–311. https://doi.org/10.1002/bdm.475

Chandler, J., Rosenzweig, C., Moss, A. J., Robinson, J., & Litman, L. (2019). Online panels in social science research: Expanding sampling methods beyond Mechanical Turk. *Behavior Research Methods*, *51*(5), 2022–2038. https://doi.org/10.3758/s13428-019-01273-7

Chen, G., Gully, S. M., & Eden, D. (2001). Validation of a New General Self-Efficacy Scale. *Organizational Research Methods*, *4*(1), 62–83. https://doi.org/10.1177/109442810141004

Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring Risk Literacy: The Berlin Numeracy Test. *Judgement and Decision Making*, *7*(1), 25–47.

Cortina, J. M. (1993). What is coefficient alpha? An examination of theory and applications. *Journal of Applied Psychology*, *78*(1), 98–104. https://doi.org/10.1037/0021-9010.78.1.98

Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, *16*, 297–334. https://doi.org/10.1007/BF02310555

Deaves, R., Lüders, E., & Luo, G. Y. (2009). An Experimental Test of the Impact of Overconfidence and Gender on Trading Activity. *Review of Finance*, *13*(3), 555–575. https://doi.org/10.1093/rof/rfn023

Ehrlinger, J., & Dunning, D. (2003). How chronic self-views influence (and potentially mislead) estimates of performance. *Journal of Personality and Social Psychology*, *84*(1), 5–17. https://doi.org/10.1037/0022-3514.84.1.5

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519–527. https://doi.org/10.1037/0033-295X.101.3.519

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, *3*(4), 552–564. https://doi.org/10.1037/0096-1523.3.4.552

Fleming, S. M. (2024). Metacognition and Confidence: A Review and Synthesis. *Annual Review of Psychology*, *75*(Volume 75, 2024), 241–268. https://doi.org/10.1146/annurev-psych-022423-032425

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00443

Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, *19*(4), 25–42. https://doi.org/10.1257/089533005775196732

Galasso, A., & Simcoe, T. S. (2011). CEO Overconfidence and Innovation. *Management Science*, *57*(8), 1469–1484. https://doi.org/10.1287/mnsc.1110.1374

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, *98*(4), 506–528. https://doi.org/10.1037/0033-295X.98.4.506

Glaser, M., Langer, T., & Weber, M. (2013). True Overconfidence in Interval Estimates: Evidence Based on a New Measure of Miscalibration. *Journal of Behavioral Decision Making*, *26*(5), 405–417. https://doi.org/10.1002/bdm.1773

Glaser, M., & Weber, M. (2007). Overconfidence and trading volume. *The Geneva Risk and Insurance Review*, *32*(1), 1–36. https://doi.org/10.1007/s10713-007-0003-3

Goodman-Delahunty, J., Granhag, P., Hartwig, M., & Loftus, E. (2010). Insightful or Wishful: Lawyers' Ability to Predict Case Outcomes. *Psychology Public Policy and Law*, *16*. https://doi.org/10.1037/a0019060

Grabman, J. H., & Dodson, C. S. (2024). Unskilled, underperforming, or unaware? Testing three accounts of individual differences in metacognitive monitoring. *Cognition*, *242*, 105659. https://doi.org/10.1016/j.cognition.2023.105659

Haran, U., Ritov, I., & Mellers, B. A. (2013). The role of actively open-minded thinking in information acquisition, accuracy, and calibration. *Judgment and Decision Making*, *8*(3), 188–201. https://doi.org/10.1017/S1930297500005921

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, *466*(7302), Article 7302. https://doi.org/10.1038/466029a

Hilton, D., Régner, I., Cabantous, L., Charalambides, L., & Vautier, S. (2011). Do positive illusions predict overconfidence in judgment? A test using interval production and probability evaluation measures of miscalibration. *Journal of Behavioral Decision Making*, *24*(2), 117–139. https://doi.org/10.1002/bdm.678

Jansen, R. A., Rafferty, A. N., & Griffiths, T. L. (2021). A rational model of the Dunning–Kruger effect supports insensitivity to evidence in low performers. *Nature Human Behaviour*, *5*(6), Article 6. https://doi.org/10.1038/s41562-021-01057-0

Jones, D. N., & Paulhus, D. L. (2014). Introducing the Short Dark Triad (SD3): A Brief Measure of Dark Personality Traits. *Assessment*, *21*(1), 28–41. https://doi.org/10.1177/1073191113514105

Jonsson, A.-C., & Allwood, C. M. (2003). Stability and variability in the realism of confidence judgments over time, content domain, and gender. *Personality and Individual Differences*, *34*(4), 559–574. https://doi.org/10.1016/S0191-8869(02)00028-4

Kahneman, D., & Lovallo, D. (1993). Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking. *Management Science*, *39*(1), 17–31. JSTOR.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, *103*(3), 582–591. https://doi.org/10.1037/0033-295X.103.3.582

Klayman, J., Soll, J. B., González-Vallejo, C., & Barlas, S. (1999). Overconfidence: It Depends on How, What, and Whom You Ask. *Organizational Behavior and Human Decision Processes*, *79*(3), 216–247. https://doi.org/10.1006/obhd.1999.2847

Klayman, J., Soll, J. B., Juslin, P., & Winman, A. (2006). Subjective confidence and the sampling of knowledge. In *Information Sampling and Adaptive Cognition*. Cambridge University Press.

Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, *100*(4), 609–639. https://doi.org/10.1037/0033-295X.100.4.609

Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for Confidence. *Journal of Experimental Psychology: Human Learning and Memory*, *6*(2). https://media.gradebuddy.com/documents/1657138/1d2520cf-4a80-4d38-be44-2b4255bc0a46.pdf

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*(6), 1121–1134. https://doi.org/10.1037/0022-3514.77.6.1121

Kruglanski, A. W., Webster, D. M., & Klem, A. (1993). Motivated resistance and openness to persuasion in the presence or absence of prior information. *Journal of Personality and Social Psychology*, *65*(5), 861–876. https://doi.org/10.1037//0022-3514.65.5.861

Larrick, R. P., Burson, K. A., & Soll, J. B. (2007). Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes*, *102*(1), 76–94. https://doi.org/10.1016/j.obhdp.2006.10.002

Lawson, M. A. (2022). *Stable Individual Differences in Overconfidence* [Doctoral]. Duke University.

Lawson, M. A., Larrick, R. P., & Soll, J. B. (2024). *The Overconfidence Test (OCT): A parsimonious measure of trait-level overconfidence* (SSRN Scholarly Paper 4893481). Social Science Research Network. https://doi.org/10.2139/ssrn.4893481

Leary, M. R., Diebels, K. J., Davisson, E. K., Jongman-Sereno, K. P., Isherwood, J. C., Raimi, K. T., Deffler, S. A., & Hoyle, R. H. (2017). Cognitive and Interpersonal Features of Intellectual Humility. *Personality and Social Psychology Bulletin*, *43*(6), 793–813. https://doi.org/10.1177/0146167217697695

Logg, J. M., Haran, U., & Moore, D. A. (2018). Is overconfidence a motivated bias? Experimental evidence. *Journal of Experimental Psychology: General*, *147*(10), 1445–1465. https://doi.org/10.1037/xge0000500

Luo, Y. L. L., Liu, Y., Cai, H., Wildschut, T., & Sedikides, C. (2016). Nostalgia and Self-Enhancement: Phenotypic and Genetic Approaches. *Social Psychological and Personality Science*, *7*(8), 857–866. https://doi.org/10.1177/1948550616660158

Macenczak, L. A., Campbell, S., Henley, A. B., & Campbell, W. K. (2016). Direct and interactive effects of narcissism and power on overconfidence. *Personality and Individual Differences*, *91*, 113–122. https://doi.org/10.1016/j.paid.2015.11.053

Malmendier, U., & Tate, G. (2004). Who Makes Acquisitions? A Test of the Overconfidence Hypothesis. *NBER Working Paper*, 73.

Malmendier, U., & Tate, G. (2005a). CEO Overconfidence and Corporate Investment. *The Journal of Finance*, *60*(6), 2661–2700. https://doi.org/10.1111/j.1540-6261.2005.00813.x

Malmendier, U., & Tate, G. (2005b). Does Overconfidence Affect Corporate Investment? CEO Overconfidence Measures Revisited. *European Financial Management*, *11*(5), 649–659. https://doi.org/10.1111/j.1354-7798.2005.00302.x

Malmendier, U., & Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics*, *89*(1), 20–43. https://doi.org/10.1016/j.jfineco.2007.07.002

Malmendier, U., Tate, G., & Yan, J. (2011). Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies. *The Journal of Finance*, *66*(5), 1687–1733. https://doi.org/10.1111/j.1540-6261.2011.01685.x

McDonald, R. P. (1999). *Test Theory: A Unified Treatment*. Psychology Press. https://doi.org/10.4324/9781410601087

Merkle, C. (2017). Financial overconfidence over time: Foresight, hindsight, and insight of investors. *Journal of Banking & Finance*, *84*, 68–87. https://doi.org/10.1016/j.jbankfin.2017.07.009

Miller, J. D., Lynam, D. R., Hyatt, C. S., & Campbell, W. K. (2017). Controversies in Narcissism. *Annual Review of Clinical Psychology*, *13*(1), 291–315. https://doi.org/10.1146/annurev-clinpsy-032816-045244

Moore, D. A. (2020). *Perfectly confident: How to calibrate your decisions wisely*. HarperCollins.

Moore, D. A., & Dev, A. S. (2018). Individual Differences in Overconfidence. In *Encyclopedia of Personality and Individual Differences.* http://doi.apa.org/getdoi.cfm?doi=10.1037/pspa0000018

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, *115*(2), 502–517. https://doi.org/10.1037/0033-295X.115.2.502

Moore, D. A., & Small, D. A. (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*(6), 972–989. https://doi.org/10.1037/0022-3514.92.6.972

Moore, D. A., & Swift, S. A. (2010). The Three Faces of Overconfidence in Organizations. In *Social Psychology of Organizations*. Taylor & Francis.

Murphy, A. H. (1973). *A New Vector Partition of the Probability Score*. https://journals.ametsoc.org/view/journals/apme/12/4/1520-0450_1973_012_0595_anvpot_2_0_co_2.xml

Nelson, T. O., & Narens, L. (1990). Metamemory: A Theoretical Framework and New Findings. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 26, pp. 125–173). Academic Press. https://doi.org/10.1016/S0079-7421(08)60053-5

Oakley, J. (2021). *SHELF: Tools to Support the Sheffield Elicitation Framework* (Version 1.8.0) [Computer software]. https://CRAN.R-project.org/package=SHELF

Oberlechner, T., & Osler, C. (2012). Survival of Overconfidence in Currency Markets. *Journal of Financial and Quantitative Analysis*, *47*(1), 91–113. https://doi.org/10.1017/S0022109012000038

Olsson, H. (2014). Measuring overconfidence: Methodological problems and statistical artifacts. *Journal of Business Research*, *67*(8), 1766–1770. https://doi.org/10.1016/j.jbusres.2014.03.002

O'Reilly, C. A., & Hall, N. (2021). Grandiose narcissists and decision making: Impulsive, overconfident, and skeptical of experts–but seldom in doubt. *Personality and Individual Differences*, *168*, 110280. https://doi.org/10.1016/j.paid.2020.110280

Oskamp, S. (1965). Overconfidence in case-study judgments. *Journal of Consulting Psychology*, *29*(3), 261–265. https://doi.org/10.1037/h0022125

Pallier, G., Wilkinson, R., Danthiir, V., Kleitman, S., Knezevic, G., Stankov, L., & Roberts, R. D. (2002). The Role of Individual Differences in the Accuracy of Confidence Judgments. *The Journal of General Psychology*, *129*(3), 257–299. https://doi.org/10.1080/00221300209602099

Paulhus, D. L., Harms, P. D., Bruce, M. N., & Lysy, D. C. (2003). The over-claiming technique: Measuring self-enhancement independent of ability. *Journal of Personality and Social Psychology*, *84*(4), 890–904. https://doi.org/10.1037/0022-3514.84.4.890

Raiffa, H. (1968). *Decision analysis: Introductory lectures on choices under uncertainty*. Addison-Wesley.

Revelle, W. (2024). *psych: Procedures for Psychological, Psychometric, and Personality Research* (Version 2.4.6.26) [Computer software]. https://cran.r-project.org/web/packages/psych/index.html

Revelle, W., & Zinbarg, R. E. (2009). Coefficients Alpha, Beta, Omega, and the glb: Comments on Sijtsma. *Psychometrika*, *74*(1), 145–154. https://doi.org/10.1007/s11336-008-9102-z

Ronis, D. L., & Yates, J. F. (1987). Components of probability judgment accuracy: Individual consistency and effects of subject matter and assessment method. *Organizational Behavior and Human Decision Processes*, *40*(2), 193–218. https://doi.org/10.1016/0749-5978(87)90012-4

Sanchez, C., & Dunning, D. (2018). Overconfidence among beginners: Is a little learning a dangerous thing? *Journal of Personality and Social Psychology*, *114*(1), 10–28. https://doi.org/10.1037/pspa0000102

Schaefer, P. S., Williams, C. C., Goodie, A. S., & Campbell, W. K. (2004). Overconfidence and the Big Five. *Journal of Research in Personality*, *38*(5), 473–480. https://doi.org/10.1016/j.jrp.2003.09.010

Scheier, M. F., Carver, C. S., & Bridges, M. W. (1994). Distinguishing optimism from neuroticism (and trait anxiety, self-mastery, and self-esteem): A reevaluation of the Life Orientation Test. *Journal of Personality and Social Psychology*, *67*(6), 1063–1078. https://doi.org/10.1037/0022-3514.67.6.1063

Sedikides, C., & Gregg, A. P. (2008). Self-Enhancement: Food for Thought. *Perspectives on Psychological Science*, *3*(2), 102–116. https://doi.org/10.1111/j.1745-6916.2008.00068.x

Soll, J. B., & Klayman, J. (2004). Overconfidence in Interval Estimates. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(2), 299–314. https://doi.org/10.1037/0278-7393.30.2.299

Soll, J. B., Palley, A. B., Klayman, J., & Moore, D. A. (2023). Overconfidence in Probability Distributions: People Know They Don't Know, but They Don't Know What to Do About It. *Management Science*. https://doi.org/10.1287/mnsc.2019.00660

Spearman, C. (1961). *'General Intelligence' Objectively Determined and Measured* (p. 73). Appleton-Century-Crofts. https://doi.org/10.1037/11491-006

Stankov, L., & Crawford, J. D. (1996). Confidence judgments in studies of individual differences. *Personality and Individual Differences*, *21*(6), 971–986. https://doi.org/10.1016/S0191-8869(96)00130-4

Stanovich, K. E., & West, R. F. (1998). Individual Differences in Rational Thought. *Journal of Experimental Psychology: General*, *127*(2), 161–188.

Stewart, N., Chandler, J., & Paolacci, G. (2017). Crowdsourcing Samples in Cognitive Science. *Trends in Cognitive Sciences*, *21*(10), 736–748. https://doi.org/10.1016/j.tics.2017.06.007

Webster, D. M., & Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, *67*(6), 1049–1062. https://doi.org/10.1037/0022-3514.67.6.1049

West, R. F., & Stanovich, K. E. (1997). The domain specificity and generality of overconfidence: Individual differences in performance estimation bias. *Psychonomic Bulletin & Review*, *4*(3), 387–392. https://doi.org/10.3758/BF03210798

Wiggins, J. S., Trapnell, P., & Phillips, N. (1988). Psychometric and Geometric Characteristics of the Revised Interpersonal Adjective Scales (IAS-R). *Multivariate Behavioral Research*, *23*(4), 517–530. https://doi.org/10.1207/s15327906mbr2304_8

Winman, A., Hansson, P., & Juslin, P. (2004). Subjective Probability Intervals: How to Reduce Overconfidence by Interval Evaluation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(6), 1167–1175. https://doi.org/10.1037/0278-7393.30.6.1167

Yates, J. F. (1982). External correspondence: Decompositions of the mean probability score. *Organizational Behavior and Human Performance*, *30*(1), 132–156. https://doi.org/10.1016/0030-5073(82)90237-9

Zinbarg, R. E., Revelle, W., Yovel, I., & Li, W. (2005). Cronbach's α, Revelle's β, and Mcdonald's ω_H: Their relations with each other and two alternative conceptualizations of reliability. *Psychometrika*, *70*(1), 123–133. https://doi.org/10.1007/s11336-003-0974-7

Appendix A: Items of the Overconfidence Test (OCT)

You're in a large group of people playing a new game, Zoob. You don't know how your ability compares to other people's. You play one game against a randomly selected person and win. What percentage of people do you think you're better than?
Imagine you're on a committee deciding between 5 different project options. There are 10 members of the committee. You propose Option A. The next three committee members propose Option B. The six other committee members will propose their choices next. When the meeting is finished, what percentage of the choices do you think will be Option A?
One-hundred people are guessing the number of jellybeans in a jar. The closest 10 guesses win $100. How likely are you to be one of the winners?