Examining the replicability of online experiments selected by a decision market

Felix Holzmeister[1], Magnus Johannesson[2], Colin F. Camerer[3], Yiling Chen[4], Teck-Hua Ho[5], Suzanne Hoogeveen[6], Juergen Huber[7], Noriko Imai[8], Taisuke Imai[8], Lawrence Jin[9], Michael Kirchler[7], Alexander Ly[10,11], Benjamin Mandl[12], Dylan Manfredi[13], Gideon Nave[13], Brian A. Nosek[14,15], Thomas Pfeiffer[16], Alexandra Sarafoglou[10], Rene Schwaiger[7], Eric-Jan Wagenmakers[10], Viking Waldén[17], Anna Dreber[1,2,*]

[1]Department of Economics, University of Innsbruck, Innsbruck, Austria. [2] Department of Economics, Stockholm School of Economics, Stockholm, Sweden. [3] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, USA. [4] John A. Paulson School of Engineering and Applied Sciences, Harvard University, Boston, USA. [5] Nanyang Technological University, Singapore. [6] Faculty of Social and Behavioural Sciences, Utrecht University, Utrecht, The Netherlands. [7] Department of Banking and Finance, University of Innsbruck, Innsbruck, Austria. [8] Institute of Social and Economic Research, Osaka University, Osaka, Japan. [9] Lee Kuan Yew School of Public Policy, National University of Singapore, Singapore. [10] Faculty of Social and Behavioural Sciences, University of Amsterdam, Amsterdam, [11] Machine Learning, Centrum Wiskunde and Informatica, Amsterdam, The Netherlands. [12] Independent researcher, Vienna, Austria. [13] Marketing Department, Wharton School, University of Pennsylvania, Philadelphia, USA. [14]Department of Psychology, University of Virginia, Charlottesville, USA. [15]Center for Open Science, Charlottesville, Virginia, USA. [16] Institute for Advanced Study, Massey University, Auckland, New Zealand. [17]Sveriges Riksbank, Stockholm, Sweden.

[*] To whom correspondence should be addressed:

Anna Dreber
Department of Economics, Stockholm School of Economics
Box 6501, SE-113 83 Stockholm, Sweden.
Email: anna.dreber@hhs.se

Abstract

In this study, we test the feasibility of using decision markets to select studies for replication and provide evidence about the replicability of online experiments. Social scientists (n = 162) traded on the outcome of close replications of 41 systematically selected MTurk social science experiments published in PNAS 2015–2018, knowing that the 12 studies with the lowest and the 12 with the highest final market prices would be selected for replication, along with two randomly selected studies. The replication rate, based on the statistical significance indicator, was 83% for the top-12 and 33% for the bottom-12 group. Overall, 54% of the studies were successfully replicated, with replication effect size estimates averaging 45% of the original effect size estimates. The replication rate varied between 54% and 62% for alternative replication indicators. The observed replicability of MTurk experiments is comparable to that of previous systematic replication projects involving laboratory experiments.

Main

Can published research findings be trusted? Unfortunately, the answer to this question is not straightforward, and the credibility of scientific findings and methods has been questioned repeatedly[1–9]. A vital tool for evaluating and enhancing the reliability of published findings is to carry out replications, which can be used to sort out likely true positive findings from likely false positives. A replication essentially updates the probability of the hypothesis being true after observing the replication outcome. A successful replication will move this probability towards 100%, while a failed replication will move it towards 0%[10,11]. In recent years, several systematic large-scale replication projects in the social sciences have been published[12–17], reporting replication rates of around 50% in terms of both the fraction of statistically significant replications and the relative effect sizes of replications. Potential factors to explain these replication rates may be low statistical power[1,18,19] in the original studies, testing original hypotheses with low priors[1,10,20], and questionable research practices[1,21,22]. Systematic replication studies led to discussions about improving research practices[23,24] and have substantially increased the interest in independent replications[25]. However, as it is time-consuming and costly to conduct replications, it has been argued that it is useful to have a principled mechanism to decide which replications to prioritize to facilitate efficient and effective usage of resources[25–37]. In this study, we test the feasibility of one potential method to select which studies to replicate. Building on previous work using prediction markets[38–40] to forecast replicability, we adapt the forecasting methodology to what is referred to as decision markets[41–44].

The decisive distinction between prediction markets and decision markets is that prediction markets elicit aggregate-level replicability forecasts on a predetermined set of studies, whereas decision market forecasts determine which studies are going to be put to a replication test. While previous studies provide evidence that prediction market forecasts are predictive of replication outcomes[10,16,17,45], prediction efficiency might not generalize to decision markets, which involve more complex procedures and incentives. The performance of decision markets as a tool for selecting which empirical claims to replicate has not been systematically examined. Note that a decision market in itself is not sufficient to provide a mechanism to select studies for replication, but it has to be combined with an objective function of which studies to replicate (an example of an objective function would be to replicate the studies with the lowest probability of replication). For decision markets to be potentially useful for selecting studies for replication, it first has to be established that the predictions of the decision markets are associated with the replication outcomes. To provide such a “proof of concept” of using a decision market as a mechanism to determine which studies to replicate, we first identified all social science experiments published in the Proceedings of the National Academy of Sciences (PNAS) between 2015 and 2018 that fulfilled our inclusion criteria for (i) the journal and period; (ii) the platform on which the experiment was performed (Amazon Mechanical Turk; MTurk); (iii) the type of design (between-subjects or within-subject treatment design); (iv) the equipment and materials needed to implement the experiment (the experiment had to be logistically feasible for us to implement); and (v) the results

reported in the experiment (at least one main or interaction effect with $p < 0.05$ reported in the main text). Based on our inclusion criteria, we identified 44 articles, three of which have been excluded due to a lack of feasibility, leaving us with a final sample of 41 articles[46–86] (see Methods for details on the inclusion criteria). For each of these articles, we identified one critical finding with $p < 0.05$ that we could potentially replicate (see Methods for details and Supplementary Table 1 for the hypotheses selected for each of the 41 studies).

We then invited social science researchers to participate as forecasters in both a prediction survey and an incentivized decision market on the 41 studies. In the survey, the forecasters independently estimated the probability of replication for the 41 studies. In the decision market, they could trade on whether the result of each of the 41 studies would replicate. Participants in the decision market received an endowment of 100 tokens corresponding to USD 50, and 162 participants made a total of 4,412 trades. Traders in the market were informed about the preregistered decision mechanism: The 12 studies with the highest and the 12 studies with the lowest market prices were to be selected for close replication; additionally, two randomly chosen studies (out of the remaining 17 studies) are replicated to ensure incentive compatibility, with participant payoffs scaled up by the inverse of their probability in the decision rule (see Methods for details). For incentive compatibility, all the 41 studies included need to have a strictly positive probability of being selected for replication, which is ensured by having at least one randomly selected study. Otherwise, traders would be incentivized to only trade on those studies that will most likely be chosen according to the decision rule.

All replication experiments, just like all original studies, were conducted on Amazon Mechanical Turk (MTurk), and the same sample restrictions and exclusion criteria as the original studies were applied, which guards against concerns about the potential moderating effects of culture differences in replications[14,87]. Replication sample sizes were determined to have 90% power to detect ⅔ of the effect size reported in the original study at the 5% significance level in a two-sided test (with the effect size estimates having been converted to Cohen's d to have a common standardized effect size measure across the original studies and the replication studies; see Methods for details). If sample size calculations led to replication sample sizes smaller than in the original study, we targeted the same sample size as in the original study. The average sample size in the replications ($n = 1{,}018$) was 3.5 times as large as the average sample size in the original studies ($n = 292$).

The replication results for the 26 MTurk experiments selected by the decision market constitute the second contribution of this project. Systematic evidence on the replicability of online experiments in the social sciences is lacking, and concerns about the quality of online experiments in general—and MTurk studies in particular—have been raised[88–94]. Needless to say, the replication results only pertain to the single focal result selected per paper, and the replication outcome does not necessarily generalize to other results reported in the original articles[95,96]. For convenience, we refer to the replications as "replication of [study reference]" though. Also, our assessment of the most central result may differ from that of the original authors.

Preregistering study protocols and analysis plans have been proposed as a means to reduce questionable research practices. While empirical evidence is still limited, some recent studies suggest that these practices enhance the credibility of published findings[97–99], although potential issues with preregistration have also been raised[100–102]. Prior to starting the survey data collection (that preceded the decision market and replications), we preregistered[103,104] an analysis plan ("replication report") for each of the 41 potential replications at OSF after obtaining feedback from the original authors (https://osf.io/sejyp). After the replications had been conducted, the replication reports of the 26 studies selected for replication were updated with the results of the replications (and potential deviations from the protocol) and were posted to the same OSF repository. We also preregistered an overall analysis plan at OSF before starting the data collection, detailing the study's design and all planned analyses and tests (https://osf.io/xsp6g). Unless explicitly stated, all analyses and tests reported in the manuscript have been preregistered and adhere exactly to our pre-registered analysis plan. The Supplementary Notes details any deviations from the planned design and analyses for the 26 replications.

We preregistered two primary replication indicators and two primary hypotheses. The two primary replication indicators are the relative effect size of the replications and the statistical significance indicator for replication (i.e., whether or not the replication results in a statistically significant effect with $p < 0.05$ in the same direction as the original effect), which was the replication outcome predicted by forecasters in the survey and the decision market.

The statistical significance indicator is a binary criterion of replication and is based on testing the hypothesis for which the original study found support using standard null hypothesis significance testing. The indicator crudely classifies replications as failed or successful depending on whether the replication study yields evidence in support of the original hypothesis at a particular significance threshold. (Critics of null hypothesis significance testing or privileging a p-value of .05 will, justifiably so, object to this crude classification; that is why it is only one of several indicators that we report.) A replication classified as failed based on this indicator, however, does not imply that the estimated replication effect size is significantly different from the original estimate (see more on this below). To keep the false negative risk at bay and to be informative, the statistical significance indicator calls for well-powered replications (as in this study).[105,106] However, a limitation of this indicator for well-powered replication studies is that it may classify a replication as successful even if the observed effect size is substantially smaller (or larger) than in the original study. While the statistical significance indicator dichotomizes replication outcomes into successful and failed, replicability may be perceived as a continuous matter of degree. This is why we also consider the relative effect size—a continuous measure of replicability—as a primary replication indicator. While the relative effect size constitutes an imprecise indicator for an individual replication study, it arguably provides an informative measurement of the extent of replicability for a group of studies as it quantifies the average degree of apparent inflation in the original effect sizes.[107] As all replication indicators have limitations, we preregistered four additional secondary replication indicators. In addition, we report the results for two non-

preregistered replication indicators, which were helpfully suggested during the review process of the manuscript.

In our two primary hypotheses, we conjecture that (i) the decision market prices positively correlate with the replication outcomes and (ii) the standardized effect sizes in the replications are lower than in the original studies. All hypotheses are evaluated using two-tailed tests, and—following Benjamin et al.[108]—we interpret results with $p < 0.005$ as "statistically significant evidence," whereas results with $0.005 \leq p < 0.05$ are considered "suggestive evidence." No adjustments were made for multiple comparisons.

Results

Replication outcomes and decision market performance. Fig. 1 and Supplementary Table 2 show the results for the decision markets where the final market price can be interpreted as the predicted replication probability. The predicted probabilities of replication range from 20.9% to 92.9% for the 41 studies, with a mean of 57.6% (sd = 23.6%). The average predicted probability for the 26 studies eventually selected for replication is 58.5%. Fig. 1 also delineates the replication outcomes based on the statistical significance indicator, which allows for gauging the relationship between the decision market prices and the replication outcomes. In Fig. 2 and Supplementary Table 3, we show the replication results for the 26 studies selected for replication. Of the 26 claims, 14 (53.8%; 95% CI [33.4%, 73.4%]) replicated successfully according to the statistical significance indicator. The point-biserial correlation between decision market prices and the binary replication outcome, testing our first primary hypothesis, is $r = 0.505$ (95% CI [0.146, 0.712]; $t(24) = 2.867$, $p = 0.008$; $n = 26$). Thus, in support of our first primary hypothesis, we find suggestive evidence of a positive association between decision market prices and replication outcomes. As a related secondary hypothesis, we test if the replication rate is lower among the 12 studies with the lowest decision market prices than for the 12 studies with the highest decision market prices. The replication rate is 33.3% (95% CI [9.9%, 65.1%]) for the studies in the "bottom-12" group and 83.3% (95% CI [51.6%, 97.9%]) for the studies in the "top-12" group, yielding suggestive evidence in support of our secondary hypothesis (Fisher's exact test; $\chi^2(1) = 6.171$, $p = 0.036$; $n = 24$). Note that Fisher's test conditions on the margin totals; hence, it is only exact for the conditional distribution and can be overly conservative if the margin totals are unknown[109,110], as is the case in our analysis. Boschloo's test[111], an exact unconditional procedure uniformly more powerful than Fisher's test, also yields suggestive evidence for the difference in proportions between the "top-12" and the "bottom-12" group (not preregistered; 95% CI [0.089, 0.799], $p = 0.017$; $n = 24$).

Relative effect sizes. The mean estimated effect size of the 26 replication studies (in Cohen's d units) is 0.253 (sd = 0.357) compared to 0.563 (sd = 0.426) for the original studies, implying a relative estimated average effect size, just dividing the two numbers, of 45.0%; the difference in estimated effect sizes is statistically significant, supporting our second primary hypothesis of

systematically smaller estimated effect sizes in the replications (Wilcoxon signed-rank test; $z = 4.203$, $p < 0.001$; $n = 26$). The relative effect size can also be estimated for each study separately (reported in Supplementary Table 3) and varies between −17.0% and 136.2%, with a mean estimate across studies of 41.1% (95% CI [24.5%, 57.7%]). For the 14 studies that replicated according to the statistical significance indicator, the first and the second relative effect size measures as defined above are 69.5% and 72.0% (95% CI [54.8%, 89.3%]), indicative of some inflation in original effect sizes even for apparent true positives. The two estimated relative effect size measures for the 12 studies that failed to replicate according to the statistical significance indicator are 3.2% and 5.0% (95% CI [−2.6%, 12.5%]), respectively. Fig. 3 illustrates the relationship between the estimated original and replication effect sizes.

Secondary replication indicators. We also preregistered four secondary replication indicators: the small-telescopes approach[112], the one-sided default Bayes factor[113], the replication Bayes factor[114], and the fixed-effects weighted meta-analytic effect size (see Methods for details). When relying on the small-telescopes approach, testing if the replication effect size is smaller than a "small effect,"[112] 15 studies (57.7%; 95% CI [36.9%, 76.6%]) are considered successful replications (Fig. 4 and Supplementary Table 4). The one-sided default Bayes factor (BF_{+0}) indicates the strength of evidence in favor of the alternative hypothesis as opposed to the null hypothesis. BF_{+0} exceeds one for the 14 studies (53.8%; 95% CI [33.4%, 73.4%]) that replicated according to the statistical significance indicator, with strong evidence ($BF_{+0} > 10$) for the tested hypothesis for nine studies (34.6%; 95% CI [17.2%, 55.7%]); BF_{+0} is below 1 for the 12 replications (46.2%; 95% CI [26.6%, 66.6%]) that failed to replicate according to the statistical significance indicator, with strong evidence ($BF_{+0} < 0.1$) for the null hypothesis for seven studies (26.9%; 95% CI [11.6%, 47.8%]) based on the evidence categories proposed by Jeffreys[115] (Fig. 5 and Supplementary Table 4). The one-sided replication Bayes factor (BF_{R0}) indicates the strength of additional evidence in favor of the alternative hypothesis as opposed to the null hypothesis, given the already acquired evidence based on the original data[114]. Replication Bayes factors lead to similar conclusions as the one-sided default Bayes factors, with $BF_{R0} > 10$ for ten studies (38.5%; 95% CI [20.2%, 59.4%]) and $BF_{R0} < 0.1$ for seven studies (26.9%; 95% CI [11.6%, 47.8%]). One exception to this is the study by Cooney et al.[56], for which the default Bayes factor exceeds one ($BF_{+0} = 8.01$) but the replication Bayes factor is below one ($BF_{R0} = 0.23$) due to the replication effect size being only about a third of the original effect size and a larger sample size in the replication compared to the original study (Fig. 5 and Supplementary Table 4). The meta-analytic effect size is statistically significant at the 5% level for 16 studies (61.5%; 95% CI [40.6%, 79.8%]) and significant at the 0.5% level for 14 studies [53.8%; 95% CI [33.4%, 73.4%]); see Fig. 6 and Supplementary Table 4. The meta-analytic effect sizes should be interpreted cautiously as original effect sizes reported as statistically significant are likely to be overestimated on average due to insufficient sample sizes and, thereby, statistical power (and potentially due to questionable research practices)[18,19]. Overall, the primary and secondary replication indicators yield the same binary conclusions for 23 of the 26 replications.

Non-preregistered replication indicators. Following the suggestion of a reviewer, we report the results for two additional replication indicators. The first alternative replication indicator is a test of whether the replication effect size is statistically significantly different from the original effect size. This indicator is closely related to the prediction interval approach[116] as the results can be illustrated as prediction intervals that the replication effect sizes are evaluated against: if the replication effect size falls outside the 95% prediction interval, the replication and original effect sizes differ at the 5% significance level, and the replication is considered a failure. The prediction interval approach, thus, yields a binary replication indicator which is complemented by a continuous replicability measure defined as the p-value of the test of a significant difference between the replication and original effect sizes. We illustrate the prediction interval results in Fig. 7 and report the z-statistics and p-values in Supplementary Table 5 (the p-values are also shown in Fig. 7). According to the prediction interval indicator, 15 studies (57.7%, 95% CI [36.9%, 76.6%]) replicate. This replication rate is close to the result for the statistical significance indicator. However, the classification of nine replication outcomes shifts: For four studies, the classification changes from successful to failed, and for five studies, the classification changes from failed to successful. These changes are due to the fact that low-powered original studies are more likely to replicate, whereas high-powered original studies are less likely to replicate based on the prediction interval indicator (as compared to evaluating replicability based on the statistical significance indicator). According to the prediction interval indicator, six studies failed to replicate among the top-12 studies in terms of decision market prices, whereas three studies failed to replicate among the bottom-12 studies.

Associations between indicators (not preregistered). To examine the relationship between the replication indicators, we estimated Kendall's rank correlations τ_b between all the replication indicators used in the study (see Supplementary Table 6). All preregistered replication indicators are strongly correlated with each other, with τ_b varying between 0.61 and 1.00 ($p < 0.005$ for all correlations). However, they are more weakly to moderately correlated with the prediction interval approach and p-values from z-tests comparing the replication and original effect sizes (with τ_b varying between 0.12 and 0.56).

Each of the various replication indicators presented in this study has its strengths and weaknesses. There is no general consensus about which indicator is most appropriate[15–17,117–119]. Therefore, we chose to report the results for a host of indicators and leave it to readers to judge the suitability of the different indicators and their degree of consensus. The overall replication rate is similar for all the binary replication indicators and varies between 14 (53.8%) and 16 (61.5%) studies. The agreement about which results are classified as successfully replicated is large between the indicators, with the exception of the prediction interval approach. The estimated average relative effect size of around 45%, which can be interpreted in terms of a replicability rate, yields an estimate in the same ballpark. The somewhat lower estimate for the relative effect size is due to the fact that not only the false positive rate but also the inflation of true positive effect sizes is factored in. Another advantage of the relative effect size indicator is that it is not affected by replication power. The three other continuous replication indicators cannot be

aggregated across studies and are thus difficult to compare to the other indicators on an aggregated level.

Replicability forecasts and indicators (not preregistered). In Supplementary Table 7, we also provide Pearson correlations between our replicability forecasts (final decision market prices and average prediction survey beliefs) and all the replication indicators. Both the decision market prices and the survey beliefs are positively correlated with all the replication indicators except the prediction interval approach (although positive, the correlations to the p-value of the test of a significant difference between the replication and original effect sizes are also close to 0). Note that these correlations should be interpreted cautiously as the forecasters predicted the replication outcomes for the statistical significance indicator but not the other replication indicators.

Survey forecasts vs. decision market predictions. We tested three additional preregistered secondary hypotheses based on the survey beliefs about replication (see Methods for details and Supplementary Table 7 for the survey results). The point-biserial correlation between average survey beliefs and the replication outcomes based on the statistical significance criterion is $r = 0.476$ (95% CI [0.107, 0.694]; $t(24) = 2.650$, $p = 0.014$; $n = 26$). The survey beliefs and the decision market prices are positively correlated with a Pearson correlation of 0.899 (95% CI [0.814, 0.944]; $t(39) = 12.830$, $p < 0.001$; $n = 41$) (Fig. 8a). The final secondary hypothesis tests if the prediction accuracy, measured in terms of the absolute prediction error and the Brier score (i.e., the squared prediction error), is higher for the decision market than the survey forecasts (Fig. 8b). The mean absolute prediction error and the mean Brier score are 0.353 and 0.188 for the decision market, and 0.421 and 0.202 for the survey, respectively, providing suggestive evidence for higher accuracy for the market forecasts based on the absolute prediction error (Wilcoxon signed-rank test: $z = 2.172$, $p = 0.030$; $n = 26$) but not the Brier score (Wilcoxon signed-rank test: $z = 1.181$, $p = 0.238$; $n = 26$). The failure to reject the null hypothesis for the Brier score does not imply that the null hypothesis is true. In the survey, we also elicited forecasters' self-rated expertise for each study. The average self-rated expertise (of participants eventually active in the markets, n = 162) for the 26 replicated studies was 2.31 (sd = 1.40; n = 4,212) on a scale from 1 ("no knowledge of the topic") to 7 ("very high knowledge of the topic"). Supplementary Figure 1 plots the absolute prediction error and the Brier score of the 26 survey and decision market forecasts over the average self-rated expertise per study. We do not find evidence for the prediction accuracy and the average self-rated expertise being significantly correlated (not preregistered; see Supplementary Figure 1 for details).

Beliefs about the Covid-19 pandemic and replicability. A potential issue raised by some original authors in giving feedback on the replication reports prior to the data collection was that the replicability of some original results might be affected by the Covid-19 pandemic (as all the original studies were conducted before the pandemic). We evaluate this possibility in a preregistered exploratory analysis, relying on the forecasters' beliefs about the impact of the pandemic on replicability. As part of the prediction survey, participants were asked to judge whether the pandemic would have affected the likelihood of successful replication, measured on

a scale from −3 (“the pandemic has definitely decreased the probability of replication”) to 3 (“the pandemic has definitely increased the probability of replication”). We test if the average response to this question differs from zero using a one-sample t-test for each of the 26 replications, and we test if the average response across all 26 studies differs from 0. We find a statistically significant result for four and a suggestive result for two replications on beliefs that Covid-19 has affected the replication probability (Supplementary Table 8). For the six studies with suggestive or statistically significant evidence, the estimate is negative for two studies and positive for four; only in two of the cases does the sign of the expectation match the eventual replication outcome. For the average belief about the impact of the pandemic on replicability across the 26 studies of 162 forecasters (who were active in the decision markets), there is suggestive evidence that the mean of 0.039 (sd = 0.190) differs from zero ($t(161) = 2.598$, $p = 0.010$; $n = 162$). Somewhat surprisingly—and in contrast to the concerns raised by some of the original authors—there is thus a tendency for forecasters to believe that the pandemic has *increased* the average likelihood that the studies will replicate. However, the magnitude of the effect is small ($d = 0.204$; 95% CI [0.049, 0.360]).

In addition, we tested, estimating the point-biserial correlation, if the average belief (per study) about the pandemic’s impact on replicability correlates with the replication outcomes based on the statistical significance indicator; we do not find a statistically significant association ($r = 0.014$, 95% CI [−0.360, 0.382]; $t(24) = 0.068$, $p = 0.946$; $n = 26$). Yet, we cannot rule out that Covid-19 has entailed effects on replicability not foreseen by scholars participating in the survey. Further work is needed to gauge whether and to which extent experimental replications—and predictions of replication success—might be sensitive to macro-historical secular change such as economic upheaval, wars, pandemics, etc. Forecasters’ beliefs about the pandemic’s impact on replicability are also neither statistically significantly correlated with the final decision market prices ($r = 0.387$, 95% CI [−0.008, 0.669]; $t(24) = 2.055$, $p = 0.051$; $n = 26$) nor the average survey belief of replication ($r = 0.347$, 95% CI [−0.053, 0.644]; $t(24) = 1.815$, $p = 0.082$; $n = 26$), although the point estimates of the correlations are quite sizeable.

Original p-value and replication (not preregistered). For comparison to previous systematic large-scale replication projects, we also report the correlation between the original p-value and the replication outcome for the statistical significance indicator. The point-biserial correlation between the original p-value and the replication outcome for the statistical significance indicator is −0.400 ($p = 0.043$; 95% CI [0.014, 0.648]) and comparable in magnitude to correlations of −0.33 in the RPP[15], −0.57 in the EERP[16], and −0.40 in the SSRP[17].

Discussion

We found suggestive evidence ($p < 0.05$) for our first primary hypothesis that final decision market prices correlate with replication outcomes ($r = 0.505$). However, the estimated effect size is somewhat smaller than the effect size of $r = 0.67$, as presumed in our a priori power calculations

(see Methods for details). The estimated correlation is within the range of previous prediction markets on systematic replication projects with correlations of 0.42 in the Replication Project: Psychology (RPP)[10,15], 0.30 in the Experimental Economics Replication Project (EERP)[16], and 0.84 in the Social Sciences Replication Project (SSRP)[17], but we expected a stronger correlation because we selected studies with the highest and the lowest prices for replication. Consistent with the primary hypothesis test, there is also suggestive evidence of a difference in the replication rate between the “top-12” (10 of 12) and “bottom-12” (4 of 12) in our secondary hypothesis test. The difference of 50 percentage points is also reflected in the difference between the forecasted replication rates of 86.6% ("top-12") vs. 29.6% ("bottom-12") in the decision market. However, the small sample size suggests caution against drawing firm conclusions about whether decision markets are appropriate for selecting studies for replication.

The pooled evidence from previous prediction market studies on replication outcomes suggests that markets are somewhat more accurate than surveys[45], although the difference tends to be small. These indications are consistent with our results, yielding suggestive evidence of higher accuracy in terms of the absolute prediction error but not in terms of the squared prediction error (although, as noted above, failing to reject the null hypothesis for the squared prediction error does not imply that the null hypothesis is true). The estimated correlation between the average survey beliefs and the replication outcomes was almost as high for the survey as the prediction market (0.476 vs. 0.505). The decision market prices and survey beliefs are also highly correlated with each other (r = 0.9). Since surveys are less resource-intensive, simple polls can be an expedient alternative to decision markets for selecting which studies to replicate, even if they should be somewhat less accurate. Another potential method for selecting which studies to replicate would be to rely on the original *p*-value for studies reporting statistically significant results[45]. Although the prediction accuracy appears to be somewhat lower for original *p*-values than market and survey forecasts[45], relying on *p*-values may well be considered a practical alternative as it does not involve any additional data collection. Another possibility would be to use predicted replication probabilities from machine learning models to select studies for replication. There has been some progress in developing such models[120–123], but evidence on whether they outperform markets or surveys is yet missing. Other potential mechanisms for selecting which studies to replicate include relying on general or study-specific characteristics (e.g., connection to theory, surprise factor, sample size, effect size, importance)[25–28], relying on cost-benefit considerations[29,30], employing Bayesian strategies[31,32], determining the “replication value”[33], adopting empirical audit and review[34], selecting studies randomly[35], or using predictions from laypeople[36,37].

Using decision markets to select the studies with the highest and lowest predicted probabilities for replication is just one of many potential selection rules for this methodology. Our goal was to test whether a decision market could distinguish findings that would replicate or not, and we aimed to maximize the statistical power of detecting an association between market prices and replication outcomes. For the practical application of decision markets, the choice of the selection mechanism will largely depend on the objective function. One selection rule would be to choose

the studies with the highest predicted false positive likelihood, i.e., the studies with the smallest market prices (in addition to at least one randomly selected study to ensure incentive compatibility). This decision mechanism would align with the objective of identifying and correcting false discoveries in the literature to facilitate an efficient allocation of resources for follow-up investigations. Another selection rule would be to replicate the studies with market predictions close to 50%, which reflects the highest possible uncertainty or disagreement regarding the likelihood of the original finding being genuinely true. Providing additional evidence on these claims could maximize the information value of replication studies, as well-powered replications will move the probability that the tested hypothesis is genuinely true towards 0% or 100%.

For our second primary hypothesis, we found strong evidence that original effect sizes are inflated on average compared to replication effect sizes, with a relative estimated average effect size of 45%. This is comparable to previous systematic replication studies, with relative average effect size estimates of 49% in the RPP[15], 59% in the EERP[14], and 54% in the *SSRP*[17]. The replication rate of 54% based on the statistical significance indicator is also similar to previous replication studies, with 36% in RPP[15], 61% in the EERP[14], and 62% in the SSRP[17]. Caution should be exercised when comparing the replication results across these studies: the number of replications in each of the projects is small, only one focal result per paper has been selected for replication, and the particular journals and time periods considered differ. However, the results of all these studies are consistent with a replication rate of about 50% for both the binary statistical significance indicator and the continuous relative effect size indicator; compatible replication results have also been observed in the Many Labs replication projects[12–14].

The ability of the statistical significance indicator to discriminate between true positives and false positives depends on replication power, and the relative average effect size of the studies that failed to replicate should be close to zero if the systematic replication study successfully separates false positives from true positives. The relative average estimated effect size of the 12 studies that failed to replicate according to the statistical significance indicator was 3.2%, which is close to zero and consistent with a successful separation between true positives and false positives. But also true positive findings can be expected to have exaggerated effect sizes in the published literature due to a lack of statistical power[18,19]. In line with this, we found an estimated average effect size of 69.5% for the 14 studies that were successfully replicated based on the statistical significance indicator. These findings are consistent with similar analyses in the SSRP[17] in which the estimated mean relative effect size among the studies that failed to replicate according to the statistical significance indicator was 0.3%, and the estimated mean relative effect size among the studies that replicated successfully was 73.1%. This illustrates how the combination of statistical significance and relative effect size can contribute to revealing possible false positives and true positives with exaggerated effect sizes.

Previous systematic replication studies have focused on laboratory experiments rather than online experiments. Concerns have been raised over data quality in online data collections using "crowd workers," as via MTurk[88–94], and part of the rationale for zeroing in on experiments

conducted via MTurk was that we tend to share these concerns. However, the results of this study do not suggest that replicability is substantively lower for experiments conducted via MTurk compared to experiments conducted in physical laboratories for studies published in top journals; more evidence is needed to draw strong conclusions. Relatedly, the predicted average replicability rate of 57.6% in the decision market—despite widespread concerns about data quality on MTurk—is within the range of replication rate forecasts in previous prediction markets of 56% in the RPP[10], 75% in the EERP[16], and 63% in the SSRP[17]. We used IP quality checks[90,124] to minimize the chances of low-quality participant data (see Methods for details), screening out participants before the random assignment into treatments. In total, across all 26 replications, 29% of the participants who accepted a "human intelligence task" (HIT) failed the IP check and were excluded (this descriptive result was not preregistered; see Methods for further details). The replication results from our study should thus not be extrapolated to MTurk experiments not using a comparable screening procedure. An important caveat is that although our IP quality checks seem to have been effective in filtering out bots, this may not be the case for artificial responses generated by large language models like ChatGPT, which could pose a challenge for collecting data online via platforms such as MTurk[125].

There are several important limitations to our study. A successful replication, on its own, does not provide valid evidence for the tested hypothesis. It goes without saying that inference in replication studies is subject to type-I and type-II errors, just as in original studies. Moreover, a finding can be replicable while being based on an invalid experimental design, leading to biased results. An example of this would be an experimental design that systematically results in more attrition in one experimental treatment, causing selection bias in favor of the tested conjecture.[88] Likewise, a failed replication, on its own, does not provide direct evidence against the tested hypothesis. A finding can be unreplicable and based on an invalid experimental design, leaving the hypothesis untested. Although the replication rate for online experiments in our study appears to be similar to previous laboratory evidence, it does not necessarily imply that online and laboratory experiments provide equally valid evidence of the tested hypotheses.

Another limitation is that we only replicate a single focal result per paper, and the replication outcome does not necessarily generalize to other results reported in the original articles. Furthermore, we only gathered data from one online population using the same experimental design as in the original study. It cannot be ruled out that the difference in timing between the replication studies and the original studies has affected the replication results as a consequence of changes in the composition of the MTurk subject pool or the tested phenomenon having changed over time. Large-scale, multi-site replication studies that collect data across various populations and settings, similar to the Many Labs replication projects[12–14], qualify as a promising method to shed light on the heterogeneity of replication effect sizes across populations and designs[126–128] in future replication work, potentially increasing the strength of evidence for whether the hypothesis supported in the original study is likely true or not. Collaborative networks such as the Psychological Science Accelerator[133] facilitate multi-site replication studies and can be a door opener to large and diverse samples.

Another caveat in interpreting our results is the lack of agreement about how to define and measure replicability. We chose to report the results for a broad set of replication indicators proposed in the literature and leave it to readers to gauge the strengths and weaknesses of the various measures. Decision markets come with the limitation of being a relatively resource-intensive tool, rendering simple polls an appealing alternative.

In our proof-of-concept investigation of using decision markets to assess replicability, decision markets show potential as a tool for selecting studies for replications, but further work is needed to draw strong conclusions. The observed replication rate of social science experiments based on data collections via MTurk published in PNAS is comparable to previous systematic replication projects of experimental studies in the social sciences, primarily based on lab experiments. However, the sample size of 26 replication studies is small, implying substantial uncertainty about both the estimated replication rate and the estimated association between the decision market prices and the replication rate. Our study is also limited to one scientific journal, and may not be representative of social science results based on MTurk samples published in other journals, or studies using other online platforms for the data collection. Thus, prudence should be exercised in generalizing our findings and comparing replication results across studies.

Methods

We preregistered an analysis plan for the project at OSF on October 7, 2021, prior to starting the survey data collection (that preceded the decision market and replications), which detailed the design of the study and the exact analyses for all planned analyses and tests (https://osf.io/xsp6g). Unless explicitly mentioned in the main text, we adhere exactly to our pre-analysis plan. The information in this Methods section follows the pre-analysis plan (PAP; with some of the information from the pre-analysis plan reported in the Supplementary Notes). Note that previous systematic replication projects like the RPP[15], the EERP[16], and the SSRP[17] did not file preregistrations of the overall study protocol and planned analyses.

Prior to starting the survey data collection, we also preregistered an analysis plan ("replication report") for each of the 41 potential replications included in the decision market at OSF after obtaining feedback from the original authors (https://osf.io/sejyp). After the replications had been conducted, the 26 replication reports of the replications selected for replication by the decision market were updated with the results of the replications and posted in the same OSF repository. Any deviations from the preregistered analysis plans for the 26 replications are detailed in the 26 "post-replication reports" and listed in Supplementary Notes. We provided all original authors the opportunity to comment on the replication results (without a particular due date) and make the comments publicly available as we receive them alongside the post-replication reports on OSF (https://osf.io/sejyp).

Below, we provide further details on the inclusion criteria, the decision market setup and the survey, the replications, and the replication rate indicators included in the study. The preregistered

analyses and tests were divided into descriptive results of the replication rate among the 26 replicated studies and hypothesis tests. The preregistered descriptive results were furthermore divided into primary replication indicators and secondary replication indicators, and the pre-registered hypothesis tests were divided into (i) primary hypotheses, (ii) secondary hypotheses, and (iii) exploratory analyses. See Supplementary Notes for more details about the preregistered hypothesis tests and exploratory analyses.

Inclusion criteria for studies

We reviewed all *PNAS* articles from 2015–2018 and searched for the terms Amazon Mechanical Turk, MTurk, and Turk. When we began planning our study at the start of 2019, we started reviewing the most recent articles published in PNAS. We then continued to look back in time, year by year, until we reached a sufficiently large number of studies to run a decision market. However, data collection was delayed after some original authors expressed concerns that the Covid-19 pandemic could affect the replication outcomes. We included all social sciences articles that fulfilled our inclusion criteria for (i) the journal and time period, (ii) the platform on which the experiment was performed (MTurk), (iii) the type of design (between-subjects or within-subject treatment design), (iv) the equipment and materials needed to implement the experiment (the experiment had to be logistically feasible for us to implement), and (v) the results reported in the experiment (that there is at least one statistically significant $p < 0.05$ main or interaction effect in the main text). Based on the inclusion criteria, we identified 44 articles. After contacting the original authors, we ended up with 41 articles (the three excluded articles[129–131] involved either software or platforms that no longer existed or methods we were unfamiliar with). In these 41 articles, we identified at least one critical finding that we could replicate. In cases where several studies in the same article fit the inclusion criteria, we randomly picked one of the studies; this was the case for 17 of the 26 replicated studies (Ames and Fiske[46], Atir and Ferguson[47], Baldwin and Lammers[48], Boswell et al.[50], Cooney et al.[56], Genschow et al.[59], Gheorghiu et al.[60], Halevy and Halali[62], Hofstetter et al.[65], John et al.[69], Jordan et al.[70], Klein and O'Brien[73], Kouchaki and Gino[74], McCall et al.[76], Rai et al.[81], Stern et al.[84], and Williams et al.[86]). In cases where the (randomly picked) study contained several conditions, we randomly picked which to compare to the control condition. After that, we looked for the central result with $p < 0.05$ for that particular study. If there were several statistically significant results, one was selected at random. The replication results thus only pertain to the single central result selected per paper, and the replication outcome does not necessarily generalize to other results reported in the original articles. For convenience, we refer to the replications as "replication of [study reference]," though.

For Cheon and Hong[54], the result chosen for replication is reported as part of a 2x2 ANOVA in the original article; since the paper does not report the main effect, the original authors kindly provided us with the corresponding estimates. For Gheorghiu et al.[60], the result to be replicated is only reported with its p-value in the paper; a precise estimate of the test statistic has been obtained from a re-analysis of the original data, which the original authors kindly provided. For

the study by Kraus et al.[75], we could not reproduce the result reported in the original article using the original data. The original authors acknowledged that there had been a reporting error in the original article. For the replication, we use the analysis described in the paper; the effect size and the test statistic reported in the original paper were replaced by the re-estimated result. For the study by Williams et al.[86], the focal hypothesis test in the replication is based on a composite score of five suites of behavior (which are tested separately in the original article) to have a single test. The original authors also report tests on composite measures in the Supporting Information of their article, and they approved the choice to investigate the replicability of the focal hypothesis using a composite score. These changes are transparently reported in the replication reports for each study (see https://osf.io/sejyp for details).

Decision market and prediction survey

We invited researchers to voluntarily participate in the decision market through public mailing lists (ESA and JDM lists) and social media (e.g., Twitter/X); we also emailed colleagues asking them to distribute the call to participants within their professional networks. Participants were required to hold a Ph.D. degree or to be a Ph.D. student currently. In the decision market, participants bet on whether or not the specific result chosen for each study would replicate based on the statistical significance indicator ($p < 0.05$ in the replication and an effect in the same direction as in the original study) as a criterion for replication (thus a binary outcome, as discussed below). Prior to the decision market, participants filled out a survey where we asked them to assign a probability of successful replication to each of the 41 results. The survey is available at https://osf.io/a24zq. Completing the survey was a prerequisite for participating in the markets. We started the recruitment of participants for the decision market on October 4 (2021), and we started sending out the prediction survey on October 8 to those who had signed up for the study (participants who signed up after October 8 received the survey invitation a few days after their registration). The deadline for registering as a participant was October 29, and the deadline for completing the survey was November 5. Overall, 289 participants signed up to participate and were forwarded the link to the survey; 193 participants started the survey, and 162 completed it by the due date. The forecasters were from the following fields of research: psychology (37.7%), economics (34.6%), management (7.4%), political science (4.9%), sociology (1.9%), and other fields (13.6%). No additional demographic information was collected.

In the survey, we asked participants to assess, for each replication study, (i) the likelihood that the hypothesis will successfully replicate (on a scale from 0% to 100%); (ii) their stated expertise for the study/the hypothesis (on a scale from 1 to 7); and (iii) whether they believe the pandemic has affected the likelihood of replication. The question about the pandemic was measured on a scale from −3 (“the pandemic has definitely decreased the probability of replication”) to 3 (“the pandemic has definitely increased the probability of replication”); the 0 midpoint implies that they do not think that the pandemic has affected the probability of replication. The survey was not incentivized.

The decision market opened on November 8 (2021) and closed after two weeks on November 22 (and before the decision market opened, participants had at least one week to complete the prediction survey). In the decision market, participants could trade (bet) on whether they expected the 41 studies to replicate. While participants had the opportunity to bet on the replication outcome of the 41 studies, we did not carry out replications for all 41 studies, but the final decision market prices determined which studies to replicate. We replicated the 12 studies that had the highest and the 12 studies that had the lowest market prices when the market closed. Additionally, two out of the remaining 17 studies were randomly selected for replication to ensure a non-zero probability for each study to be replicated (i.e., we replicated 12 + 12 + 2 = 26 studies in total). Since payoffs are only determined based on forecasts of studies that were eventually replicated, payoffs were scaled up by the inverse of their probability of being selected for replication in the decision rule (see below for details). This incentive scheme encourages trading based on traders' true beliefs, even though some studies will not be replicated. Consequently, participants have the incentive to buy shares of a particular study whenever they believe that the likelihood of replication is higher than the current market price; likewise, participants have an incentive to (short) sell shares whenever they believe that the likelihood of replication is lower than the current market price. Thus, as long as the market price differs from the predicted likelihood of replication for a participant, the participant has an incentive to buy or (short) sell shares of a particular study and realizing a trade according to a trader's belief will move the market price in the direction of the trader's belief. The decision rule for which studies to replicate was based on final market prices and was common knowledge to the market participants; the instructions (provided to participants who completed the prediction survey by the due date) are available at https://osf.io/a24zq/.

We chose 12 studies with the lowest predicted probability and 12 studies with the highest predicted probability based on a power calculation using the pooled data from our previous prediction market studies[45]. The power calculations were conducted by randomly sampling 41 studies from the dataset described in Gordon et al.[45] in a simulation with 10,000 iterations and then selecting the forecasts and outcomes from the 12 studies with the lowest predicted probability, the 12 studies with the highest predicted probability, and two random studies. We failed to set a random seed for the simulation when the study was conducted, implying that the pre-registered power estimates could not be numerically reproduced when we wrote up the study results. For full transparency, we report the power estimates included in the PAP in parentheses below for full transparency. The median point-biserial correlation coefficient across the 10,000 runs is 0.671 (reported as 0.66 in the PAP), and we have 91.0% power (reported as >90% in the PAP) to detect a statistically significant correlation (n = 26) between the decision market prices and the replication outcomes at the 0.5% level and 99.4% power (reported as >95% in the PAP) to detect a statistically significant correlation at the 5% level, which is our first primary hypothesis test. As a secondary hypothesis test, we test if the fraction of studies that successfully replicate differs between the 12 studies with the highest and the 12 studies with the lowest predicted replication probabilities using Fisher's exact test. Applying the same sampling approach as for primary hypothesis 1, the median difference in replication rates between the 12 studies with the

highest and the 12 studies with the lowest market prices is 0.663; the secondary test (n = 24) has 66.5% power at the 0.5% level (reported as 66% in the PAP) and 94.9% power at the 5% level (reported as 95% in the PAP). The code for the power simulations is available at https://osf.io/47drs.

Implementation of the decision market. We used a web-based trading platform, similar to the ones used in Camerer et al.[16,17] and identical to the one used in Botvinik-Nezer et al.[6]. The trading platform involves two main views: (i) the market overview and (ii) the trading page. The market overview listed the 41 assets (i.e., one corresponding to each study) in tabular format, including information on the current price for buying a share and the number of shares held (separated for long and short positions). Via the trading page, which was shown after clicking on any of the assets, participants could make investment decisions (i.e., buy or sell shares) and view price developments in graphical format for the particular asset.

Trading and incentivization. Decision market participants received an endowment of 100 tokens corresponding to USD 50. Once the markets opened, market participants could use the tokens to trade shares of the assets available in the market. An automated market maker, implementing a logarithmic market scoring rule[132], determined the assets' prices. At the beginning of the markets, all assets were valued at 0.50 tokens per share. The market maker calculated the share price for each infinitesimal transaction and updated the price based on the scoring rule. With this mechanism, participants had incentives to invest according to their beliefs[43,44]. With the logarithmic scoring rule, the price p for an infinitesimal trade is determined as $p = e^{s/b} \div (e^{s/b} + 1)$, where s denotes the net sales (shares held − shares borrowed) that the market maker has done so far in a market; the liquidity parameter b determines how strongly the market price is affected by trade and was set to b = 100, implying that by investing ten tokens, traders could move the price of a single asset from 0.50 to about 0.55. We opted for the same value of b as the one employed in the prediction markets in SSRP[17], which appears intuitively sensible in terms of striking a good balance between price sensitivity and liquidity. Notwithstanding, it is worth noting that it is unclear whether or not our results are sensitive to the choice of this parameter. Decision market participants were paid only for studies chosen for replication (based on their final holdings). Participants received one token per correct share for the replications with the 12 lowest and 12 highest final market prices. For the two randomly selected replications, participants received 17 ÷ 2 = 8.5 tokens for each share; for replications that were not chosen for replication, participants received no compensation for their holdings. We followed this procedure to keep information revelation in the decision market incentive-compatible, with the increased payouts for the randomly selected studies compensating for the "voided" shares in studies not chosen for replication. Participants were paid after all 26 replications had been completed.

Participation. A total of 193 participants completed the prediction survey (a prerequisite to participate in market trading) after providing consent to participate and were subsequently invited to trade on the decision market. Of these 193 participants, 162 (83.9%) traded in the market at least once. During the two-week trading period, a total of 4,412 transactions were recorded. On

663 average, each trader prompted 27.2 transactions (sd = 30.7; min = 1, max = 185). The average
664 number of traders per hypothesis was 65.1 (sd = 15.3; min = 35, max = 98); the average number
665 of transactions recorded per hypothesis was 107.6 (sd = 35.2; min = 56, max = 213). See
666 Supplementary Table 2 for descriptive statistics on the trading activity for each market.
667

668 ## Replications

669 We carry out close replications[107] as closely as possible following the experimental design,
670 sample restrictions, exclusion criteria, and analysis as used in the original studies and carried out
671 in the same population (Amazon Mechanical Turk) as the original studies. The replications started
672 in January 2022 and were completed in October 2023. The replications were planned and pre-
673 registered by five replication teams: a team at CalTech, LMU, and Wharton; a team at the
674 Stockholm School of Economics; a team at the National University of Singapore; a team at the
675 University of Amsterdam; and a team at the University of Innsbruck.

676 **Participants in replication studies.** All replications were carried out at Amazon Mechanical Turk
677 as in the original studies. We ensured that participants could only participate once using the same
678 account in a specific study. If the original study had not specified a HIT approval rate, we recruited
679 participants with a HIT approval rate of at least 95%; if the original study had specified a higher
680 approval rate, we applied the same requirement as used in the original study.

681 To ward off concerns about impaired data quality due to low-attention participants and bots[88–94],
682 we implemented several "quality filters." Particularly, before redirecting participants to each study,
683 we forwarded the IP addresses to https://www.ipqualityscore.com/ for a quality check to minimize
684 the chances of low-quality participant data (we initially planned to use this filter ex-post but during
685 the data collection of the first two replication studies of Klein & O'Brien[73] and Halevy & Halali[62] we
686 decided to set it up so that the IP address quality check happened before participants got
687 redirected to the study). Participants for whom one or more of the following was true could not
688 proceed with participating in the study: fraud score ≥ 85; TOR = True; VPN = True; Bot = True;
689 abuse velocity = high. This means that, for example, participants were not allowed to use a virtual
690 private network (VPN) or Tor connections or participate if they had IP addresses that had recently
691 engaged in automated bot activity (the VPN exclusions were made ex-ante, i.e., before
692 participants were redirected to the study, for four studies and ex-post for 22 studies). After that,
693 in all replications, participants were first shown a Captcha and then provided informed consent.
694 After this, we included an attention check that participants had to pass to proceed to the study
695 (with the exception of Reeck et al.[82]; see Section 4 in the Supplementary Information for details).
696 The attention check was implemented in addition to any other potential attention check(s) used
697 in the original study. All these exclusions based on the "quality filters" were preregistered, but the
698 pre-analysis plan did not specify if participants would be excluded before or after participating in
699 the study.

700 The individual replication studies sometimes also used additional exclusion criteria that are
701 detailed in the preregistered replication report for each replication (we tried to use the same

exclusion criteria for the replications as used in the original studies as much as possible). The replication sample sizes defined below are the sample sizes after any exclusions of participants.

Replication sample sizes. The replications were carried out with high statistical power. Replication sample sizes were based on having 90% power to detect ⅔ of the effect size reported in the original study (with the effect size converted to Cohen's d to have a common standardized effect size measure across the original studies and the replication studies). See Supplementary Notes for more details about the power calculations and replication sample sizes. The criteria for replication were an effect in the same direction as the original study and a p-value < 0.05 (in a two-sided test). In cases where this power estimation led to a sample size smaller than the original one, we used the same sample size as in the original study. The average replication sample ($\bar{n}$ = 1,018) size was 3.5 times as large as the average sample size of the original studies ($\bar{n}$ = 292). We continued the data collection for each replication until we reached at least the preregistered sample size after exclusions for that replication, and this led to slightly larger replication sample sizes than preregistered in all replications except one (as it is not possible with exclusion criteria to get an exact sample size as the number of exclusions is not known ex-ante).

Conversion of effect sizes to Cohen's d. We converted the effect sizes of all the original studies and all the replication studies to Cohen's d to have a standardized effect size (the effect size in the original study was always assigned a positive sign; the effect size in the replication study was assigned a positive sign if the effect was in the same direction as in the original study and a negative sign if the effect was in the opposite direction of the original study). See Supplementary Notes for details about the conversion of effect sizes to Cohen's d.

Replication reports. For each of the 41 studies, we prepared a pre-replication plan/report stating the hypothesis we had chosen from each paper and how we planned to proceed with the replication study. These reports were shared with the original authors for feedback, and at least one original author from each paper replied. These pre-replication reports were posted at OSF (https://osf.io/sejyp at the same time as the pre-analysis plan and prior to the start of the prediction survey (that preceded the decision markets and the replication data collections). For those studies that were selected for replication, we have updated the replication reports with the replication results after the replications were completed. After sharing them with the original authors for feedback, we have posted the updated replication reports at OSF as well (https://osf.io/sejyp). Additionally, we reached out to the original authors for their comments on the replication reports and results. We promised to make their comments available along with the replication reports, and any comments received can be found at https://osf.io/sejyp.

Incentivization in the replication experiments. We standardized payments across all replications such that studies had a certain show-up fee depending on the expected length of the study. In particular, we paid an hourly fee of USD 8.00 for all studies, and we calculated the show-up fee for each study based on the expected length of the study. For all studies, we implemented a minimum payoff of USD 1.00. For studies with incentive payments, we used the same incentive

payment as in the original study, paid in addition to the show-up fee. If we faced problems in recruiting participants, we increased the show-up fee, which happened for two studies[61,65].

Replication indicators

Statistical significance criterion (primary indicator). The first primary replication indicator was the statistical significance criterion – i.e., whether the replication resulted in an effect size in the same direction as the original study and a two-sided p-value less than 0.05. Unless otherwise stated above, we used the same statistical test as in the original study. We report the replication rate (i.e., the fraction of the 26 studies that replicated according to this criterion) and the 95% Clopper-Pearson CI of this fraction in the Results section. We also report the 95% CI of the replication effect size for each of the 26 replication studies in Fig. 2 and Supplementary Table 3.

Relative effect sizes (primary indicator). As a second primary replication indicator, we used relative effect sizes. Relative effect sizes were estimated in two different ways. We report the mean effect size of all 26 replications and compare it to the mean effect size of the 26 original studies (see also primary hypothesis test 2 below). We furthermore estimate the relative effect size of each replication (the replication effect size divided by the original effect size) and estimate the mean of this variable for the 26 replication studies and the 95% CI of this mean (based on a one-sample t-test). We report both of these measures of the relative effect size separately for the replications that replicate and those that do not. These results are reported in the Results section, Fig. 3, and Supplementary Table 3.

Small-telescopes approach (secondary indicator). We also used the small-telescopes approach[112]. For this indicator, we estimated whether the replication effect size was significantly smaller (using a one-sided test at the 5% level) than a “small effect,” defined as the effect size the original study would have had 33% power to detect. For studies using t-tests (or F-tests converted to a t-test statistic), we based “the small effect size” on the effect size that a t-test had 33% power to detect (at the 5% level in a two-sided test); for studies using z-test statistics (or chi-square tests converted to a z-test statistic), we based “the small effect size” on the effect size that a z-test had 33% power to detect (at the 5% level in a two-sided test). To test whether the replication effect size was significantly smaller than “the small effect size” in a one-sided test at the 5% level, we estimated a 90% CI of the replication effect size. We tested if the 90% CI overlapped the small effect size with CIs constructed as described in Supplementary Notes. If the effect size in the replication was significantly smaller than this “small effect size,” the result was considered a failed replication; otherwise, it was considered successful. We report the fraction of studies that replicate according to this criterion and the 95% Clopper-Pearson CI of this fraction. The small-telescopes results are reported in Fig. 4 and Supplementary Table 4.

Bayes factors (secondary indicators). We also compute the one-sided default Bayes factors on the replication data, allowing us to obtain the strength of evidence in favor of the hypothesis that stipulates an effect in the direction of the original experiment (where a default prior in terms of a truncated Cauchy distribution with scale 0.707 was assigned to the size of the effect) versus

the null hypothesis that stipulates the effect to be absent[113]. In addition, we also computed (one-sided) replication Bayes factors, which quantifies the additional evidence for the hypothesis given the evidence already provided by the original study[114]. (We are counting the one-sided default and replication as Bayes factors as two separate indicators, which they are.) These results are reported in Fig. 5 and Supplementary Table 4. We use the evidence categories proposed by Jeffreys[115] to interpret the Bayes Factors. A detailed report on the estimation of the Bayes factors is available at https://osf.io/47drs/.

Meta-analytic effect sizes (secondary indicator). We estimated the meta-analytic estimate of the effect size by combining the original result and the replication result in a fixed-effect meta-analysis. We report the fraction of the 26 studies that replicated according to the 0.05 and the 0.005 significance threshold and the 95% Clopper-Pearson CI of these fractions. We also use the stricter 0.005 significance threshold as a replication indicator for the meta-analytic effect sizes because this is similar to observing two studies (an original study and a replication study) that are significant at the 0.05 level. We report these results in the Results section, Fig. 6, and Supplementary Table 4.

Data Availability

The data reported in this paper is tabulated in Supplementary Tables 1–8. The replication reports (both the pre-replication and the post-replication versions), the pre-analysis plan, the data from the survey and the decision market, and the data for each of the 26 replications are available at the project's OSF repository (https://osf.io/sk82q).

Code Availability

The analysis scripts, generating all results, figures, and tables reported in the main text and the Supplementary Information, are available at the project's OSF repository (https://osf.io/sk82q).

Acknowledgments

We thank Alexander Andevall for helping with the data collection and programming of experiments and Robb Willer for helpful advice on defining the IP address check and exclusion criteria used to exclude individuals from participating to minimize low-quality participant data. For financial support, we thank the Austrian Science FWF (grant SFB F63 to J.H. and M.K.), Jan Wallander and Tom Hedelius Foundation (grants P21-0091 and P23-0098 to A.D.), Knut and Alice Wallenberg Foundation and Marianne and Marcus Wallenberg Foundation (Wallenberg Scholar grant to A.D.), and Riksbankens Jubileumsfond (grant P21-0168 to M.J.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. One author (V.W.) is currently employed by Sveriges Riksbank but did this work before being employed by Sveriges Riksbank; the opinions expressed in this article are the sole responsibility of the authors and should not be interpreted as reflecting the views of Sveriges Riksbank. We sought ethical approval from the Swedish Ethical Review Authority who had no ethical objections to the decision market part of the project and judged the replication part of the project to not be covered by the Swedish ethical review law (Dnr 2019-06501).

Author Contributions

A.D., F.H., J.H., M.J., M.K., B.A.N., and T.P. designed the study; A.D., F.H., and M.J. managed the study; Y.C., A.D., F.H., M.J., and T.P. designed and implemented the decision market; A.D., F.H., M.J., B.M., and V.W. selected articles and critical findings for (potential) replication; A.D., C.F.C., F.H., T.-H.H., S.H., J.H., N.I., T.I., L.J., M.J., M.K., B.M., D.M., G.N., A.S., R.S., E.-J.W., and V.W., designed the replications and collected replication data; F.H, S.H., T.I., L.J., A.S., R.S., and V.W. conducted the preregistered statistical tests on the individual replications; A.L.

computed the Bayes factors; F.H. conducted all analyses reported in the manuscript; A.D., F.H., and M.J. wrote the paper; all authors reviewed and approved the final manuscript.

Competing Interest Statement

The authors declare no competing interests.

Figure legends

Fig. 1. Decision market prices for the 41 included studies. Plotted are the decision market prices for the 41 MTurk social science experiments published in PNAS between 2015 and 2018. The small gray dots indicate the market prices after each market transaction; the larger dots indicate the final market price. The studies are ordered based on the final decision market prices, which can be interpreted as the market's probability forecast of successful replication. The 12 studies with the highest decision market prices and the 12 studies with the lowest decision market prices were selected for replication; in addition, two of the remaining 17 studies were selected for replication at random to ensure that the decision market is incentive compatible. The replication outcomes for the statistical significance indicator are also illustrated for the 26 replicated studies. The point-biserial correlation between the decision market prices and the replication outcomes in primary hypothesis 1 is $r = 0.505$ (95% CI [0.146, 0.712], $t(24) = 2.867$, $p = 0.008$; $n = 26$, two-sided test).

Fig. 2. Replication results. Plotted are the point estimates and the 95% CIs (standardized to Cohen's d units) of the 26 replications (d_R) and original studies. Studies within each of the three panels (top-12, random, bottom-12) are sorted based on the decision market prices as in Fig. 1. There is a statistically significant effect ($p < 0.05$) in the same direction as the original study for 14 out of 26 replications (53.8%; 95% CI [33.4%, 73.4%]). For the 12 studies with the highest decision market prices, there is a statistically significant effect ($p < 0.05$) in the same direction as the original study for 10 out of 12 replications (83.3%; 95% CI [51.6%, 97.9%]). For the 12 studies with the lowest decision market prices, there is a statistically significant effect ($p < 0.05$) in the same direction as the original study for 4 out of 12 replications (33.3%; 95% CI [9.9%, 65.1%]). Our secondary hypothesis test provides suggestive evidence that the difference in replication rates between the top-12 and the bottom-12 group is different from zero (Fisher's exact test; $\chi^2(1) = 6.171$, $p = 0.036$; $n = 24$, two-sided test). The error bars denote the 95% confidence intervals (CIs) of the original and the replication effect size estimates. The numbers of observations used to estimate the 95% CIs are the original and replication sample sizes noted on the right as n_O and n_R.

Fig. 3. Relationship between estimated original and replication effect sizes. Plotted are the estimated original and replication effect sizes for each of the 26 replication studies (the estimated effect sizes of both the original and replication studies are standardized to Cohen's d units). The 95% confidence intervals for the original and replication effect size estimates are illustrated in Fig.

2 and tabulated in Supplementary Table 3. The mean estimated effect size of the 26 replication studies is 0.253 (sd = 0.357) compared to 0.563 (sd = 0.426) for the original studies, resulting in a relative estimated average effect size of 45.0%, confirming our second primary hypothesis (Wilcoxon signed-rank test; $z = 4.203$, $p < 0.001$; $n = 26$, two-sided test). The estimated relative effect size of the 13 replications that have been successfully replicated according to the statistical significance indicator is 69.5%, and the estimated relative effect size of the 13 studies that did not replicate is 3.2%. The box plots show the median, the interquartile range, and the 5th and 95th percentile of the effect size estimates in the 26 original studies and the 26 replication studies.

Fig. 4. Replication results based on the small-telescopes approach (a secondary replication indicator). Plotted are the 90% CIs of replication effect sizes in relation to small-effect sizes as defined by the small-telescopes approach[112] (the effect size that the original study would have had 33% power to detect). Studies within the three panels (top-12, random, bottom-12) are sorted based on the decision market prices as in Fig. 1. A study is defined as failing to replicate if the 90% CI is below the small effect. According to the small-telescopes approach, 15 out of 26 studies (57.7%; 95% CI [36.9%, 76.6%]) replicate. The error bars denote the 90% confidence intervals (CIs) of the estimated replication effect sizes. The numbers of observations used to estimate the 90% CIs are the replication sample sizes noted on the right as n_R.

Fig. 5. Replication results based on Bayes factors (secondary replication indicators). The figure plots the one-sided default Bayes factor (BF_{+0}) and the replication Bayes factor (BF_{R0}) for the 26 replications[113]. $BF_{+0} > 1$ favors the hypothesis of an effect in the direction of the original paper, whereas $BF_{+0} < 1$ favors the null hypothesis of no effect. BF_{R0} quantifies the additional evidence provided by the replication results on top of the original evidence. $BF_{R0} > 1$ indicates additional evidence in favor of the alternative over the null, whereas $BF_{R0} < 1$ indicates additional evidence for the null instead. The evidence categories proposed by Jeffreys[115] are also shown (from extreme support for the null hypothesis to extreme support for the original hypothesis). Studies within the three panels (top-12, random, bottom-12) are sorted based on the decision market prices as in Fig. 1. The BF_{+0} is above one for all 14 replication studies that successfully replicated according to the statistical significance indicator and below one for all 12 replication studies that failed to replicate according to the statistical significance indicator. The BF_{R0} is above one for 13 of the 14 replication studies that replicated according to the statistical significance indicator and below one for Cooney et al.[56] whose estimated relative effect size of 0.36 is the lowest among these 14 studies; the BF_{R0} is below one for all of the 12 replication studies that failed to replicate according to the statistical significance indicator. The numbers of observations used to estimate BF_{+0} and BF_{R0} are the original and replication sample sizes noted on the right as n_O and n_R.

Fig. 6. Meta-analytic estimated effect sizes combining the original and the replication estimated effect sizes (a secondary replication indicator). The figure plots the point estimates and 95% and 99.5% CIs of the fixed-effects weighted meta-analytic effect sizes, combining the original and the replication studies (standardized to Cohen's d units). Studies within the three

907 panels (top-12, random, bottom-12) are sorted based on the decision market prices as in Fig. 1.
908 As preregistered, we report the significance of the estimated meta-analytic effect sizes for both
909 the 0.05 significance threshold and the 0.005 significance threshold (based on a two-sided z-test).
910 Sixteen out of 26 (61.5%; 95% CI [40.6%, 79.8%]) studies replicated according to the statistical
911 significance indicator using the 0.05 significance threshold, and 14 out of 26 (53.8%; 95% CI
912 [33.4%, 73.4%]) studies replicated using the 0.005 significance threshold. The error bars denote
913 the 95% confidence intervals (CIs) of the estimated meta-analytic effect sizes. The number of
914 observations used to estimate the 95% CIs are the sums of the original and replication sample
915 sizes noted next to the study identifier on the Y-axis as n_{O+R}.

916 **Fig. 7. Replication results based on prediction intervals (not preregistered).** Plotted are the
917 95% prediction intervals[116] for the standardized original effect sizes (Cohen's d). Studies within
918 the three panels (top-12, random, bottom-12) are sorted based on the decision market prices as
919 in Fig. 1. Fifteen replications out of 26 (57.7%; CI [36.9%, 76.6%]) are within the 95% prediction
920 interval and replicate according to this indicator. The p-values reported on the right are based on
921 two-sample z-tests for a difference between the replication effect size and the original effect size.
922 The gray lines denote the 95% prediction intervals, and the small circles denote the mean
923 replication effect sizes. All tests are two-sided. The numbers of observations used to estimate the
924 95% prediction intervals are the original and replication sample sizes noted next to the study
925 identifier on the y-axis as n_O and n_R.

926 **Fig. 8. Relationship between decision market prices and mean survey beliefs and**
927 **forecasting accuracy. a,** Plotted are the decision market prices and the mean survey beliefs
928 about replication for the 41 studies included in the decision market and the survey; the color
929 coding highlights the replication outcomes for the 26 replicated studies. The decision market
930 prices and the mean survey beliefs about replication are highly correlated with a Pearson
931 correlation of r = 0.899 (95% CI [0.814, 0.944]; t(39) = 12.830, p = $1.4e^{-15}$; n = 41, two-sided test).
932 **b**, Plotted are the absolute prediction errors and the Brier scores (the squared prediction errors)
933 for the decision market and the prediction survey for the 26 replicated studies. There is suggestive
934 evidence of higher prediction accuracy for the decision market in terms of the absolute prediction
935 error (0.353 for the decision markets and 0.421 for the survey; Wilcoxon signed-rank test:
936 z = 2.172, p = 0.030; *n* = 26, two-sided test), but not in terms of the Brier score (0.188 for the
937 decision markets and 0.202 for the survey; Wilcoxon signed-rank test: z = 1.181, p = 0.238;
938 n = 26, two-sided test). The box plots show the median, the interquartile range, and the 5th and
939 95th percentile of the absolute prediction errors and Brier scores for the survey and decision
940 market predictions of the 26 replication studies**.**

References

1. Ioannidis, J. P. A. Why most published research findings are false. *PLoS Med.* **2**, e124 (2005).
2. Leamer, E. E. Let's take the con out of econometrics. *Am. Econ. Rev.* **73**, 31–43 (1983).
3. Begley, C. G. & Ellis, L. M. Raise standards for preclinical cancer research. *Nature* **483**, 531–533 (2012).
4. McNutt, M. Reproducibility. *Science* **343**, 229–229 (2014).
5. Gertler, P., Galiani, S. & Romero, M. How to make replication the norm. *Nature* **554**, 417–419 (2018).
6. Botvinik-Nezer, R. *et al.* Variability in the analysis of a single neuroimaging dataset by many teams. *Nature* **582**, 84–88 (2020).
7. Breznau, N. *et al.* Observing many researchers using the same data and hypothesis reveals a hidden universe of uncertainty. *Proc. Natl. Acad. Sci.* **119**, e2203150119 (2022).
8. Delios, A. *et al.* Examining the generalizability of research findings from archival data. *Proc. Natl. Acad. Sci.* **119**, e2120377119 (2022).
9. Huber, C. *et al.* Competition and moral behavior: A meta-analysis of forty-five crowd-sourced experimental designs. *Proc. Natl. Acad. Sci.* **120**, e2215572120 (2023).
10. Dreber, A. *et al.* Using prediction markets to estimate the reproducibility of scientific research. *Proc. Natl. Acad. Sci.* **112**, 15343–15347 (2015).
11. Maniadis, Z., Tufano, F. & List, J. A. To replicate or not to replicate? Exploring reproducibility in economics through the lens of a model and a pilot study. *Econ. J.* **127**, F209–F235 (2017).
12. Klein, R. A. *et al.* Investigating variation in replicability: A "many labs" replication project. *Soc. Psychol.* **45**, 142–152 (2014).
13. Ebersole, C. R. *et al.* Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *J. Exp. Soc. Psychol.* **67**, 68–82 (2016).
14. Klein, R. A. *et al.* Many Labs 2: Investigating variation in replicability across samples and settings. *Adv. Methods Pract. Psychol. Sci.* **1**, 443–490 (2018).
15. Open Science Collaboration. Estimating the reproducibility of psychological science. *Science* **349**, aac4716 (2015).
16. Camerer, C. F. *et al.* Evaluating replicability of laboratory experiments in economics. *Science* **351**, 1433–1436 (2016).
17. Camerer, C. F. *et al.* Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nat. Hum. Behav.* **2**, 637–644 (2018).

18. Ioannidis, J. P. A. Why most discovered true associations are inflated. *Epidemiology* **19**, 640 (2008).

19. Button, K. S. *et al.* Power failure: Why small sample size undermines the reliability of neuroscience. *Nat. Rev. Neurosci.* **14**, 365–376 (2013).

20. Wegener, D. T., Fabrigar, L. R., Pek, J. & Hoisington-Shaw, K. Evaluating research in personality and social psychology: Considerations of statistical power and concerns about false findings. *Pers. Soc. Psychol. Bull.* **48**, 1105–1117 (2022).

21. Simmons, J. P., Nelson, L. D. & Simonsohn, U. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol. Sci.* **22**, 1359–1366 (2011).

22. John, L. K., Loewenstein, G. & Prelec, D. Measuring the Prevalence of Questionable Research Practices With Incentives for Truth Telling. *Psychol. Sci.* **23**, 524–532 (2012).

23. Finkel, E. J., Eastwick, P. W. & Reis, H. T. Replicability and other features of a high-quality science: Toward a balanced and empirical approach. *J. Pers. Soc. Psychol.* **113**, (2017).

24. Flake, J. K., Davidson, I. J., Wong, O. & Pek, J. Construct validity and the validity of replication studies: A systematic review. *Am. Psychol.* **77**, 576–588 (2022).

25. Pittelkow, M.-M. *et al.* The process of replication target selection in psychology: What to consider? *R. Soc. Open Sci.* **10**, 210586 (2023).

26. Makel, M. C., Plucker, J. A. & Hegarty, B. Replications in psychology research: How often do they really occur? *Perspect. Psychol. Sci.* **7**, 537–542 (2012).

27. Lindsay, D. S. Replication in psychological science. *Psychol. Sci.* **26**, 1827–1832 (2015).

28. Block, J. & Kuckertz, A. Seven principles of effective replication studies: Strengthening the evidence base of management research. *Manag. Rev. Q.* **68**, 355–359 (2018).

29. Coles, N. A., Tiokhin, L., Scheel, A. M., Isager, P. M. & Lakens, D. The costs and benefits of replication studies. *Behav. Brain Sci.* **41**, e124 (2018).

30. Alipourfard, N. *et al.* Systematizing Confidence in Open Research and Evidence (SCORE). Preprint at https://doi.org/10/hn4g (2021).

31. Hardwicke, T. E., Tessler, M. H., Peloquin, B. N. & Frank, M. C. A Bayesian decision-making framework for replication. *Behav. Brain Sci.* **41**, e132 (2018).

32. Field, S. M., Hoekstra, R., Bringmann, L. & van Ravenzwaaij, D. When and why to replicate: As easy as 1, 2, 3? *Collabra Psychol.* **5**, 46 (2019).

33. Isager, P. M. *et al.* Deciding what to replicate: A decision model for replication study selection under resource and knowledge constraints. *Psychol. Methods* **28**, 438–451 (2023).

34. O'Donnell, M. *et al.* Empirical audit and review and an assessment of evidentiary value in

research on the psychological consequences of scarcity. *Proc. Natl. Acad. Sci.* **118**, e2103313118 (2021).

35. Kuehberger, A. & Schulte-Mecklenbeck, M. Selecting target papers for replication. *Behav. Brain Sci.* **41**, e139 (2018).

36. Hoogeveen, S., Sarafoglou, A. & Wagenmakers, E.-J. Laypeople can predict which social-science studies will be replicated successfully. *Adv. Methods Pract. Psychol. Sci.* **3**, 267–285 (2020).

37. Marcoci, A. *et al.* Predicting the replicability of social and behavioural science claims from the COVID-19 Preprint Replication Project with structured expert and novice groups. *Nat. Hum. Behav.* (forthcoming).

38. Wolfers, J. & Zitzewitz, E. Prediction markets. *J. Econ. Perspect.* **18**, 107–126 (2004).

39. Arrow, K. J. *et al.* The promise of prediction markets. *Science* **320**, 877–878 (2008).

40. Tziralis, G. & Tatsiopoulos, I. Prediction markets: An extended literature review. *J. Predict. Mark.* **1**, 75–91 (2012).

41. Hanson, R. Decision markets. *IEEE Intell. Syst.* **14**, 16–19 (1999).

42. Hanson, R. Combinatorial information market design. *Inf. Syst. Front.* **5**, 107–119 (2003).

43. Chen, Y., Kash, I., Ruberry, M. & Shnayder, V. Decision markets with good incentives. in *Internet and Network Economics* (eds. Chen, N., Elkind, E. & Koutsoupias, E.) 72–83 (Springer, Berlin, 2011). doi:10/b4vtjj.

44. Wang, W. & Pfeiffer, T. Securities based decision markets. in *Distributed Artificial Intelligence* (eds. Chen, J., Lang, J., Amato, C. & Zhao, D.) vol. 13170 79–92 (Springer, Shanghai, China, 2022).

45. Gordon, M., Viganola, D., Dreber, A., Johannesson, M. & Pfeiffer, T. Predicting replicability—Analysis of survey and prediction market data from large-scale forecasting projects. *PLoS One* **16**, e0248780 (2021).

46. Ames, D. L. & Fiske, S. T. Perceived intent motivates people to magnify observed harms. *Proc. Natl. Acad. Sci.* **112**, 3599–3605 (2015).

47. Atir, S. & Ferguson, M. J. How gender determines the way we speak about professionals. *Proc. Natl. Acad. Sci.* **115**, 7278–7283 (2018).

48. Baldwin, M. & Lammers, J. Past-focused environmental comparisons promote proenvironmental outcomes for conservatives. *Proc. Natl. Acad. Sci.* **113**, 14953–14957 (2016).

49. Bear, A., Fortgang, R. G., Bronstein, M. V. & Cannon, T. D. Mistiming of thought and perception predicts delusionality. *Proc. Natl. Acad. Sci.* **114**, 10791–10796 (2017).

50. Boswell, R. G., Sun, W., Suzuki, S. & Kober, H. Training in cognitive strategies reduces

eating and improves food choice. *Proc. Natl. Acad. Sci.* **115**, E11238–E11247 (2018).

51. Caruso, E. M., Burns, Z. C. & Converse, B. A. Slow motion increases perceived intent. *Proc. Natl. Acad. Sci.* **113**, 9250–9255 (2016).

52. Casella, A., Kartik, N., Sanchez, L. & Turban, S. Communication in context: Interpreting promises in an experiment on competition and trust. *Proc. Natl. Acad. Sci.* **115**, 933–938 (2018).

53. Chao, M. Demotivating incentives and motivation crowding out in charitable giving. *Proc. Natl. Acad. Sci.* **114**, 7301–7306 (2017).

54. Cheon, B. K. & Hong, Y.-Y. Mere experience of low subjective socioeconomic status stimulates appetite and food intake. *Proc. Natl. Acad. Sci.* **114**, 72–77 (2017).

55. Clarkson, J. J. *et al.* The self-control consequences of political ideology. *Proc. Natl. Acad. Sci.* **112**, 8250–8253 (2015).

56. Cooney, G., Gilbert, D. T. & Wilson, T. D. When fairness matters less than we expect. *Proc. Natl. Acad. Sci.* **113**, 11168–11171 (2016).

57. Côté, S., House, J. & Willer, R. High economic inequality leads higher-income individuals to be less generous. *Proc. Natl. Acad. Sci.* **112**, 15838–15843 (2015).

58. Flesch, T., Balaguer, J., Dekker, R., Nili, H. & Summerfield, C. Comparing continual task learning in minds and machines. *Proc. Natl. Acad. Sci.* **115**, E10313–E10322 (2018).

59. Genschow, O., Rigoni, D. & Brass, M. Belief in free will affects causal attributions when judging others' behavior. *Proc. Natl. Acad. Sci.* **114**, 10071–10076 (2017).

60. Gheorghiu, A. I., Callan, M. J. & Skylark, W. J. Facial appearance affects science communication. *Proc. Natl. Acad. Sci.* **114**, 5970–5975 (2017).

61. Guilbeault, D., Becker, J. & Centola, D. Social learning and partisan bias in the interpretation of climate trends. *Proc. Natl. Acad. Sci.* **115**, 9714–9719 (2018).

62. Halevy, N. & Halali, E. Selfish third parties act as peacemakers by transforming conflicts and promoting cooperation. *Proc. Natl. Acad. Sci.* **112**, 6937–6942 (2015).

63. Handley, I. M., Brown, E. R., Moss-Racusin, C. A. & Smith, J. L. Quality of evidence revealing subtle gender biases in science is in the eye of the beholder. *Proc. Natl. Acad. Sci.* **112**, 13201–13206 (2015).

64. Hoffman, K. M., Trawalter, S., Axt, J. R. & Oliver, M. N. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proc. Natl. Acad. Sci.* **113**, 4296–4301 (2016).

65. Hofstetter, R., Rüppell, R. & John, L. K. Temporary sharing prompts unrestrained disclosures that leave lasting negative impressions. *Proc. Natl. Acad. Sci.* **114**, 11902–11907 (2017).

66. Horne, Z., Powell, D., Hummel, J. E. & Holyoak, K. J. Countering antivaccination attitudes. *Proc. Natl. Acad. Sci.* **112**, 10321–10324 (2015).

67. Isley, S. C., Stern, P. C., Carmichael, S. P., Joseph, K. M. & Arent, D. J. Online purchasing creates opportunities to lower the life cycle carbon footprints of consumer products. *Proc. Natl. Acad. Sci.* **113**, 9780–9785 (2016).

68. Jachimowicz, J. M., Chafik, S., Munrat, S., Prabhu, J. C. & Weber, E. U. Community trust reduces myopic decisions of low-income individuals. *Proc. Natl. Acad. Sci.* **114**, 5401–5406 (2017).

69. John, L. K., Barasz, K. & Norton, M. I. Hiding personal information reveals the worst. *Proc. Natl. Acad. Sci.* **113**, 954–959 (2016).

70. Jordan, J. J., Hoffman, M., Nowak, M. A. & Rand, D. G. Uncalculating cooperation is used to signal trustworthiness. *Proc. Natl. Acad. Sci.* **113**, 8658–8663 (2016).

71. Jun, Y., Meng, R. & Johar, G. V. Perceived social presence reduces fact-checking. *Proc. Natl. Acad. Sci.* **114**, 5976–5981 (2017).

72. KC, R. P., Kunter, M. & Mak, V. The influence of a competition on noncompetitors. *Proc. Natl. Acad. Sci.* **115**, 2716–2721 (2018).

73. Klein, N. & O'Brien, E. People use less information than they think to make up their minds. *Proc. Natl. Acad. Sci.* **115**, 13222–13227 (2018).

74. Kouchaki, M. & Gino, F. Memories of unethical actions become obfuscated over time. *Proc. Natl. Acad. Sci.* **113**, 6166–6171 (2016).

75. Kraus, M. W., Rucker, J. M. & Richeson, J. A. Americans misperceive racial economic equality. *Proc. Natl. Acad. Sci.* **114**, 10324–10331 (2017).

76. McCall, L., Burk, D., Laperrière, M. & Richeson, J. A. Exposure to rising inequality shapes Americans' opportunity beliefs and policy support. *Proc. Natl. Acad. Sci.* **114**, 9593–9598 (2017).

77. Morris, A., MacGlashan, J., Littman, M. L. & Cushman, F. Evolution of flexibility and rigidity in retaliatory punishment. *Proc. Natl. Acad. Sci.* **114**, 10396–10401 (2017).

78. Mummolo, J. Militarization fails to enhance police safety or reduce crime but may harm police reputation. *Proc. Natl. Acad. Sci.* **115**, 9181–9186 (2018).

79. Payne, B. K., Brown-Iannuzzi, J. L. & Hannay, J. W. Economic inequality increases risk taking. *Proc. Natl. Acad. Sci.* **114**, 4643–4648 (2017).

80. Phillips, J. & Cushman, F. Morality constrains the default representation of what is possible. *Proc. Natl. Acad. Sci.* **114**, 4649–4654 (2017).

81. Rai, T. S., Valdesolo, P. & Graham, J. Dehumanization increases instrumental violence, but not moral violence. *Proc. Natl. Acad. Sci.* **114**, 8511–8516 (2017).

82. Reeck, C., Wall, D. & Johnson, E. J. Search predicts and changes patience in intertemporal choice. *Proc. Natl. Acad. Sci.* **114**, 11890–11895 (2017).

83. Schilke, O., Reimann, M. & Cook, K. S. Power decreases trust in social exchange. *Proc. Natl. Acad. Sci.* **112**, 12950–12955 (2015).

84. Stern, C., West, T. V. & Rule, N. O. Conservatives negatively evaluate counterstereotypical people to maintain a sense of certainty. *Proc. Natl. Acad. Sci.* **112**, 15337–15342 (2015).

85. Vacharkulksemsuk, T. *et al.* Dominant, open nonverbal displays are attractive at zero-acquaintance. *Proc. Natl. Acad. Sci.* **113**, 4009–4014 (2016).

86. Williams, K. E. G., Sng, O. & Neuberg, S. L. Ecology-driven stereotypes override race stereotypes. *Proc. Natl. Acad. Sci.* **113**, 310–315 (2016).

87. Schimmelpfennig, R. *et al.* The moderating role of culture in the generalizability of psychological phenomena. *Adv. Methods Pract. Psychol. Sci.* **7**, 25152459231225163 (2024).

88. Zhou, H. & Fishbach, A. The pitfall of experimenting on the web: How unattended selective attrition leads to surprising (yet false) research conclusions. *J. Pers. Soc. Psychol.* **111**, 493–504 (2016).

89. Chmielewski, M. & Kucker, S. C. An MTurk crisis? Shifts in data quality and the impact on study results. *Soc. Psychol. Personal. Sci.* **11**, 464–473 (2020).

90. Aguinis, H., Villamor, I. & Ramani, R. S. MTurk research: Review and recommendations. *J. Manag.* **47**, 823–837 (2021).

91. Brodeur, A., Cook, N. & Heyes, A. We need to talk about Mechanical Turk: What 22,989 hypothesis tests tell us about publication bias and p-hacking in online experiments. Preprint at https://doi.org/10/nd9h (2022).

92. Peer, E., Rothschild, D., Gordon, A., Evernden, Z. & Damer, E. Data quality of platforms and panels for online behavioral research. *Behav. Res. Methods* **54**, 1643–1662 (2022).

93. Webb, M. A. & Tangney, J. P. Too good to be true: Bots and bad data from Mechanical Turk. *Perspect. Psychol. Sci.* 174569162211200 (2022) doi:10/gq7nw4.

94. Douglas, B. D., Ewell, P. J. & Brauer, M. Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *PLoS One* **18**, e0279720 (2023).

95. Abelson, R. P. *Statistics as Principled Argument*. (Psychology Press, New York, 1995). doi:10/gf5svk.

96. Macdonald, R. R. Statistical inference and Aristotle's Rhetoric. *Br. J. Math. Stat. Psychol.* **57**, 193–203 (2004).

97. Scheel, A. M., Schijen, M. R. M. J. & Lakens, D. An excess of positive results: Comparing the standard psychology literature with registered reports. *Adv. Methods Pract. Psychol. Sci.* **4**, 25152459211007467 (2021).

98. Soderberg, C. K. *et al.* Initial evidence of research quality of registered reports compared with the standard publishing model. *Nat. Hum. Behav.* **5**, 990–997 (2021).

99. Brodeur, A., Cook, N., Hartley, J. & Heyes, A. Do pre-registration and pre-analysis plans reduce p-hacking and publication bias? Evidence from 15,992 test statistics and suggestions for improvement. *J. Polit. Econ. Microecon.* **2**, (2024).

100. Yamada, Y. How to crack pre-registration: Toward transparent and open science. *Front. Psychol.* **9**, (2018).

101. Flis, I. The function of literature in psychological science. *Rev. Gen. Psychol.* **26**, 146–156 (2022).

102. Rubin, M. Questionable metascience practices. *J. Trial Error* **4**, (2024).

103. Wagenmakers, E.-J., Wetzels, R., Borsboom, D., van der Maas, H. L. J. & Kievit, R. A. An agenda for purely confirmatory research. *Perspect. Psychol. Sci.* **7**, 632–638 (2012).

104. Nosek, B. A., Ebersole, C. R., DeHaven, A. C. & Mellor, D. T. The preregistration revolution. *Proc. Natl. Acad. Sci.* **115**, 2600–2606 (2018).

105. Maxwell, S. E., Lau, M. Y. & Howard, G. S. Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *Am. Psychol.* **70**, 487–498 (2015).

106. Shrout, P. E. & Rodgers, J. L. Psychology, science, and knowledge construction: Broadening perspectives from the replication Crisis. *Annu. Rev. Psychol.* **69**, 487–510 (2018).

107. Dreber, A. & Johannesson, M. A framework for evaluating reproducibility and replicability in economics. *Econ. Inq.* 1–19 (2024) doi:10/gt3vmw.

108. Benjamin, D. J. *et al.* Redefine statistical significance. *Nat. Hum. Behav.* **2**, 6–10 (2018).

109. Barnard, G. A. Significance tests for 2x2 tables. *Biometrika* **34**, 123–138 (1947).

110. Mehrotra, D. V., Chan, I. S. F. & Berger, R. L. A cautionary note on exact unconditional inference for a difference between two independent binomial proportions. *Biometrics* **59**, 441–450 (2003).

111. Boschloo, R. D. Raised conditional level of significance for the 2×2-table when testing the equality of two probabilities. *Stat. Neerlandica* **24**, 1–9 (1970).

112. Simonsohn, U. Small Telescopes: Detectability and the Evaluation of Replication Results. *Psychol. Sci.* **26**, 559–569 (2015).

113. Ly, A., Verhagen, J. & Wagenmakers, E.-J. Harold Jeffreys’s default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *J. Math. Psychol.*

72, 19–32 (2016).

114. Ly, A., Etz, A., Marsman, M. & Wagenmakers, E.-J. Replication Bayes factors from evidence updating. *Behav. Res. Methods* **51**, 2498–2508 (2019).

115. Jeffreys, H. *The Theory of Probability*. (Oxford University Press, Oxford, 1961).

116. Patil, P., Peng, R. D. & Leek, J. T. What should researchers expect when they replicate studies? A statistical view of replicability in psychological science. *Perspect. Psychol. Sci.* **11**, 539–544 (2016).

117. Gelman, A. & Stern, H. The difference between "significant" and "not significant" is not Itself statistically significant. *Am. Stat.* **60**, 328–331 (2006).

118. Cumming, G. Replication and p intervals: p values predict the future only vaguely, but confidence intervals do much better. *Perspect. Psychol. Sci.* **3**, 286–300 (2008).

119. Muradchanian, J., Hoekstra, R., Kiers, H. & van Ravenzwaaij, D. How best to quantify replication success? A simulation study on the comparison of replication success metrics. *R. Soc. Open Sci.* **8**, 201697 (2021).

120. Altmejd, A. *et al.* Predicting the replicability of social science lab experiments. *PLoS One* **14**, e0225826 (2019).

121. Yang, Y., Youyou, W. & Uzzi, B. Estimating the deep replicability of scientific findings using human and artificial intelligence. *Proc. Natl. Acad. Sci.* **117**, 10762–10768 (2020).

122. Rajtmajer, S. *et al.* A synthetic prediction market for estimating confidence in published work. *Proc. AAAI Conf. Artif. Intell.* **36**, 13218–13220 (2022).

123. Youyou, W., Yang, Y. & Uzzi, B. A discipline-wide investigation of the replicability of psychology papers over the past two decades. *Proc. Natl. Acad. Sci.* **120**, e2208863120 (2023).

124. Agley, J., Xiao, Y., Nolan, R. & Golzarri-Arroyo, L. Quality control questions on Amazon's Mechanical Turk (MTurk): A randomized trial of impact on the USAUDIT, PHQ-9, and GAD-7. *Behav. Res. Methods* **54**, 885–897 (2022).

125. Veselovsky, V., Ribeiro, M. H. & West, R. Artificial artificial artificial intelligence: Crowd workers widely use large language models for text production tasks. Preprint at https://doi.org/10/gsc3gr (2023).

126. Olsson-Collentine, A., Wicherts, J. M. & van Assen, M. A. L. M. Heterogeneity in direct replications in psychology and its association with effect size. *Psychol. Bull.* **146**, 922–940 (2020).

127. Linden, A. H. & Hönekopp, J. Heterogeneity of research results: A new perspective from which to assess and promote progress in psychological science. *Perspect. Psychol. Sci.* **16**, 358–376 (2021).

128. Holzmeister, F. *et al.* Heterogeneity in effect size estimates. *Proc. Natl. Acad. Sci.* **121**, e2403490121 (2024).

129. Epstein, R. & Robertson, R. E. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proc. Natl. Acad. Sci.* **112**, E4512–E4521 (2015).

130. Gallo, E. & Yan, C. The effects of reputational and social knowledge on cooperation. *Proc. Natl. Acad. Sci.* **112**, 3647–3652 (2015).

131. Li, V., Michael, E., Balaguer, J., Herce Castañón, S. & Summerfield, C. Gain control explains the effect of distraction in human perceptual, cognitive, and economic decision making. *Proc. Natl. Acad. Sci.* **115**, E8825–E8834 (2018).

132. Hanson, R. Logarithmic market scoring rules for modular combinatorial information aggregation. *J. Predict. Mark.* **1**, 3–15 (2007).

133. Moshontz, H. *et al.* The Psychological Science Accelerator: Advancing psychology through a distributed collaborative network. *Adv. Methods Pract. Psychol. Sci.* **1**, 501–515 (2018).

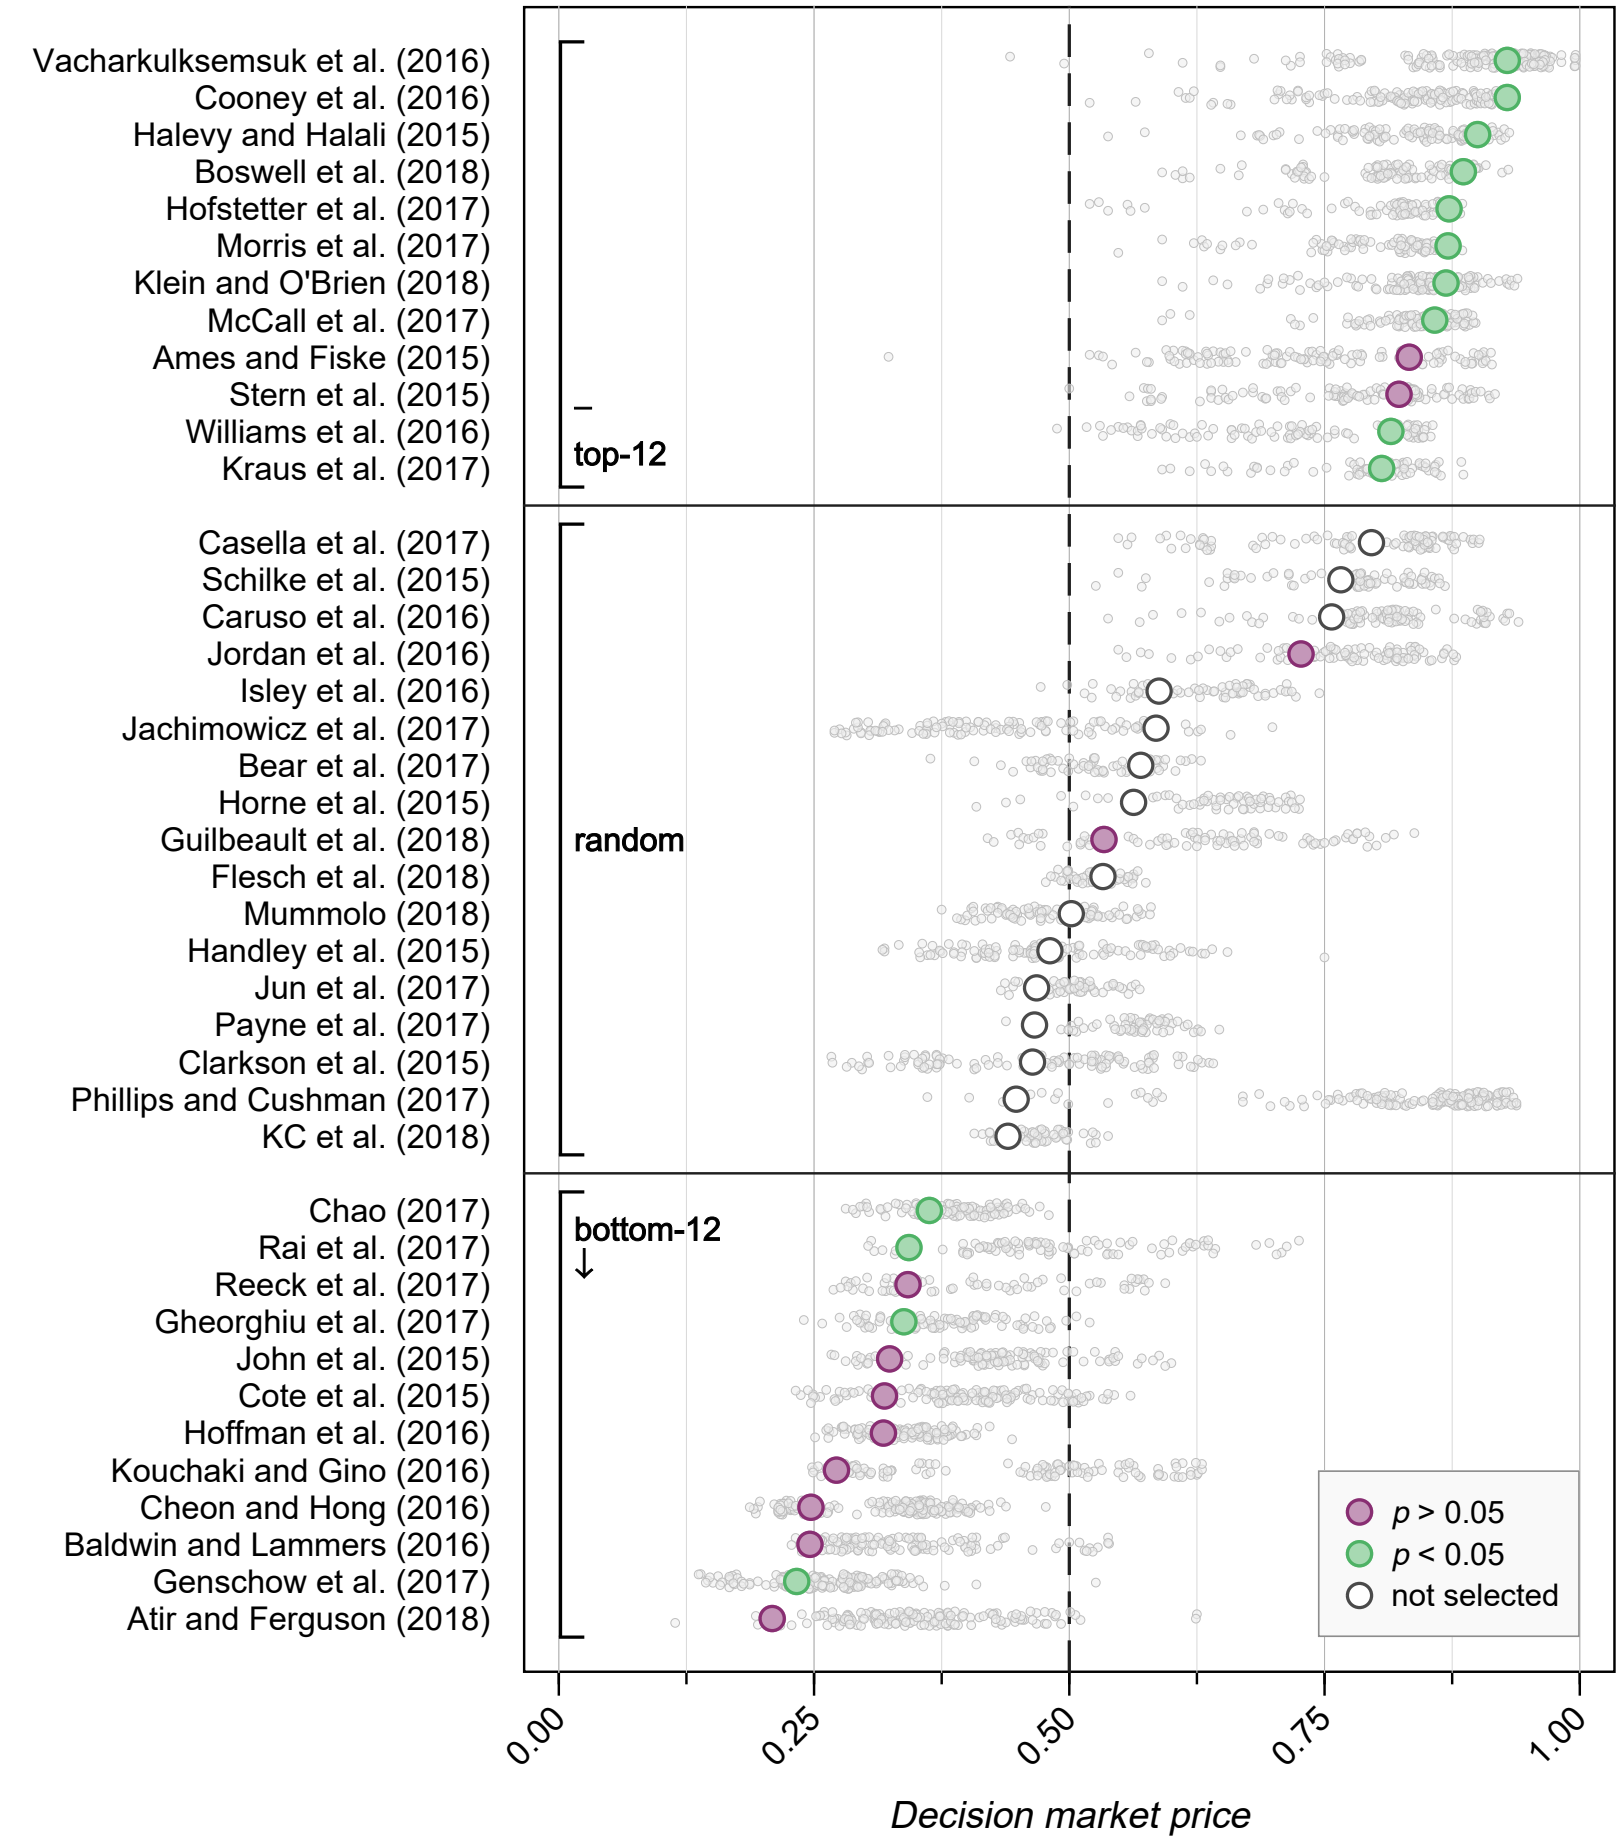
Vacharkulksemsuk et al. (2016)
Cooney et al. (2016)
Halevy and Halali (2015)
Boswell et al. (2018)
Hofstetter et al. (2017)
Morris et al. (2017)
Klein and O'Brien (2018)
McCall et al. (2017)
Ames and Fiske (2015)
Stern et al. (2015)
Williams et al. (2016)
Kraus et al. (2017)
top-12
Casella et al. (2017)
Schilke et al. (2015)
Caruso et al. (2016)
Jordan et al. (2016)
Isley et al. (2016)
Jachimowicz et al. (2017)
Bear et al. (2017)
Horne et al. (2015)
Guilbeault et al. (2018)
Flesch et al. (2018)
Mummolo (2018)
Handley et al. (2015)
Jun et al. (2017)
Payne et al. (2017)
Clarkson et al. (2015)
Phillips and Cushman (2017)
KC et al. (2018)
random
Chao (2017)
Rai et al. (2017)
Reeck et al. (2017)
Gheorghiu et al. (2017)
John et al. (2015)
Cote et al. (2015)
Hoffman et al. (2016)
Kouchaki and Gino (2016)
Cheon and Hong (2016)
Baldwin and Lammers (2016)
Genschow et al. (2017)
Atir and Ferguson (2018)
bottom-12
↓
p > 0.05
p < 0.05
not selected
0.00
0.25
0.50
0.75
1.00
Decision market price

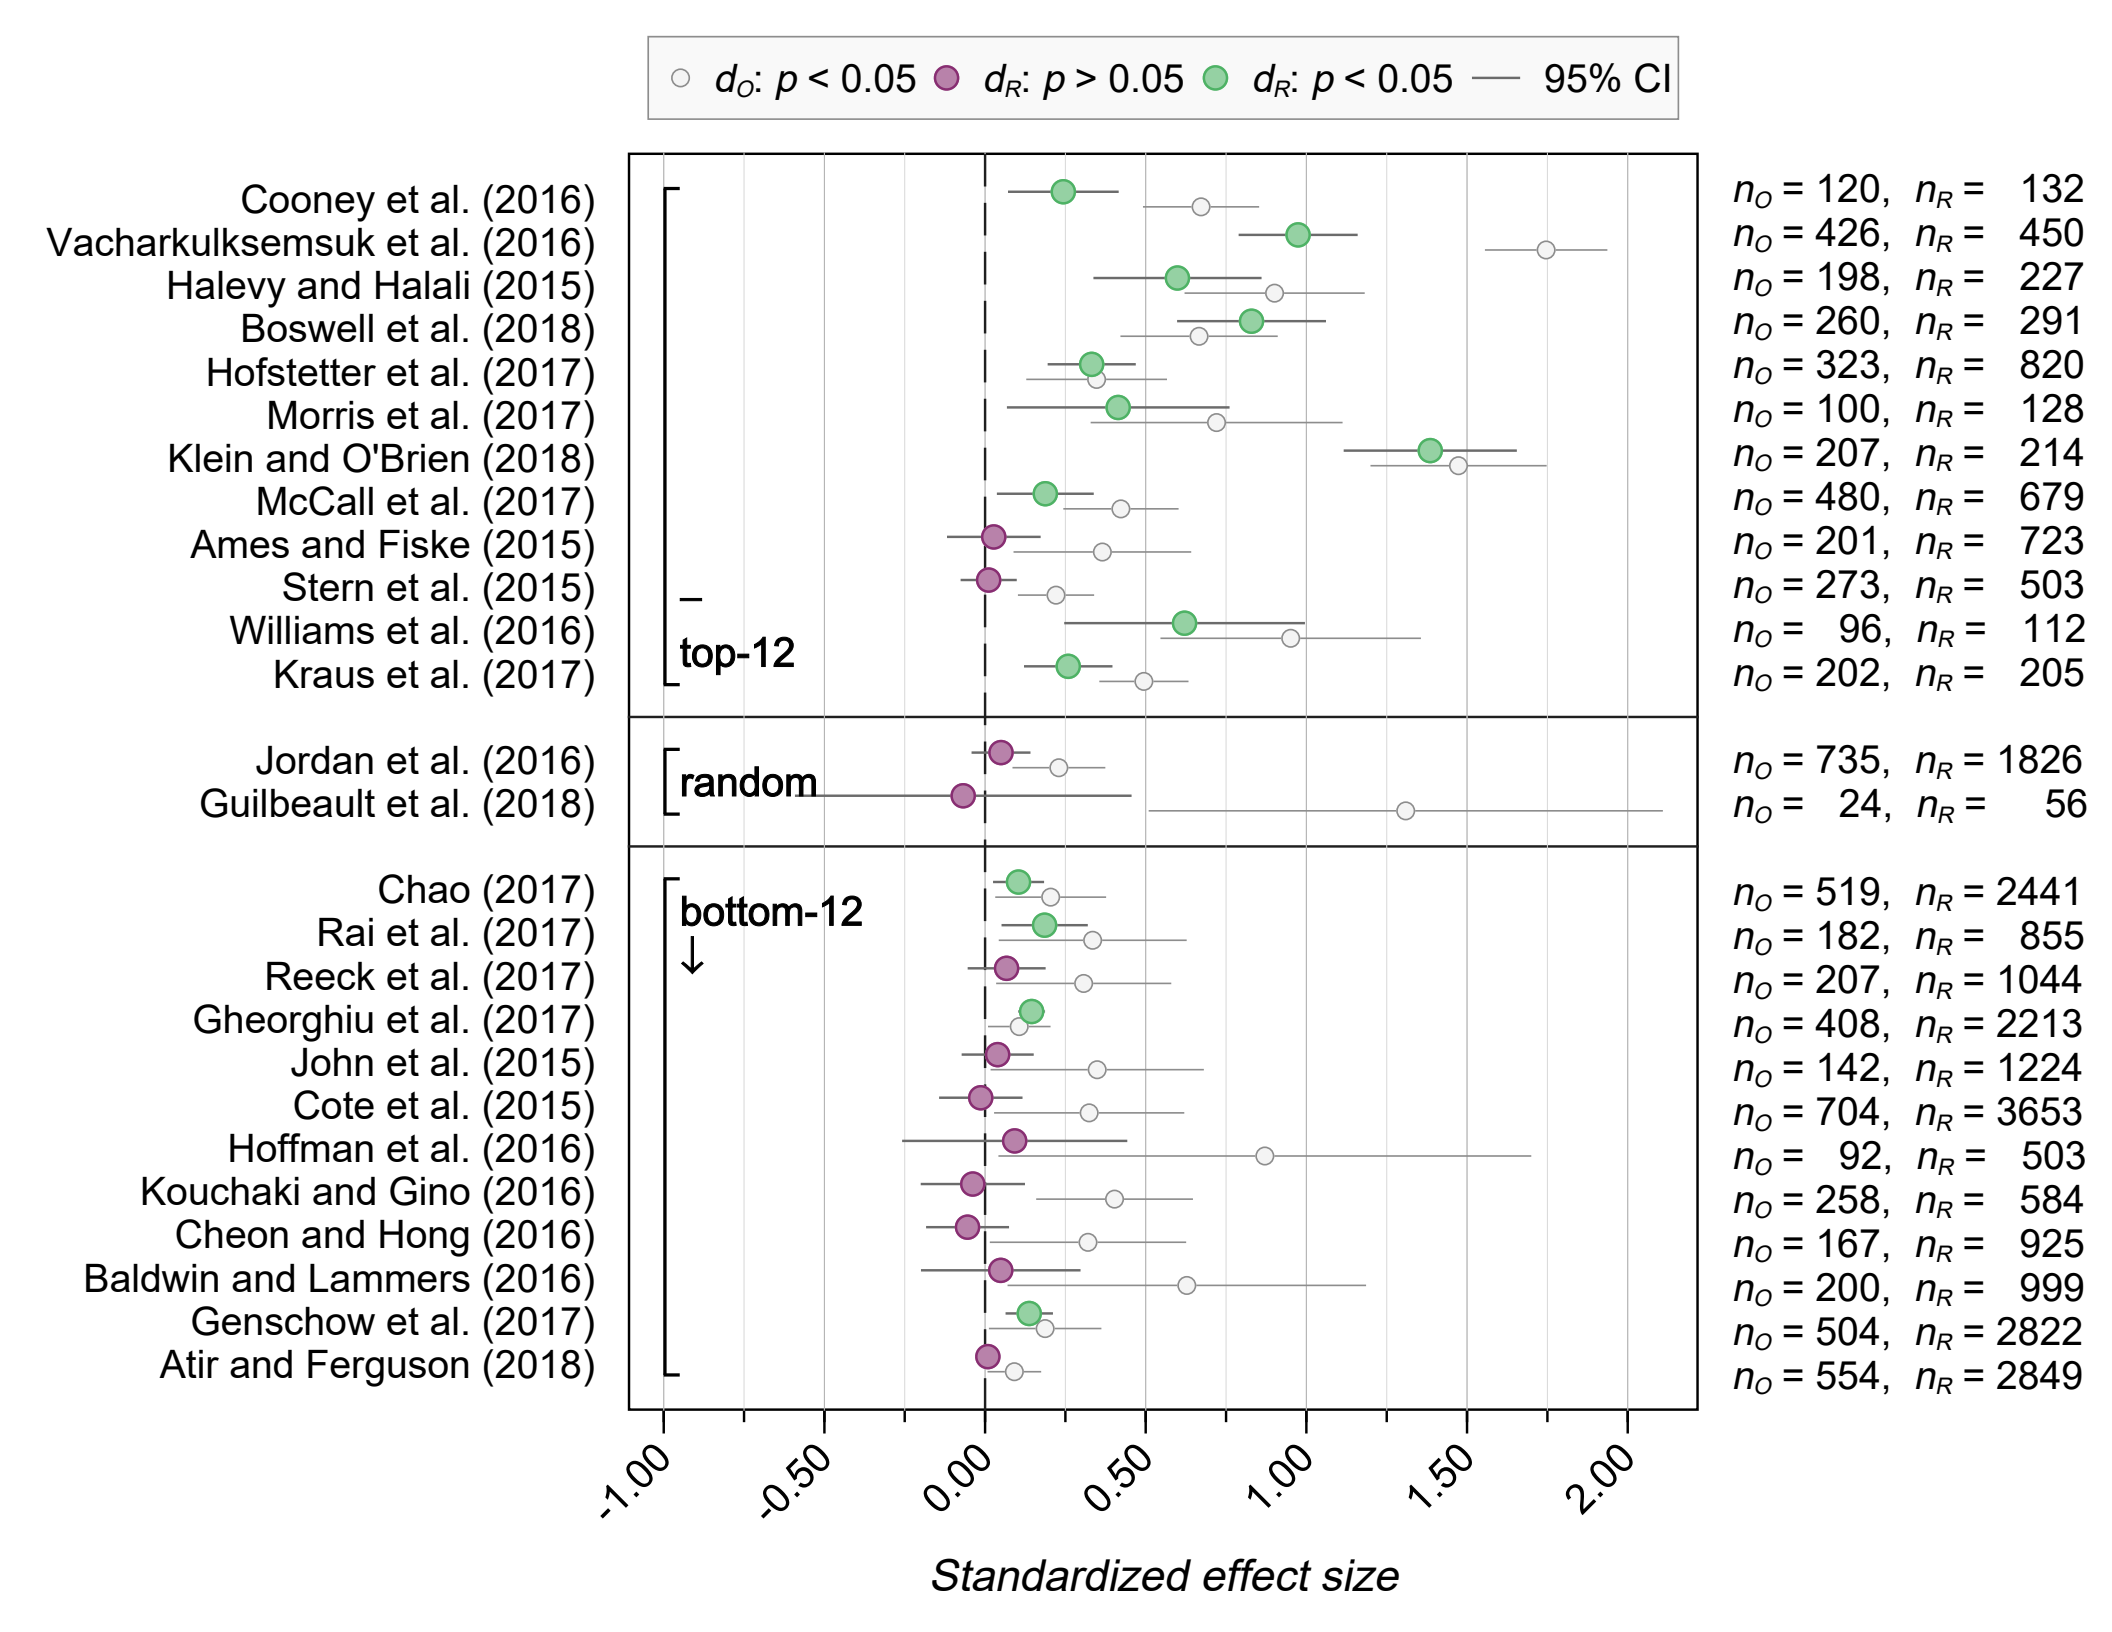
d_O: $p < 0.05$
d_R: $p > 0.05$
d_R: $p < 0.05$
95% CI
Cooney et al. (2016)
Vacharkulksemsuk et al. (2016)
Halevy and Halali (2015)
Boswell et al. (2018)
Hofstetter et al. (2017)
Morris et al. (2017)
Klein and O'Brien (2018)
McCall et al. (2017)
Ames and Fiske (2015)
Stern et al. (2015)
Williams et al. (2016)
Kraus et al. (2017)
top-12
Jordan et al. (2016)
Guilbeault et al. (2018)
random
Chao (2017)
Rai et al. (2017)
Reeck et al. (2017)
Gheorghiu et al. (2017)
John et al. (2015)
Cote et al. (2015)
Hoffman et al. (2016)
Kouchaki and Gino (2016)
Cheon and Hong (2016)
Baldwin and Lammers (2016)
Genschow et al. (2017)
Atir and Ferguson (2018)
bottom-12
n_O = 120, n_R = 132
n_O = 426, n_R = 450
n_O = 198, n_R = 227
n_O = 260, n_R = 291
n_O = 323, n_R = 820
n_O = 100, n_R = 128
n_O = 207, n_R = 214
n_O = 480, n_R = 679
n_O = 201, n_R = 723
n_O = 273, n_R = 503
n_O = 96, n_R = 112
n_O = 202, n_R = 205
n_O = 735, n_R = 1826
n_O = 24, n_R = 56
n_O = 519, n_R = 2441
n_O = 182, n_R = 855
n_O = 207, n_R = 1044
n_O = 408, n_R = 2213
n_O = 142, n_R = 1224
n_O = 704, n_R = 3653
n_O = 92, n_R = 503
n_O = 258, n_R = 584
n_O = 167, n_R = 925
n_O = 200, n_R = 999
n_O = 504, n_R = 2822
n_O = 554, n_R = 2849
-1.00
-0.50
0.00
0.50
1.00
1.50
2.00
Standardized effect size

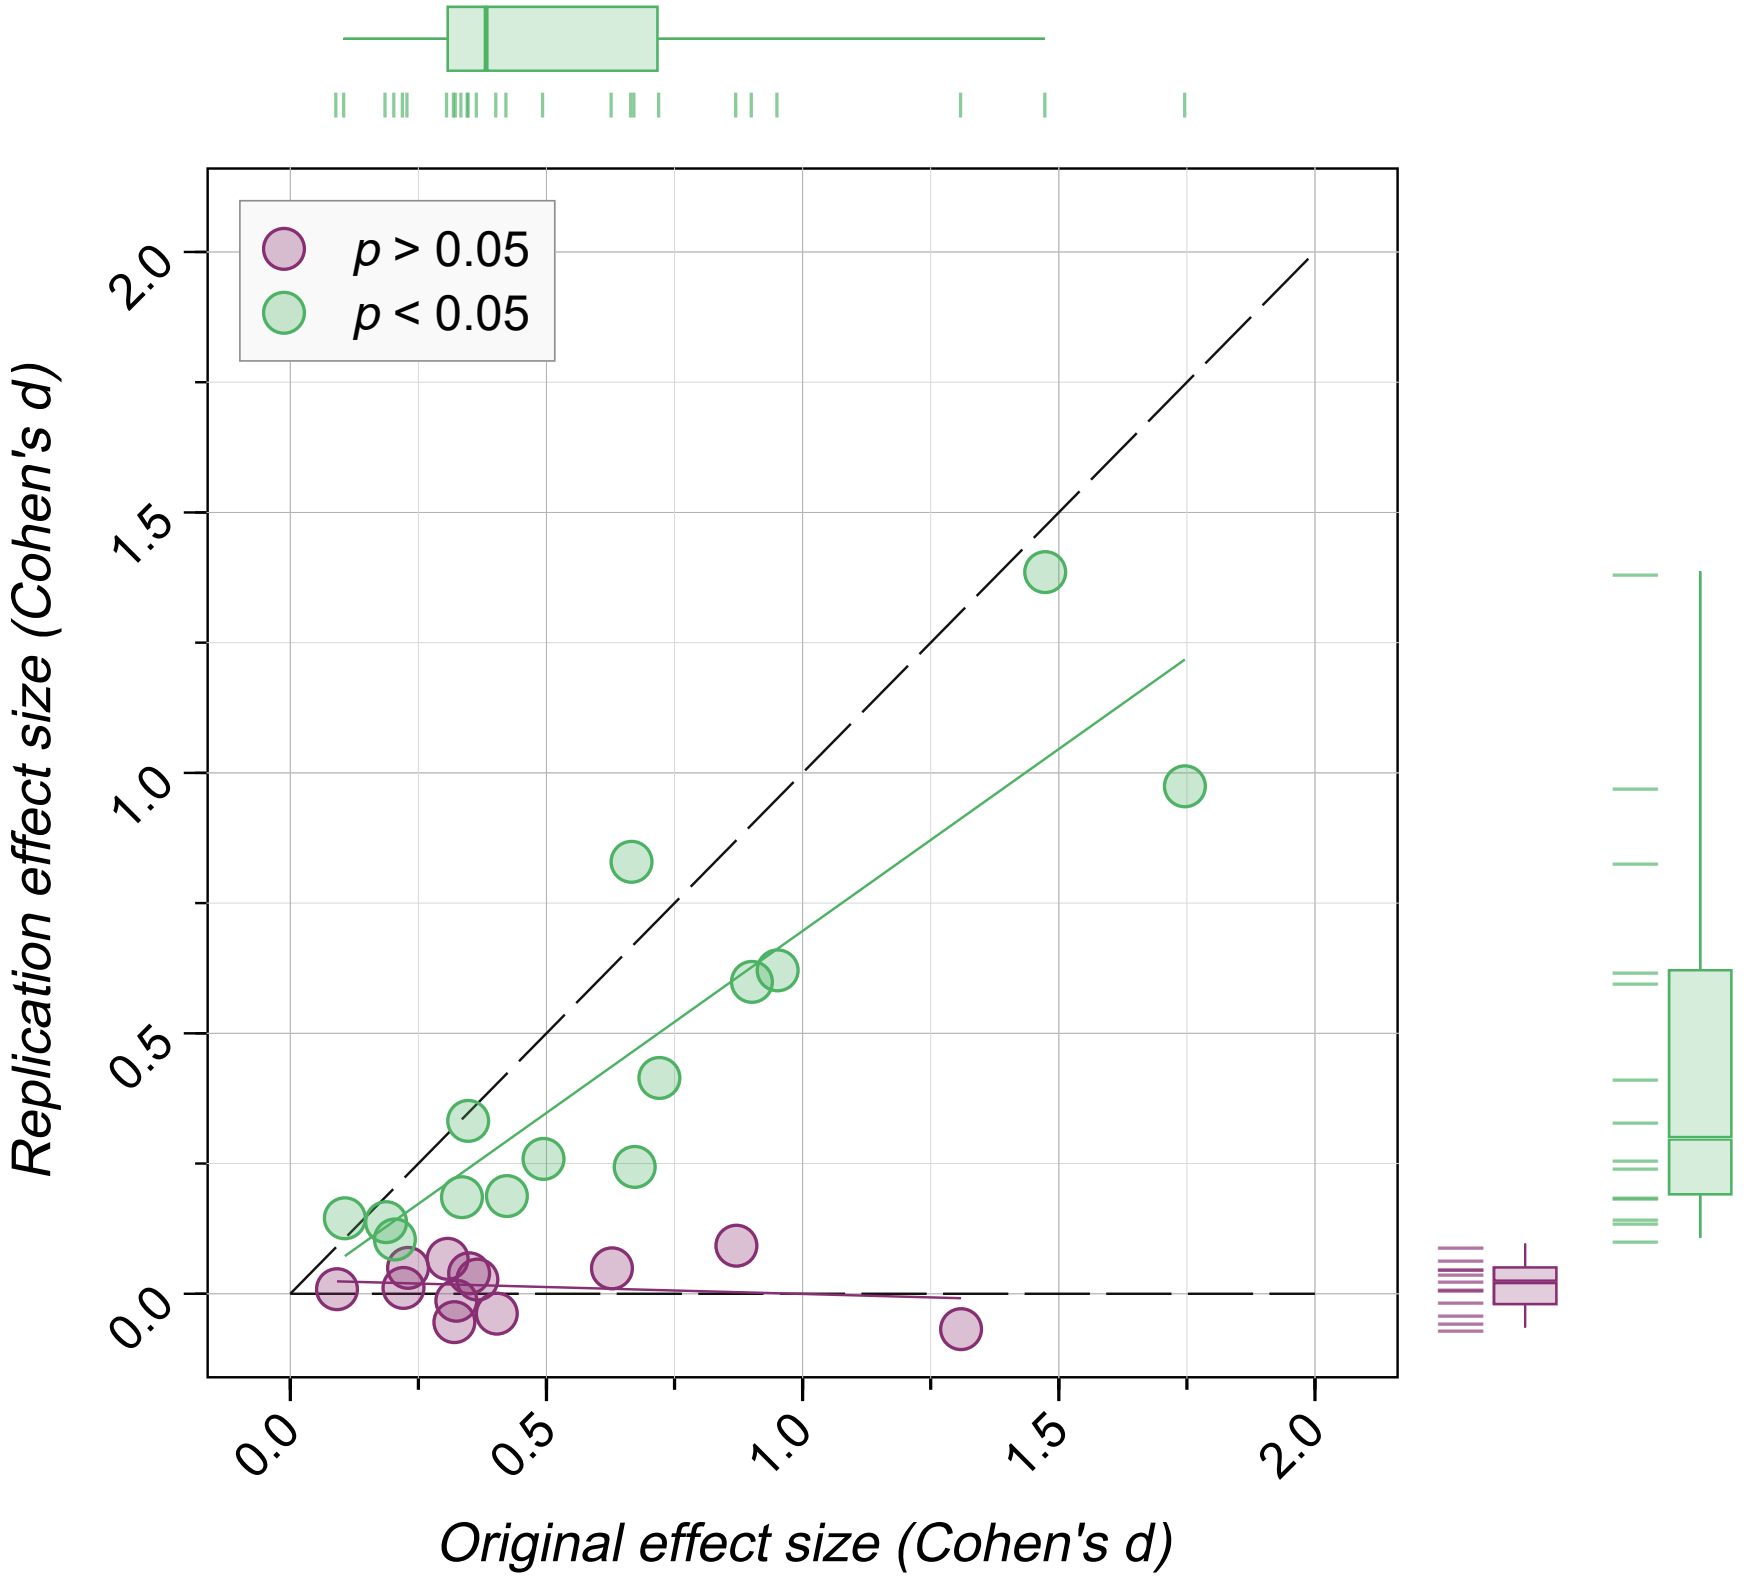
p > 0.05
p < 0.05
Replication effect size (Cohen's d)
Original effect size (Cohen's d)
0.0
0.5
1.0
1.5
2.0

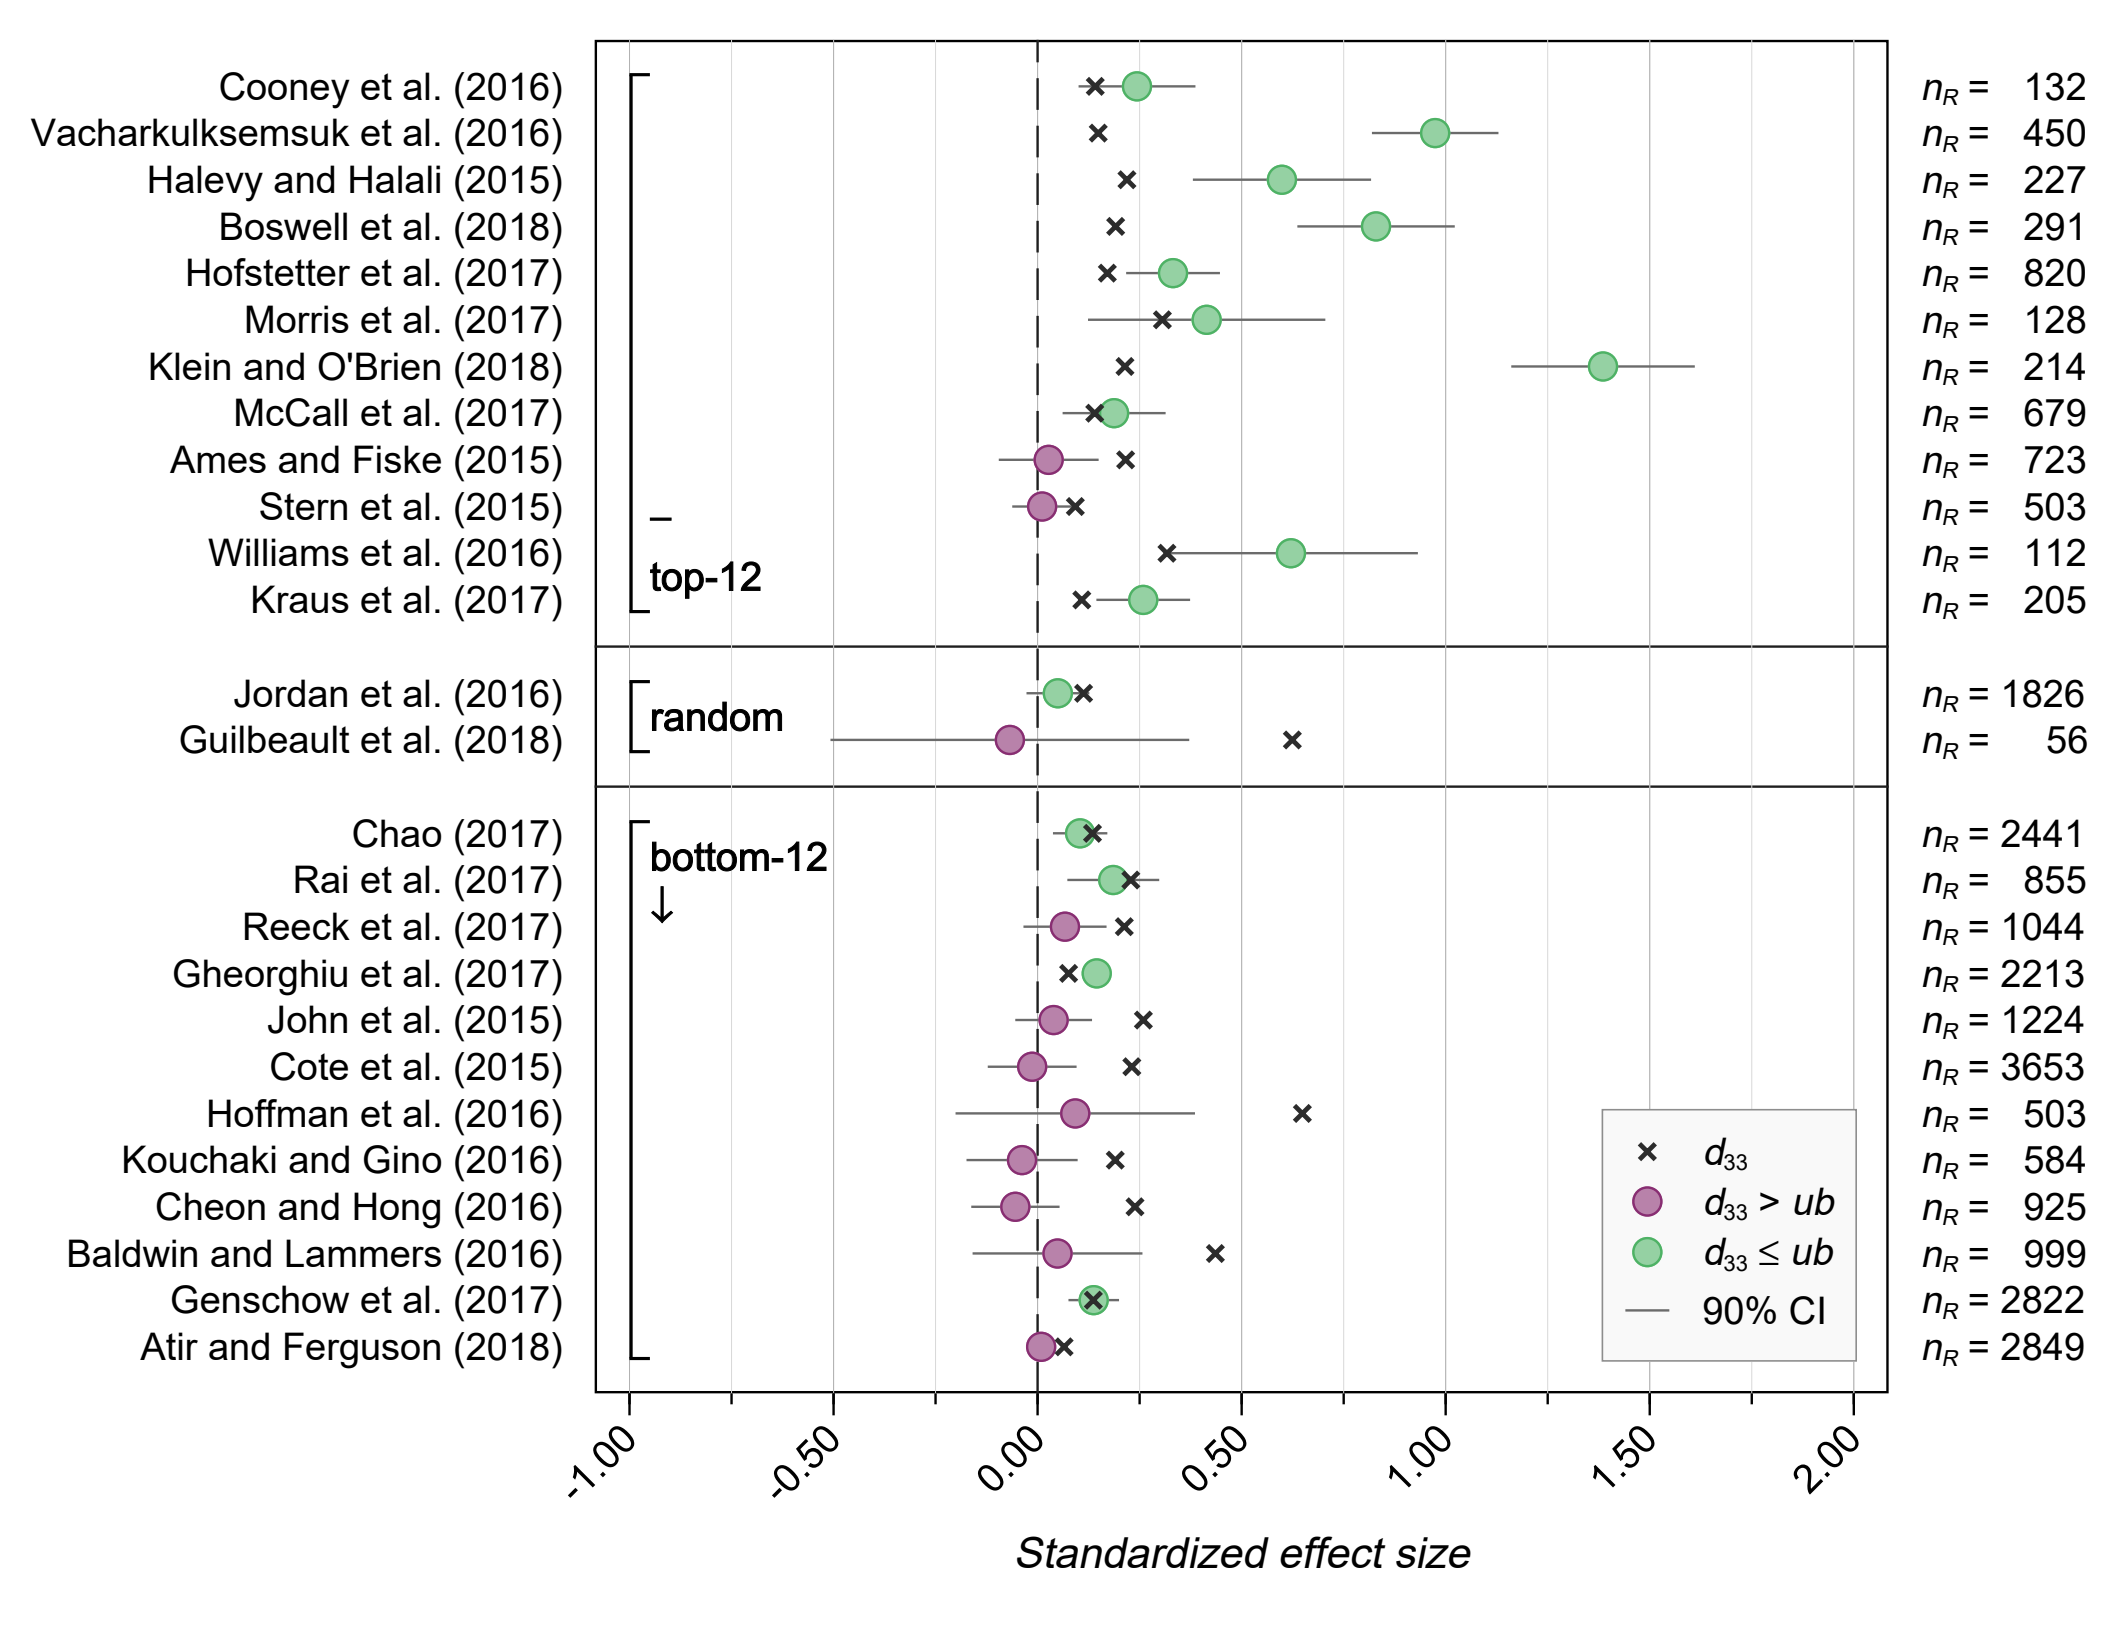
Cooney et al. (2016)
Vacharkulksemsuk et al. (2016)
Halevy and Halali (2015)
Boswell et al. (2018)
Hofstetter et al. (2017)
Morris et al. (2017)
Klein and O'Brien (2018)
McCall et al. (2017)
Ames and Fiske (2015)
Stern et al. (2015)
Williams et al. (2016)
Kraus et al. (2017)
top-12
Jordan et al. (2016)
Guilbeault et al. (2018)
random
Chao (2017)
Rai et al. (2017)
Reeck et al. (2017)
Gheorghiu et al. (2017)
John et al. (2015)
Cote et al. (2015)
Hoffman et al. (2016)
Kouchaki and Gino (2016)
Cheon and Hong (2016)
Baldwin and Lammers (2016)
Genschow et al. (2017)
Atir and Ferguson (2018)
bottom-12
n_R = 132
n_R = 450
n_R = 227
n_R = 291
n_R = 820
n_R = 128
n_R = 214
n_R = 679
n_R = 723
n_R = 503
n_R = 112
n_R = 205
n_R = 1826
n_R = 56
n_R = 2441
n_R = 855
n_R = 1044
n_R = 2213
n_R = 1224
n_R = 3653
n_R = 503
n_R = 584
n_R = 925
n_R = 999
n_R = 2822
n_R = 2849
d_{33}
$d_{33} > ub$
$d_{33} \leq ub$
90% CI
-1.00
-0.50
0.00
0.50
1.00
1.50
2.00
Standardized effect size

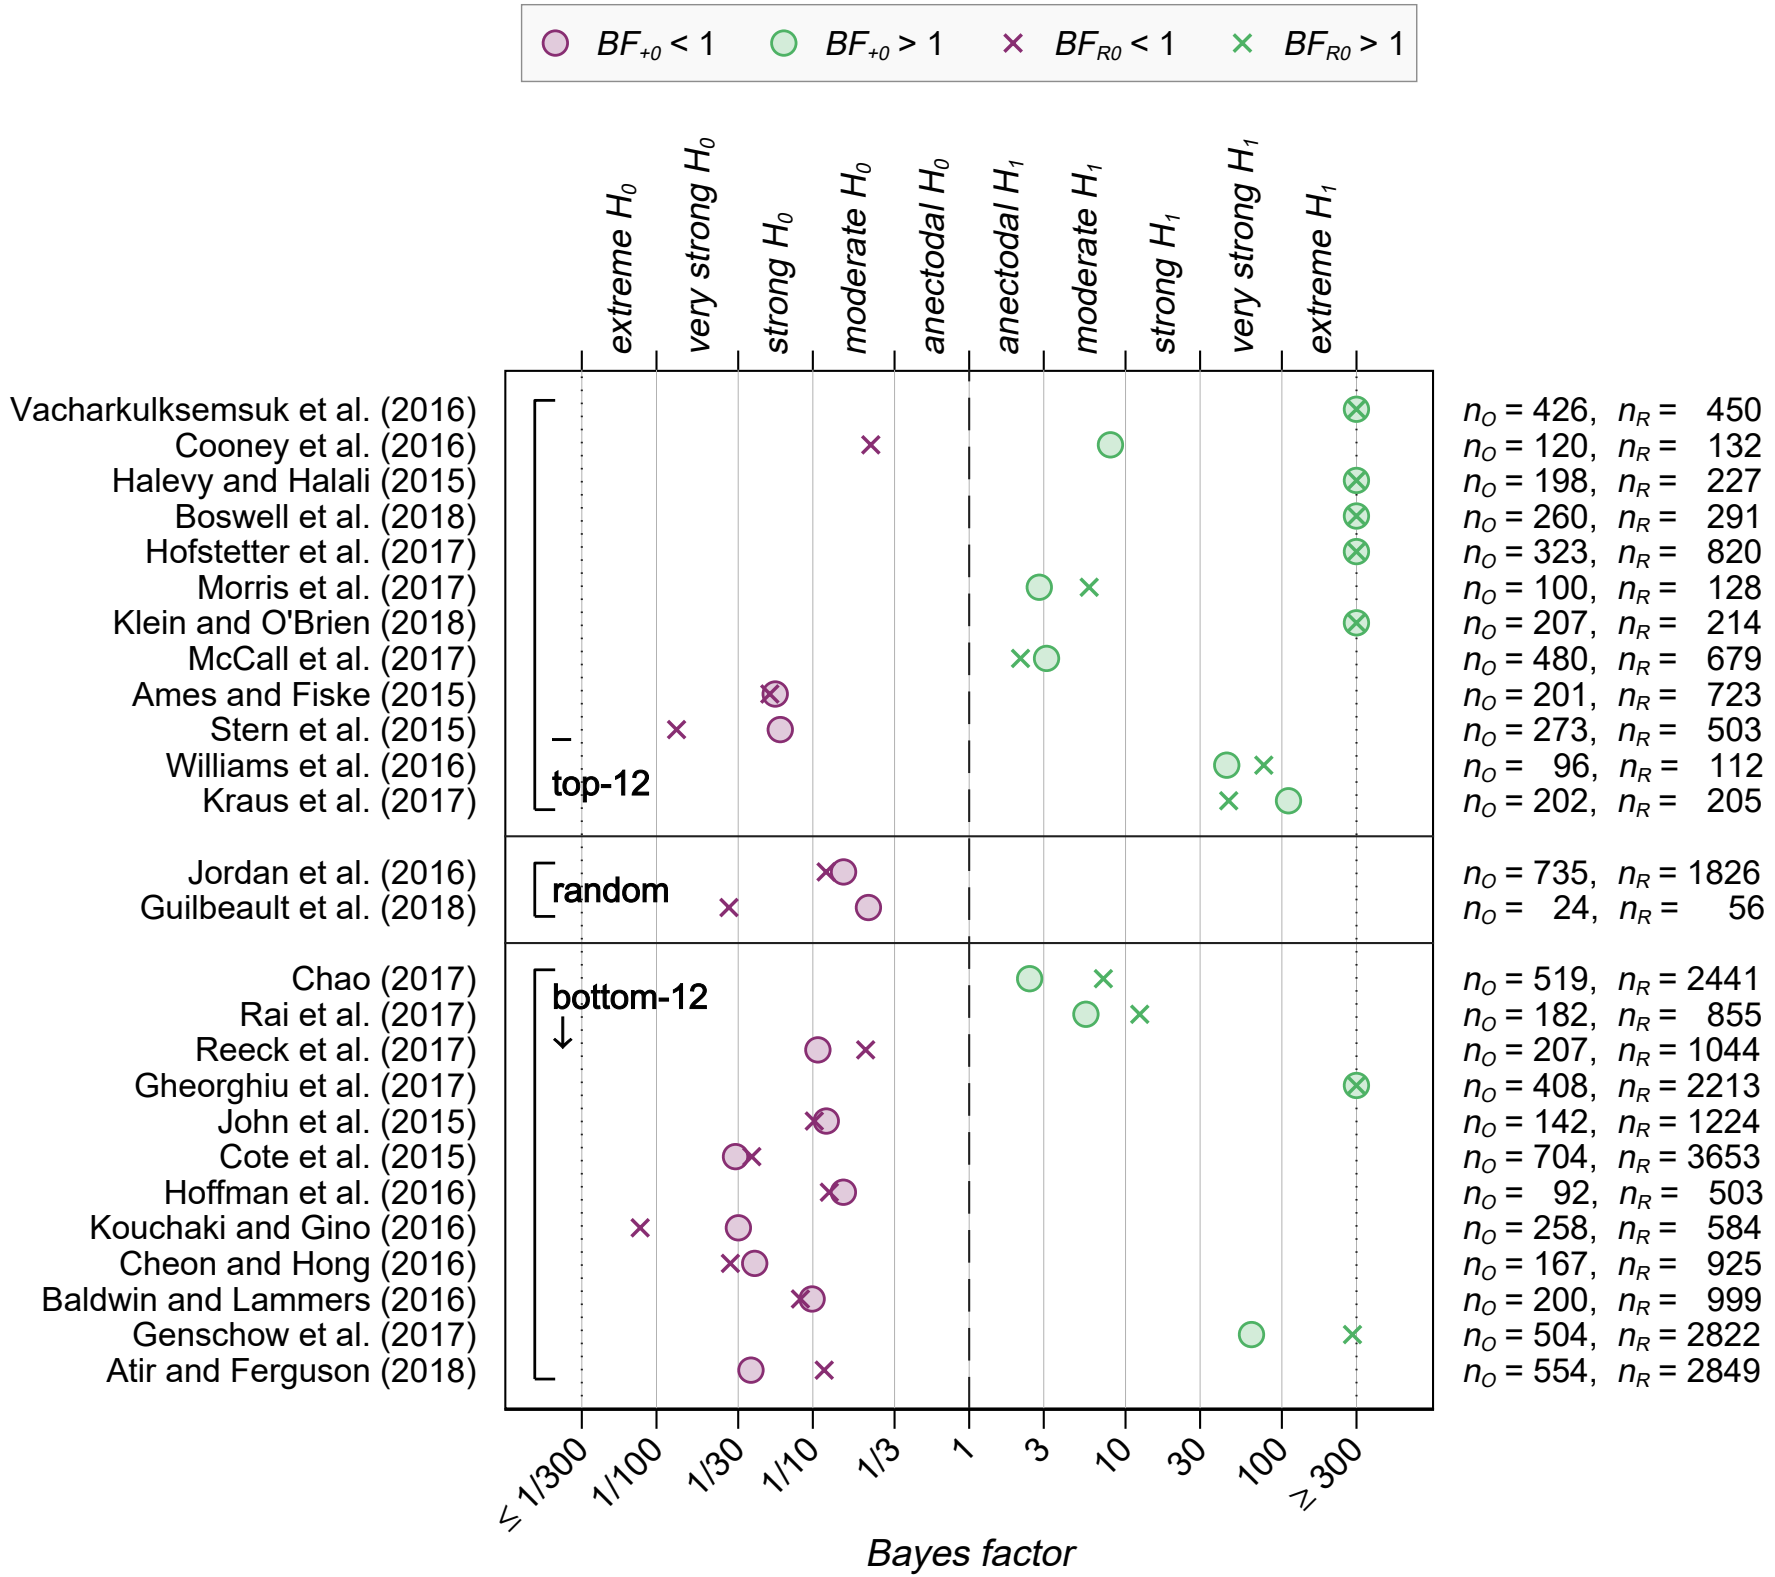
BF+0 < 1
BF+0 > 1
BFR0 < 1
BFR0 > 1
extreme H0
very strong H0
strong H0
moderate H0
anectodal H0
anectodal H1
moderate H1
strong H1
very strong H1
extreme H1
Vacharkulksemsuk et al. (2016) nO = 426, nR = 450
Cooney et al. (2016) nO = 120, nR = 132
Halevy and Halali (2015) nO = 198, nR = 227
Boswell et al. (2018) nO = 260, nR = 291
Hofstetter et al. (2017) nO = 323, nR = 820
Morris et al. (2017) nO = 100, nR = 128
Klein and O'Brien (2018) nO = 207, nR = 214
McCall et al. (2017) nO = 480, nR = 679
Ames and Fiske (2015) nO = 201, nR = 723
Stern et al. (2015) nO = 273, nR = 503
Williams et al. (2016) nO = 96, nR = 112
Kraus et al. (2017) nO = 202, nR = 205
top-12
Jordan et al. (2016) nO = 735, nR = 1826
Guilbeault et al. (2018) nO = 24, nR = 56
random
Chao (2017) nO = 519, nR = 2441
Rai et al. (2017) nO = 182, nR = 855
Reeck et al. (2017) nO = 207, nR = 1044
Gheorghiu et al. (2017) nO = 408, nR = 2213
John et al. (2015) nO = 142, nR = 1224
Cote et al. (2015) nO = 704, nR = 3653
Hoffman et al. (2016) nO = 92, nR = 503
Kouchaki and Gino (2016) nO = 258, nR = 584
Cheon and Hong (2016) nO = 167, nR = 925
Baldwin and Lammers (2016) nO = 200, nR = 999
Genschow et al. (2017) nO = 504, nR = 2822
Atir and Ferguson (2018) nO = 554, nR = 2849
bottom-12
≤ 1/300
1/100
1/30
1/10
1/3
1
3
10
30
100
≥ 300
Bayes factor

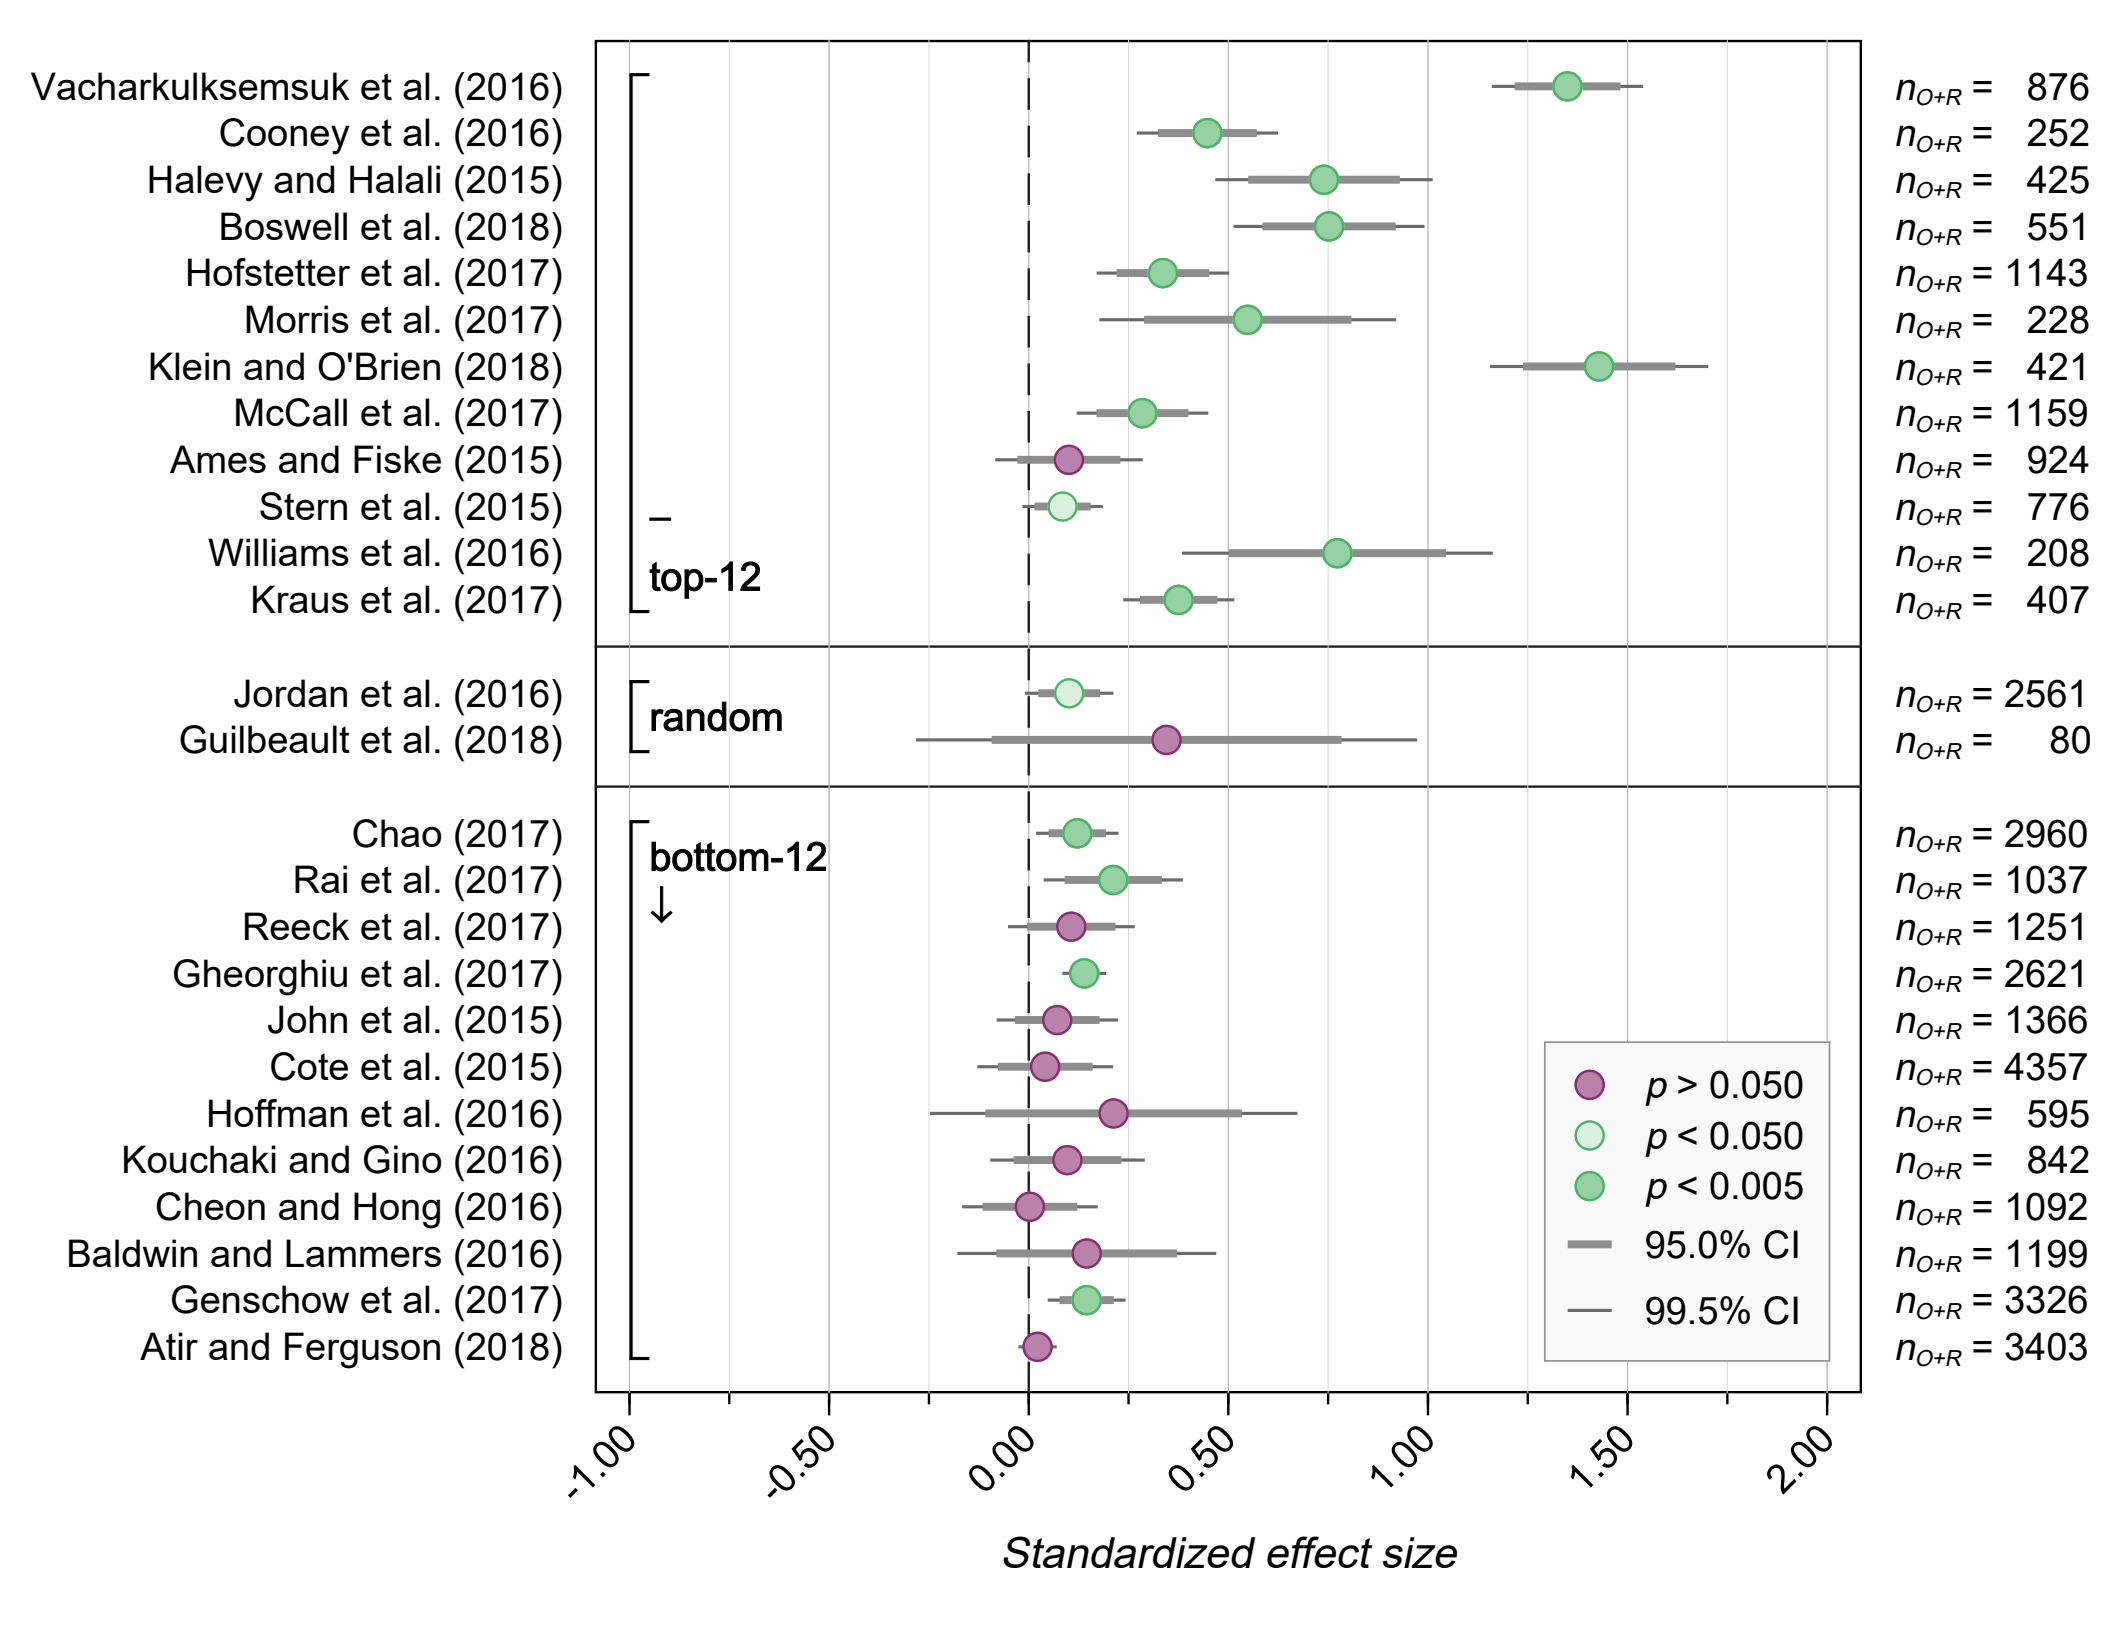
Vacharkulksemsuk et al. (2016)
Cooney et al. (2016)
Halevy and Halali (2015)
Boswell et al. (2018)
Hofstetter et al. (2017)
Morris et al. (2017)
Klein and O'Brien (2018)
McCall et al. (2017)
Ames and Fiske (2015)
Stern et al. (2015)
Williams et al. (2016)
Kraus et al. (2017)
top-12
Jordan et al. (2016)
Guilbeault et al. (2018)
random
Chao (2017)
Rai et al. (2017)
Reeck et al. (2017)
Gheorghiu et al. (2017)
John et al. (2015)
Cote et al. (2015)
Hoffman et al. (2016)
Kouchaki and Gino (2016)
Cheon and Hong (2016)
Baldwin and Lammers (2016)
Genschow et al. (2017)
Atir and Ferguson (2018)
bottom-12
n_{O+R} = 876
n_{O+R} = 252
n_{O+R} = 425
n_{O+R} = 551
n_{O+R} = 1143
n_{O+R} = 228
n_{O+R} = 421
n_{O+R} = 1159
n_{O+R} = 924
n_{O+R} = 776
n_{O+R} = 208
n_{O+R} = 407
n_{O+R} = 2561
n_{O+R} = 80
n_{O+R} = 2960
n_{O+R} = 1037
n_{O+R} = 1251
n_{O+R} = 2621
n_{O+R} = 1366
n_{O+R} = 4357
n_{O+R} = 595
n_{O+R} = 842
n_{O+R} = 1092
n_{O+R} = 1199
n_{O+R} = 3326
n_{O+R} = 3403
$p > 0.050$
$p < 0.050$
$p < 0.005$
95.0% CI
99.5% CI
-1.00
-0.50
0.00
0.50
1.00
1.50
2.00
Standardized effect size

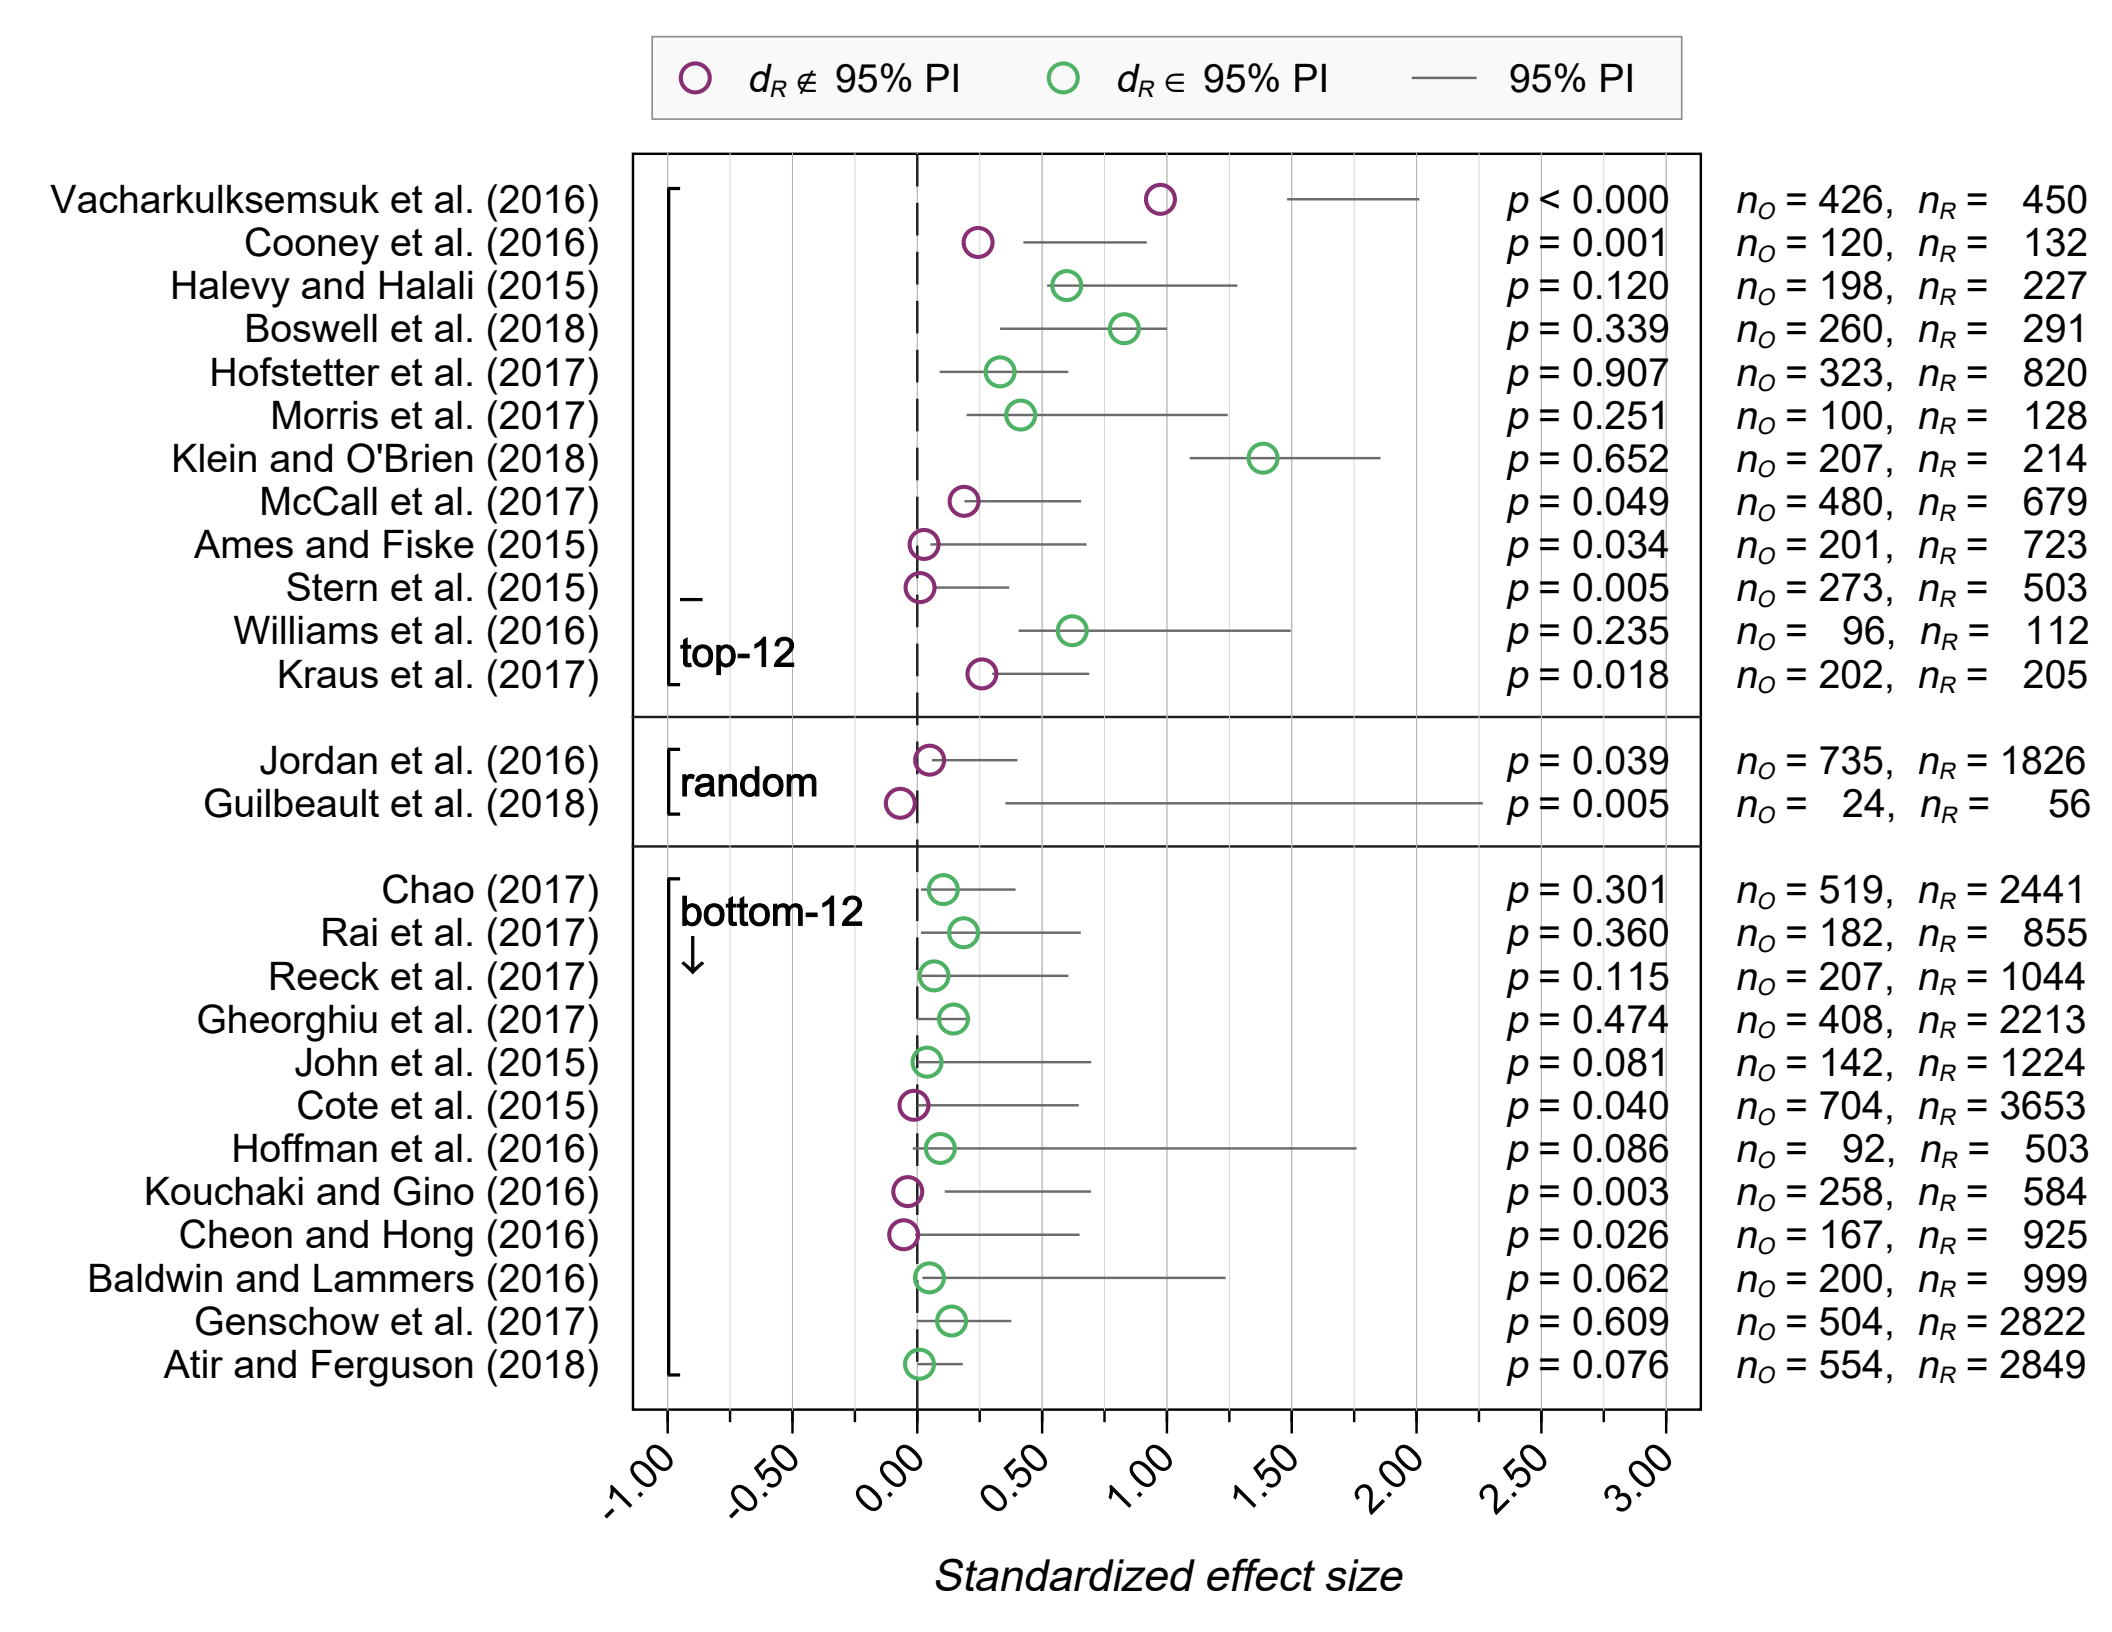
$d_R \notin$ 95% PI
$d_R \in$ 95% PI
95% PI
Vacharkulksemsuk et al. (2016) p < 0.000 n_O = 426, n_R = 450
Cooney et al. (2016) p = 0.001 n_O = 120, n_R = 132
Halevy and Halali (2015) p = 0.120 n_O = 198, n_R = 227
Boswell et al. (2018) p = 0.339 n_O = 260, n_R = 291
Hofstetter et al. (2017) p = 0.907 n_O = 323, n_R = 820
Morris et al. (2017) p = 0.251 n_O = 100, n_R = 128
Klein and O'Brien (2018) p = 0.652 n_O = 207, n_R = 214
McCall et al. (2017) p = 0.049 n_O = 480, n_R = 679
Ames and Fiske (2015) p = 0.034 n_O = 201, n_R = 723
Stern et al. (2015) p = 0.005 n_O = 273, n_R = 503
Williams et al. (2016) p = 0.235 n_O = 96, n_R = 112
Kraus et al. (2017) p = 0.018 n_O = 202, n_R = 205
top-12
Jordan et al. (2016) p = 0.039 n_O = 735, n_R = 1826
Guilbeault et al. (2018) p = 0.005 n_O = 24, n_R = 56
random
Chao (2017) p = 0.301 n_O = 519, n_R = 2441
Rai et al. (2017) p = 0.360 n_O = 182, n_R = 855
Reeck et al. (2017) p = 0.115 n_O = 207, n_R = 1044
Gheorghiu et al. (2017) p = 0.474 n_O = 408, n_R = 2213
John et al. (2015) p = 0.081 n_O = 142, n_R = 1224
Cote et al. (2015) p = 0.040 n_O = 704, n_R = 3653
Hoffman et al. (2016) p = 0.086 n_O = 92, n_R = 503
Kouchaki and Gino (2016) p = 0.003 n_O = 258, n_R = 584
Cheon and Hong (2016) p = 0.026 n_O = 167, n_R = 925
Baldwin and Lammers (2016) p = 0.062 n_O = 200, n_R = 999
Genschow et al. (2017) p = 0.609 n_O = 504, n_R = 2822
Atir and Ferguson (2018) p = 0.076 n_O = 554, n_R = 2849
bottom-12
-1.00
-0.50
0.00
0.50
1.00
1.50
2.00
2.50
3.00
Standardized effect size

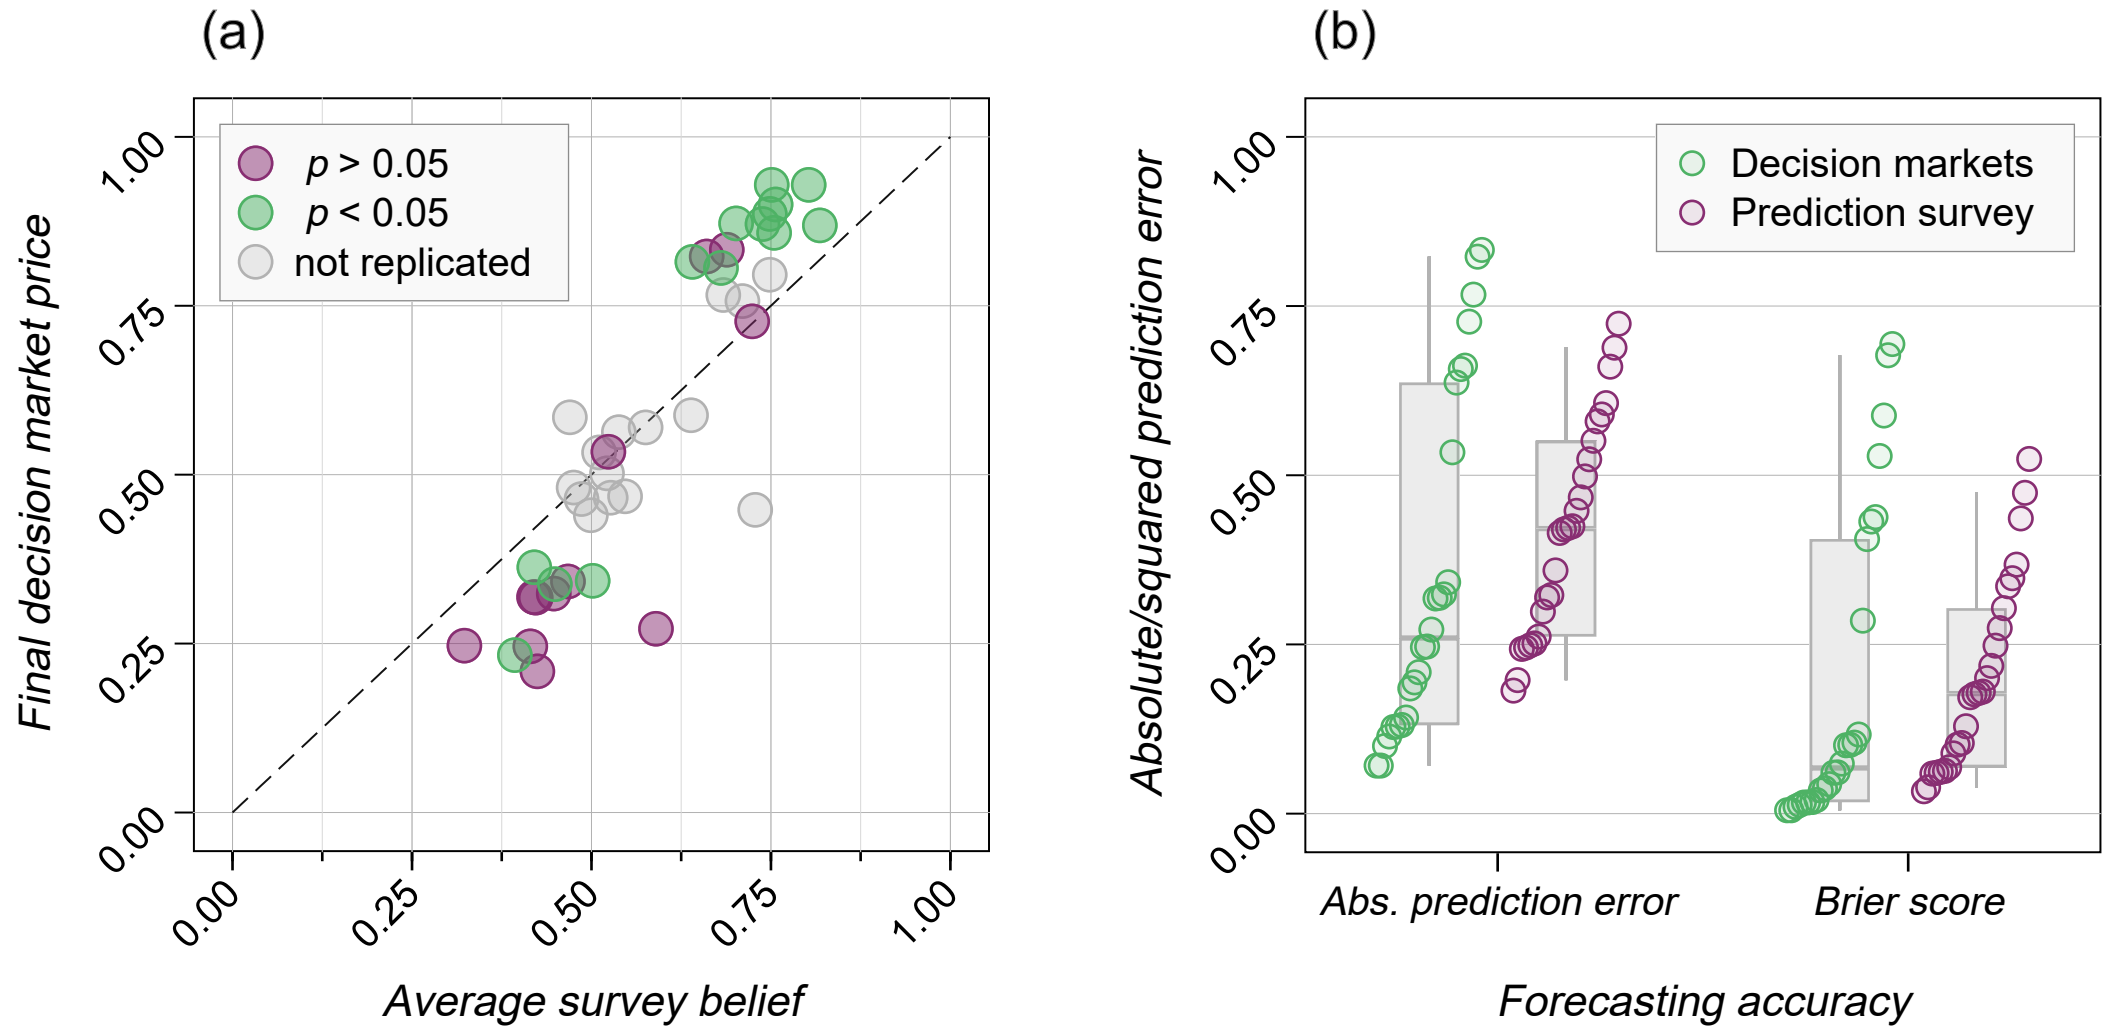
(a)
p > 0.05
p < 0.05
not replicated
Final decision market price
Average survey belief
0.00
0.25
0.50
0.75
1.00
(b)
Decision markets
Prediction survey
Absolute/squared prediction error
Abs. prediction error
Brier score
Forecasting accuracy