

The Implicit Honesty Premium:
Why Honest Advice is More Persuasive than Highly Informed Advice

Supplementary Online Material: Additional Measures and Results

Experiment 2

In this experiment, we took the following measures, in addition to those reported in the manuscript:

1. In each of the six estimation tasks:
 - a. Confidence that the initial estimate is accurate (on a 0%-100% scale).
 - b. Confidence that the final estimate is accurate (on a 0%-100% scale).
2. Post-task attitudes toward the advisor, rated on a 1-9 scale.
 - a. The extent to which participants thought the advisor wanted to hurt them.
 - b. The extent to which participants thought the advisor wanted to help them.
 - c. The degree to which participants thought they were harmed by the advisor.
 - d. The degree to which participants trusted the advisor.
 - e. The degree to which participants were willing to use the advisor's advice.

Results

Advice taking. To supplement the advice taking results reported in the main article, Table S1 summarizes the rates of advice taking, remaining with one's initial estimate and active distrust for all experimental conditions.

Table S1.

Rates of advice taking, remaining with the initial estimate and active distrust by advice quality and reason for suspicion in Experiment 2.

Outcome	Advice quality	Reason for suspicion	Mean (SD)	Effect of reason for suspicion
Advice taking	Advisor A (6 accurate)	Error	.87 (.22)	$t(58) = 1.61, p = .11,$ $d = 0.42$
		Intentional bias	.78 (.21)	
	Advisor B (5 accurate)	Error	.77 (.26)	$t(58) = 2.19, p = .03,$ $d = 0.58$
		Intentional bias	.62 (.27)	
	Advisor C (4 accurate)	Error	.74 (.26)	$t(49) = 1.94, p = .06,$ $d = 0.55$
		Intentional bias	.59 (.31)	
	Advisor D (3 accurate)	Error	.65 (.35)	$t < 1$
		Intentional bias	.57 (.29)	
Remaining with initial estimate	Advisor A (6 accurate)	Error	.11 (.21)	$t(58) = 1.40, p = .17,$ $d = 0.37$
		Intentional bias	.18 (.19)	
	Advisor B (5 accurate)	Error	.19 (.26)	$t(58) = 1.43, p = .16,$ $d = 0.38$
		Intentional bias	.29 (.28)	
	Advisor C (4 accurate)	Error	.23 (.25)	$t(49) = 1.54, p = .13,$ $d = 0.44$
		Intentional bias	.34 (.27)	
	Advisor D (3 accurate)	Error	.28 (.33)	$t < 1$
		Intentional bias	.33 (.30)	
Active distrust	Advisor A (6 accurate)	Error	.00 (.00)	N/A
		Intentional bias	.00 (.19)	
	Advisor B (5 accurate)	Error	.006 (.03)	$t(58) = 1.65, p = .10,$ $d = 0.47$
		Intentional bias	.03 (.09)	
	Advisor C (4 accurate)	Error	.00 (.00)	$t(49) = 2.83, p = .007,$ $d = 0.81$
		Intentional bias	.04 (.08)	
	Advisor D (3 accurate)	Error	.01 (.04)	$t(48) = 2.08, p = .04,$ $d = 0.81$
		Intentional bias	.05 (.10)	

Confidence. We conducted univariate ANOVAs on confidence levels associated with the final estimates and on the change in confidence between initial and final estimates. We observed a significant effect of advisor quality on both measures, $F_s(3,214) \geq 12.99$, $p_s < .001$, $\eta^2 \geq 0.15$, but no significant main effect of reason for suspicion, $F_s < 1$ or interaction effects, $F_s \leq 1.63$, $p_s \geq .18$, $\eta^2 \leq .02$.

Post-task attitudes. A confirmatory factor analysis produced two Promax-rotated factors, corresponding to the five post-task attitudes, which account for 80.48% of the variance. We grouped the items into these two variables: negative attitudes and positive attitudes. See factor loadings in Table S2. A multivariate analysis of variance (MANOVA) on the two reaction types with reason for suspicion and advisor quality as between-subjects factors revealed main effects of reason for suspicion, $F(2,213) = 12.66$, $p < .001$, $\eta^2 = .11$, and advisor, $F(6,428) = 17.72$, $p < .001$, $\eta^2 = .20$ and a significant interaction, $F(6,428) = 2.59$, $p = .02$, $\eta^2 = .04$. Descriptive statistics are reported in Table S3.

Table S2.

Factor loadings and reliability scores of post-task attitudes in Experiment 2.

	Factor Loadings	
	Factor 1	Factor 2
Negative (alpha = .68)		
The extent to which:		
The advisor wanted to hurt you	.864	
You were harmed by the advisor	.876	
Positive (alpha = .90)		
The extent to which:		
The advisor wanted to help you		.875
You trusted the advisor		.934
You were willing to use the advisor's advice		.920

Table S3.

Negative and positive post-task attitudes by advisor quality and reason for suspicion in Experiment 2. Standard deviations are in parentheses.

	Negative attitudes		Positive attitudes	
	Error	Intentional bias	Error	Intentional bias
Advisor A (6 accurate)	1.88 (1.23)	2.20 (0.97)	7.51 (1.75)	7.53 (1.23)
Advisor B (5 accurate)	2.74 (1.52)	3.18 (1.40)	6.92 (1.27)	5.91 (1.89)
Advisor C (4 accurate)	2.89 (1.72)	4.54 (1.61)	6.42 (1.69)	4.45 (1.45)
Advisor D (3 accurate)	3.07 (1.77)	3.34 (1.51)	4.84 (1.43)	4.20 (1.96)

Willingness to pay for advice and its use. In the article, we report measures of participants' willingness to pay for advice from each advisor and their use of advice after receiving it. We found the price participants paid for advice predicted their use of it, such that the more participants paid for the advice, the greater weight they ascribed to it in their final estimates, $B = .381, p < .001$. The result is consistent with earlier findings (Gino, 2008) and was true in cases of both suspected error, $B = .389, p < .001$, and intentional bias, $B = .371, p < .001$.

Experiment 4

In addition to the measures reported in the manuscript, we collected the following measures:

1. Post-task attitudes toward the advisor. Participants rated on a 1-5 scale:
 - a. The degree to which they used the advice they received in round 2.
 - b. The extent to which the advisor's previous errors [lying] influenced their use of the advice in round 2.
 - c. How angry they were at the advisor.
 - d. Their estimates of the advisor's capability.
 - e. Their estimates of the advisor's morality.

Results

A series of 2 (reason for suspicion: error vs. intentional bias) \times 2 (direction of inaccuracy: random vs. non-random) ANOVAS found significant effects of the reason for suspicion on participants' subjective ratings of the advisor and of their own emotion and behavior. Receiving intentionally biased advice in round 1 was associated with providing lower ratings of the advisor's capabilities, lower ratings of the advisor's morality, greater reported influence of the advisor's behavior, lower reported use of the advice in round 2, and higher levels of anger toward the advisor, $F_s(1,196) \geq 4.22, p \leq .04, \eta^2 \geq .02$. The inaccuracy direction manipulation had no effect on any of these measures, $F_s \leq 2.29, p_s > .13, \eta^2 \leq .01$. Note that although all of these measures were affected by the suspicion manipulation, this pattern in participants' subjective reports was not consistent with their actual use of advice. The means and standard deviations of all items and results of the analyses are presented in Table S4.

Table S4.

Descriptive statistics and results of factorial ANOVAs of subjective advisor evaluations in Experiment 4.

Measure	Mean (SD) Random		Mean (SD) Bias		Main effect: suspicion	Main effect: direction	Interaction
	Error	Intentional bias	Error	Intentional bias			
Use of advice	3.08 (1.07)	2.76 (0.77)	3.04 (0.82)	2.77 (0.96)	*	ns	ns
Effect of prior error/bias	2.60 (1.09)	3.24 (0.89)	2.85 (1.15)	3.44 (1.13)	***	ns	ns
Anger toward the advisor	1.14 (0.41)	1.68 (1.06)	1.21 (0.58)	1.67 (0.88)	***	ns	ns
Advisor's capability	3.34 (0.82)	3.02 (0.65)	3.19 (0.61)	3.10 (0.72)	*	ns	ns
Advisor's morality	3.38 (0.64)	2.74 (0.83)	3.19 (0.49)	2.63 (0.77)	***	ns	ns

Note: ns $p \geq .1$, $^\dagger p < .1$, * $p < .05$, ** $p < .01$, *** $p < .001$.

Experiment 5

In addition to the measures reported in the manuscript, we collected the following measures:

1. In each of the six estimation tasks:
 - a. Confidence that the initial estimate is accurate (on a 0%-100% scale).
 - b. Confidence that the final estimate is accurate (on a 0%-100% scale).
2. Post-task attitudes toward the advisor. Participants rated on a 1-9 scale:
 - a. The extent to which participants were harmed by advisor(s).
 - b. The extent to which participants believed the advisor(s) meant to hurt them.
 - c. Participants' willingness to use the advice.
 - d. Participants' trust in the advisor(s).
 - e. The degree to which participants felt the advisor(s) meant to help them.

Results

Confidence. We conducted 2 (reason for suspicion: error vs. intentional bias) \times 2 (advice setting: one advisor vs. randomly-alternating advisors) univariate ANOVAs on participants' average levels of confidence in the accuracy of their initial and final estimates. The analyses found no significant effects on either measure, $F_s(1, 201) < 1.16$, $p_s > .28$, $\eta^2 < .01$.

Post-task attitudes. An exploratory factor analysis on post-task attitudes produced two Promax-rotated factors, representing 72.96% of the variance. We grouped the items into two factors. See Table S5.

Table S5.

Factor loadings and reliability scores of post-task attitude items in Experiment 5.

	Factor Loadings	
	Factor 1	Factor 2
<i>Emotional reaction</i> (alpha = .59)		
Harm suffered by the advisor	.841	
The advisor wanted to hurt you	.841	
<i>Evaluation of the advisor</i> (alpha = .82)		
Willingness to use the advice		.860
Trust in the advisor		.913
The advisor wanted to help you		.805

A multivariate ANOVA on these two factors revealed marginally significant main effects of advice setting on both emotional reaction, $F(1, 201) = 2.92, p = .09, \eta^2 = .01$, and advisor evaluation, $F(1, 201) = 3.71, p = .06, \eta^2 = .02$. Participants perceived the harm done to them by the two randomly-alternating advisors as lower and evaluated them more favorably than those who received all of the advice from one advisor. The reason for suspicion did not significantly affect any measure (emotional reaction: $F(1, 201) = 2.65, p = .10, \eta^2 = .01$; evaluation: $F < 1$), and no significant interactions were observed, $F_s \leq 1.16, p_s \geq .28, \eta^2 \leq .01$. Regression analyses show that, consistent with findings of Experiment 2, advisor evaluation was a significant predictor of advice use, $B = .296, p < .001$, but emotional reaction was not, $B = .009, p = .90$. Controlling for these attitude measures, however, did not attenuate the observed effects of our manipulations on use of advice.

Experiment S1

We conducted an additional study, which was not included in the main manuscript. This experiment tested the replicability of the results of Experiment 2, which found that the type of suspicion affects advice use, but that this effect was not consistent across all levels of advice quality. When quality was low, such that the advice was as likely to be inaccurate as accurate, the effect of the reason to suspect it weakened considerably. This could be due to a floor effect or an artefact of the experimental setting in Experiment 2. To test the robustness of this pattern, we conducted another experiment, which varied both advice quality and the reason for suspicion and measured advice use.

Method. One hundred seventy-three undergraduate management students at an Israeli university ($M_{age} = 24.25$; 133 females) participated in a lab study. For their participation, they received course credit and a chance to win a performance-based monetary bonus. We determined our sample size based on a power analysis of a medium-size effect ($f = 0.25$; $d = 0.5$) with 90% power, which suggested a sample of 171 participants.

Participants completed 12 estimation tasks, divided into two rounds. The tasks were in the same format as the ones used in Experiment 2. In each task, participants saw a series of nine numbers that appeared on the screen for two seconds and provided an initial estimate of their average. Next, they

received advice from an advisor, who had participated in a pilot study for the experiment and saw the numbers for 10 seconds, and provided a final estimate. The order of the tasks was randomized within each round.

We employed a 2 (round 1 advice quality: high vs. low) \times 2 (reason for suspicion: error vs. intentional bias) between-subjects design. We used the advice quality levels which yielded different effects in Experiment 2. In round 1, the high-quality group received accurate advice (i.e., within 10 units of the correct answer) for four estimates, and inaccurate advice that missed the correct answer by more than 20 units for two estimates. The low-quality group received accurate advice for three of the six estimates. Participants learned of these quality levels after the first round, when they received feedback about the advice. We also used the feedback to manipulate the reason for suspicion: in the error group, the feedback referred to the inaccurate advice in round 1 as “erroneous”, whereas in the intentional bias condition, the inaccurate advice was described as “untruthful”.

Round 2 also consisted of six estimation tasks. Participants learned they would receive advice from the same advisor as before. In this round, the advice was the same for participants in all conditions. The participant whose final estimates in this round were the most accurate (i.e., had the lowest average error) won a bonus of 40 shekels (~\$11.50). After completing the tasks, participants estimated the quality of the advice in round 2 by estimating its average error and the number of estimates for which they received inaccurate advice.

We also took several additional measures. In each of the twelve estimation task, we elicited participants’ confidence in the accuracy of their initial and final estimates, on a 0%-100% scale. After the task, we elicited participants’ attitudes toward the advisor. On a 1-9 scale, participants rated:

- a. How angry participants felt toward the advisor.
- b. The extent to which participants felt the advisor meant to harm them.
- c. The degree to which participants thought they were harmed by the advisor.
- d. The degree to which participants wanted to harm the advisor.
- e. The degree to which participants trusted the advisor.

- f. The extent to which participants thought the advisor meant to help them.
- g. The degree to which participants were willing to use the advisor's advice after receiving the information about the advisor's prior performance.
- h. How good participants thought the advisor was at estimating.
- i. How good participants thought the advisor was at giving advice.

Results. We used the same inclusion criteria as in Experiment 2. We removed individual invalid estimates, and included participants who provided at least three valid estimates in the analyses. Save for one participant, who did not provide any valid estimates, all participants met the inclusion criterion, resulting in a final sample of 172 participants.

Advice Use. Average advice use in round 1 was 0.41 ($SD = 0.20$) in the high-quality condition and 0.39 ($SD = 0.21$) in the low-quality condition, with no significant differences between groups, $ts < 1$. In round 2, average advice use in the high-quality condition declined to 0.33 ($SD = 0.22$). We conducted an ANOVA on advice use in round 2, using round 1 advice use as a covariate, among participants who received high-quality advice. The analysis found a significant effect of the reason for suspicion, $F(1,87) = 5.05, p = .03, \eta^2 = .06$. As Figure S1 shows, these participants used advice less after learning that the advisor's prior inaccuracies were intentional than when these inaccuracies were due to error. In the low-quality condition, average advice use in round 2 was 0.30 ($SD = 0.30$). Consistent with Experiment 2, the reason for suspicion did not affect subsequent advice use in this condition, $F < 1$. An alternative analysis measured advice use in round 1 and round 2 as repeated measures. The analysis yielded a marginally-significant Round \times Reason for Suspicion interaction in the high quality condition $F(1, 88) = 3.36, p = .07, \eta^2 = .04$ and no significant interaction in the low quality condition, $F < 1$.

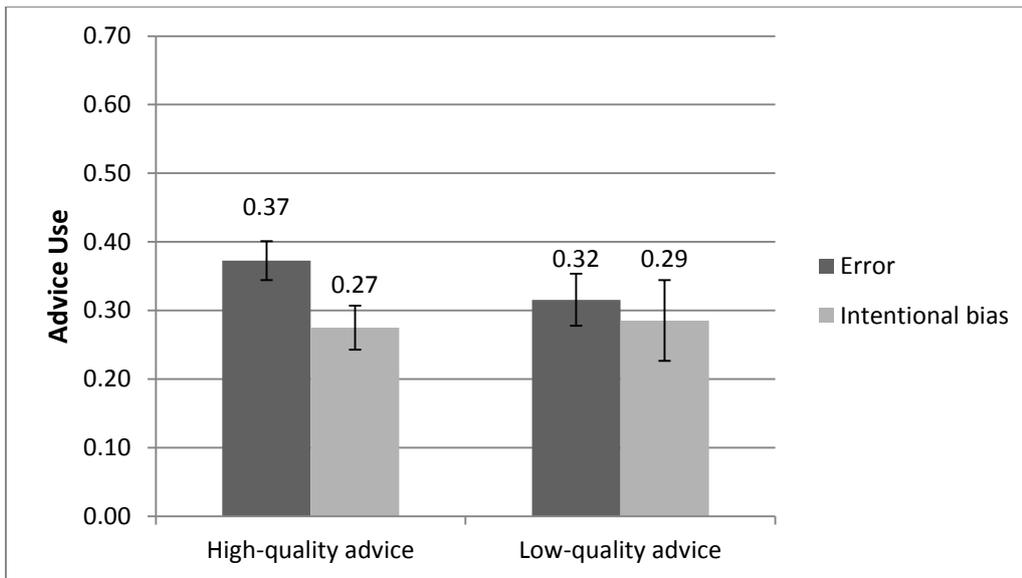


Figure S1. Average rates of advice use by prior advice quality and reason for suspicion in Experiment 3. Error bars represent ± 1 SEM.

Post-task Advice Quality Evaluation. To measure the perceived quality of the advice, we elicited participants' estimates of (a) the number of times they received inaccurate advice, and (b) average advice error. We tested the effect of suspicion on these two measures in both the high-quality and low-quality conditions using independent samples t-tests. As we found in the previous two experiments, suspecting intentional bias did not reduce estimates of advice quality, relative to suspecting error. Participants who had received intentionally biased advice in round 1 did not estimate a higher frequency of inaccurate advice in round 2 (high quality: $M = 1.70$, $SD = 1.24$; low quality: $M = 2.33$, $SD = 1.29$) than did those whose advice in round 1 was prone to error (high quality: $M = 1.85$, $SD = 1.25$; low quality: $M = 2.87$, $SD = 1.20$), low quality: $t(80) = 1.94$, $p = .06$, $d = 0.43$; high quality: $t < 1$. No differences were observed in participants' estimates of the average error in the advice they received in round 2, both after receiving low quality advice (error: $M = 20.03$, $SD = 43.84$; intentional bias: $M = 23.57$, $SD = 26.23$) and after receiving high quality advice (error: $M = 15.13$, $SD = 10.89$; intentional bias: $M = 16.51$, $SD = 23.22$), $ts < 1$. Participants who suspected intentional bias did not think the advice they received was inferior to that received by those who suspected error, but nevertheless gave the advice less weight in their judgments.

Confidence. We conducted univariate ANOVAs on the confidence participants expressed in the accuracy of their initial and final estimates. We found no significant main effects, $F_s < 1$, or significant interactions, $F(1,168) \leq 2.42$, $p_s > .12$, $\eta^2 \leq .01$.

Attitudes toward the advisor. An exploratory factor analysis produced two Promax-rotated factors, corresponding to the nine post-task attitudes, which account for 61.13% of the variance. We grouped the items into these two variables: negative attitudes and positive attitudes. See factor loadings in Table S6. A two-way MANOVA, with prior advice use as a covariate, reveals main effects of reason for suspicion and round 1 advice quality on participants' negative attitudes (reason for suspicion: $F(1,167) = 12.80$, $p < .001$, $\eta^2 = .07$; advice quality: $F(1,167) = 5.02$, $p = .03$, $\eta^2 = .03$) as well as on their positive attitudes (reason for suspicion: $F(1,167) = 8.24$, $p = .005$, $\eta^2 = .05$; advice quality: $F(1,167) = 12.17$, $p = .001$, $\eta^2 = .07$). Regression analyses reveal positive attitudes were a significant predictor of advice use, $B = .218$, $p = .005$, but negative attitudes were not, $B = .042$, $p = .58$.

Table S6.

Factor loadings and reliability scores of post-task attitudes in Experiment 2.

	Factor Loadings	
	Factor 1	Factor 2
Negative (alpha = .73)		
The extent to which:		
You felt angry at the advisor	.790	
The advisor meant to harm you	.747	
You were harmed by the advisor	.670	
You want to harm the advisor	.760	
Positive (alpha = .85)		
The extent to which:		
The advisor wanted to help you		.678
You trusted the advisor		.863
You were willing to use the advisor's advice		.779
The advisor was good at estimating		.806
The advisor was good at giving advice		.824

Discussion

Experiment S1 replicates the findings of Experiment 2, whereby suspecting intentional bias leads to greater discounting of advice than does suspecting unintentional error. As in Experiment 2, the effect

occurred when the advice was generally expected to be accurate, but when advice quality was low (i.e., not more likely to be accurate than inaccurate), it received equally low weight in both conditions.

The lower weight ascribed to advice in the intentional bias condition was not reflected in participants' assessments of its quality, which participants evaluated similarly regardless of the reason for their suspicion. This result is consistent with Experiment 2, which found similar levels of willingness to pay for advice, which implies that advice seekers are not willing to pay a premium for honesty, if it does not objectively increase advice quality. Together, these findings demonstrate that the detrimental effect of suspicion on advice use has to do with how the received information is processed, rather than with overt perceptions of the quality of the advice.