Online Appendix for

Enabling or Limiting Cognitive Flexibility? Evidence of Demand for Moral Commitment

Silvia Saccardo and Marta Serra-Garcia

September 2022, updated September 2023

Table of Contents

A Theoretical Model: Proofs, Additional Results and Example

A.1 Main Model: Effects of Information Order

We discuss the strategies and equilibria when the advisor is exogenously assigned to see quality first ($f = q$) and when the advisor is exogenously assigned to see the incentive first ($f = i$). Let us denote the game as G^f.

Proposition A.1. *If $m_0 > \iota$, the PBE of G^f, with $f \in \{i, q\}$, is characterized by Self 0 not suppressing the signal of quality. If $m_0 \leq \iota$, in any PBE of G^f,*

$$p_s^* = \min\{\frac{(1-\lambda^f)(\iota - \phi M)}{\lambda^f \phi(M-\iota)}, 1\}$$

Hence, suppression is more likely when the advisor sees the incentive first.

Proof. The expected utility of Self 0 given p_s is:

$$E(U_0) = \lambda^f \Big((1-\phi)\iota + \phi((1-p_s)((\iota - m_0)\frac{\iota}{M}) + p_s(\iota - m_0)q)\Big) + (1-\lambda^f)(\iota - \phi m_0)q.$$

Self 1 recommends the incentivized option if she receives $\hat{\sigma} = \varnothing$ with certainty if $\frac{\iota}{Mr(\varnothing)} \geq 1$. Self 1 is uncertain about whether the signal of quality is empty because of Self 0 suppressed it or because it was not encoded to begin with. Using Bayes' Rule,

$$\frac{(p_s\lambda^f + (1-\lambda^f))\phi}{\lambda^f p_s \phi + (1-\lambda^f)} \leq \frac{\iota}{M}$$

which implies

$$p_s \leq \frac{(1-\lambda^f)(\iota - \phi M)}{\lambda^f \phi(M-\iota)}$$

Hence, a selfish Self 0 prefers to suppress as often as possible and hence chooses,

$$p_s^* = \min\{\frac{(1-\lambda^f)(\iota - \phi M)}{\lambda^f \phi(M-\iota)}, 1\}$$

Since p_s^* is decreasing with λ^f,

$$\frac{\partial p_s^*}{\partial \lambda^f} = \frac{\iota - \phi M}{\phi(M-\iota)} \frac{-\lambda^f - (1-\lambda^f)}{(\lambda^f)^2} < 0,$$

it follows that suppression is more likely when the incentive information is shown first, than when the signal of quality is shown first.

Ex-ante, if Self 0 is selfish, her expected payoff in G^f is:

$$U_0(G^f) = \lambda^f \Big((1-\phi)\iota + \phi((1-p_s^*)((\iota - m_0)\frac{\iota}{M}) + p_s^*(\iota - m_0))\Big) + (1-\lambda^f)(\iota - \phi m_0).$$

If Self 0 is moral, her expected payoff in G^f is:

$$U_0(G^f) = \lambda^f \Big((1-\phi)\iota + \phi((\iota - m_0)\frac{\iota}{M}) \Big) + (1-\lambda^f)(\iota - \phi m_0).$$

We also consider the case in which the advisor is naive, in the sense that she believes that her attention is imperfect under both information orders.

Proposition A.2. *If the advisor believes she encodes quality signals with the same probability for both information orders, she exhibits the same suppression under both information orders. Since encoding is actually less likely when the information is shown first, she still recommends the incentivized product more often when the incentive is shown first.*

Proof. The advisor's belief is $\hat{\lambda}^q = \hat{\lambda}^i = \lambda^i < 1$. Since the advisor's belief is correct when the incentive is shown first, the same prediction holds as for Proposition A.1. When the signal of quality is shown first, Self 0 and Self 1 both believe that there is a probability $1 - \hat{\lambda}^q$ that the signal is not encoded to begin with. Hence, Self 0 suppresses signals that are in conflict with the incentive with probability p_s^* believing the likelihood of encoding a signal is λ^i. Self 1 updates using the same belief about attention. Therefore, the same behavior as in Proposition A.1. would arise. Since in actuality the signal would be encoded less often when the incentive is shown first, there would still be more suppression in that case.

A.2 Main Model: Preferences for information order

Below we provide the proof for Proposition 2 (from the main text).

Proposition 2.

- *If Self 0 is selfish ($m_0 \leq \iota$), she chooses to see the incentive first ($f^* = i$). This order increases the likelihood that Self 1 recommends the incentivized product when the signal is in conflict with the incentive.*
- *If Self 0 is moral ($m_0 > \iota$), she chooses to see quality first ($f^* = q$), which decreases the likelihood that Self 1 recommends the incentivized product when the signal is in conflict with the incentive.*

Proof. If $m_0 > \iota$, Self 0's utility increases as the likelihood that Self 1 recommends the incentivized product when it is in conflict with the incentive decreases. By choosing $f = q$, the likelihood that the incentivized product is recommended is lowered to $\frac{\iota}{M}$. If

$m_0 \leq \iota$, Self 0's utility increases as the likelihood that Self 1 recommends the incentivized product when it is in conflict with the incentive increases. For any p_s^*, the likelihood is higher when $f = i$ because $\lambda^i < \lambda^q$.

The proof of Proposition 3 follows directly from the fact that the advisor does not believe that information order differentially affects her likelihood of encoding a quality signal. Given that she does not anticipate a difference in behavior, she is not willing to pay for any information order.

In the main text advisors are assumed to receive their preferred order. In the experiment, however, there is a 25% chance that the advisor's preference is not implemented. In that case, a selfish Self 0 would be assigned to see quality first, $f = q$, while a moral Self 0 would see the incentive first, $f = i$. Since a selfish Self 0 would suppress with p_s^*, given λ^q, Self 1 would receive a signal in conflict with the incentive with probability $\lambda^q(1 - p_s^*)$. By contrast, since a moral Self 0 does not suppress, Self 1 would receive a signal in conflict with the incentive with likelihood λ^i. This implies that whether Self 0s who desire but are not assigned to see the incentive first recommend the incentivized option when there is a conflict of interest more or less often than Self 0s who desire but are not assigned to see quality first is ambiguous and depends on the exact value of λ^i and λ^q.

Numerical Example for NoChoice and Choice. In what follows we use a simple numerical example to illustrate the differences between the predicted behavior of Self 0 and Self 1, when the incentive information is shown first ($f = i$) and when the quality signal is shown first ($f = q$). Consider the case where advisors who see the incentive first encode the quality signal 50% of the time ($\lambda^i = 0.5$), while advisors who see the quality signal first encoded it 70% of the time ($\lambda^q = 0.7$). The incentive ι is 0.15, while the prior likelihood that the signal is in conflict with the incentive is $\phi = 0.5$. The range of moral costs is from 0 to $M = 0.25$.

Given these parameter values, and applying the formula for p_s^* in Proposition A.1., if the incentive is shown first, the optimal likelihood of suppression is

$$p_s^* = \frac{(1 - 0.5)(0.15 - 0.5 \cdot 0.25)}{0.5 \cdot 0.5 \cdot (0.25 - 0.15)} = 0.5$$

Similarly, if the quality signal is shown first, the optimal likelihood of suppression is 0.21. Both if the incentive is shown first and if the signal of quality is shown first, the posterior belief of Self 1 after receiving an empty signal ($\hat{\sigma} = \varnothing$) is 0.5. What differs between both information orders is how often an empty signal is received.

Given Self 0's suppression strategy, what signal distribution does Self 1 receive? We consider a Self 0 who is selfish and suppresses with the likelihoods shown above. We

calculate (a) how often Self 1 receives a signal that is in conflict with the incentive, given that the signal was of conflict, and (b) how often Self 1 receives a signal that is not in conflict with the incentive, given that there is no conflict.

Consider the case where the signal of quality is shown first. Then a signal is encoded with a likelihood of 0.7, and conditional on it being of conflict, it is suppressed with a likelihood of 0.21. Therefore, if there is a conflict with the incentive, Self 1 receives the signal with a 0.55 likelihood $(0.7 \cdot (1 - 0.21))$. If there is no conflict, Self 1 receives the signal with a 0.7 likelihood.

Consider the case where the incentive is shown first. Then a signal is encoded with a likelihood of 0.5, and conditional on it being of conflict, it is suppressed with a likelihood of 0.5. If there is a conflict with the incentive, Self 1 receives the signal with a 0.25 likelihood $(0.5 \cdot (1 - 0.5))$. If there is no conflict, Self 1 receives the signal with a 0.5 likelihood.

Therefore, seeing the incentive first has a (large) first-order effect: it decreases the likelihood that Self 1 learns the actual signal, both if it is in conflict or not in conflict with the incentive. There is a second effect: seeing the incentive first increases the difference in the likelihood that a signal of conflict is received relative to a signal that is not in conflict.

Recommendations would reflect these differences. Consider the case where Self 0s feature low moral costs ($m_0 < i$) and are classified as selfish in 50% of the cases, while Self 0s would exhibit high moral costs in the remaining 50% of the cases. Selfish Self 0s would prefer to see the incentive first, while moral Self 0s would prefer to see quality first. Given their suppression strategies, how often would Self 1 recommend the incentivized product when the signal is in conflict with the incentive? Given the independent draw of moral costs for Self 0 and Self 1, each Self 0 would have a 50% chance to be matched with a selfish Self 1. Consider a selfish Self 0 who chooses to see the incentive first. Then, Self 1 sees a signal in conflict with the incentive with a 0.25 chance, and receives an empty signal with a 0.75. Since Self 1 is selfish with a 0.5 chance, Self 1 would recommend the incentivized product with a 0.875 chance $(0.5 + 0.5 \cdot 0.5 + 0.5 \cdot 0.5 \cdot 0.25)$. If, by contrast, Self 0 would choose to see the signal of quality first, and suppress optimally, Self 1 would recommend the incentivized product with a 0.7235 chance $(0.3 + 0.7 \cdot 0.21 + 0.7 \cdot 0.79 \cdot 0.5)$.

Consider by contrast a moral Self 0 who chooses to see quality first. Since she prefers not to suppress, Self 1 would recommend the incentivized product with a 0.65 chance $(0.3 + 0.7 \cdot 0.5)$. If, by contrast, Self 0 would choose to see the incentive first, due to the lower attention, Self 1 would recommend the incentivized product with a 0.75 chance $(0.5 + 0.5 \cdot 0.5)$.

If advisors are sophisticated and can choose their preferred order (Choice Experiment),

this numerical example of the model would predict recommendations of the incentivized product to occur in 87.5% of the cases, among those who choose to see the incentive first, and 65% of the cases, among those who choose to see quality first. This 22.5% percentage point gap is similar, though slightly larger, than the gap we observe in the Choice experiment (19.5% percentage points). Comparing cases in which the advisors receive their preferred information order and cases in which they do not, conditional on preference, this numerical example would predict a 10 to 15 percentage point gap in recommendations, which is close to the 10 to 12 percentage point gap observed in the Choice experiment.

If advisors cannot choose their preferred order, those assigned to see the incentive first would be selfish Self 0s in 50% of the cases, and moral Self 0s in the remaining cases. Combining the behavior of these two types, the numerical example would predict that 81.25% of advisors recommend the incentivized product when the signal is in conflict with the incentive. Among those assigned to see quality first, the numerical example would predict that 68.7% of advisors would recommend the incentivized product when the signal is in conflict with the incentive. This predicted gap of 12.55 percentage points is qualitatively similar, but somewhat smaller than the gap of 17 percentage points observed in NoChoice.

A.3 Incentives and preferences for information order

In the Choice Stakes experiment, we vary the advisor's incentive ι, while keeping it costly to see the incentive first. Instead of having $\iota = 0.15$, the Low Commission treatment has $\iota = 0.01$ and the High Commission treatment is $\iota = 0.30$. We examine how increasing or decreasing the incentive affects the utility of seeing the incentive first, assuming advisors are sophisticated.

Corollary A.1. *If the advisor's incentive to recommend the incentivized option decreases, the demand to see the incentive first decreases. If the advisor's incentive to recommend the incentivized option increases, the demand to see the incentive first may increase or decrease.*

Proof. When ι decreases, as in the Low Commission treatment, it is more likely that $m_0 > \iota$, and the advisor is more likely to prefer to see quality first ($f = q$). Further, for advisors who still prefer to see the incentive first, the likelihood of suppression decreases when the incentive decreases. Specifically, since $\iota > \phi M$, then p_s^* decreases with ι since:

$$\frac{\partial p_s^*}{\partial \iota} = \frac{1 - \lambda^f}{\lambda^f \phi} \frac{(M - \iota) + (\iota - \phi M)}{(M - \iota)^2} > 0.$$

By contrast, if the incentive increases, we have that p_s^* increases. This can increase the utility from seeing the incentive first, as long as $\iota < M$. However, as the incentive becomes higher, it can become higher than the highest moral cost $\iota > M$. Then, the potential for conflict between Self 0 and Self 1 disappears, and Self 0 no longer strictly prefers to see the incentive first.

B Detailed Experimental Design and Procedures

B.1 The Experiments

Table B.1 reports all the data collected for this paper, their corresponding pre-registration, recruitment platform and incentives for advisors and clients. We pre-registered the design, sample sizes, exclusion criteria, and analyses of all Amazon Mechanical Turk (AMT) experiments on aspredicted.org. The experiment on professionals was not pre-registered. The design of NoChoice experiment and the Choice Stakes Experiment are described in full in the main text.

NoChoice Experiment. The NoChoice experiment was conducted on AMT. All details about the experiment are reported in the main text.

Choice Experiment. As displayed in Table B.1, the Choice experiment on Amazon Mechanical Turk (AMT) was conducted in three different waves (AMT-1, AMT-2, AMT-3). The first wave of the experiment, AMT-1, was conducted in 2019 and randomized participants in the *ChoiceFree* and *See Incentive First treatment*. The second wave of the experiment, AMT-2, was conducted in 2020, and collected additional data for the *ChoiceFree* treatment. As part of this wave of data collection, we randomized whether incentives were identical to those in the first wave of the experiment or probabilistic as in the professional sample. In the treatments with probabilistic incentives, the products were urns containing payoff balls that were worth either \$0 or \$20 (rather than \$0 and \$2 as in the regular treatments), and the commission for the advisor was \$15 to 1 out of 100 participants (instead of paying \$0.15 as in the regular treatments). One out of 100 advisors was selected for payment, and their recommendations were sent to a client. The experiment also counterbalanced whether the incentivized product was presented on the left side or the right side of the screen. These two factors varied independently between subjects (2x2 design). In 2020 we also collected data for the sample of professionals in the ChoiceFree treatment (ChoiceFree-Professionals). The third wave of the experiment, AMT-3, was conducted in 2021. The majority of the participants recruited for this wave (80%) were randomized into the *ChoiceFree* treatment, the *See Incentive First* treatment and the *See Quality First*. Since we already had data on the latter treatments from prior waves, participants were randomized to those treatments at a 1:1:5 ratio. The remaining 20% of participants was randomized (at a 2:1 ratio) into one of two robustness treatments that increased the incentives for both the advisor and the client incentives in the ChoiceFree treatment. The goal of these robustness treatments was to test whether the effects of information order on recommendations documented in the *ChoiceFree* treatment are specific to the small stakes used in the experiment, or whether

Table B.1: The Experiments

Sample-Wave	Aspredicted pre-reg #	Treatment	Advisor/DM Commission	Client's Payoff Balls	Matching with Client	Year	N
Main Text: NoChoice Experiment							
AMT	22709	See Incentive First	\$0.15	\$0 or \$2	1 advisor out of 10	2019	152
		See Quality First	\$0.15	\$0 or \$2	1 advisor out of 10	2019	147
Main Text: Choice Experiment							
AMT-1	23272	Choice Free	\$0.15	\$0 or \$2	1 advisor out of 10	2019	1308
		Incentive First Costly	\$0.15	\$0 or \$2	1 advisor out of 10	2019	1347
AMT-2	42246	Choice Free	\$0.15	\$0 or \$2	1 advisor out of 10	2020	511
		Choice Free	\$15 to 1/100	\$0 or \$20	1 to 1*	2020	542
Professionals	NA	ChoiceFree Professionals	\$15 to 1/100	\$0 or \$20	1 to 1*	2020	712
AMT-3	70817	Choice Free	\$0.15	\$0 or \$2	1 advisor out of 10	2021	213
		Quality First Costly	\$0.15	\$0 or \$2	1 advisor out of 10	2021	1067
		Incentive First Costly	\$0.15	\$0 or \$2	1 advisor out of 10	2021	215
		ChoiceFree Highx10	\$1.5	\$0 or \$20	1 advisor out of 10	2021	275
		ChoiceFree Highx100	\$15	\$0 or \$20	1 to 1	2021	110
Main Text: Choice Stakes Experiment							
AMT	76771	Low Incentive	\$0.01	\$0 or \$2	1 advisor out of 10	2021	483
		Intermediate Incentive	\$0.15	\$0 or \$2	1 advisor out of 10	2021	511
		High Incentive	\$0.30	\$0 or \$2	1 advisor out of 10	2021	478
Main Text: Information Architect Experiment							
AMT	76771	IA-Advisor	\$0.15	\$0 or \$2	1 advisor out of 10	2021	245
		IA-Client	\$0.15	\$0 or \$2	1 advisor out of 10	2021	253
Appendix: Choice Deterministic							
AMT	82298	ChoiceFree - Replication	\$0.15	\$0 or \$2	1 advisor out of 10	2021	385
AMT		ChoiceFree - Deterministic	\$0.15	\$0 or \$2	1 advisor out of 10	2021	369
Appendix (Additional Exp): NoChoice Simultaneous							
AMT	79521**	See Incentive First	\$0.15	\$0 or \$2	1 advisor out of 10	2021	70
AMT		See Quality First	\$0.15	\$0 or \$2	1 advisor out of 10	2021	78
AMT		Simultaneous	\$0.15	\$0 or \$2	1 advisor out of 10	2021	128
Appendix (Additional Exp): Predictions							
AMT	37081	Prediction	-	\$0 or \$2		2020	288

Notes. This table presents all the experiments we conducted for this paper, their corresponding sample, wave of data collection, pre-registration, treatments, incentive features for advisors and clients in the experiment, matching between advisors and clients, and sample sizes after excluding inattentive participants and participants with inconsistent responses in the MPL measure of moral costs, as pre-registered.

* In these studies, only 1 out of 100 advisors were selected for payment. These advisors were all matched with a client.

** In this study, due to higher rates of inattention than in other studies, we updated the pre-registration to increase the size of the recruited sample, see Aspredicted #82164.

they persist when advisors conflict of interests that have higher stakes. In the *High Stakes - 10 fold* treatment, we increased the incentives by a factor of 10. We paid each advisor a \$1.50 commission if she recommended the incentivized product. In this treatment, one out of 10 advisors was then matched to a client, who received either \$0 or \$20. In the *High Stakes - 100 fold* treatment, we increased the incentives by a factor of 100, increasing the commission of the advisor to \$15, and matched each advisor with a client, who received either \$0 or \$20.

Choice Stakes Experiment. In the Choice Stakes experiment, all instructions were

identical to those of the ChoiceFree experiment, but we varied the size of the commission while keeping the incentives for the clients the same. All details about this experiment are reported in the main text.

Information Architect Experiment. In the IA-Experiment, all instructions were identical to those of the ChoiceFree experiment, except that participants were assigned to the role of Information Architects (IAs). That is, participants were informed that they would be matched with an advisor and a client and that they will have to make a decision about how the advisor receive information. The Instructions for this experiment are reported in Online Appendix G.2. IAs received information about their incentive and were asked to choose an information order for the advisor they were matched with. Importantly, the IA did not receive information about the product that yielded the advisors a commission nor the signal of quality directly, but only determined the order with which advisors receive such information. We subsequently recruited 498 advisors, and presented them with the information order chosen by the IA. We informed these advisors that the order of information was chosen by the IA. Advisors were not informed about IAs' incentives.

B.2 Sample, Recruitment Procedures and Exclusion Criteria

Sample. All experiments were conducted via the CloudResearch platform (Litman, Robinson, and Abberbock, 2016), which we used to recruit high quality subjects from Amazon Mechanical Turk (AMT), except for the ChoiceFree Professionals treatment, which was conducted on both CloudResearch and Prolific Academic (Palan and Schitter, 2018) targeting participants who self-report to work in two industries in which advice is very frequent: finance and insurance, and legal services. Prolific has their own sample of participants, and we recruited as many professionals as possible within the UK, the US, and Canada. CloudResearch draws professionals from AMT, and again we recruited as many professionals based in the US as possible.

Out of 712 professionals, 677 (95.1%) provided job descriptions that could be used by our independent raters to judge whether their position was fiduciary or not. Two independent raters were asked to classify each job title as fiduciary or not fiduciary, based on the description provided by the participant. They were provided the following information regarding what is defined as fiduciary: *"According to Investopedia, a fiduciary is "a person or organization that acts on behalf of another person or persons, putting their clients' interest ahead of their own, with a duty to preserve good faith and trust. Being a fiduciary thus requires being bound both legally and ethically to act in the other's best interests. A fiduciary may be responsible for the general well-being of another (e.g. a*

child's legal guardian), but often the task involves finances; managing the assets of another person, or a group of people, for example. Money managers, financial advisors, bankers, insurance agents, accountants, executors, board members, and corporate officers all have fiduciary responsibility"." This definition is broad and may include examples of job titles that in some cases have fiduciary duties and others not (e.g., financial advisors may or may not be fiduciaries). It hence aims to broadly capture potential conflicts of interest, given the limited information provided by participants.

The raters agreed on their classification of fiduciary duty in 87% of the cases (interrater agreement κ=0.85). In 63% of the cases the job title was considered as fiduciary by at least one rater. Focusing on the cases with agreement, 58% of the job titles were considered as fiduciary. Job titles frequently found in the data included the word analyst (financial, actuarial, etc., in 9% of the cases), accountant or account manager (12% of the cases), and lawyer or paralegal (in 7% of the cases). In their job titles, 14% of participants included the word "manager." Our raters were also asked to classify the job titles into industry (finance and insurance or legal, or neither if it was not clear from the job title). Prolific provides this information for some of our participants, but it was missing in 156 of 496 cases. The agreement between raters regarding industry classification was high for CloudResearch ($\kappa = 0.80$) and somewhat lower for the missing cases on Prolific ($\kappa = 0.65$). Overall, for cases in which there is an agreement (636 out of 712), we find that 72% of professionals work in the finance and insurance industry, 19% in legal service, and for the remaining 9% the industry is unknown.

Recruitment and Procedures. We recruited participants in the role of advisors to a 5-7 minutes study on decision-making and compensated them with $0.50 for completing the study and providing a recommendation to a participant in the role of client. Professionals and participants in the third wave of the Choice experiment (AMT-3), the IA experiment and the NoChoice-Simoultaneous experiment were instead paid $1 for completing the study. Participants had to be located in the US and have an approval rating of at least 90%. All experiments were conducted online, using Qualtrics surveys. All experiments conducted in the convenience sample leveraged the CloudResearch platform, which enabled us to recruit high quality participants, limiting the number of BOTs, duplicate participants, and inattentive participants.

Upon being recruited, participants were assigned to the role of advisors and, in almost all treatments, informed that one of ten advisors would be matched with a client, as described in Table B.1. Participants were presented with several understanding questions about the products while reading the instructions. Before randomizing participants to treatments, we included one question that participants had to answer correctly in order to

continue in the study (i.e., the attention check). Those who failed to answer it correctly, were disqualified from participation and were not randomized to treatments. Advisors were then provided additional information about the experiment, and then moved to a screen where they were given the choice of information order. Advisors were informed that they would receive a commission for recommending one of two products, A or B. In the AMT-3 wave of the Choice experiment, we clarified that the commission was determined at random by the computer. In the Choice Experiments, advisors were prompted to make a choice between information order. A summary of the treatments we ran in each wave is presented in Table B.1. After receiving information about the incentive and the signal of quality, advisors were asked to provide their recommendation to the client. We then collected measures of beliefs, selfishness, and, in wave AMT-3, preferences for blinding.

B.3 Additional measures

In all experiments (except for the Prediction experiment) we collected additional measures.

Beliefs. As explained in the main text, we elicit advisors' beliefs about the likelihood that the quality of Product B was low by asking advisors i) to choose one of ten options, where Option 1 ranged between 0% and 10% and Option 10 ranged between 91% and 100%, and ii) to indicate the exact likelihood by entering a number from 0 to 100. The first measure was incentivized: in most treatments, advisors received \$0.15 for a guess in the correct range. In the ChoiceFree-Professionals and the ChoiceFree-Probabilistic treatment in the AMT-2 wave of the Choice experiment, this payment was \$15 to 1 out of 100 advisors. In the High Stakes - 100 fold treatment, this payment was \$15.

Moral costs. After the belief measure, we asked participants to complete one additional advice task, aimed at measuring advisors' selfishness/morality using a multiple price list. We informed advisors that they would be asked to make a second recommendation to an advisee –a participant different than the one who received their first recommendation. Advisors were told they would need to make a series of recommendations to another participant (an advisee), choosing between two products, X and Y. Product Y had the same payoffs of product B in the main experiment. Product X varied across 5 different decisions. It paid \$2 with probabilities 1, 0.8, 0.6, 0.4, and 0 respectively, and \$0 otherwise. Advisors were incentivized to recommend Y, with a \$0.15 commission, and received a signal of quality of Product Y that indicated that a \$0 had been drawn from Y. Given the payoffs of X, recommending Y (the incentivized product) harmed the client if X paid \$2 with a probability of 0.6 or higher. We use this elicitation to measure the advisor's

selfishness, as the number of times the advisor chose to recommend Y, and standardize it within each experiment.

In the ChoiceFree-Probabilistic treatment in the AMT-2 wave of the Choice experiment, these products were scaled up to paying either \$0 or \$20; the commission to the advisor was \$15 to 1 out of 100 participants. At the end of the experiment, we randomly selected one out of 10 advisors, we randomly picked one of the 5 recommendations, and showed them to an advisee. For this purpose, we recruited a total of 866 advisees.

Blinding. In the AMT-3 wave of the choice experiment, we measure preferences for blinding in an additional advice task. This additional task was conducted after participants took part in the main experiment, and after the additional task designed to elicit moral costs. In the task, participants learn that they are matched with another participants, a different client from the one of the main task and of the Selfishness task. We present the advisors with two products, 1 and 2, which yield the same expected payoff of products A and B in the main experiment. As in the main experiment, advisor know that, before making their recommendation they will receive a signal about the quality of Product 2. Advisors learn that they will receive a \$0.15 commission depending on their recommendation. The commission can be either for Product 1 or Product 2, determined at random, and advisors are notified that they will learn for which product the commission is before the end of the study. We then ask advisors to choose whether to learn for which product is the commission *before* receiving a signal about the quality of product and making the recommendation, or *after* learning the quality of Product 2 and making the recommendation. That is, in this task, advisors can either learn their incentive before making the recommendation or after the recommendation is made. By choosing the latter option, advisors can ensure that their recommendation is blind to incentive information. At the end of the study, we recruited $N = 188$ advisees and sent them the advisors' recommendation in this task.

Explanation of Advisors' Choices. In the second wave of data collection of the ChoiceFree experiment, we added an open-ended question asking participants to explain how they made their decision about order of information. The question was "When you had to decide between learning about your commission Before or After getting information about the quality of Product B [A, if the order was flipped], how did you make this decision?". Two independent raters, who were blind to advisors' choices, coded the responses of advisors from the AMT-2 wave of the experiment and the advisors from the sample of professionals. They classified their responses into four categories, which apply to 91% of the open-ended responses. The remaining 9% consists of empty or unrelated comments. The first category was "limiting bias" and was assigned to messages that ex-

plicitly stated that the reason for their preference was to be less biased in the evaluation and to want what is best for the client. This category was meant to capture preference for commitment to accurate beliefs and moral behavior. The second category, "does not matter," captured indifference—whether advisors stated that information order did not matter. The third category, "commission," was for advisors who indicated explicitly that they cared only about their own commission. The fourth category, "other reasons," captured whether advisors indicated that gut feeling, curiosity, or other reasons guided their preference. We did not expect advisors to openly express wanting cognitive flexibility in their comments. Consistent with this, we find no such comments in the data. We allowed coders to indicate multiple categories, though this was rarely done (in less than 3% of the cases). We analyze the relationship between these categories of motives and advisor preferences in Online Appendix C.2.4.

Demographics. We collected information on the participants' gender, age, their first language, ethnicity, and difficulty in understanding the instructions.

B.4 Exclusion Criteria

In all of the experiments, participants who failed to answer the attention check correctly were not randomized into treatments and therefore, as pre-registered, they were excluded from completing the experiment. Further, we pre-registered that we would exclude participants who provide non-monotone responses to the multiple price list used to measure selfishness. Nonmonotone participants are $N = 28$ (8.6%), in the NoChoice experiment, $N = 676$ (10.8%) across the three waves of the Choice experiment, $N = 209$ (12.43%) in Choice Stakes experiment, and $N = 51$ (9.3%) in the Information Architect experiment. For the analyses in the main text, we apply this exclusion criteria across studies with the exception of the ChoiceFree-Professionals sample for whom we did not collect this measure. In Online Appendix C.4, we repeat the main analyses including these participants. On top of this, all of our studies systematically exclude duplicate responses and participants classified as bots by Qualtrics bot detection feature.[1] Finally, since in all regressions we control for gender and age, participants with missing information for these variables are dropped from the analyses ($N = 11$ in the Choice Experiment, $N = 1$ in the Choice Stakes Experiment; we did not have missing demographics in the other experiments). In the Prediction experiment, as an extra measure of attention, we ask participants to

[1]In particular, we exclude participants with a Q_RecaptchaScore score lower than 0.5 on a scale from 0-1, which indicates a high probability that a given response comes from a bot, see https://www.qualtrics.com/support/survey-platform/survey-module/survey-checker/fraud-detection/#BotDetection.

give a reason for their predictions by writing one sentence. As pre-registered, we exclude participants who provide answers to this question that are unrelated to the experiment, as determined by a research assistant.

C Additional Analyses

C.1 NoChoice Experiment

The tables below show regression analyses for advisors in the NoChoice treatment. As pre-registered, Table C.1 focuses on advisors who gave consistent answers in the elicitation of selfishness. Table C.2 includes all advisors.

Table C.1: Recommendations

	(1) **Conflict**	(2) **No Conflict**	(3) **Both**
See Incentive First	0.137**	0.026	0.143**
	(0.058)	(0.079)	(0.057)
No Conflict			0.265***
			(0.073)
See Incentive First * No Conflict			-0.120
			(0.096)
Incentive for B	-0.182***	-0.091	-0.169***
	(0.056)	(0.085)	(0.047)
Selfishness	0.189***	0.041	0.148***
	(0.030)	(0.045)	(0.024)
Constant	0.656***	0.933***	0.666***
	(0.108)	(0.169)	(0.094)
Observations	213	86	299
R^2	0.247	0.037	0.209

Notes: This table displays the estimated coefficients from linear probability models on the advisors' recommendations. See Incentive first is a binary indicator coded as 1 for participants who were randomly assigned to see the incentive first. Selfishness is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task aimed at measuring moral costs. The sample includes attentive participants who did not switch multiple times in this elicitation. The regression includes individual controls for the advisor's gender and age. Robust standard errors (HC3) in parentheses

Table C.2: Recommendations including Inattentive

	(1) **Conflict**	(2) **No Conflict**	(3) **Both**
See Incentive First	0.172***	0.096	0.173***
	(0.059)	(0.083)	(0.059)
No Conflict			0.206***
			(0.076)
See Incentive First * No Conflict			-0.082
			(0.099)
Incentive for B	-0.217***	-0.089	-0.181***
	(0.062)	(0.085)	(0.050)
Constant	0.735***	0.827***	0.707***
	(0.108)	(0.180)	(0.095)
Observations	232	95	327
R^2	0.091	0.030	0.088

Notes: This table displays the estimated coefficients from linear probability models on the advisors' recommendations. See Incentive first is a binary indicator coded as 1 for participants who were randomly assigned to see the incentive first. The sample includes all participants, including those who switched multiple times in the MPL task to measure moral costs. The regression includes individual controls for the advisor's gender and age. Robust standard errors (HC3) in parentheses

C.2 Choice experiment

C.2.1 Preferences

Table C.3 breaks down preferences for information order by treatment and wave of the Choice Experiment. In Tables C.4 and C.5 we report the correlation between preferences to see the incentive first and preferences for blinding.

Table C.3: Preferences for Information Order by Wave

(1)	(2)	(3)	(4)
Wave	Treatment	N	Demand to See Incentive First
AMT-1	Choice Free	1308	0.563
AMT-1	Incentive First Costly	1347	0.422
AMT-2	Choice Free ($0.15 commission)	511	0.628
AMT-2	Choice Free ($15 for 1/100 commission)	542	0.565
AMT-3	Choice Free	213	0.451
AMT-3	Choice Free - Highx10	275	0.556
AMT-3	Choice Free - Highx100	110	0.600
AMT-3	Quality First Costly	1067	0.619
AMT-3	Incentive First Costly	215	0.340
Professionals	ChoiceFree—Professionals	712	0.480

Table C.4: Preferences for Blindness and Preferences for Information Order

	(1)	(2)	(3)
	Advisor Preference to Blind		
	Assigned Pref.	Not Assigned Pref.	Both
Prefer Incentive First	-0.217***	-0.268***	-0.216***
	(0.030)	(0.051)	(0.030)
No Conflict	-0.044	-0.092*	-0.057**
	(0.032)	(0.053)	(0.027)
Not Assigned Preference			0.010
			(0.045)
Prefer Incentive First X Not Assigned Preference			-0.054
			(0.058)
See Incentive First Costly	0.039	-0.101	-0.003
	(0.057)	(0.095)	(0.049)
Assess Quality First Costly	0.000	-0.070	-0.022
	(0.043)	(0.074)	(0.037)
Constant	0.550***	0.517***	0.545***
	(0.064)	(0.109)	(0.056)
Observations	1121	363	1484
R^2	0.053	0.104	0.060

Notes: This table displays the coefficient estimates of OLS regressions on the advisor's preferences to blind themselves to incentives information in the Blinding task. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.5: Preferences for Blindness, Information Order & Selfishness

	(1)	(2)	(3)
	Advisor Preference to Blind		
	Assigned Pref.	Not Assigned Pref.	Both
Prefer Incentive First	-0.198***	-0.245***	-0.198***
	(0.030)	(0.052)	(0.030)
No Conflict	-0.055*	-0.098*	-0.066**
	(0.031)	(0.052)	(0.027)
Not Assigned Preference			0.005
			(0.045)
Prefer Incentive First X Not Assigned Preference			-0.048
			(0.057)
See Incentive First Costly	0.044	-0.111	-0.001
	(0.056)	(0.092)	(0.048)
Assess Quality First Costly	-0.000	-0.077	-0.024
	(0.043)	(0.071)	(0.037)
Selfishness	-0.095***	-0.098***	-0.096***
	(0.015)	(0.026)	(0.013)
Constant	0.558***	0.550***	0.560***
	(0.063)	(0.108)	(0.055)
Observations	1121	363	1484
R^2	0.088	0.144	0.097

Notes: This table displays the coefficient estimates of OLS regressions on the advisor's preferences to blind themselves to incentives information in the Blinding task, controlling for selfishness. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.2.2 Recommendations

Figure C.1 displays recommendations for the cases in which advisors did not face a conflict of interest.

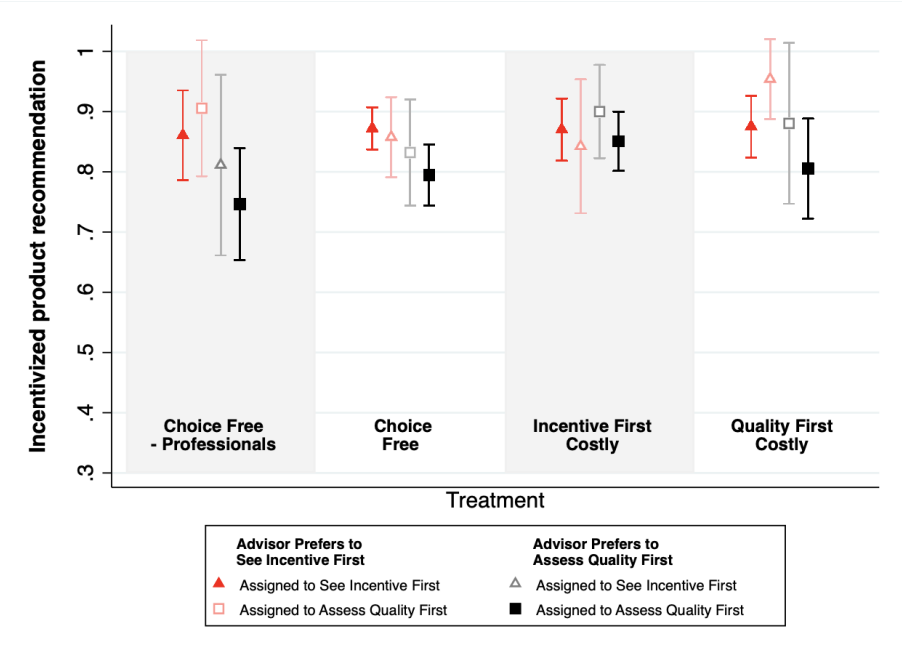

Notes: This figure presents the covariate-adjusted recommendations of the incentivized product for cases in which there was no conflict of interest. Error bars indicate ± 1 SE.

Figure C.1: Advisor Recommendations - No Conflict

Table C.6 reports the regression results for recommendations and looks at the relationship between recommendations and selfishness. In Tables C.7 and C.8, we break down the results for recommendations by product.

Table C.6: Advisor Recommendations

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.158***	-0.015	0.145***
	(0.017)	(0.030)	(0.016)
Not Assigned Preference			0.049**
			(0.022)
Prefer to See Incentive First X Not Assigned Pref.			-0.127***
			(0.027)
No Conflict	0.259***	0.220***	0.243***
	(0.021)	(0.034)	(0.019)
No Conflict X Prefer to See Incentive First	-0.125***	-0.005	-0.094***
	(0.026)	(0.046)	(0.023)
No Conflict X Not Assigned Preference			0.021
			(0.026)
See Incentive First Costly	0.025	0.021	0.023
	(0.017)	(0.031)	(0.015)
Assess Quality First Costly	0.005	0.093*	0.028
	(0.029)	(0.050)	(0.025)
Incentive for B	-0.158***	-0.179***	-0.163***
	(0.013)	(0.023)	(0.011)
Selfishness	0.100***	0.115***	0.104***
	(0.007)	(0.012)	(0.006)
Female	-0.008	-0.018	-0.011
	(0.013)	(0.023)	(0.011)
Age	-0.001**	-0.002*	-0.001***
	(0.001)	(0.001)	(0.000)
Constant	0.740***	0.825***	0.749***
	(0.028)	(0.049)	(0.025)
Observations	3915	1281	5196
R^2	0.150	0.146	0.147

Notes: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.7: Advisor Recommendations: Incentive for A

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.180***	-0.008	0.172***
	(0.020)	(0.035)	(0.019)
Not Assigned Preference			0.073***
			(0.027)
Prefer to See Incentive First X Not Assigned Pref.			-0.161***
			(0.033)
No Conflict	0.202***	0.170***	0.182***
	(0.027)	(0.038)	(0.025)
No Conflict X Prefer to See Incentive First	-0.112***	0.035	-0.077***
	(0.032)	(0.050)	(0.027)
No Conflict X Not Assigned Preference			0.048*
			(0.029)
See Incentive First Costly	0.022	0.044	0.027
	(0.023)	(0.041)	(0.020)
Assess Quality First Costly	0.002	-0.008	-0.002
	(0.039)	(0.066)	(0.033)
Female	0.017	-0.001	0.013
	(0.016)	(0.030)	(0.014)
Age	-0.002**	-0.001	-0.002***
	(0.001)	(0.001)	(0.001)
Constant	0.735***	0.785***	0.730***
	(0.035)	(0.061)	(0.032)
Observations	2242	725	2967
R^2	0.074	0.048	0.065

Notes: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option, focusing on the cases in which advisors were incentivized to recommend product A. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.8: Advisor Recommendations: Incentive for B

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.216***	0.018	0.197***
	(0.027)	(0.048)	(0.025)
Not Assigned Preference			0.047
			(0.034)
Prefer to See Incentive First X Not Assigned Pref.			-0.122***
			(0.039)
No Conflict	0.295***	0.229***	0.275***
	(0.028)	(0.049)	(0.026)
No Conflict X Prefer to See Incentive First	-0.157***	-0.011	-0.120***
	(0.038)	(0.067)	(0.033)
No Conflict X Not Assigned Preference			0.008
			(0.038)
See Incentive First Costly	0.049*	-0.000	0.035
	(0.027)	(0.047)	(0.023)
Assess Quality First Costly	0.003	0.201**	0.055
	(0.045)	(0.082)	(0.040)
Female	-0.009	-0.024	-0.015
	(0.019)	(0.034)	(0.017)
Age	-0.002***	-0.004***	-0.003***
	(0.001)	(0.001)	(0.001)
Constant	0.562***	0.740***	0.599***
	(0.041)	(0.074)	(0.037)
Observations	2206	735	2941
R^2	0.097	0.087	0.088

Notes: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option, focusing on the cases in which advisors were incentivized to recommend product B. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.9 examines the effect of selection on recommendations as a function of preference, conditioning on receiving information on a given order. Focusing on individuals who receive information about the incentive first (Column 1), we see that advisors who prefer to see the incentive first are 10.6 percentage points (pp) more likely to provide the incentivized product recommendations than those who prefer to see quality first. Similarly, focusing on individuals who receive information about the quality first (Column 2) we see that those who prefer to see the incentive first are 9.2 pp more likely to recommend the incentivized product than those who prefer to see quality first. In column (3) of Table C.9, we pool all advisors. We find that, among those assigned to see quality first, those who prefer to see the incentive first are more likely to recommend the incentivized product (9.6 pp). Among those who prefer to see quality first, the effect of being assigned to see the incentive first is 8.9 pp. Hence, focusing on these differences, we find that, while selection matters, there is a similarly important role for experience.

Table C.9: Advisor Recommendations - Role of Selection and Experience

	(1)	(2)	(3)	(4)	(5)
	Recommend incentivized product				
Sample:	Assigned to See:		Prefer to See:		Full Sample
	Incentive First	Quality First	Incentive First	Quality First	
Prefer See Inc. First	0.106***	0.092***			0.096***
	(0.023)	(0.024)			(0.021)
Assigned See Inc. First			0.100***	0.086***	0.089***
			(0.022)	(0.024)	(0.021)
Prefer X Assigned See Inc. First					0.008
					(0.026)
No Conflict	0.194***	0.260***	0.209***	0.261***	0.261***
	(0.033)	(0.020)	(0.030)	(0.020)	(0.018)
No Conflict X Prefer See Inc. First	-0.073**	-0.044			-0.059**
	(0.036)	(0.036)			(0.025)
No Conflict X Assigned See Inc. First			-0.089***	-0.066*	-0.078***
			(0.034)	(0.038)	(0.025)
Choice Free–Professionals	-0.004	-0.006	0.019	-0.023	-0.006
	(0.029)	(0.033)	(0.028)	(0.034)	(0.022)
See Inc. First Costly	-0.004	0.066***	-0.018	0.075***	0.032**
	(0.020)	(0.023)	(0.020)	(0.023)	(0.015)
Assess Quality First Costly	0.004	0.059	-0.029	0.081**	0.028
	(0.036)	(0.037)	(0.035)	(0.038)	(0.026)
Incentive for B	-0.169***	-0.180***	-0.162***	-0.188***	-0.175***
	(0.014)	(0.017)	(0.014)	(0.017)	(0.011)
Female	-0.003	-0.001	0.003	-0.005	-0.000
	(0.014)	(0.017)	(0.014)	(0.017)	(0.011)
Age	-0.003***	-0.001*	-0.002***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Constant	0.883***	0.698***	0.847***	0.740***	0.748***
	(0.034)	(0.035)	(0.032)	(0.036)	(0.025)
Observations	3038	2870	3102	2806	5908
R^2	0.083	0.090	0.073	0.094	0.098

Note: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) and (2) focus on participants assigned to experience a given information order. Column (1) focuses on individuals who are assigned to see the incentive first, column (2) focuses on individuals who are assigned to see quality first. Columns (3) and (4) focus on individuals' who prefer to be assigned to a given order, with Column (3) focusing on those who prefer to see the incentive first, and Column (4) focusing on those who prefer to see quality first. These groups are merged in column (5). Prefer See Inc. First is an indicator of the advisor's preference to see her incentive first, and Assigned See Inc. First is a indicator for whether advisors are assigned to see their incentive first. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. The same analysis including a measure of advisor's selfishness are shown in Online Appendix C. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

C.2.3 Beliefs

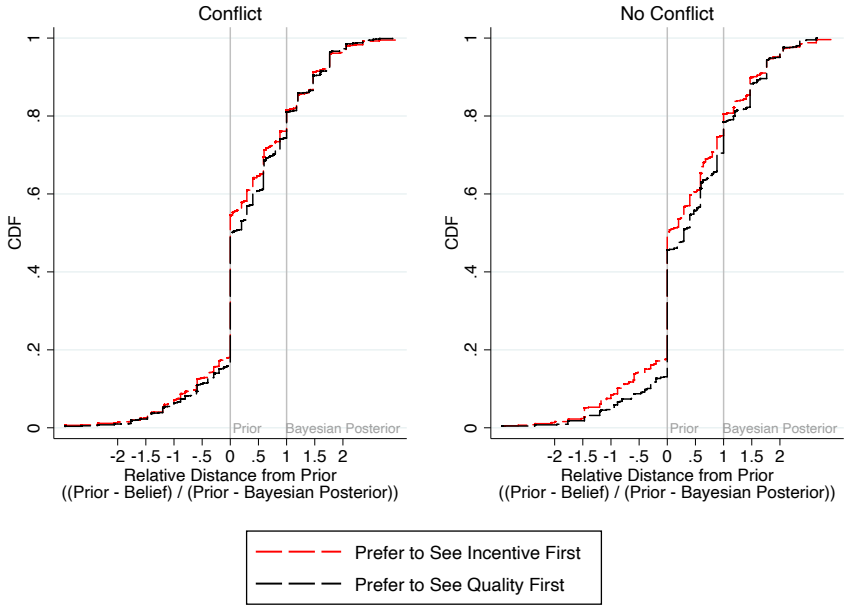

Notes: This figure presents the cumulative distribution of the extent of belief updating by advisors. This measure is the ratio of the difference between the advisor's belief and the prior, divided by the difference between the Bayesian posterior and the prior. The figure focuses on cases in which advisors are assigned to their preferred information order.

Figure C.2: Beliefs - Assigned Preferred Information Order

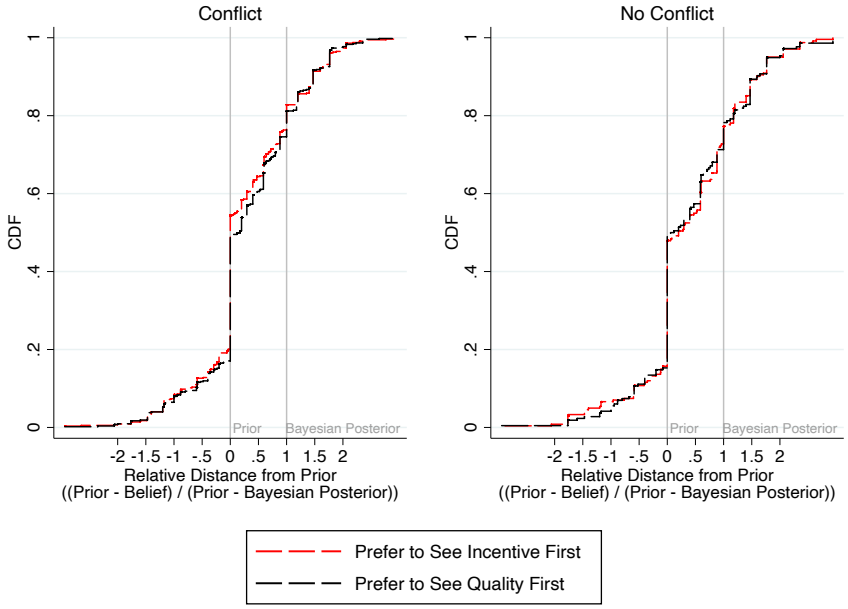

Notes: This figure presents the cumulative distribution of the extent of belief updating by advisors. This measure is the ratio of the difference between the advisor's belief and the prior, divided by the difference between the Bayesian posterior and the prior. The figure focuses on cases in which advisors are not assigned to their preferred information order.

Figure C.3: Beliefs - Not Assigned Preferred Information Order

Next, we conduct additional analyses on belief updating. Throughout we focus on analysis in which advisors who updated in the wrong direction are excluded. In Panel B of Table 4 of the main text, when focusing on participants who update in the correct direction (column 3) we saw that, when individuals face a conflict of interest there is a marginally significant difference between those who receive information about the incentive first and those who see quality first ($p = 0.078$) whereas this difference is not significant for the cases of no conflict ($p = 0.417$). Here, we also compare updating in response to signals of conflict (versus no conflict) separately for those who see information about incentive first and those who see information about quality first. Our results reveal evidence of asymmetric updating in both cases, though the gap between conflict and no conflict is slightly larger for cases in which advisors see the incentive first ($\beta_C^{f=q} = 0.574$, $\beta_{NC}^{f=q} = 0.6644$, $p = 0.0217$ for quality first; $\beta_C^{f=q} = 0.525$, $\beta_{NC}^{f=q} = 0.626$, $p = 0.0089$ for incentive first).

In Tables C.10 and C.11, we separate out these analyses by signal, looking at belief updating in response to signals of conflict and no conflict separately when the signal received was a \$0 ball (Table C.10) and when it was a \$2 ball (Table C.11). For the \$0 signal (Table C.10), Panel A shows that, in the aggregate, there is evidence of suppression: advisors update less in response of signals of conflict as compared to signals of no conflict. In Panel B, we see that individuals update less in response of signal of conflict when they see information about their incentive first (vs. quality first); this is not the case for signals of no conflict (Column 3). When advisors pursue and are assigned to see the incentive first, the difference in updating between signals of conflict and signals of no conflict is large and statistically significant ($\beta_C^{f=i} = 0.419$, $\beta_{NC}^{f=i} = 0.552$, $p = 0.0084$). This difference is smaller and not statistically significant when advisors pursue and are assigned to information about quality first ($\beta_C^{f=i} = 0.479$, $\beta_{NC}^{f=i} = 0.555$, $p = 0.1429$). Taken together, these results are suggestive that, when advisors received a \$0 signal, they engaged in additional suppression of signals of conflict when presented with information about incentive first (as opposed to information about quality first).

For the \$2 signal (Table C.11), we see a somewhat different updating pattern. First, we see that, regardless of information order, advisors update at a greater extent when presented with \$2 signals for both signals of conflict and signals of no conflict, and we see no asymmetry in updating in the aggregate (Panel A, column 3). In Panel B, we observe a directional difference between the cases in which advisor see the incentive first and those in which they see quality first both for cases of conflict ($\beta_C^{f=i} = 0.702$ versus $\beta_C^{f=q} = 0.752$),and for cases of no conflict ($\beta_{NC}^{f=i} = 0.724$ versus $\beta_{NC}^{f=q} = 0.789$), though none of the differences are statistically significant ($p = 0.276$ and $p = 0.339$, respectively). We do not see an asymmetry in updating in response to signals of conflict (vs. no conflict)

either when the information about incentive is presented first ($p = 0.7011$), or when quality information is presented first ($p = 0.5256$).

Table C.10: Belief Updating when Signal is \$0

	(1)	(2)	(3)	(4)
			Log-odds Belief	
Assignment:	Assigned Pref.	Not Assigned Pref.	Assigned Pref.	Not Assigned Pref.
Data:		All	Excl. update in wrong direction	
Panel A: Pooled				
β_C	0.217***	0.259***	0.449***	0.506***
	(0.020)	(0.036)	(0.017)	(0.031)
β_{NC}	0.309***	0.287***	0.553***	0.538***
	(0.038)	(0.064)	(0.032)	(0.052)
Panel B: By Choice of Information Order				
$\beta_C^{f=i}$	0.173***	0.265***	0.419***	0.517***
	(0.028)	(0.049)	(0.024)	(0.042)
$\beta_C^{f=q}$	0.262***	0.253***	0.479***	0.495***
	(0.029)	(0.052)	(0.024)	(0.045)
$\beta_{NC}^{f=i}$	0.285***	0.361***	0.552***	0.562***
	(0.052)	(0.086)	(0.045)	(0.069)
$\beta_{NC}^{f=q}$	0.339***	0.194**	0.555***	0.503***
	(0.056)	(0.096)	(0.046)	(0.080)
Observations	1765	569	1428	453
$\beta_C^{f=q} = \beta_C^{f=i}$	0.026	0.872	0.076	0.727
$\beta_{NC}^{f=q} = \beta_{NC}^{f=i}$	0.476	0.193	0.964	0.576

Notes: The outcome in all regressions is the log belief ratio, when the advisors sees a \$0 ball for product B. β_C^f and β_{NC}^f are the estimated effects of the log likelihood ratio for conflict and no conflict signals, respectively, for advisors who prefer order ($f = i$ indicates a preference to see the incentive first, and $f = q$ indicates a preference to see quality first). Columns(1) and (2) include all advisors. Columns (3) and (4) exclude advisors who updated in the wrong direction. Columns (1) and (3) include only advisors who were assigned their preference, while columns (2) and (4) include only advisors who were not assigned their preference. Robust standard errors (HC3) in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Tables C.12 and Table C.13, we make use of the incentivized belief data where advisors had to pick one of 10 belief bins. Table C.12 focuses on posterior beliefs, showing that advisors who prefer to see the incentive first and are assigned their preference are less likely to pick the correct choice bin as their posterior. Table C.13 focuses on the likelihood of sticking to prior of 0.50, showing that that advisors who prefer to see the incentive first and are assigned their preference are more likely to stick to the bin containing the prior. In Table C.14, we repeat the analysis from Table C.13 using the continuous beliefs data and find consistent results.

Table C.11: Belief Updating when Signal is $2

	(1)	(2)	(3)	(4)
	Log-odds Belief			
Assignment:	Assigned Pref.	Not Assigned Pref.	Assigned Pref.	Not Assigned Pref.
Data:	All		Excl. update in wrong direction	
Panel A: Pooled				
β_C	0.471^{***}	0.410^{***}	0.726^{***}	0.506^{***}
	(0.026)	(0.043)	(0.023)	(0.031)
β_{NC}	0.473^{***}	0.484^{***}	0.757^{***}	0.538^{***}
	(0.039)	(0.066)	(0.034)	(0.052)
Panel B: By Choice of Information Order				
$\beta_C^{f=i}$	0.436^{***}	0.363^{***}	0.702^{***}	0.661^{***}
	(0.035)	(0.060)	(0.031)	(0.053)
$\beta_C^{f=q}$	0.511^{***}	0.460^{***}	0.752^{***}	0.732^{***}
	(0.038)	(0.063)	(0.034)	(0.055)
$\beta_{NC}^{f=i}$	0.380^{***}	0.460^{***}	0.724^{***}	0.826^{***}
	(0.058)	(0.106)	(0.048)	(0.085)
$\beta_{NC}^{f=q}$	0.573^{***}	0.508^{***}	0.789^{***}	0.700^{***}
	(0.053)	(0.080)	(0.048)	(0.072)
Observations	2620	878	2246	740
$\beta_C^{f=q} = \beta_C^{f=i}$	0.149	0.265	0.276	0.352
$\beta_{NC}^{f=q} = \beta_{NC}^{f=i}$	0.014	0.715	0.339	0.257

Notes: The outcome in all regressions is the log belief ratio, when the advisors sees a $2 ball for product B. β_C^f and β_{NC}^f are the estimated effects of the log likelihood ratio for conflict and no conflict signals, respectively, for advisors who prefer order ($f = i$ indicates a preference to see the incentive first, and $f = q$ indicates a preference to see quality first). Columns(1) and (2) include all advisors. Columns (3) and (4) exclude advisors who updated in the wrong direction. Columns (1) and (3) include only advisors who were assigned their preference, while columns (2) and (4) include only advisors who were not assigned their preference. Robust standard errors (HC3) in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.12: Belief Updating: Correct Choice

	(1)	(2)	(3)	(4)
	Belief Correct			
Assignment:	Assigned Pref.	Not Assigned Pref.	Assigned Pref.	Not Assigned Pref.
Data:	All		Excl. update in wrong direction	
Prefer to See Incentive First	-0.039***	-0.032	-0.042***	-0.037
	(0.013)	(0.024)	(0.016)	(0.028)
No Conflict	0.018	0.012	0.015	0.005
	(0.019)	(0.033)	(0.022)	(0.038)
No Conflict X Prefer to See Incentive First	0.012	0.022	0.019	0.037
	(0.025)	(0.045)	(0.029)	(0.052)
See Incentive First Costly	-0.002	-0.003	-0.001	-0.007
	(0.015)	(0.027)	(0.018)	(0.032)
Assess Quality First Costly	0.036	-0.046	0.042	-0.051
	(0.025)	(0.047)	(0.030)	(0.052)
Incentive for B	-0.001	0.014	-0.002	0.017
	(0.011)	(0.021)	(0.013)	(0.024)
Female	-0.029**	-0.016	-0.033**	-0.011
	(0.011)	(0.020)	(0.013)	(0.024)
Age	-0.000	-0.001	0.000	-0.001
	(0.000)	(0.001)	(0.001)	(0.001)
Constant	0.196***	0.192***	0.221***	0.218***
	(0.024)	(0.044)	(0.028)	(0.050)
Observations	4448	1460	3700	1224
R^2	0.007	0.009	0.008	0.010

Notes: This table displays the estimated coefficients from linear probability models on the advisor's beliefs that the quality of product B is low measured via their choice of one out of 10 possible belief bins (ranging from 0 to 100, in steps of 10). Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Columns (3) and (4) exclude individuals who chose a bin that is consistent with updating in the incorrect direction. Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.13: Belief Updating: Likelihood of Sticking to the Prior Belief - Incentivized Elicitation

	(1)	(2)	(3)	(4)
	Belief Bin Containing the Prior of 50%			
Data:	All		Excl. update in wrong direction	
Prefer to See Incentive First	0.026	0.019	0.039**	0.031
	(0.018)	(0.031)	(0.020)	(0.034)
No Conflict	-0.008	0.027	-0.023	0.019
	(0.023)	(0.042)	(0.026)	(0.046)
No Conflict X Prefer to See Incentive First	-0.031	-0.077	-0.030	-0.080
	(0.032)	(0.056)	(0.036)	(0.063)
See Incentive First Costly	0.007	0.027	0.011	0.025
	(0.020)	(0.036)	(0.023)	(0.040)
Assess Quality First Costly	0.006	-0.044	0.001	-0.051
	(0.034)	(0.059)	(0.038)	(0.063)
Incentive for B	-0.015	-0.073***	-0.022	-0.093***
	(0.015)	(0.026)	(0.017)	(0.029)
Female	0.056***	0.008	0.070***	0.028
	(0.015)	(0.026)	(0.016)	(0.029)
Age	-0.002***	-0.001	-0.001**	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.412***	0.440***	0.467***	0.495***
	(0.031)	(0.055)	(0.035)	(0.061)
Observations	4448	1460	3700	1224
R^2	0.009	0.017	0.011	0.020

Notes: This table displays the estimated coefficients from linear probability models on the advisor's likelihood to stick to the bin containing the prior belief (50%) in the incentivized belief elicitation. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Columns (3) and (4) exclude individuals who chose a bin that is consistent with updating in the incorrect direction. Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.14: Belief Updating: Likelihood of Sticking to the Prior Belief - Continuous Elicitation

	(1)	(2)	(3)	(4)
	Belief at Prior of 50%			
Data:	All		Excl. update in wrong direction	
Prefer to See Incentive First	0.033*	0.016	0.054***	0.034
	(0.017)	(0.030)	(0.020)	(0.035)
No Conflict	-0.012	0.027	-0.027	0.024
	(0.023)	(0.041)	(0.025)	(0.046)
No Conflict X Prefer to See Incentive First	-0.040	-0.047	-0.040	-0.060
	(0.031)	(0.055)	(0.035)	(0.063)
See Incentive First Costly	-0.011	0.030	-0.009	0.032
	(0.020)	(0.035)	(0.023)	(0.040)
Assess Quality First Costly	-0.042	-0.024	-0.063*	-0.024
	(0.034)	(0.055)	(0.038)	(0.062)
Incentive for B	0.011	-0.024	0.027*	-0.024
	(0.014)	(0.025)	(0.016)	(0.030)
Female	0.072***	0.008	0.084***	0.017
	(0.014)	(0.025)	(0.016)	(0.029)
Age	-0.001	-0.002*	-0.000	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.351***	0.401***	0.382***	0.455***
	(0.031)	(0.054)	(0.035)	(0.061)
Observations	4448	1460	3712	1201
R^2	0.011	0.009	0.015	0.011

Notes: This table displays the estimated coefficients from linear probability models on the advisor's likelihood to stick to the bin containing the prior belief (50%) in the incentivized belief elicitation. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Columns (3) and (4) exclude individuals who chose a bin that is consistent with updating in the incorrect direction. Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

C.2.4 Explanations for Choices

In the second wave of the Choice experiment (AMT-2) and the Choice Free Professional treatment of the Choice Experiment, we asked advisors to explain their choices of information order. A total of $N = 1,747$ advisors ($N = 1,035$ from AMT-2 and $N = 712$ professionals) from the main sample answered this question. We then had two independent raters code the reasons for choosing to see the incentive first or quality first. The two raters agreed in over 82% of their classifications, leading to an interrater agreement κ of 0.76. We average their ratings to examine how advisors' explanations vary with their preference of information order.

Table C.12 below reports the raters' coding of the 91% of the data ($N = 1556$ advisors, of which $N = 660$ were professionals) of advisors who provided an explanation.

Table C.15: Advisors' Explanations: Detailed Results

		Advisors' Explanations of Preference (Categories)			
		Limiting Bias	Indifference	Commission	Other reasons
Sample: All (N=1,556)					
Prefer to:	Assess Quality First	47.1%	10.8%	7.8%	37.4%
	See Incentive First	5.9%	7.1%	36.4%	55.1%
Sample: AMT (N=896)					
Prefer to:	Assess Quality First	41.4%	11.6%	10.7%	39.3%
	See Incentive First	5.1%	7.4%	36.5%	55.0%
Sample: Professionals (N=660)					
Prefer to:	Assess Quality First	53.1%	10%	4.8%	35.4%
	See Incentive First	7.2%	6.6%	36.4%	55.2%

Note: This table displays the average rating of advisors whose explanation to see their incentive first or assess quality first was classified into each category. This classification excludes answers that were blank or unrelated to the choice.

C.3 Comparing the Choice and NoChoice Experiments

In this section, we compare the recommendation decisions in the Choice experiment to those in the NoChoice experiment. For this comparison, we focus on the ChoiceFree treatment, which uses the same sample and incentives of the NoChoice experiment, and cases in which the advisor's incentives were in conflict with the quality signal.

When advisors are assigned to see the Incentive first in the NoChoice experiment, we estimate a 16.9pp increase in recommendations as compared to advisors who are randomly assigned to see Quality first (Table **??** column (1)). Instead, when advisors

prefer and are assigned to see the incentive first, we estimated a 23.5pp increase in recommendations, as compared to advisors who prefer and are assigned to see quality first (Table **??** column (2)). This increase is 7.6 percentage points larger, albeit not statistically different, than the increase in recommendations observed when advisors see the incentive first in NoChoice.

In addition to examining the difference in recommendations between advisors who choose (Choice experiment) or are assigned (NoChoice Experiment) an information order, we also conduct exploratory comparisons of the levels. Table C.17 reports covariate-adjusted recommendation rates for the See Incentive First and See Quality first treatments of the NoChoice experiment, and for assignment to see the Incentive first (vs. see Quality First) conditional on preferences for the Choice experiment.

We first focus on advisors who choose or are assigned to seeing the incentive first. In the NoChoice experiment, advisors who are assigned to "See Incentive First" recommend the incentivized option in 79% of the cases. In the Choice experiment, advisors who prefer to see the incentive first and are assigned to see the incentive first recommend the incentivized option in 81% of the cases. When advisors prefer to see quality first but are assigned to see the incentive first, they instead recommend the incentivized option in 67.9% of the cases. This difference in recommendations is in line with advisors who prefer to see the incentive first being more selfish. In the experiment, 55% of advisors prefer to see the incentive first whereas 45% of advisors prefer to see quality first. Weighting by their preference, we can estimate what the recommendation rate would have been if advisors in the Choice experiment were not asked to make a choice, and were rather randomly assigned to see the incentive first as in the NoChoice Experiment. In that case, the predicted rate of recommendations is 76.6% (from 0.55 X 0.810 + 0.45 X 0.679). This recommendation rate is 2.4pp smaller than the rate of 79% we observe in the No-Choice experiment, but falls within its 95% confidence interval. This small decrease in recommendation could be due to a some advisors being limited in their ability to justify self-serving recommendations or due to noise in the data.

Next, we consider the preference to see quality first. In the NoChoice experiment, advisors assigned to this preference order recommend the incentivized option in 62% of the cases. In the Choice experiment, advisors who prefer and are assigned to see quality first recommend the incentivized option in 56.9% of the cases, while advisors who prefer to see the incentive first but are assigned to see quality first recommend the incentivized option in 71.3% of the cases. Weighting by advisors' preferences, we estimate that the recommendation rate if advisors had been randomly assigned to see quality first as in the NoChoice experiment would have been 63%. This estimate is very close to the 62% of recommendations we estimate in the NoChoice experiment.

Taken together, the results observed in the Choice experiment are consistent with those observed in the NoChoice experiment. Advisor recommendations in the Choice experiment are still significantly affected by assignment to the advisors' preferred order, which indicates that their active choice did not remove the scope for self-deception. At the same time, the 2.4pp difference between our prediction from the Choice data and the results of the NoChoice experiment in our exploratory analyses, might potentially indicate that the scope for self-deception may have been directionally restricted.

Table C.16: Advisor Recommendations

	(1)	(2)	(3)
	Recommend incentivized product		
Sample:	NoChoice.	Choice	Both
See Incentive First	0.1686***	0.2352***	0.2359***
	(0.0575)	(0.0226)	(0.0226)
No Choice			0.0471
			(0.0444)
See IncentiveFirst X NoChoice			-0.0764
			(0.0605)
No Conflict	0.2784***	0.2710***	0.2717***
	(0.0770)	(0.0313)	(0.0313)
No Choice X No Conflict			0.0006
			(0.0809)
See Incentive First X No Conflict	-0.1418	-0.1748***	-0.1752***
	(0.1063)	(0.0411)	(0.0412)
See Incentive First X No Choice X No Conflict			0.0361
			(0.1125)
Incentive for B	-0.1634***	-0.1514***	-0.1527***
	(0.0495)	(0.0191)	(0.0178)
Constant	0.6918***	0.7456***	0.7317***
	(0.0921)	(0.0371)	(0.0353)
Observations	299	1931	2230
R^2	0.093	0.117	0.113

Note: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) focuses on the NoChoice Experiment, while column (2) focuses on the Choice Experiment (ChoiceFree Treatment only) and on individuals who are assigned their preference. Both groups are merged in column (3). See Incentive First is an indicator for whether advisors are randomly assigned to see the incentive first in NoChoice, and whether, conditional on preferring to see the incentive first, they are assigned to see the incentive first in Choice. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. All regression models include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.17: Recommendations in the NoChoice and Choice Experiment

Experiment	Treatment	Mean	95% CI
NoChoice	See Incentive First	79%	[71%-87%]
	See Quality First	62%	[53%-71%]
Choice	Prefer to See Incentive First & Assigned to See Incentive First	81%	[78%-84%]
	Prefer to See Incentive First & Assigned to See Quality First	71%	[66%-77%]
	Prefer to See Quality First & Assigned to See Incentive First	68%	[61%-74%]
	Prefer to See Quality First & Assigned to See Quality First	57%	[53%-61%]
Choice - Predicted	See Incentive First	77%	-
	See Quality First	63%	-

Note: This table displays the proportion of incentivized recommendations in the NoChoice experiment, the covariate-adjusted estimates of frequency of recommendations of the incentivized product by treatment and assignment in the ChoiceFree treatment of the Choice experiment, obtained via OLS regression, and the predicted recommendations in the Choice experiment, weighted by average preferences to see the incentive first/assess quality first.

C.4 The Higher Incentives Treatments

In this section, we report the results from the two robustness treatments that we collected as part of the AMT-3 wave of the choice experiment. In these treatments, we scale the incentives by a factor of 10 (High Stakes - 10-fold incentives) or 100 (High Stakes - 100-fold incentives). As part of that wave, we also collected data for our regular version of the Choice Free treatment with low incentives (a \$0.15 commission and products that yielded \$0 or \$2 to the client). As shown in Table C.3, the share of advisors who choose to see the incentive first is larger when advisors face larger incentives (45% with regular incentives as opposed to 55% with 10-fold incentives and 60% with 100-fold incentives; $p = 0.02$ and $p = 0.01$, respectively). This data shows that despite the substantially higher incentives, the fraction of advisors who prefers to assess quality first remains substantial.

As shown in Table C.18, when looking at preferences for information order in our full sample using OLS regressions and controlling for wave, we see that advisors in the treatments with higher incentives are 9 (High Stakes - 10 fold incentives) and 13 (High Stakes - 100 fold incentives) percentage points more likely to choose to see the incentives first than participants who were presented with smaller incentives.

In Table C.19, we report the results for recommendations. As displayed in the table, advisors in these treatments are directionally more likely to recommend the incentivized product than those who faced smaller incentives in the Choice Free treatment. Importantly, the coefficient for the interaction between preferring to see the incentive first and the indicator for these treatments is not statistically significant (directionally, it is positive). Taken together, these results suggest that the effect of information order is robust to increasing the stakes.

Table C.18: Preference for Information Order: Including Incentives Treatments

	(1)	(2)	(3)
	Prefer to See Incentive First		
See Incentive First Costly	-0.14***	-0.14***	-0.14***
	(0.02)	(0.02)	(0.02)
Assess Quality First Costly	0.15***	0.15***	0.15***
	(0.03)	(0.03)	(0.03)
Choice Free – Professionals	-0.10***	0.00	0.00
	(0.03)		
High Stakes (10-fold incentives)	0.09**	0.08**	0.08*
	(0.04)	(0.04)	(0.04)
High Stakes (100-fold incentives)	0.13**	0.12**	0.12**
	(0.05)	(0.05)	(0.05)
Selfishness		0.06***	0.08***
		(0.01)	(0.01)
See Incentive First Costly X Selfishness			-0.05***
			(0.02)
See Quality First Costly X Selfishness			-0.03
			(0.02)
Female	-0.03**	-0.03**	-0.03**
	(0.01)	(0.01)	(0.01)
Age	-0.00***	-0.00***	-0.00***
	(0.00)	(0.00)	(0.00)
Constant	0.68***	0.66***	0.66***
	(0.02)	(0.02)	(0.02)
Observations	6293	5581	5581
R^2	0.034	0.049	0.051

Notes: This table displays the estimated coefficients from linear probability models on the preference to see the incentive first. See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. The regression models in columns (2) and (3) include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.19: Advisor Recommendations: Including Incentives Treatments

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.196***	0.002	0.182***
	(0.016)	(0.029)	(0.015)
Not Assigned Preference			0.061***
			(0.021)
Prefer to See Incentive First X Not Assigned Pref.			-0.140***
			(0.026)
No Conflict	0.256***	0.201***	0.236***
	(0.020)	(0.033)	(0.018)
No Conflict X Prefer to See Incentive First	-0.137***	0.013	-0.098***
	(0.025)	(0.045)	(0.022)
No Conflict X Not Assigned Preference			0.019
			(0.025)
See Incentive First Costly	0.035**	0.020	0.031**
	(0.017)	(0.032)	(0.015)
Assess Quality First Costly	0.004	0.091*	0.027
	(0.030)	(0.052)	(0.026)
Incentive for B	-0.168***	-0.182***	-0.171***
	(0.012)	(0.022)	(0.011)
Female	0.006	-0.031	-0.004
	(0.012)	(0.022)	(0.011)
Age	-0.002***	-0.003***	-0.002***
	(0.001)	(0.001)	(0.000)
High Stakes (10-fold incentives)	0.119*	0.133	0.133**
	(0.064)	(0.100)	(0.061)
High Stakes (100-fold incentives)	0.131	0.291	0.142
	(0.096)	(0.180)	(0.093)
Prefer to See Incentive First X High Stakes (10-fold)	0.038	0.013	0.041
	(0.067)	(0.131)	(0.064)
Prefer to See Incentive First X High Stakes (100-fold)	0.044	-0.403	0.052
	(0.104)	(0.248)	(0.101)
Constant	0.734***	0.867***	0.753***
	(0.027)	(0.047)	(0.024)
Observations	4743	1550	6293
R^2	0.110	0.085	0.101

Notes: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

C.5 Including Inattentive Participants

Our main sample exclude all participants who give non-monotone responses to the multiple price list questions that aims to measure selfishness and classify participants into moral types. A total of 1355 participants switched multiple times in the MPL, and, as pre-registered were therefore excluded from the main analyses. Here, we repeat the analyses for preferences and recommendations from the main text (Tables 2 and 3) but include participants who switch multiple times in the multiple price list to measure selfishness. Column 1 includes only attentive participants from the main sample, and Column 2 includes all participants (including the inattentive ones).

Table C.20: Preference for Information Order—Including Inattentive

	(1)	(2)
	Prefer to See Incentive First	
See Incentive First Costly	-0.139***	-0.139***
	(0.018)	(0.018)
Assess Quality First Costly	0.152***	0.152***
	(0.029)	(0.029)
Choice Free – Professionals		-0.095***
		(0.026)
Female	-0.026*	-0.029**
	(0.014)	(0.013)
Age	-0.002***	-0.003***
	(0.001)	(0.001)
Constant	0.668***	0.674***
	(0.025)	(0.024)
Observations	5196	5908
R^2	0.037	0.034

Note: This table displays the estimated coefficients from linear probability models on the advisor's preference to see the incentive first. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.21: Advisor Recommendations—Including Inattentive

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.195***	0.003	0.181***
	(0.016)	(0.029)	(0.015)
Not Assigned Preference			0.060***
			(0.021)
Prefer to See Incentive First X Not Assigned Preference			-0.140***
			(0.026)
No Conflict	0.256***	0.202***	0.236***
	(0.020)	(0.033)	(0.018)
No Conflict X Prefer to See Incentive First	-0.137***	0.012	-0.098***
	(0.025)	(0.045)	(0.022)
No Conflict X Not Assigned Preference			0.019
			(0.025)
Choice Free–Professionals	-0.026	0.051	-0.006
	(0.025)	(0.044)	(0.022)
See Incentive First Costly	0.035**	0.020	0.031**
	(0.017)	(0.031)	(0.015)
Assess Quality First Costly	0.004	0.093*	0.027
	(0.030)	(0.052)	(0.026)
Incentive for B	-0.171***	-0.187***	-0.175***
	(0.013)	(0.023)	(0.011)
Female	0.005	-0.015	-0.001
	(0.013)	(0.023)	(0.011)
Age	-0.002***	-0.003***	-0.002***
	(0.001)	(0.001)	(0.000)
Constant	0.737***	0.864***	0.755***
	(0.027)	(0.048)	(0.025)
Observations	4448	1460	5908
R^2	0.106	0.083	0.097

Note: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Choice Free-Professionals, See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. All regression models include controls for each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

C.6 The Choice Stakes Experiment: Additional Results

Table C.22: Preference for Information Order

	(1)	(2)
	Prefer to See Incentive First	
Low Incentive	-0.276***	-0.278***
	(0.027)	(0.027)
High Incentive	0.029	0.028
	(0.031)	(0.032)
Selfishness		0.009
		(0.012)
Female	-0.039	-0.038
	(0.024)	(0.024)
Age	-0.001	-0.001
	(0.001)	(0.001)
Constant	0.477***	0.474***
	(0.044)	(0.044)
Observations	1471	1471
R^2	0.088	0.089

Note: This table displays the estimated coefficients from linear probability models on the preference to see the incentive first. Low Incentive and High Incentive are indicator variables that take value 1 in the respective treatment, 0 otherwise. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table C.23: Advisor Recommendations

	(1)	(2)	(3)
	Recommend incentivized product		
Assignment:	Assigned Pref.	Not Assigned Pref.	Both
Prefer to See Incentive First	0.1363***	0.0456	0.1498***
	(0.0512)	(0.1023)	(0.0471)
Not Assigned Preference			0.0887**
			(0.0402)
Prefer to See Incentive First X Not Assigned Preference			-0.1515***
			(0.0582)
No Conflict	0.2842***	0.2456***	0.2836***
	(0.0344)	(0.0603)	(0.0329)
No Conflict X Prefer to See Incentive First	-0.1307**	-0.1912*	-0.1383***
	(0.0558)	(0.1093)	(0.0498)
No Conflict X Not Assigned Preference			-0.0484
			(0.0555)
Low Incentive	-0.1451***	-0.0561	-0.1249***
	(0.0403)	(0.0719)	(0.0351)
Low Incentive X Prefer to See Incentive First	0.0121	0.0775	0.0134
	(0.0855)	(0.1704)	(0.0764)
High Incentive	0.0241	0.0535	0.0323
	(0.0442)	(0.0751)	(0.0379)
High Incentive X Prefer to See Incentive First	0.0748	0.0067	0.0565
	(0.0636)	(0.1197)	(0.0560)
Incentive for B	-0.1269***	-0.0869*	-0.1186***
	(0.0272)	(0.0518)	(0.0239)
Female	-0.0317	-0.1215**	-0.0535**
	(0.0273)	(0.0499)	(0.0239)
Age	-0.0017	-0.0004	-0.0014
	(0.0012)	(0.0021)	(0.0010)
Constant	0.7145***	0.7399***	0.6987***
	(0.0555)	(0.1023)	(0.0497)
Observations	1104	367	1471
R^2	0.121	0.063	0.104

Notes: This table displays the estimated coefficients from linear probability models on the advisor's decision to recommend the incentivized option. Column (1) focuses on individuals who are assigned their preference, while column (2) focuses on individuals who are not assigned their preference. Both groups are merged in column (3). Prefer to See Incentive First is an indicator of the advisor's preference, and Not Assigned Preference is an indicator for not receiving the preferred order. No Conflict is an indicator for the cases in which the signal of quality is not in conflict with the advisor's commission. Low Incentive and High Incentive are indicator variables that take value 1 in the respective treatment, 0 otherwise. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

C.7 The Information Architect Experiment: Additional Results

In the Information Architect experiment, we investigate preferences for information order of a third party who determines how advisors receive information. The sample is comprised by 498 attentive participants. An additional 51 participants switched multiple times in the task that measured selfishness. As preregistered, these participants are dropped from the main analysis, but for robustness, we repeat the analysis including these participants. Table C.25 presents regression results comparing IA preferences in IA-Advisor, relative to IA-Client (omitted category), controlling for the IA's gender and age.

	(1)	(2)
	DM Choice to See Incentive First	
Sample:	Main Sample	Including Inattentive
DMAdvisor		
IA-Advisor	0.148^{***}	0.142^{***}
	(0.045)	(0.042)
Constant	0.334^{***}	0.353^{***}
	(0.082)	(0.078)
Observations	498	549
R^2	0.033	0.031

Notes: This table displays the coefficient estimates of OLS regressions on the Information Architect's preferences to have the advisor see the incentive first for the main sample (Column 1) and the sample that includes inattentive participants who switched multiple times in the selifishness measure. DM-Advisor is an indicator for whether advisors have an incentive to receive information about their incentive first. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As part of the experiment, we then recruited $N = 498$ advisors, presented with the information order selected by the Information Architect and had them make their recommendation to the client. For this purpose, we recruited $N = 50$ clients for the main task; of these 86% followed the recommendations. We also recruited an additional $N = 50$ advisees for the MPL task that measured advisors' moral costs and matched them with 1 out of 10 Information Architects, and an additional $N = 50$ advisees for the same task and matched them with 1 out of 10 advisors. Of these, 86% and 80% of advisees followed the recommendation.

D The NoChoiceSimultaneous Experiment

D.1 Experimental Design

To test how advisors behave when they receive information about their own incentive and information about the quality of the product (i.e., the quality signal) simultaneously, we conducted an additional wave of the NoChoice experiment. The experiment replicates the design and procedures of the NoChoice experiment. On top of the See Incentive First and See Quality First treatments, this wave of data collection added an additional treatment (Simultaneous) where information about incentive and the quality signal were presented to participants on the same screen. Participants were assigned to the three treatments at a 1:1:3 ratio, as we planned to merge the data with those of the original NoChoice experiment for the analyses. By comparing the rates of recommendations of the incentivized product in the Simultaneous treatment to those in the See Incentive First and See Quality First treatment, we can investigate how receiving information simultaneously affects recommendation in case of a conflict of interest.

Procedures. The experiment was conducted on Amazon Mechanical Turk (AMT); the design and analyses plan were pre-registered on aspredicted.org (#79521 and #82164). Participants received $1 payment for taking part in the experiment and for making a recommendation. On top of that, they received a $.15 commission for recommending either Product A or Product B. Advisors were informed that one out of 10 advisors would be matched with a client, another AMT participant, and their advice was delivered to them.

For the See Incentive First and See Quality First treatments, all procedures were identical to those of the NoChoice Experiment, with some small modifications. In particular, to address potential concerns about demand effect, whereby participants may assume that the order of information is determined by the experimenter thereby leading participants to "react" to the experimenter decisions, we informed participants that the order of information in the experiment was randomly determined by the computer. Further, we also informed participants that whether the commission was for Product A or Product B was randomly determined by the computer. In the See Incentive First treatment, participants first saw information about what product yielded a commission and then received further information about the quality of the product. In the See Quality First treatment, participants first learned about the quality of the product and then received information about their incentive. In the Simultaneous treatment, the information about the incentive and the quality signal appeared on the same screen. We counterbalanced whether the information about the incentive appeared on the top or the bottom of the

screen, to control for the potential effect of position on the screen on attention. Then, participants were prompted to make a recommendation. We further collected additional measures of beliefs and selfishness using a the same measures used in the NoChoice Experiment. At the end of the experiments, we randomly selected 1 out of 10 advisors and sent their recommendation to a client.

D.2 Results

As pre-registered, we merge the data from the NoChoiceSimultaneous experiment with the data collected for the NoChoice experiment, and control for the wave in which the data was collected. The main sample comprises of a total of 276 attentive participants from the NoChoiceSimultaneous experiment and 298 attentive participants from the original wave of the NoChoice experiment, for a total of 574 attentive participants. However, in this experiment, overall, we had much lower data quality that in the prior wave of the NoChoice experiment as well as all the prior experiments. Among those who completed the NoChoiceSimultaneous experiment, 50.63% ($N = 283$) of participants switched multiple times in the multiple price list measure of selfishness, one of our exclusion criteria in the pre-registration. This fraction is much larger than the fraction of inconsistent participants in any of the other study we ran.[2] Given these differences in data quality, we analyze the data both including and excluding participants who switch more than once in the measure of selfishness.

As shown in Table D.1, participants in the Simultaneous treatment, who received both the information about the incentive and the quality signal on the same screen, were more likely to recommend the incentivized product in cases of conflict than participants in the Assess Quality First treatment. As shown in the table, these participants behaved similarly to those in the See Incentive First treatment. The results are similar both if we include (Column 3) and exclude (Columns 1-2) inattentive participants.

The 276 attentive participants were matched with $N = 28$ clients for the main task; of these 96% followed the recommendation. They were also matched with $N = 28$ clients for the MPL task; of these 79% followed the recommendation.[3]

[2]At the time we ran he experiment, Cloudresearch changed some of the features it used to filter participants (https://www.cloudresearch.com/resources/blog/cloudresearch-is-retiring-the-block-low-quality-participants-option/) In particular, CloudResearch removed their "Block Low Quality Participants" which is what we have used in all prior experiments. This change resulted in data quality issues as, at the time we ran the study, we could not filter out inattentive participants/BOTs as well as before.

[3]We also recruited advisees ($N = 28$ for the main task and $N = 28$ for the MPL task that measured moral costs) for the $N = 283$ inattentive participants who swicthed multiple times in the MPL.

Table D.1: Advisor Recommendations - No Choice (Simultaneous)

	(1)	(2)	(3)
	Main Sample		**Including Inattentive**
See Incentive First	0.167***	0.149***	0.151***
	(0.051)	(0.049)	(0.043)
No Conflict	0.249***	0.218***	0.137**
	(0.062)	(0.060)	(0.058)
See Incentive First * No Conflict	-0.156*	-0.126	-0.109
	(0.083)	(0.081)	(0.077)
Simultaneous	0.172**	0.167**	0.120**
	(0.067)	(0.066)	(0.051)
Simultaneous X No Conflict	-0.267**	-0.219**	-0.083
	(0.104)	(0.096)	(0.083)
Incentive for B	-0.149***	-0.155***	-0.163***
	(0.037)	(0.035)	(0.031)
Selfishness		0.134***	
		(0.018)	
Constant	0.745***	0.730***	0.745***
	(0.077)	(0.074)	(0.067)
Observations	574	574	883
R^2	0.069	0.161	0.053

Note: This table displays the estimated coefficients from linear probability models on the advisors' recommendations. See Incentive first is a binary indicator coded as 1 for participants who were randomly assigned to see the incentive first. Simultaneous is a binary indicator coded as 1 for participants who saw all infornation at the same time. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. The regression models in columns (1) and (2) restrict the analyses to participants who did not switch multiple times in the MPL. Column (3) includes all participants. The regression includes individual controls for the advisor's gender and age, and a binary indicator for the wave in which participants took part in the experiment. Robust standard errors (HC3) in parentheses

E The Choice Deterministic Experiment

E.1 Experimental Design

The goal of this experiment is to establish whether the behavior of participants in the Choice experiment is affected by our design choice in the main experiment of assigning individuals to their preferred information order with 75% chance. While this design choice allowed us to separate selection from the effect of actually getting flexibility or commitment, it is possible that the presence of uncertainty may have provided participants with an additional excuse to behave self-servingly, affecting both information preferences and subsequent behavior.

In the Choice Deterministic experiment, we replicate the Choice Free treatment from the Choice Experiment and randomly assign participants to one of two treatments that vary whether assignment to the preferred information order occur with 75% chance as in the original experiment, or is certain. We then add the *ChoiceFree-Deterministic* treatment in which advisors know that they will receive information in their desired order with certainty. Comparing these two treatments allows us to understand whether the presence of uncertainty with respect to how advisors received information, conditional on preferring a given order, affected their recommendation behavior.

Procedures. The experiment was conducted on Amazon Mechanical Turk (AMT) and was pre-registered on aspredicted.org (#82298). Participants received $1 payment for taking part in the experiment and for making a recommendation. On top of that, they received a $.15 commission for recommending either Product A or Product B. Advisors were informed that one out of 10 advisors would be matched with a client, another AMT participant, and their advice was delivered to them. In the Choice Free-Probabilistic experiment, the procedures were identical to those in the Choice Free treatment of the Choice Experiment. In particular, advisors knew that there was a 75% chance that their preference would be implemented. After making the choice, advisors learned whether their choice was implemented, and then proceeded to see either the commission followed by the signal, or the signal followed by the commission, with the order depending on whether their choice was implemented. In the Choice Free-Deterministic experiment, participants were not told that there was a 75% chance that their preference would be implemented. Instead, upon making their choice, advisors proceeded to receive information in their desired order. Upon making their recommendations, we collected additional measures of beliefs and morality using a the same measures used in the Choice Experiment. At the end of the experiments, we randomly selected 1 out of 10 advisors and sent their recommendation to a client.

E.2 Results

A total of 952 attentive participants completed the experiment; of these 198 participants (20.8%) gave inconsistent responses to the multiple price list measure of morality. Following our pre-registration, we conduct analyses excluding these participants as our main analyses ($N = 385$ participants in the Choice Free-Replication treatment and $N = 36$ in the Choice Free-Deterministic treatment), and also including these participants.

We find that 59.7% of advisors demand to see the incentive first in the Choice Free-Probabilistic treatment and 53.4% in the Choice Free - Deterministic treatment. The decrease in demand is marginally significant (χ^2-test= 3.09, $p = 0.078$), and consistent with prior work showing that self-serving behavior increases when there is uncertainty (e.g., Haisley and Weber, 2010; Exley, 2015). Including inattentive participants, demand to see the incentive first is 62.9% and 58.7%, respectively, and the difference is not significant (χ^2-test= 1.66, $p = 0.197$).

Figure E.1. shows advisors' recommendation decisions, when there is a conflict between the signal of quality and the advisor's incentive. The difference in recommendations, depending on advisors' preferences (and when assigned their preference) is similar in the Choice Free-Deterministic and the Choice Free-Probabilistic treatments. A total of 70% of the $N = 223$ participants who preferred and got assigned to see the incentive first recommended the incentivized option in the Choice Free-Replication treatment; this fraction was 76% (out of $N = 281$ participants) in the Choice Free - Deterministic treatment. For those who preferred and got assigned to assess quality first, 55% (out of $N = 132$) and 53% out of $N = 197$) of participants recommended the incentivized option. Further, the figure shows that only 52% (out of $N = 75$) participants who preferred but were not assigned to see the incentive first in the Choice Free -Replication treatment recommended the incentivized option; and 52% (out of $N = 33$) participants who preferred but did not get assigned to see quality first recommended the incentivized option.

Advisors who prefer to see the incentive first (and are assigned their preferred information order) are, on average, 21 percentage points more likely to recommend the incentivized product, as shown in Table E.1. Interactions between the Deterministic treatment and preferences as well as the presence of conflict are not significant. Hence, the results show that recommendation decisions are robust to the probabilistic implementation of advisors' preferences for information order.

At the end of the experiment, we recruited $N = 75$ clients and matched them with 1 out of 10 advisors for the main task; of these 87% followed the recommendation. Advisors were also matched with $N = 75$ additional advisees for the MPL task that measured moral costs; of these 84% followed the recommendation.

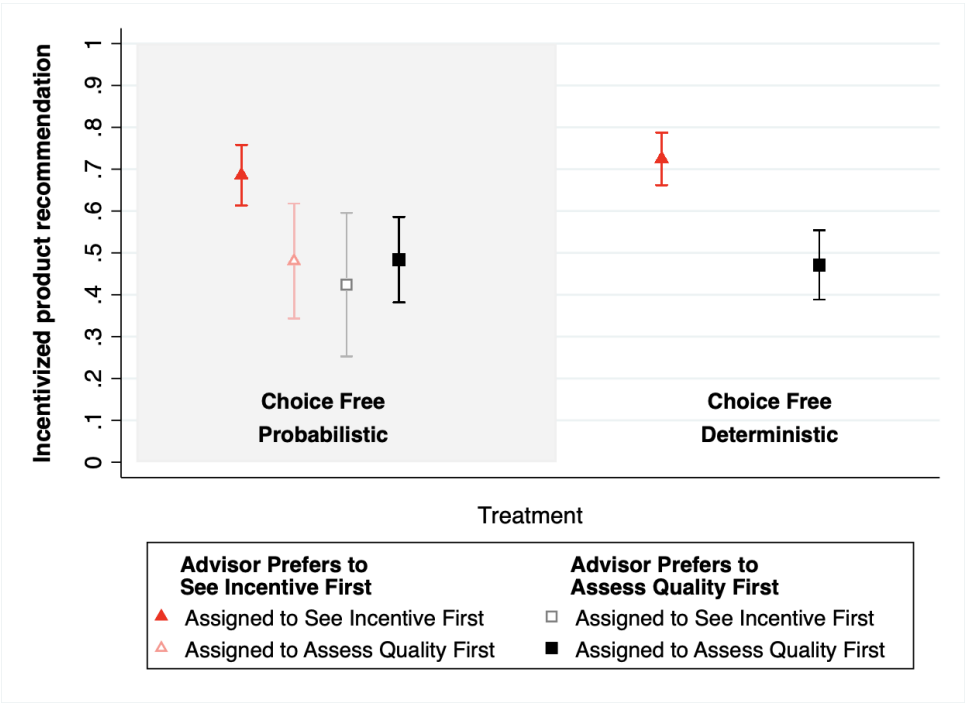

Notes: This figure presents the covariate-adjusted recommendations of the incentivized product when there is a conflict between the signal of quality and the advisor's incentive

Figure E.1: Advisors' Recommendations

Table E.1: Recommendations: Assigned Preferences

	(1)	(2)	(3)
	Main Sample		**Including Inattentve**
Prefer to See Incentive First	0.211***	0.191***	0.200***
	(0.069)	(0.066)	(0.064)
No Conflict	0.253***	0.263***	0.253***
	(0.095)	(0.093)	(0.091)
Prefer to See Incentive First * No Conflict	-0.099	-0.083	-0.153
	(0.120)	(0.115)	(0.113)
Deterministic	-0.015	-0.017	-0.012
	(0.071)	(0.069)	(0.067)
Deterministic X No Conflict	0.040	0.041	-0.016
	(0.124)	(0.123)	(0.118)
Deterministic X Prefer to See Incentive First	0.102	0.113	0.053
	(0.090)	(0.087)	(0.083)
Deterministic X Prefer to See Incentive First x No Conflict	-0.089	-0.140	0.053
	(0.154)	(0.151)	(0.145)
Incentive for B	-0.145***	-0.150***	-0.118***
	(0.036)	(0.034)	(0.033)
Female	-0.016	-0.005	0.014
	(0.036)	(0.035)	(0.032)
Age	-0.002	-0.001	-0.001
	(0.001)	(0.001)	(0.001)
Selfishness		0.106***	
		(0.018)	
Constant	0.616***	0.603***	0.582***
	(0.082)	(0.079)	(0.076)
Observations	656	656	832
R^2	0.113	0.163	0.080

Note: This table displays the estimated coefficients from linear probability models on the advisors' recommendations. Deterministic is a binary indicator coded as 1 for participants in the Deterministic treatment. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. The regression model in column (3) extends the analyses to included advisors who switched multiple times in the multiple price list eliciting selfishness. The regression includes individual controls for the advisor's gender and age. Robust standard errors (HC3) in parentheses

F Additional Data: Predictions

In the Prediction experiment, we recruited forecasters on AMT and asked them to read a summary description of the recommendation decisions advisors made in the Incentive First Costly treatment of the Choice Experiment. A link to the original instruction was provided to participants. We then asked forecasters to predict the recommendation decisions of advisors who choose to see the incentive first. In particular, forecasters were told to consider the recommendation decisions of advisors who chose to see their incentives first. They were asked to estimate the recommendations of advisors who were either assigned to see the incentives first or assigned to assess quality first.

To aid participants in making their predictions, and following the approach of DellaVigna and Pope (2018), participants received information about the counterfactual—the fraction of recommendations of the incentivized product for cases in which advisors were assigned to receive information in the opposite order. For this purpose, we provided forecasters with the fraction of incentivized recommendations in the See Incentive First Costly treatment (AMT-1) of the Choice experiment. Then, we first ask forecasters to predict the direction of the effect (more, equal or fewer recommendations of the incentivized product), and then to provide their estimated fraction of recommendations. If participants anticipate that seeing the incentive first gives advisors more flexibility to provide self-serving recommendations, then we would expect to see a positive and significant gap between the two information sequences, with participants predicting a higher fraction of recommendations of the incentivized product when advisors see their incentive first. Forecasters were paid \$1 and received an additional \$2 bonus if their predictions laid within 5 percentage points of the true value.

In order to interpret advisors' preferences to see their incentive first or, on the contrary, assess quality first, as evidence that individuals actively pursue or constrain cognitive flexibility, it is important to test whether individuals anticipate that the order of information will affect their recommendations. To investigate this question, we turn to the Prediction experiment, in which a group of forecasters predicted the difference in recommendations between the two information orders for the case in which seeing the incentive first is costly.

Figure F.1 shows the cumulative distribution function of forecasts, as well as the average predicted effect and the average actual effects of seeing the incentive first. The predicted effect of seeing the incentive first—relative to seeing quality first—is 6.2 percentage points (SE=0.13, $N = 288$). This is significantly different from zero ($t - stat = 4.76$, p <0.001). It is not significantly different from the actual effect of 11.36 percentage points ($t - stat = 1.34$, $p = 0.181$), which we documented in the See Incentive First

Costly (AMT-1) of the Choice experiment. As shown in Figure F.1, the majority of participants expect a positive effect of seeing the incentive first (51.4%), while 24.0% predict no effect and 24.7% predict a negative effect.

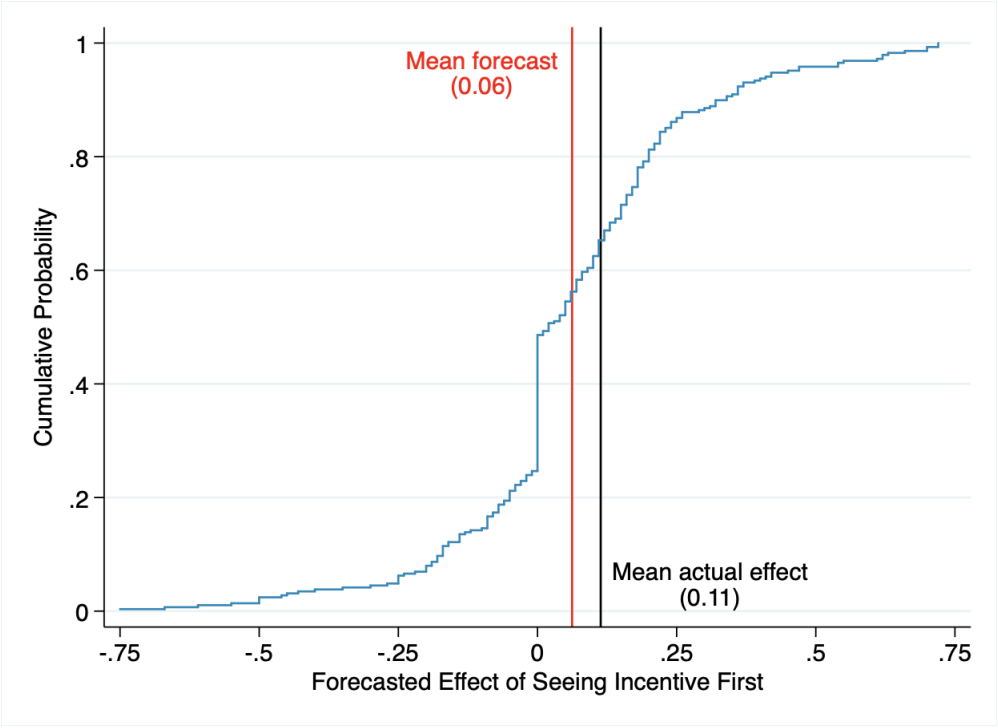

Notes: This figure displays the distribution of forecasts regarding the effect of seeing the incentive first on recommendations of the incentivized product in the Choice experiment, for advisors who prefer to see the incentive first when seeing it is costly, the average forecast (from the Predictions experiment) and the average actual effect (from the Choice experiment).

Figure F.1: Predicted and Actual Effect of Seeing the Incentive First on Recommendations

This experiment therefore provides some evidence that individuals evaluating the task of advisors can anticipate the effects of seeing the incentive first, although on average they may somewhat underestimate the magnitude of those effects. This result is consistent with the interpretation that the choice to see the incentive first or assess quality first is at least in part driven by the anticipated effect of this information order on recommendations.

G Experimental Instructions

Below we present instructions for the Choice experiment and the IA experiment.

G.1 Choice Experiment

Below we present the screenshots that advisors were presented with in the Choice experiment.

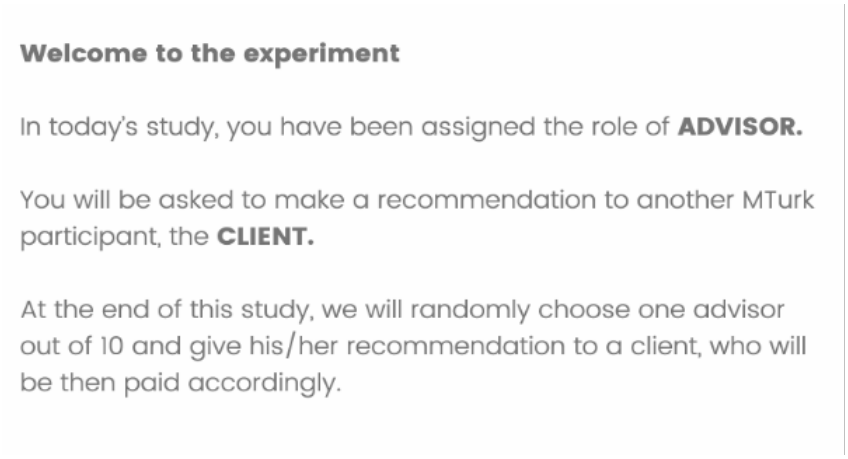

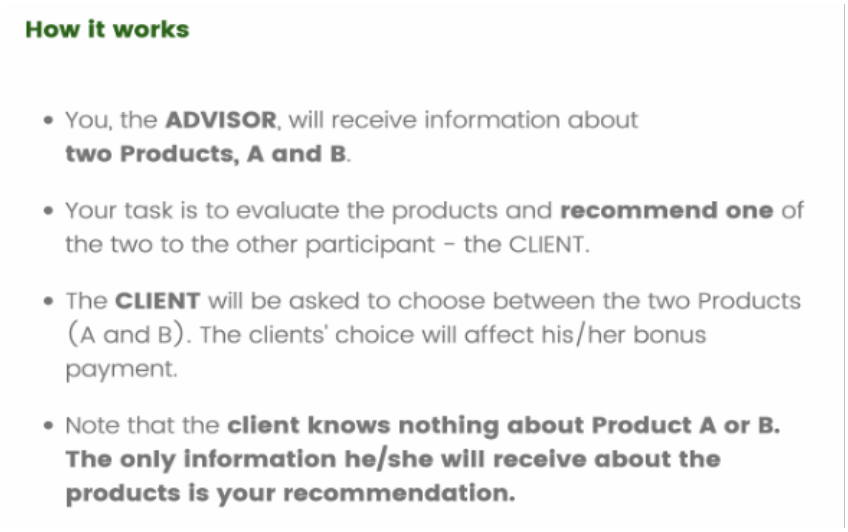

The Products

Product A and Product B are urns containing 5 payoff balls each. The payoff balls are either blue or red. Blue balls are worth $2 and red balls are worth $0. The combination of balls is different for the two Products, as described below:

- **Product A** is an urn with **3 blue ($2)** balls and **2 red ($0)** balls

- **Product B** is an urn that is either **high or low quality**, with equal chance.
  - If the Urn is **high quality** (50% chance), it has **four blue ($2) balls** (more than Product A).
  - If the Urn is **low quality** (50% chance), it has only **two blue ($2) balls** (fewer than Product A).
  - The quality of the urn was determined by the computer at random. **You do not know whether Product B is high or low quality for sure, but will soon receive information that will help you infer the quality of Product B.**

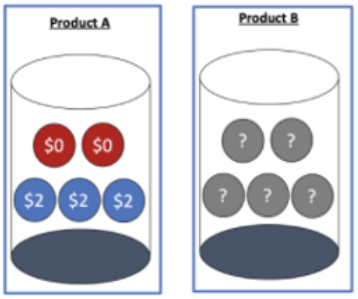

After receiving your recommendation, the client will choose one product between A and B. He/she will then randomly draw one ball from the urn. The payoff ball he/she draws will determine his/her payment.

Before you proceed, just a few questions to help you go over the instructions.

1. How much is a **red ball** worth?

- ○ $0.15
- ○ $1
- ○ $0
- ○ $2

2. How much is a **blue ball** worth?

- ○ $0.15
- ○ $1
- ○ $0
- ○ $2

3. How many **blue balls** are there in Product A?

- ○ 3 out of 5 balls are blue
- ○ 2 out of 5 balls are blue
- ○ 5 out of 5 balls are blue
- ○ 1 out of 5 balls is blue

4. The **quality of Product B** is **high** with...

- ○ 75% chance
- ○ 25% chance
- ○ 30% chance
- ○ 50% chance

5. Which of the following statements is correct? **Product B...**

- ○ ...is an urn with **4 blue balls ($2)** and **1 red ball ($0)** if its quality is **HIGH**, and it is an urn with **2 blue balls ($2)** and **3 red balls ($0)** if its quality is **LOW**
- ○ ...is an urn with **3 blue balls ($2)**and **2 red balls ($0)** if its quality is **HIGH**, and it is an urn with **3 blue balls ($2)** and **2 red balls ($0)** if its quality is **LOW**
- ○ ...is an urn with **5 blue balls ($2)** and **0 red balls ($0)** if its quality is **HIGH**, and it is an urn with **0 blue balls ($2)** and **5 red balls ($0)** if its quality is **LOW**

Before you proceed, **make sure you read these instructions carefully**. On the next screen, there will be one more question to verify that you paid attention. If you don't answer that question correctly, you will not be eligible to receive a bonus for this study.

A Question for You

Before proceeding with your task, please answer the question below.

Imagine the client chooses **Product A**. What is the **chance** he/she gets **$2** (a **blue** ball)?

- ○ 1 in 5, because 1 out of 5 balls in Product A is **blue** ($2)
- ○ 2 in 5, because 2 out of 5 balls in Product A are **blue** ($2)
- ○ 3 in 5, because 3 out of 5 balls in Product A are **blue** ($2)
- ○ 5 in 5, because 5 out of 5 balls in Product A are **blue** ($2)

What You Know

- **You will soon receive more information that will help you gain some insights on the quality of Product B.**

- **The client does not know anything about Product A and B.** He/she will choose a Product after receiving your recommendation. The computer will then randomly draw a ball from the Product chosen by the advisor. The advisor will then will be paid accordingly.

Advisor's choice in See Quality First Costly (adjusted accordingly for Choice Free and See Incentive First Costly).

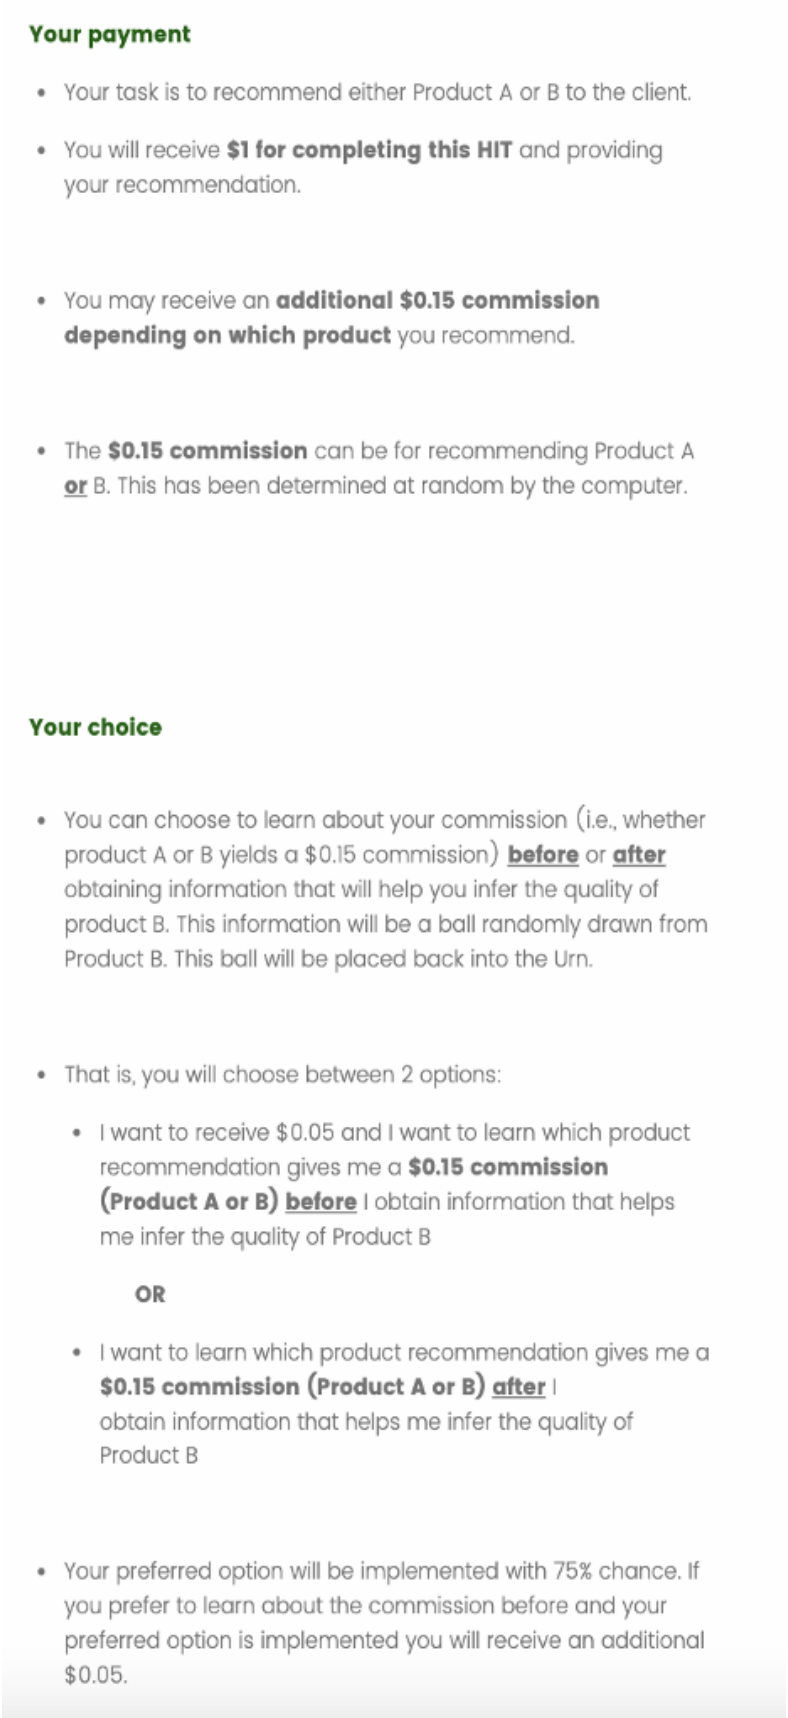

Your payment

- Your task is to recommend either Product A or B to the client.
- You will receive **$1 for completing this HIT** and providing your recommendation.
- You may receive an **additional $0.15 commission depending on which product** you recommend.
- The **$0.15 commission** can be for recommending Product A **or** B. This has been determined at random by the computer.

Your choice

- You can choose to learn about your commission (i.e., whether product A or B yields a $0.15 commission) **before** or **after** obtaining information that will help you infer the quality of product B. This information will be a ball randomly drawn from Product B. This ball will be placed back into the Urn.
- That is, you will choose between 2 options:
  - I want to receive $0.05 and I want to learn which product recommendation gives me a **$0.15 commission (Product A or B) before** I obtain information that helps me infer the quality of Product B

    OR

  - I want to learn which product recommendation gives me a **$0.15 commission (Product A or B) after** I obtain information that helps me infer the quality of Product B
- Your preferred option will be implemented with 75% chance. If you prefer to learn about the commission before and your preferred option is implemented you will receive an additional $0.05.

Choice screen

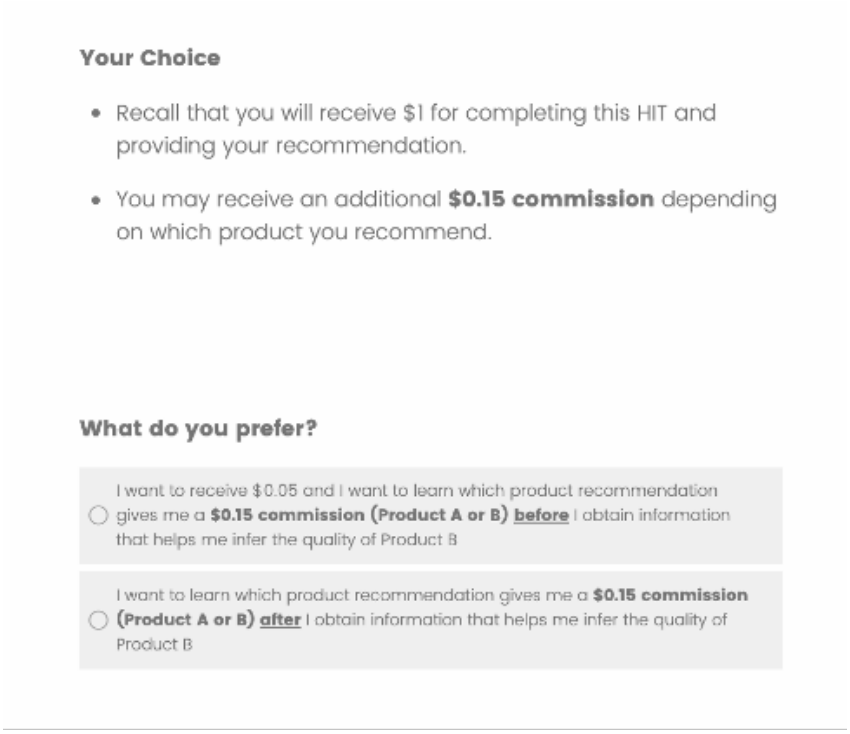

Assignment screen (adjusted accordingly depending on the assignment)

Information for advisors who see the incentive after (incentive shown earlier if assignment is to before)

Next, you will obtain information that will help you infer the quality of Product B.

As a reminder, you know what Product A is. Instead, you don't know whether Product B is the low or high-quality Urn. Product B could have either High or Low quality with equal chance. The combination of blue and red balls for both cases is depicted in the picture below. The quality of Product B was determined at random by the computer at the beginning of the study.

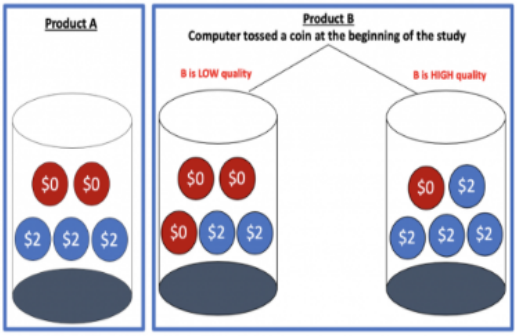

- **On the next screen, you will gain some insights on the quality of Product B. That is, we will randomly draw a payoff ball from Product B and display it on the next screen.**

- **After seeing the ball, you will be asked to choose which product, A or B, to recommend to the CLIENT.**

Quality signal

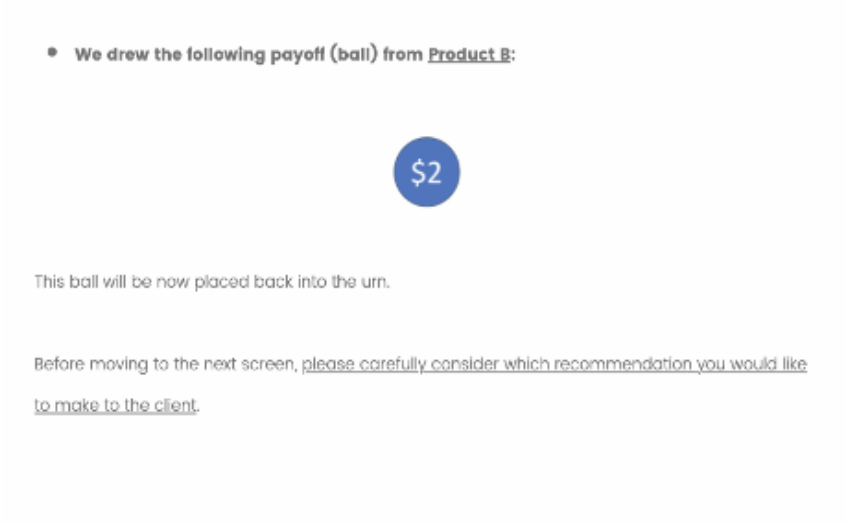

Recommendation decision

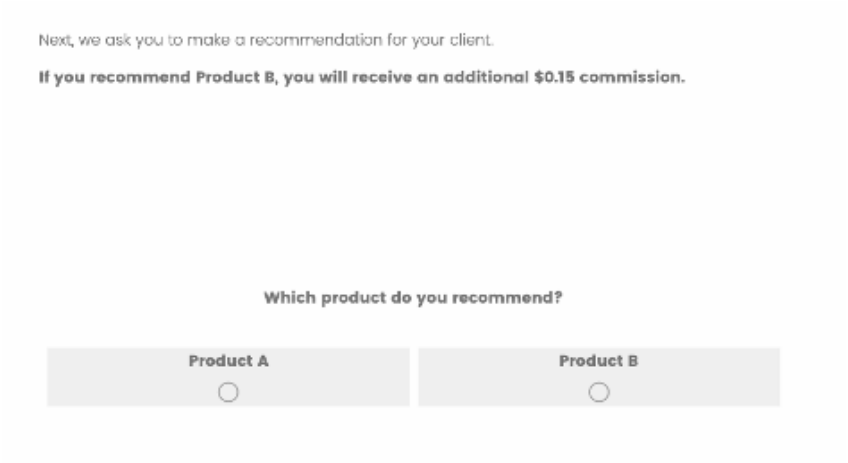

Additional measures of advisors' beliefs and preferences

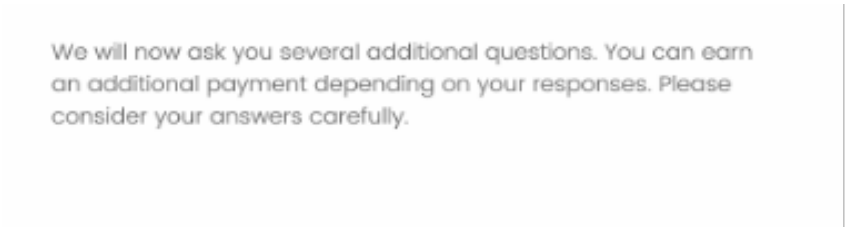

After observing the ball, what do you think is the likelihood that the quality of Product B is LOW? *(0% means that Product B is extremely unlikely to be of LOW quality, whereas 100% means that Product B is extremely likely to be of LOW quality)*

Please indicate your estimated likelihood by choosing one of the ten options below.

Bonus payment. There are 10 options. If your answer is correct (i.e., if you choose the right range), you will receive an additional **$0.15 payment.**
To get the bonus, consider your answer carefully.

- ○ The likelihood that **Product B** is of low quality is between **0% and 10%**
- ○ The likelihood that **Product B** is of low quality is between **11% and 20%**
- ○ The likelihood that **Product B** is of low quality is between **21% and 30%**
- ○ The likelihood that **Product B** is of low quality is between **31% and 40%**
- ○ The likelihood that **Product B** is of low quality is between **41% and 50%**
- ○ The likelihood that **Product B** is of low quality is between **51% and 60%**
- ○ The likelihood that **Product B** is of low quality is between **61% and 70%**
- ○ The likelihood that **Product B** is of low quality is between **71% and 80%**
- ○ The likelihood that **Product B** is of low quality is between **81% and 90%**
- ○ The likelihood that **Product B** is of low quality is between **91% and 100%**

Now that you have chosen a bin, what do you believe is the **exact likelihood**? Please enter a number from 0 to 100.

Next, you will complete 2 additional tasks that will ask you to make two additional recommendations. Each of the 2 tasks will have specific instructions and payments. Please read them carefully.

Moral costs

Task #1

- Next, we will show you **5 new sets of Products**.
- For each set, we will ask you to recommend one of two products, X or Y, to an **advisee**.
- The **advisee** is a **different MTurk Worker** than the CLIENT (for whom you made your earlier recommendation).
- At the end of the study, we will randomly select 1 of the 5 sets of products and 1 out of 10 advisors to send the advisee the corresponding recommendation.
- When the advisee receives your recommendation, he/she will be asked to choose between Product A and Product B, without having any information other than your recommendation.

In each of the following 5 sets, **Product X** will change whereas **Product Y** will always stay the same.

In decisions 1 through 5, if you recommend Product Y you will receive an additional $0.15 payment.

1. If you recommend Product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

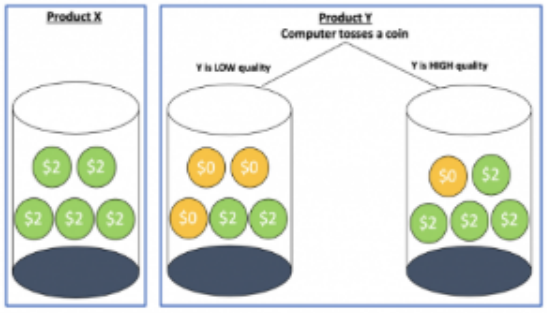

We drew the following payoff (ball) $0 **from Product Y.**

Which product do you recommend to the advisee?

- ○ Product X
- ○ Product Y

2. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

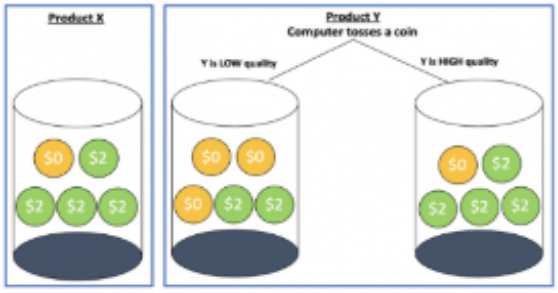

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

3. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

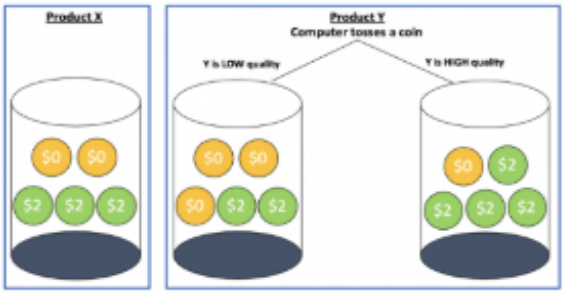

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

4. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

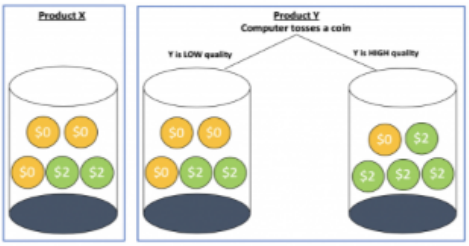

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

5. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

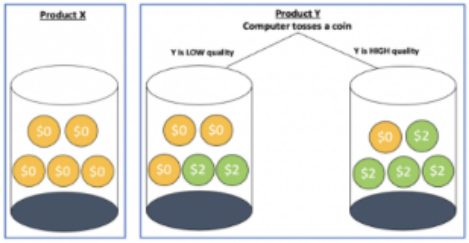

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

Blinding

Task #2.

- For this task, you again will be asked to make a recommendation between two products to another participant, **an ADVISEE**.
- This ADVISEE (another MTurk worker) is different from the one in TASK 1. **At the end of the study, we will randomly select 1 out of 10 advisors to send the advisee the corresponding recommendation.**

The two products, **1 and 2,** are displayed below.

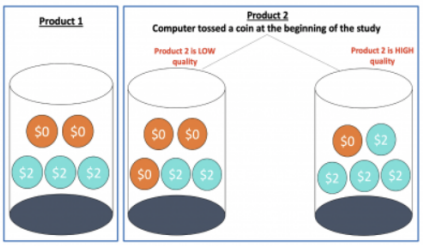

- Product 2 can have **LOW or HIGH quality**, with 50% chance but you don't know if the quality is LOW or HIGH.

To help with your recommendation, you will receive information about the quality of product 2.

- That is, you will be shown a ball ($0 or $2), randomly drawn from Product 2, **which will help you infer its quality**. The ball will then be placed in the urn and you willl be asked to make your recommendation.

◯ I understand that I will soon receive information about the quality of Product 2

You could receive a **$0.15 commission**, depending on your recommendation.

- The commission can be for **product 1 OR product 2**, determined at random, and you will learn it before the end of the HIT.

You have the option to learn whether the $0.15 commission is for product 1 or product 2 **after** you have made your recommendation or **before** making your reccomendation.

- **If you choose to see the commission after,** in the next screen you will learn about the quality of Product 2 and will then be asked to make your recommendation. In the following screen, you will learn the commission.
- **If you choose to see the commission before**, in the next screen you will learn about the quality of Product 2 and whether the commission is for product 1 or 2. You will then be asked to make your recommendation.

Which option do you prefer?

○ **Learn whether the commission is for product 1 or 2 AFTER making my recommendation**, i.e., in the next screen I will receive information about the quality of Product 2 and then make my recommendation. I will learn the commission in the following screen.

○ **Learn whether the commission is for product 1 or 2 BEFORE making my recommendation**, i.e., in the next screen I will receive information about the quality of Product 2, about whether the commission is for Product 1 or 2, and then make my recommendation.

Recommendation decision screen if blinded

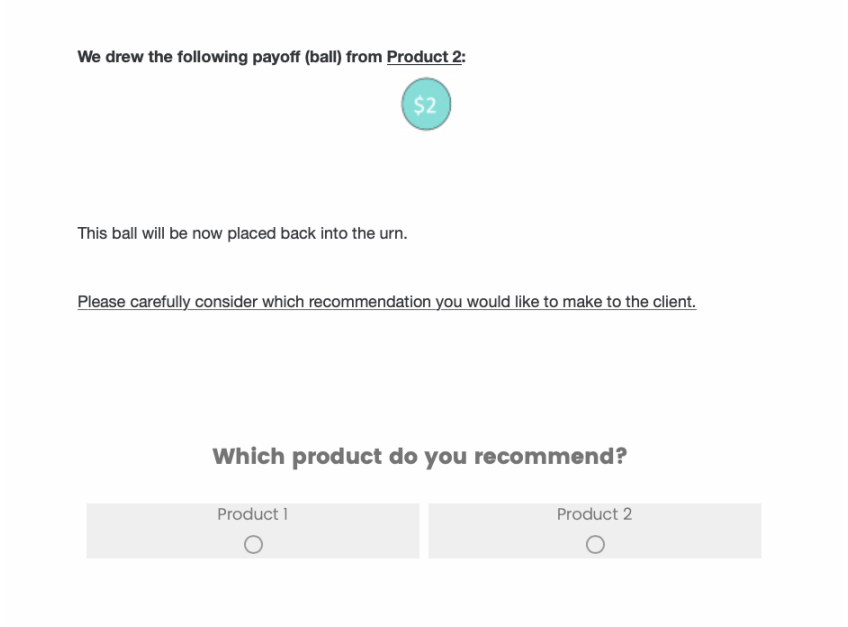

Incentive information screen (shown after recommendation for advisors who chose to blind)

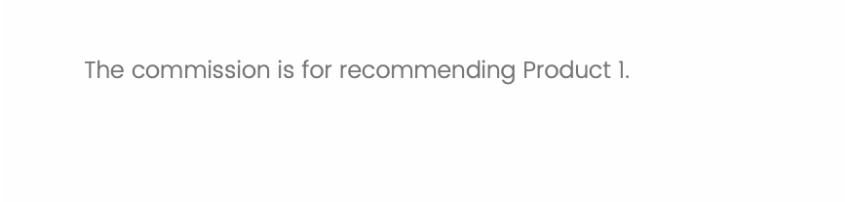

G.2 Information Architect experiment

Welcome to the experiment

In today's study, you have been assigned the role of **DECISION-MAKER.**

There are two other participants in the role of **ADVISOR** and **CLIENT**.

The **ADVISOR** will be asked to make a recommendation to the **CLIENT.** At the end of this study, we will randomly choose 1 advisor out of 10 to give his/her recommendation to a client, who will be then paid accordingly.

Your decision today may affect the ADVISOR's recommendation and the CLIENT's payment. Your task today is to protect the CLIENT's interest.

How it works

- The **ADVISOR** will receive information about **two Products, A and B**.
- The **ADVISOR's** task is to evaluate the products and **recommend one** of the two to the other participant - the CLIENT.
- The **CLIENT** will be asked to choose between the two Products (A and B). The clients' choice will affect his/her bonus payment.
- Note that the **CLIENT knows nothing about Product A or B. The only information he/she will receive about the products is the ADVISOR's recommendation.**

- **You**, the **DECISION-MAKER**, can **decide how an ADVISOR will receive information that helps them evaluate the two Products**, A and B.

Understanding Question: What is your role and that of others in this study?

- ◯ I am the DECISION-MAKER, and in my role I provide advice to another participant in the role of CLIENT.
- ◯ I am the DECISION-MAKER, and in my role I assign a product to the ADVISOR (Product A or B)
- ◯ I am the DECISION-MAKER, and in my role I will make a decision that may affect the ADVISOR's recommendation and the CLIENT's payment

The Products

Product A and Product B are urns containing 5 payoff balls each. The payoff balls are either blue or red. Blue balls are worth $2 and red balls are worth $0. The combination of balls is different for the two Products, as described below:

- **Product A** is an urn with **3 blue ($2)** balls and **2 red ($0)** balls

- **Product B** is an urn that is either **high or low quality**, with equal chance.
  - If the Urn is **high quality** (50% chance), it has **four blue ($2) balls** (more than Product A).
  - If the Urn is **low quality** (50% chance), it has only **two blue ($2) balls** (fewer than Product A).
  - The quality of the urn was determined by the computer at random. The advisor does not know whether Product B is high or low quality for sure, **but will receive information that will help him/her infer the quality of Product B**.

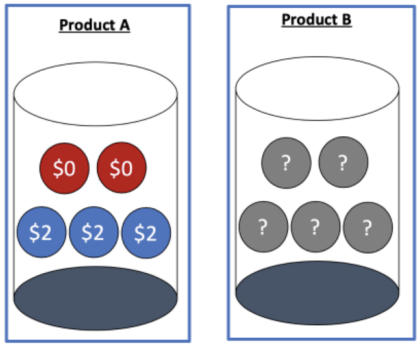

After receiving the advisor's recommendation, the client will choose one product between A and B. He/she will then randomly draw one ball from the urn. The payoff ball he/she draws will determine his/her payment.

Before you proceed, just a few questions to help you go over the instructions.

1. How much is a **red ball** worth?

- ◯ $0.15
- ◯ $1
- ◯ $0
- ◯ $2

2. How much is a **blue ball** worth?

- ◯ $0.15
- ◯ $1
- ◯ $0
- ◯ $2

3. How many **blue balls** are there in Product A?

- ◯ 3 out of 5 balls are blue
- ◯ 2 out of 5 balls are blue
- ◯ 5 out of 5 balls are blue
- ◯ 1 out of 5 balls is blue

4. The **quality of Product B** is **high** with...

- ◯ 75% chance
- ◯ 25% chance
- ◯ 30% chance
- ◯ 50% chance

5. Which of the following statements is correct? **Product B...**

- ◯ ...is an urn with **4 blue balls ($2)** and **1 red ball ($0)** if its quality is **HIGH**, and it is an urn with **2 blue balls ($2)** and **3 red balls ($0)** if its quality is **LOW**
- ◯ ...is an urn with **3 blue balls ($2)**and **2 red balls ($0)** if its quality is **HIGH**, and it is an urn with **3 blue balls ($2)** and **2 red balls ($0)** if its quality is **LOW**
- ◯ ...is an urn with **5 blue balls ($2)** and **0 red balls ($0)** if its quality is **HIGH**, and it is an urn with **0 blue balls ($2)** and **5 red balls ($0)** if its quality is **LOW**

Before you proceed, **make sure you read these instructions carefully**. On the next screen, there will be one more question to verify that you paid attention. If you don't answer that question correctly, you will not be eligible to receive a bonus for this study.

A Question for You

Before proceeding with your task, please answer the question below.

Imagine the client chooses **Product A**. What is the **chance** he/she gets **$2** (a **blue** ball)?

- ◯ 1 in 5, because 1 out of 5 balls in Product A is **blue** ($2)
- ◯ 2 in 5, because 2 out of 5 balls in Product A are **blue** ($2)
- ◯ 3 in 5, because 3 out of 5 balls in Product A are **blue** ($2)
- ◯ 5 in 5, because 5 out of 5 balls in Product A are **blue** ($2)

What You Know

- **The advisor will soon receive more information that will help them gain some insights on the quality of Product B.**
- **The client does not know anything about Product A and B.** He/she will choose a Product after receiving your recommendation. The computer will then randomly draw a ball from the Product chosen by the advisor. The client will then will be paid accordingly.

Information for IAs in the IA-client treatment (adjusted accordingly for IA-advisor).

Advisor's payment

- The ADVISOR's task is to recommend either Product A or B to the client.
- In addition to their fixed fee for completing this HIT, the advisor may receive an additional payment depending on their decisions.
- **Depending on which product the advisor recommends, they may receive an additional $0.15 commission as a bonus payment**.
- The **$0.15 commission** can be for recommending Product A **or** B. This has been determined at random by the computer.

Client's payment

- The CLIENT will receive a product recommendation from the ADVISOR.
- **After receiving the advisor's recommendation, the client will choose one product between A and B.** He/she will then randomly draw one ball from the urn. The payoff ball he/she draws will determine his/her payment.

Your payment

- You will be paid $1 for completing the study.
- You will receive **an additional $0.15 payment if the advisor recommends the best product for the CLIENT.**
- That is, if the advisor recommends **the product with more $2 balls,** you receive **a $0.15 payment.**

Your choice

As Decision-Maker, you can decide how the advisor will receive information.

- You can choose for the adviosr to learn about his/her commission (i.e., whether product A or B yields a $0.15 commission) **before** or **after** obtaining information that will help the advisor infer the quality of Product B. This information will be a ball randomly drawn from Product B. This ball will be placed back into the Urn.

- That is, you will choose between 2 options:
  - I want the advisor to learn which product recommendation gives them (the advisor) a **$0.15 commission (Product A or B) before** the advisor obtains information that helps them infer the quality of Product B

    OR

  - I want the advisor to learn which product recommendation gives them (the advisor) a **$0.15 commission (Product A or B) after** the advisor obtains information that helps them infer the quality of Product B

- Your preferred option will be implemented with 75% chance.

Question: What is your payment today?

- ◯ I am paid a $1 fixed fee.
- ◯ I am paid a $1 fixed fee. On top of that, I receive an additional $0.15 if the advisor recommends the product that yields them (the advisor) the commission.
- ◯ I am paid a $1 fixed fee. On top of that, I receive an additional $0.15 if the advisor recommends the best product to the client (i.e, the product with more $2 balls)

Your Choice

- Recall that you will receive $1 for completing this HIT and making this choice.
- You will receive **an additional $0.15 payment if the advisor recommends the best product to the client** (i.e., the product with more $2 balls).

What do you prefer?

- ◯ I want the advisor to learn which product recommendation gives them (the advisor) a **$0.15 commission (Product A or B) before** the advisor obtains information that helps them infer the quality of Product B
- ◯ I want the advisor to learn which product recommendation gives them (the advisor) a **$0.15 commission (Product A or B) after** the advisor obtains information that helps them infer the quality of Product B

Next, you will complete an additional task that will ask you to make some recommendations. The task will have specific instructions and payments. Please read them carefully.

Extra Task

- Next, we will show you **5 new sets of Products**.
- For each set, we will ask you to recommend one of two products, X or Y, to an **advisee**. That is, you will be an advisor.
- The **advisee** is a **different MTurk Worker** than the ADVISOR and CLIENT (for whom you made a decision earlier).
- At the end of the study, we will randomly select 1 of the 5 sets of products and 1 out of 10 participants to send the advisee the corresponding recommendation.
- When the advisee receives your recommendation, he/she will be asked to choose between Product A and Product B, without having any information other than your recommendation.

In each of the following 5 sets, **Product X** will change whereas **Product Y** will always stay the same.

In decisions 1 through 5, if you recommend Product Y you will receive an additional $0.15 payment.

1. If you recommend Product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

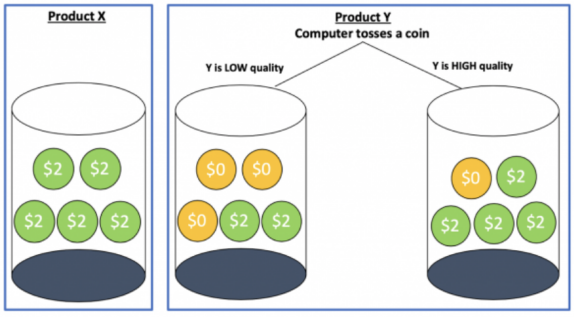

We drew the following payoff (ball)  **from Product Y.**

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

2. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

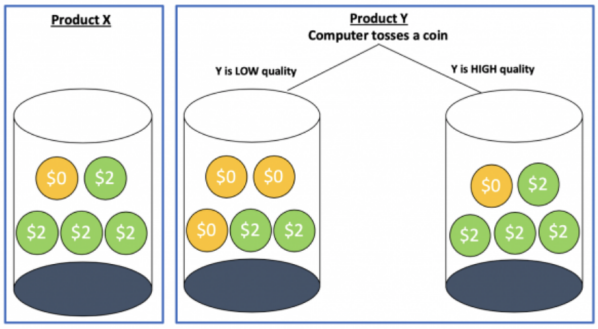

We drew the following payoff (ball) $0 **from Product Y.**

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

3. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

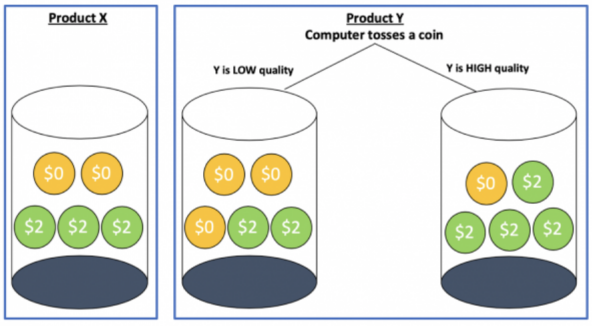

We drew the following payoff (ball) $0 **from Product Y.**

Which product do you recommend to the advisee?

- ◯ Product X
- ◯ Product Y

4. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

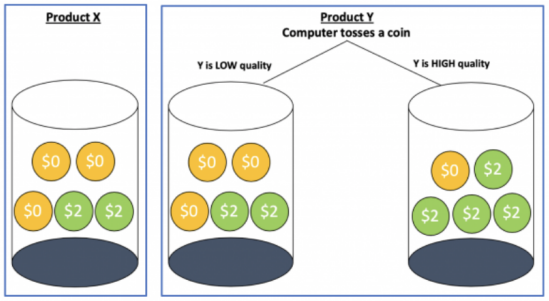

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ○ Product X
- ○ Product Y

5. If you recommend product Y for the randomly selected set that is sent to the advisee, you will receive $0.15.

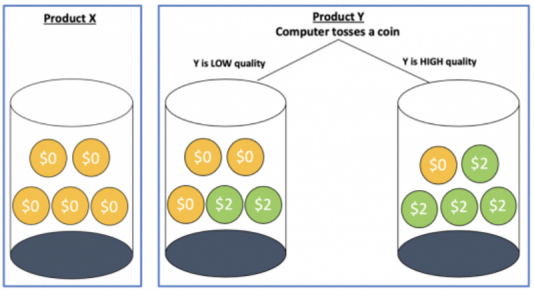

We drew the following payoff (ball) $0 from Product Y.

Which product do you recommend to the advisee?

- ○ Product X
- ○ Product Y

H Corrigendum, September 29, 2023

The authors were made aware of a coding error on August 16, 2023, reported by Hsu et al. (2023), which computationally reproduced this paper. The coding error affects the variable *Selfishness* and the report shows that "After correcting the erroneous coding, we find stronger support for the authors' main conclusion regarding *Selfishness* driving incentive information avoidance with double effect size." The coding error does not qualitatively affect any of the results reported in the paper. In the main text, correcting the coding error results only affects columns 2 and 3 of Table 2. Table H.2 reports the corrected version of Table 2, which is reproduced here as Table H.1.

Differences between Table H.2 (corrected Table 2) and Table H.1. (original Table 2). As can be seen from comparing Table H.1 and Table H.2, the main consequence of the error is that the coefficient for the variable *Selfishness* increases in size. The coefficient is 0.028 and 0.039 in columns (2) and (3), respectively, of the Original Table 2 (Table H.1). It becomes 0.056 and 0.075 in the Corrected Table 2 (Table H.2).

The main text in the paper (p.416) states "Table 2 shows the determinants of the preference to see the incentive first, and columns 2 and 3 investigate its relationship with advisor selfishness. In line with Hypothesis 2 (i), advisors who make more selfish choices in the task designed to measure advisors' moral costs prefer to see the incentive first significantly more often." This text remains correct and the conclusion is unchanged.

Implications for additional results reported in the Online Appendix. The variable *Selfishness* is included as a control variable in seven tables of the Online Appendix: Tables C.1, C.5, C.6, C.18, C.22, D.1, E.1. In all cases, there is no qualitative change in the interpretation of the main results. As for the variable *Selfishness*, the significance of the coefficients remains unchanged, except for Table C.22 where *Selfishness* is no longer marginally significant. In this updated appendix, we have corrected the tables in the text above. Below we detail the list of changes for each table.

Table H.1: Original Table 2 - Preference for Information Order

	(1)	(2)	(3)
	Prefer to See Incentive First		
See Incentive First Costly	-0.139***	-0.140***	-0.140***
	(0.018)	(0.018)	(0.018)
Assess Quality First Costly	0.152***	0.152***	0.152***
	(0.029)	(0.029)	(0.029)
Choice Free – Professionals	-0.095***		
	(0.026)		
Selfishness		0.028***	0.039***
		(0.007)	(0.009)
See Incentive First Costly X Selfishness			-0.022
			(0.016)
See Quality First Costly X Selfishness			-0.021
			(0.018)
Female	-0.029**	-0.024*	-0.023*
	(0.013)	(0.014)	(0.014)
Age	-0.003***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)
Constant	0.674***	0.662***	0.661***
	(0.024)	(0.025)	(0.025)
Observations	5908	5196	5196
R^2	0.034	0.040	0.040

Notes: This table displays the estimated coefficients from linear probability models on the preference to see the incentive first. See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. The regression models in columns (2) and (3) include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Table H.2: Corrected Table 2 - Preference for Information Order

	(1)	(2)	(3)
	Prefer to See Incentive First		
See Incentive First Costly	-0.139***	-0.141***	-0.142***
	(0.018)	(0.018)	(0.018)
Assess Quality First Costly	0.152***	0.150***	0.150***
	(0.029)	(0.029)	(0.029)
Choice Free – Professionals	-0.095***		
	(0.026)		
Selfishness		0.056***	0.075***
		(0.007)	(0.010)
See Incentive First Costly X Selfishness			-0.049***
			(0.016)
See Quality First Costly X Selfishness			-0.024
			(0.018)
Female	-0.029**	-0.028**	-0.028**
	(0.013)	(0.014)	(0.014)
Age	-0.003***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)
Constant	0.674***	0.655***	0.656***
	(0.024)	(0.025)	(0.025)
Observations	5908	5196	5196
R^2	0.034	0.049	0.051

Notes: This table displays the estimated coefficients from linear probability models on the preference to see the incentive first. See Incentive First Costly and Assess Quality First Costly are indicator variables that take value 1 in the respective treatment, 0 otherwise. Selfishness was elicited at the end of the experiment, using a multiple price list (MPL) with 5 decisions. The variable is a standardized measure of the number of times the advisor chose to recommend the incentivized product in the MPL task. The regression models in columns (2) and (3) include individual controls for the advisor's gender and age, each wave of the experiment, whether incentives were probabilistic, the position of the products on the screen and the interaction between these two variables. Robust standard errors (HC3) in parentheses. * p<.10; ** p<.05; *** p<.01

Online Appendix with updated *Selfishness* variable. The specific changes that resulted are the following:

- In Table C.1. (Recommendations), the coefficients for the variable *Selfishness* increase in magnitude: 0.189 (s.e. 0.030) instead of 0.108 (s.e. 0.028) in column (1), 0.041 (s.e. 0.045) instead of -0.026 (s.e. 0.035) in column (2), and 0.148 (s.e. 0.024) instead of 0.076 (s.e. 0.028) in column (3). The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 4 decimal points).

- In Table C.5. (Preferences for Blindness, Information Order & Selfishness), the coefficients for the variable *Selfishness* become more negative. In the original Table the coefficients are: -0.055 (s.e. 0.015), -0.050 (s.e. 0.025), and -0.052 (s.e. 0.013), in columns (1), (2), and (3), respectively. In the updated Table they are: -0.095 (s.e. 0.015), -0.098 (s.e. 0.026), and -0.096 (s.e. 0.013), respectively. The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 4 decimal points), with the exception of the coefficient for *No Conflict* which becomes marginally significant, from -0.046 (s.e. 0.032) to -0.055 (s.e. 0.031).

- In Table C.6. (Advisor Recommendations), the coefficients for the variable *Selfishness* increase in magnitude: 0.100 (s.e. 0.007) instead of 0.054 (s.e. 0.006) in column (1), 0.115 (s.e. 0.012) instead of 0.036 (s.e. 0.011) in column (2), and 0.104 (s.e. 0.006) instead of 0.049 (s.e. 0.005) in column (3). The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 3 decimal points), with the exception of the coefficient for *See Incentive First Costly* which is no longer marginally significant, 0.025 (s.e. 0.017) compared to 0.033 (s.e. 0.017) before.

- In Table C.18 (Preference for Information Order: Including Incentives Treatments), the coefficients for *Selfishness* increase in magnitude: 0.06 (s.e. 0.01) instead of 0.03 (s.e. 0.01) and 0.08 (s.e. 0.01) instead of 0.04 (s.e. 0.01). The coefficient for the interaction variable *See Incentive First Costly X Selfishness* becomes statistically significant. In the corrected table it is -0.05 (s.e. 0.02), while it was previously -0.02 (s.e. 0.02). The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 3 decimal points), with the exception of the coefficient for *Female* which becomes statistically significant in columns (2) and (3), -0.03 (s.e. 0.01) in both cases, compared to -0.02 (s.e. 0.01) before.

- In Table C.22 (Preferences for Information Order in the Choice Stakes Experiment), the coefficient for *Selfishness* is no longer marginally significant. It is 0.009 (s.e.

0.012) in the corrected table, compared to 0.022 (s.e. 0.012) before. The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 4 decimal points).

- In Table D.1. (Advisor Recommendations - No Choice (Simultaneous)), the coefficient for *Selfishness* increases in magnitude. It is 0.134 (s.e. 0.018) in column (2) of the corrected table, compared to 0.083 (s.e. 0.018) before. The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 4 decimal points).

- In Table E.1. (Recommendations: Assigned Preferences) the coefficient for *Selfishness* increases in magnitude. It is 0.106 (s.e. 0.018) in column (2) of the corrected table, compared to 0.078 (s.e. 0.017) before. The significance of all other coefficients remains unchanged and the magnitude changes are small (less than 4 decimal points).

References

Hsu, D., Wang, J.T., Weng, W., and Yang, G.C. (2023). "UCSD Replication Game Report: A Comment on Saccardo and Serra-Garcia (2023)". Manuscript.