Preferences for information in a strategic setting: a prediction market experiment*

Zakaria Babutsidze† Nobuyuki Hanaki‡ Adam Zylbersztejn§

March 15, 2019

Abstract

We experimentally investigate people's preferences towards different sources of information on a prediction market. The incentivized experimental task is to predict an unknown second-mover's behavior in an earlier hidden game experiment. In Experiment 1, we vary the source of information about that second-mover (picture, neutral video, video containing strategic content). Observed prediction accuracy rates serves as a measure of the empirical value of each source of information. In Experiment 2, we elicit the subjective value of information using the stated preferences method (WTA). The main result is that preferences towards information in a strategic setting are rational in the sence that they are closely aligned with its empirical value.

Keywords: Prediction, information, stated preferences, experiment
JEL Code: C72, D83

*The authors acknowledge the support from the following programs operated by the French National Research Agency (Agence Nationale de Recherche): UCA^{JEDI} (ANR-15-IDEX-01) and LABEX CORTEX (ANR-11-LABX-0042) being part of Université de Lyon (ANR-11-IDEX-007). Ynès Bouamoud, Imen Bouhlel, Maxime David, Carla Lapointe, Ismaël Rafaï, Mira Toumi, Dorian Vincileoni, and Charlène Wojerz have provided quality research assistance. Quentin Thévenet has provided valuable assistance with software programming.

†SKEMA Business School, Université Côte d'Azur (GREDEG) and OFCE, Sciences Po Paris
‡Universitè Côte d'Azur, CNRS, GREDEG
§Université Lumière Lyon 2, GATE L-SE UMR 5824, 69130 Ecully, France

1 Introduction

(under construction)

2 Experimental design and procedures

We run two sets of experiments (henceforth Experiment 1 and Experiment 2) with a total of $N = 287$ participants.

Background data for the prediction task. For implementing the prediction task, we use the data previously reported by Babutsidze, Hanaki, and Zylbersztejn (2019). That study is based on the classic hidden action game by Charness and Dufwenberg (2006) presented in Figure 1. All payoffs are in Euros. The game is played between two parties: the principal (trustor, or player A) and the agent (trustee, or player B). Player A may either choose an outside option *Out* which yields 5 to both players and ends the interaction, or go *In*. Then, player B may either choose to *Roll* a die (which yields 12 to A and 10 to B with the probability of 5/6, and 0 to A and 10 to B with the probability of 1/6), or not to *Roll* (yielding 0 to A and 14 to B with certainty). This game provides a simple setting for studying principal-agent relationships with moral hazard: incentives are not aligned between the two parties, and earning 0 is not perfectly informative for player A about player B's action.

Like Charness and Dufwenberg (2006), we simultaneously elicit both players' decisions. Namely, the player B makes a decision without knowing player A's decision. B's decision is only implemented had player A gone *In*. The game is preceded by a pre-play, face-to-face communication stage in which player B delivers a message to player A. Further details are provided in Appendix A.

In addition to the player Bs' decisions from the experimental game, our dataset contains several recordings of each of those players. Following van Leeuwen, Noussair, Offerman, Suetens, van Veelen, and van de Ven (2017), upon arrival to the laboratory and before learning about the rules of the hidden action game, each subject acting as player B is invited to a separate room for a mugshot picture and a short (about 30 seconds), standardized video recording (reading a short extract from a printer instruction manual, while keep a neutral face expression). These two sources of information are used, respectively, in our PHOTO and VIDNE ("neutral video") treatments. In addition, player Bs are video recorded while making a statement in the pre-play communication stage of the hidden action game. We use this information in our VIDLO ("loaded video") treatment.

The resulting dataset on trustees contains 41 observations. Below, we describe how this rich source of information is used in the prediction tasks in Experiments 1 and 2.

Experiment 1. Participants make a series of twenty decisions, each consisting in predicting a trustee's behavior in an earlier hidden action game (i.e., whether that person rolled a die or

not).[1] Each time, a trustee is randomly drawn without replacement from the main sample of 41 observations. A correct (an incorrect) prediction is worth 10 (2) euros, no feedback is provided from one prediction to the other, and two rounds out of ten are randomly drawn for payoff at the end of the session. Our experimental treatments progressively enrich the set of cues about the trustee that becomes available to the prediction-maker before the prediction: either that trustee's mugshot picture (PHOTO; $N = 44$), or a neutral video recording (with sound) showing that trustee making a non-strategic statement that has been recorded before (and indepently of) the experimental hidden action game (VIDNE; $N = 43$), or a loaded video recording (with sound) in which the trustee makes a strategic pre-play statement in front of the trustors (VIDLO; $N = 45$).
Experiment 2. The second experiment ($N = 145$) is based on the one-shot version of the prediction task used in the first experiment. We rely on the stated preferences approach to eliciting participants' subjective valuation of the different sources of information. We use a within-subject design to elicit individual "willingess to accept" (WTA) via the classic Becker-DeGroot-Marschak (BDM) method. Importantly, our design guarantees that the elicited WTA reflects the strategic value of information in a prediction exercise rather than other confounding motives, such as the prediction-maker's curiosity about the trustee.

First, we ask each participant for a certainty equivalent they would be willing to accept for not having to make a prediction in our three main environments of interest, i.e. PHOTO, VIDNE, VIDLO, as well as an additional environment where no specific information is available about the trustee from the previous hidden action game. The latter feature captures the subjective value of the prediction task *per se* (i.e., with no additional information about the trustee) which we use as a baseline for estimating the subjective value of the additional information provided in PHOTO, VIDNE, and VIDLO conditions. Note that the participants are only informed about the type of information they could receive in each of the four environments, but do not inspect any specific content.

In the next stage (and without prior notice), we elicit subject's beliefs about the prediction accuracy rates previously observed in Experiment 1. This procedure is based on Schlag and van der Weele (2015). For each of the three conditions, a participant is asked how many predictions out of 100 were accurate, and to choose one of the following intervals: ± 5, ± 10, or ± 15. The task is incentivized as follows. An answer is considered as correct if the actual prediction accuracy rate from Experiment 1 lies within the chosen band from the stated belief. A participant is only rewarded for providing three correct answers. Giving three correct answers with interval ± 5 is worth 10 euros, and whenever interval ± 10 (± 15) is chosen instead of ± 5, this amount decreases by 1.50 (3) euros.

Finally, one of the four environments is chosen at random. The stated value corresponding to

[1] Due to the technical glitches in the laboratory – visual content not being displayed properly, or sound being muted – we lost the data from 4 individual predictions (involving 3 subjects) in VIDLO and 8 individual predictions (involving 6 subjects) in VIDNE.

Figure 1: Experimental hidden action game

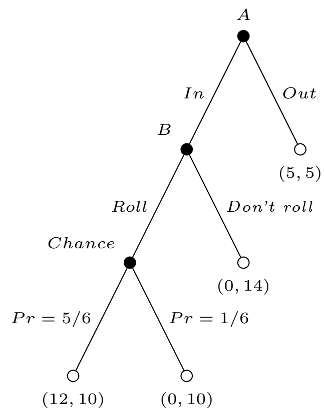

that environment is transformed by the BDM procedure into one of the following outcomes: either the prediction-maker inspects the informational content of a given environment and then makes a prediction, or receives an amount randomly generated by the BDM mechanism and inspect the informational content without being asked to make a prediction.

2.1 Procedure

First, we have collected the dataset for the direct face-to-face (F2F) communication condition.
Implementation and experimental sample.

3 Results

The goal of the first experiment is to estimate the empirical value of the different sources of information in a prediction task. In the second experiment, we elicit the subjective value of those sources of information using the stated preferences approach. Our main result is that stated preferences closely follow the empirical value of information.

3.1 Experiment 1: the objective value of information

We first draw link between the predictions of behavior and the actual behavior. For each treatment, we regress an indicator variable $1[PredictionRoll]$ (set to 1 if one predicts that the second-mover rolled a die in the previous experiment, and to 0 otherwise) on another indicator variable $1[ActualRoll]$ (set to 1 if the second-mover actually rolled a die in the previous experiment, and to 0 otherwise), and report the obtained estimates in Table 1. Coefficient α_0 captures the aggregate likelihood of prediction "Roll" made for those second-movers that did not roll a die. This

Table 1: Predicted and actual behavior: regression analysis

Treatment:	VIDLO		VIDNE		PHOTO	
	coeff. (SE)	p	coeff. (SE)	p	coeff. (SE)	p
Intercept (α_0)	0.491	<0.000	0.476	<0.000	0.471	<0.000
	(0.039)		(0.038)		(0.037)	
$1[ActualRoll]$ (α_1)	0.088	0.005	-0.019	0.532	-0.019	0.529
	(0.030)		(0.031)		(0.030)	
N of obs./clusters	896/45		852/43		880/44	

Note. Results of OLS regression models of the individual predictio (indicator variable $1[PredictionRoll] = 1$ if one predicts that the second-mover rolled a die in the previous experiment; 0 otherwise) on the indicator variable $1[ActualRoll]$ (set to 1 if the second-mover actually rolled a die in the previous experiment, and to 0 otherwise). Observations are clustered for each individual, stanard errors are cluster-robust.

Comparing prediction accuracy across treatments yields similar evidence. Figure 2a summarizes the aggregate rates of accurate predictions which equal: 55.4% in VIDLO, 48.4% in VIDNE, and 48.2% in PHOTO. Figure 2b disaggregates those data and shows the distributions of individual accuracy rates. Both figures clearly show that the accuracy is similar in VIDNE and PHOTO, and improves in VIDLE. For those data, the Epps-Singleton test (Epps and Singleton, 1986) rejects the null hypothesis of identical distributions when comparing VIDLO with either VIDNE ($p = 0.007$) or PHOTO ($p = 0.014$), but does not reject it when comparing VIDNE with PHOTO ($p = 0.255$).[2] The parametric analysis reported in Table 2 corroborates these findings.

coefficient is similar (slightly below 0.5) in all treatments. Coefficient α_1, in turn, captures how this likelihood varies once we consider those second-movers that did roll a die. This coefficient is found to be close to zero and insignificant for PHOTO and VIDNE, suggesting that these sources of information do not suffice to distinguish between the two types of second-movers. In VIDLO, in turn, α_1 is positive and significant, meaning that in this treatment predictions become adjusted to the actual type of the second-mover.

In conclusion, we formulate the following result:

Result 1 [Empirical value of information]
Prediction accuracy does not vary between different forms of non-strategic information provided in PHOTO and VIDNE. Strategic information deliverd in VIDLO improves prediction accuracy as compared to any non-strategic content.

3.2 Experiment 2: the subjective value of information

Figure 3a summarizes the aggregate mean WTA in the four informational environments of Experiment 2: $WTA_{VIDLO} = 71.48$, $WTA_{VIDNE} = 57.41$, $WTA_{PHOTO} = 55.01$, and $WTA_{Baseline} =$

Figure 2: Information and prediction accuracy in Experiment 1

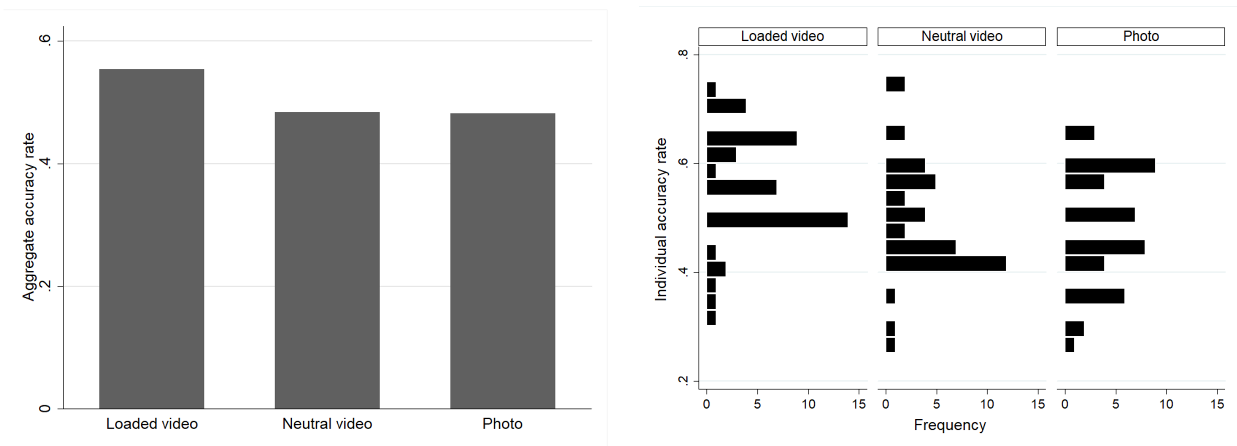

(a) Aggregate accuracy rates

(b) Histograms of individual accuracy rates

Table 2: Individual accuracy across treatments: regression analysis

	coeff. (SE)	p
Intercept (β_0)	0.553	<0.001
	(0.016)	
$1[VIDNE]$ (β_1)	-0.069	0.002
	(0.022)	
$1[PHOTO]$ (β_2)	-0.071	0.002
	(0.022)	
Additional tests:		
$H_0 : \beta_1 = \beta_2$, $p = 0.252$		

Note. Results of OLS regression of the individual accuracy rate on treatment indicator variables ($1[VIDNE] = 1$ for VIDNE, $= 0$ otherwise; $1[PHOTO] = 1$ for PHOTO, $= 0$ otherwise). The intercept provides the average individual accuracy rate in VIDLO, and the coefficient β_1 (β_2) provides the difference between the average individual accuracy rates in VIDLO and VIDNE (VIDLO and PHOTO). All p-values correspond to two-sided t-test. $N = 132$, $R^2 = 0.094$.

48.24. We capture the subjective value of information provided in a given environment – either VIDLO, or VIDNE, or PHOTO – by the difference between the respective WTA and one's stated WTA in the no-information baseline environment. This relative measure is denoted ΔWTA, and its distributions are given in Figure 3b. In aggregate, we observe the following mean values: $\Delta WTA_{PHOTO} = 6.77$, $\Delta WTA_{VIDNE} = 9.17$, $\Delta WTA_{VIDLO} = 23.14$. One-sample t-test

Figure 3: Stated preferences in Experiment 2

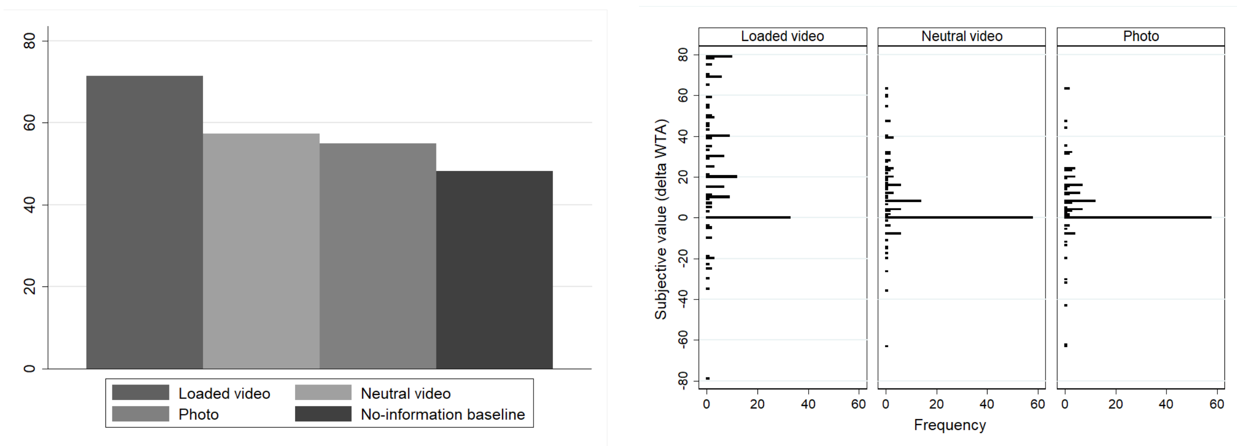

(a) Average absolute WTA

(b) Histograms of relative WTA

shows that each of these figures is significantly different to zero ($p < 0.001$).[3] This suggests that the experimental subjects have a general preference for receiving additional information about the trustee whose decision they are about to predict. Finally, paired t-test does not detect a significant difference between ΔWTA_{VIDNE} and ΔWTA_{PHOTO} ($p = 0.144$), but does so once we compare VIDLO to either PHOTO or VIDNE (both $p < 0.001$).[4] This leads us to the main results of the second experiment:

Result 2. [Subjective value of information]
The subjective values of the different sources of information follow the same pattern as their empirical values: equivalence between PHOTO and VIDNE, both of which are dominated by VIDLO.

To gain more insight into the formation of the subjective value of information, we exploit the data on how individuals perceive the empirical efficiency of the three sources of information. First, Table 3 summarizes the individual beliefs about accuracy rates generated by the three sources of information in Experiment 1, as well as the related confidence measurement based on interval choice. Comparing the mean beliefs from Table 3 to the actual accuracy rates depicted in Figure 2a, we find that participants' beliefs are fairly aligned with the actual accuracy rates in PHOTO and VIDNE, and tend to be overstated in VIDLO. Moving to pairwise comparisons, we report that mean stated beliefs are higher in VIDLO than in either VIDNE or PHOTO (both comparisons

[3] Analogous nonparametric signrank test also yields $p < 0.001$.

[4] Analogous signrank test confirms these findings, yielding $p = 0.109$ for the first comparison and $p < 0.001$ for the two remaining ones.

Table 3: Beliefs about prediction accuracy rates in Experiment 1 and interval measures of confidence

		Interval choice			
Condition	Mean belief	Mean	share of "±5"	share of "±10"	share of "±15"
VIDLO	69.68%	7.66%	44.84%	41.14%	4.14%
VIDNE	53.40%	9.41%	24.83%	62.07%	13.10%
PHOTO	47.14%	8.97%	30.34%	60.00%	9.66%

Note.

yield $p < 0.001$ using two-sided t-test). The difference between VIDNE and PHOTO small yet significant ($p < 0.001$). The corresponding interval measure of confidence, in turn, reveals a somewhat different pattern. Participants are most confident in their beliefs in VIDLO ($p < 0.001$ compared to either VIDNE or PHOTO, two-sided t-test). This time, however, the average interval chosen in PHOTO is slightly lower than in VIDNE, although this difference falls short of attaining statistical significance at the 5% ($p = 0.118$). Altogether, these data clearly show that:

Result 3a. [Individual beliefs and confidence]
The source of information that generates the highest WTA in Experiment 2 – VIDLO – also clearly dominates the two other conditions in terms of both subjects' expectations about the accuracy rate in Experiment 1 and their confidence in this expectation.

Second, we study correlation between ΔWTA and beliefs. Spearman's rank correlation coefficients are as follows: $\rho_{VIDLO} = 0.207$ (significantly different to zero with $p = 0.012$), $\rho_{VIDNE} = 0.006$ ($p = 0.945$), $\rho_{PHOTO} = -0.042$ ($p = 0.619$). We interpret these results as consistent with the data previously reported in Table 3: the existence of a positive relationship between one's beliefs about the empirical value of information and one's subjective value of information requires confidence in those beliefs. This leads us to our final result:

Result 3b. [Individual beliefs and subjective value of information]
VIDLO generates a signficant and positive correlation between the subjective value of information (ΔWTA) and the beliefs about the empirical value of information. We do not observe such a relationship in the remaining conditions that are characterized by a lower degree of confidence in beliefs.

4 Conclusion

(under construction)

A Instructions

B Sample characteristics

References

Babutsidze, Z., N. Hanaki, and A. Zylbersztejn (2019): "Digital Communication and Swift Trust," .

Charness, G., and M. Dufwenberg (2006): "Promises and Partnership," Econometrica, 74(6), 1579–1601.

Epps, T., and K. J. Singleton (1986): "An omnibus test for the two-sample problem using the empirical characteristic function," Journal of Statistical Computation and Simulation, 26(3-4), 177–203.

Forsythe, R., J. L. Horowitz, N. E. Savin, and M. Sefton (1994): "Fairness in simple bargaining experiments," Games and Economic behavior, 6(3), 347–369.

Goerg, S. J., and J. Kaiser (2009): "Nonparametric testing of distributions–The Epps-Singleton two-sample test using the empirical characteristic function," Stata Journal, 9(3), 454.

Schlag, K. H., and J. J. van der Weele (2015): "A method to elicit beliefs as most likely intervals," Judgment and Decision Making, 10(5), 456–469.

van Leeuwen, B., C. N. Noussair, T. Offerman, S. Suetens, M. van Veelen, and J. van de Ven (2017): "Predictably Angry – Facial Cues Provide a Credible Signal of Destructive Behavior," Management Science, 64, 3352–3364.

A Implementation of the hidden action game

Each experimental session involves 6 player As and 6 player Bs. All player As remain in one room during the whole experiment. They are seated in a single row, isolated one from another by separators, and not allowed to talk. The space in front of them is left open and used by a player B to make a brief statement. Player Bs enter the room one by one, so that player As play six rounds of the game (which is common knowledge). Each time, player B faces the center of player As' row, and all player As have a clear view on the speaker. Player B also has a clear, unobstructed view on all six player As. After making a statement, player B is invited to a separate room where s/he privately decides whether to *Roll* a die or not. Then, s/he is asked to leave the laboratory and wait outside until the end of the experiment. At the same time, each player A makes a decision whether to go *In* or stay *Out*. All decisions are made on a sheet of paper, which is then put in an envelope, sealed, and collected by the laboratory staff after each round. In addition, once player B has made a decision and left the separate room, a laboratory staff member rolls a die in private and marks the outcome on player B's sealed envelope. At the end of the experiment, player As and Bs are randomly and anonymously matched in pairs. The outcome of the game for each pair is based on the payoff structure described in Figure 1 and defined by the decision made by player A after player B's statement, as well as the decision made by player B in a private room had A chosen to go *In*. For B's decision to *Roll*, the outcome of the die roll is also taken into account.

For the sake of logistics and efficient time management, player Bs arrive 30 minutes prior to player As. First, they are asked to take up several computerized tasks that measure their preferences and characteristics. Then, they are all led to a waiting room. To avoid any communication or subjects overhearing what others are saying or doing, each participant is seated in a separate cubicle, puts on a headphone and listens to a classical music until further notice. Then, they are taken one by one to a separate room for a mugshot picture and a short, standardized video recording.[5]

Then, each subject is seated back in his cubicle with headphones on. He now listens to an audio file containing the experimental instructions (paper version is also provided). There is a brief comprehension quiz assisted by a laboratory staff member. Finally, he receives additional paper instructions about the upcoming statement in front of player As, as well as a pen and an empty sheet of paper, and is given approximately two minutes to prepare his message.[6] After that, a player B is invited to player As' room where he delivers a statement, leaves for another room, and the game proceeds as explain above in Section 2. The average duration of a message is 26.39 seconds (SD 2.09). Player Bs' statements are recorded using a small, non-intrusive video camera set up in the middle of player As' row, right in front or player Bs' zone, so that the perspective in the video camera recording resembles the one of a player A. The camera is always adjusted to the height of player B (so as to capture head, shoulders, and thorax), and to the luminosity in the room. The sake of the quality of the video recordings, the background in player Bs' zone is covered with light canvas. While making a statement, each player B also has a portable microphone attached below their face. The distance between player As and a player B is set to 2.50 meters.

Upon their arrival to the laboratory, player As also take up the set of preliminary questionnaires. Then, they receive and read paper instructions for the experimental game, and finally they fill in a short comprehension quiz. A laboratory staff member then reads aloud all the questions from the quiz along with the corrects answers, and answers any remaining questions. Finally, player As are asked to wait for the arrival of the first player B.

[5]Like in van Leeuwen, Noussair, Offerman, Suetens, van Veelen, and van de Ven (2017), subjects are asked to read neutral content (a short extract from a printer instruction manual) and keep a neutral face expression. The recording takes about 30 seconds. This information is not part of the present investigation and is not reported in the paper.

[6]Those additional instructions remind the subject about his role in the game; emphasize the fact that the message may affect player A's decisions and, consequently, the subject's gain from the experiment; instruct the subject to avoid making a visual or verbal contact with the experimenter, to aim at communicating with all player As, and not to introduce oneself or give any details about one's own identity.

We have conducted 7 sessions. However, one player B in session 6 decided to quit after the preliminary measurements and before receiving the instruction of the hidden action game, and was replaced by a research assistant unknown to player As. To avoid any contamination of player As' behavior, that research assistant acted as player B in the final round of the experimental game. The data from that round were dismissed and our dataset from that session only covers 5 player Bs, and thus 41 player Bs in total.