Information Valuation and Processing in Performance Contexts with Noisy Feedback: Experimental Evidence

by

Ergün Kotan

A Thesis Submitted to the

Graduate School of Social Sciences & Humanities

in Partial Fulfillment of the Requirements for

the Degree of

Master of Arts

in

Economics

Koç University

August 2012

Koç University

Graduate School of Social Sciences & Humanities (GSSSH)

This is to certify that I have examined this copy of a master's thesis by

Ergün Kotan

and have found that it is complete and satisfactory in all respects,

and that any and all revisions required by the final

examining committee have been made.

Committee Members:

Asst. Prof. Seda Ertaç

Assoc. Prof. Levent Koçkesen

Asst. Prof. Burak Dindaroğlu

Institute of Social Sciences, Director

Date: ____

To nothing in particular.

Abstract

This study explores (1) initial beliefs, (2) the processing of information, (3) the willingness to pay for information, where individuals receive noisy feedback about their own performance (a *self-relevant* case) as well as the performance of others (a *self-irrelevant* case) in an experimental task. The experiment constructs a setting where individuals have monetary incentives to hold accurate beliefs, but where the information received or purchased can also be relevant to self-image. In terms of beliefs, we find that although individuals are aware of overconfidence in others, they do not correct for this bias in their self-assessments, displaying significant overconfidence. Regarding information processing, there is evidence for conservatism and asymmetry in both contexts, with individuals being more responsive to positive information. Analyses on the demand for information reveal the existence of overbidding for information in both contexts. However, when individuals assess their own performance, they: (1) are more likely to interpret negative feedback in a self-serving way, (2) overvalue a further piece of information significantly more after receiving negative than positive feedback, potentially pointing to ego-related motives. Finally, we find some evidence supporting a decline in overbidding when information acquisition is delegated to disinterested parties.

Keywords: Overconfidence, performance feedback, belief updating, information processing, information valuation, self-image.

Özet

Bu çalışma, bireylerin gerek kendi performansları gerekse başkalarının performansları hakkında doğruluğu kesin olmayan geri bildirimler aldıkları bir deney ortamında, sırasıyla (1) bilgi-öncesi kanaatler (2) bilgi işleme (3) bilgiye değer biçme olgularını araştırmaktadır. Deney tasarımı, bireylerin doğru kanaatlere sahip olmalarının parasal açıdan teşvik edildiği, ancak aldıkları bilgilerin kişisel yetenekleri ile de ilgili olabildiği bir ortam kurgulamaktadır. Kişinin kendi performansıyla ilgili bilgi aldığı durum, başkaları hakkında bilgi aldığı durumla karşılaştırılmaktadır. Bilgi-öncesi kanaatlerin incelenmesi sonucunda elde edilen bulgular, bireylerin başka insanlardaki aşırı özgüvenin farkında olduklarını, ancak kendi performanslarını değerlendirdiklerinde aynı yanılgıdan kurtulamayarak yüksek özgüven gösterdiklerini ortaya koymaktadır. Bilgi işleme konusunda ise, hem kişinin kendisi ile ilgili bilgi aldığı durumda hem de başkalarıyla ilgili bilgi aldığı durumda, bilginin kanaatler üzerindeki etkisinin sınırlı olması (muhafazakarlık) ve pozitif bilginin kanaatleri daha fazla etkilemesi (asimetri) olgularının var olduğu gözlenmiştir. Bilgi talebi üzerine yapılan analizler, yine her iki bağlamda da, bilgiye teorik değerinden anlamlı ölçüde fazla paha biçildiğine işaret etmektedir. Diğer taraftan, bireylerin kendi performanslarını değerlendirdikleri durumlarda (1) olumsuz geri bildirimleri kendi öz benliklerine hizmet edecek şekilde ve optimist bir yaklaşımla yorumladıkları, (2) olumsuz bir geri bildirimden sonra bir bilgi daha almak için biçtikleri fiyatın olumlu geri bildirim alınmış olduğu duruma göre teorik değerlerden anlamlı ölçüde daha fazla saptığı gözlemlenmiştir. Bu bulguların öz benlikle alakalı güdülerin bilgi kullanımı ve talebinde önemli olabileceğine işaret ettiği düşünülmektedir. Son olarak, bilgi edinme süreci bilginin içeriği konusunda tarafsız kişilere ihale edildiğinde, bilgiye ederinden fazla fiyat verme eğiliminin azaldığı yönünde bulgular edinilmiştir.

Anahtar Kelimeler: Yüksek özgüven, performans geri bildirimi, kanaat güncelleme, bilgi işleme, bilgi değerleme, öz farkındalık

Acknowledgements

I would like to express my deepest gratitude to Asst. Prof. Seda Ertaç for her patience, continuous support and exceptional guidance throughout all stages of this research. It has been a great experience for me to study under her supervision.

I would also like to thank Assoc. Prof. Levent Koçkesen and Asst. Prof. Burak Dindaroğlu for their interest to this study and valuable comments.

I owe very special thanks to my friends Emrah Yılmaz and Neslihan Sakarya for their friendship and support throughout my graduate study. I am also indebted to my colleague Gökmen Ökmen for his encouragement during the writing process of this thesis.

Lastly, I am grateful to Koç University Graduate School of Social Sciences and Humanities and the European Commission (through Dr. Ertac's grant FP-7 IRG 239529) for their financial support for conducting the experiments, and TÜBİTAK and Vehbi Koç Foundation for their financial assistance during my graduate study.

Table of Contents

List of Figures

List of Tables

I. Introduction

In many cases people make decisions under uncertainty. Therefore, their beliefs on the realization of each possible outcome determine how they act. Standard economic theory predicts that people incorporate new piece of information into existing beliefs according to the well-known Bayes' rule. Moreover, it predicts that context of information does not affect information processing. That is, people update beliefs like perfect Bayesian agents no matter what the underlying quantities are. There is already a vast literature, supported by experimental evidence, showing significant deviations from Bayes' rule (e.g. Tversky and Kahneman (1974), Grether (1980)). However, the scope of this study is beyond that. This paper, additionally, questions whether the self-relevance of information affects information processing or not. Moreover, this research investigates how demand for information varies under different contexts.

The motivation of this research is that the results might have some policy implications. Many economic decisions we make through our lives involve self-relevance, therefore it is very important to know how we assess information about ourselves. For instance, a student who is considering pursuing a higher educational degree should have already formed the accurate beliefs about herself to make the right decision. However, since the underlying quantities are her own attributes like intelligence, discipline, diligence, etc., she might have assessed good and bad feedback differently through her life. Thus, asymmetries in the processing of good and bad feedback might lead her to hold overconfident/underconfident posteriors and lead her to make the wrong decision. Similar judgmental errors might be effective in mate selection, price bubbles, etc. In addition, individuals are usually faced with situations where they have the option of how much information to collect before making a decision. While extra information always has a non-negative value for decision-making from an economic perspective, psychological mechanisms such as ego-preservation might be in effect in cases where the information to be received is relevant to the self.[1] That is, in addition to the way information is processed, the way information is collected can also be potentially influenced by self-relevance. If experimental evidence supports that people process and demand information differently when the information is self-relevant and self-irrelevant, the delegation of information acquisition and processing to disinterested parties can be considered as a policy (Ertac (2011)).

[1] See Köszegi (2006) for a theoretical model of information collection with ego-utility.

There are already a few papers examining how information processing is affected by the underlying context. The papers closest to the current study are Möbius et al. (2011), Eil and Rao (2011), Ertac (2011), and Grossman and Owens (2010). This paper attempts to retain many of the attractive features of each.

The most similar paper to this study in design is Möbius et al. (2011), in which subjects perform an IQ test and are fed with four noisy signals on their relative performance rankings. Before and after each signal, they are asked to report their subjective probabilities for being in the top half of the session. They use a novel belief elicitation mechanism called *crossover*, which was introduced by Karni (2009). This mechanism is superior to its alternatives since it works under all risk-preferences and it does not generate perverse incentives to hedge performance on the quiz. The main findings in Möbius et al. (2011) are the existence of both asymmetry and conservatism in belief updating when the underlying context is ego-relevant. On the other hand, Eil and Rao (2011) find support for only asymmetry but not for conservatism. Similarly, they ask subjects to update their beliefs in response to new information on their IQ/Beauty scores. They use a different signal structure such that each subject learns if she is better/worse than a randomly selected other participant. This bilateral comparison results are repeated for three times, and subjects update their beliefs for each ranking (1-10) after each message. One common procedure in these two papers is that at the end subjects are given the chance to purchase complete information about their position in the performance stage. However, as one can easily notice, this information has no instrumental value to improve the decision making process, since subjects can get it only at the end. The current study contributes to the literature by analyzing the demand for information when the information is still valuable from a decision-making perspective, in addition to being self-relevant.

Ertac (2011) is another paper, which tries to answer if self-relevance affects information processing. In her design, subjects perform in two different tasks (addition task, GRE verbal test) and then are asked to report their beliefs for which third of the distribution they occupy. Her signal structure is different than ours such that subjects learn whether they are in the top partition or not (and in some sessions, whether they are in the bottom partition or not). Unlike previous papers, Ertac (2011) finds asymmetry in opposite direction, that is, subjects overweigh bad information in the self-relevant contexts, and end with pessimistic posteriors. Like Eil and Rao (2011), she uses the well-known *quadratic scoring rule* to elicit beliefs. However, this mechanism is problematic, since it works under only risk-neutrality assumption and is vulnerable to moral hazard issues. One last study to note is Grossman and Owens (2010), which is very different from previous ones in terms of both design and findings. In Grossman and Owens (2010), subjects update their beliefs on their absolute performances rather than relative performances, after getting noisy feedback. They find neither conservatism nor asymmetry in self-relevant information processing. However, they explain the overconfident posteriors in self-relevant case with the existence of overconfident priors in this case. Nevertheless, their design has an important similarity with ours, in the sense of using a *randomly selected other participant* as a control treatment.

As noted before, this paper aims to retain the attractive features of the above papers, and add several new attributes that are not present in them. In the basic design, subjects perform in an IQ test for a particular time and are requested to report their prior beliefs for being in the top partition of that session, and then are asked to update their beliefs according to a series of noisy feedback.[2] While the first signal is for free, the second one is subject to a fee. Unlike previous research, in our design information has an instrumental value, since it helps to improve decision making, and therefore monetary payoffs. To isolate any curiosity effect, subjects are told that they will have the opportunity to learn their states at the end of the experiment, free of charge. In the control treatment we ask subjects to process information on the performance of some other subject, whose prior belief in the first round is declared.

We find the following main results. First, people are aware of overconfidence in others; however they keep showing the same bias in their self-assessments. Second, people are both conservative and asymmetric belief updaters independent of self-relevance of information. In other words, they update their beliefs less than what Bayes' rule predicts, and they weigh positive feedback more than negative feedback by contrast with the theoretical predictions. Third, people's offers for information remarkably exceed theoretical values no matter whom the information is about. Fourth, after a negative feedback in a self-relevant context, people are more likely to process information in a self-serving way in comparison to a positive signal case. Additionally, in the same case, the severity of overbidding seems to be increasing compared to a situation after getting good news.

The rest of the paper is organized as follows. Section 2 gives a detailed explanation of experimental design. Section 3 provides theoretical predictions and hypotheses. Section 4 describes data and presents results, and Section 5 concludes.

[2] Top state corresponds to the best 50% of subjects in terms of performance, and Bottom to the rest.

II. Experimental Design

The experiment mainly consists of the following four parts.

In the first part, subjects perform in an IQ test during six minutes.[3] Most of the questions are multiple-choice, and every two mistakes in multiple choice questions eliminate a true answer.[4] On the other hand, mistakes in open-ended questions do not affect the score. Therefore, the final score of an agent is the number of correct answers minus the half of number of wrong answers in multiple-choice questions. In addition, subjects have the opportunity to skip a question, when they do not want to answer that one. There is sufficient number of questions such that a subject who wants to skip the current question can find a new one on her screen instantly.[5] Earnings for the quiz are score multiplied by 20 ECU (the exchange rate is 0.03 for all rounds of the experiment).

As soon as the quiz part is completed and subjects are ranked according to their relative performances, the second part (*Self*) starts. In this part, subjects report their beliefs for being in the top half three times (*prior, interim,* and *posterior*). However, these reports are framed as investment decisions, such that earnings from belief elicitation stages correspond to a substantial part of total earnings from the experiment.[6] While prior beliefs are reported without any feedback, the others are reported after a noisy feedback is acquired. Since the first feedback is free of charge, everyone is expected to report their interim beliefs. However, the final information, which has a larger accuracy, is subject to a fee. Hence, only subjects who purchase information are allowed to update their beliefs. For the remaining subjects, interim beliefs are accepted as posterior beliefs, as well.

In the third part, each subject is asked to repeat the same procedures for some other participant (we call this subject *SOP – selected other person*), whose prior belief in the *Self* round is given as starting information. However, subjects are free to report whatever they want as their priors. The rest of this part is similar to the second one, only this time, success of each investment decision depends on the performance of the other person. Also note that a subject only knows that her SOP is an anonymous person in the same session. What she does not know is that this person is selected according to some criteria. For any subject, her SOP is

[3] The questions in the IQ test are taken from Phillip Carter's "IQ GYM – Grow Your Mind" with the permission of Arcturus Publishing.

[4] While 22 out of 30 questions are multiple-choice, the others are open-ended.

[5] Only 1 subject reached the end of questions. The second largest number of questions monitored by a subject was 18.

[6] The reward is 200 ECU at each round.

the person, who is in the same half with her and has reported a prior belief in the *Self* round closest to her prior. In addition to that, as long as a subject has the chance to get information, we feed her with the same signal sequence as her SOP in the *Self* stage in order to make this paired comparison as effective as possible.

The fourth part proceeds just like the third one, but differs only in the way that other person is selected. In this part, *SOP* is assigned in a complete random way. This way, we aim to observe how people process information about other people with different confidence levels. Lastly, we have to underline that the third and the fourth parts will be mentioned as *Other* rounds in common during the rest of this paper.

Also, note that before the main treatment starts, subjects go through a quiz, checking whether they understood the instructions on how each mechanism works or not. In addition, we run a practice round before the main treatment starts. Also, it is important to underline that instructions on each part are read just before that part starts. Additionally, we run a Holt-Laury task to elicit subjects' risk preferences.[7] Lastly, at the end of the main treatment, subjects go through a survey that contains questions on some personal information, and assessment on task difficulty.[8]

Signal Accuracy: The first signal is accurate with 75% probability. That is, while a person in bottom half gets a *Top* signal with 25% probability, a person in the top half gets the same signal with 75% probability. On the other hand, the second signal is accurate with 90% probability. That is, while a person in bottom half gets a *Bottom* signal with 90% probability, a person in the top half gets the same signal with 10% probability.

Belief Elicitation: We use the crossover mechanism as in Möbius et al. (2011) to elicit subjects' beliefs, since we believe that this mechanism is superior to its alternatives in many ways.

In this mechanism, subjects are given two different lotteries to earn the reward**.** The first lottery gives the reward if it turns out that the agent is in top half in the performance stage. The second lottery gives the reward with probability X. We simply ask subjects to report the minimum amount of X to choose the second lottery. Under this mechanism, a rational agent is

[7] See Appendix 1 for the payoffs at the Holt-Laury task.
[8] See Appendix 5 for survey questions and summary statistics on them.

expected to report X as the subjective probability she assigns to being in the top half, no matter what kind of risk-preferences she has.

Crossover mechanism is elegant since unlike *QSR*, it works under any kind of risk-preferences. Moreover, this mechanism does not generate perverse incentives to hedge performance in quiz.

Lastly, note that at each stage only one investment decision among three is chosen randomly for payment in order to prevent any hedging behavior.

Information Purchasing: After subjects report their interim beliefs under the light of free signal, they are given the chance to purchase information with 90% accuracy to improve their decision making (investment decision as reporting posterior beliefs). At this point, the well-known BDM mechanism (Becker et al. (1964)) is used to elicit subjects' maximum willingness to pay for information. According to this mechanism, subjects report their max WTPs in an interval and computer determines a random price in that interval.[9] If a subject's max WTP is more than the price, she purchases the information at that price. Otherwise she cannot buy information. Therefore, under this mechanism, a rational agent is better off by reporting her exact max WTP for information. Also, note that in our design, subjects are given an endowment which they can use to purchase information.[10] They are free to spend as much as they want from this endowment to purchase information, and keep the remaining. However, all saving and consumption decisions given at this stage is relevant only if this investment decision (reporting posteriors) is chosen as the paid belief elicitation stage.

[9] The price of information is an element of the set {0, 1, 2, … 49, 50}.

[10] The amount of the endowment is 50 ECU at each round.

III. Theoretical Predictions

For the experimental design detailed in the previous section, standard economic theory makes the following predictions. These predictions can be presented as a summary of the questions that are tried to be answered via this research.

i. There is no initial underconfidence/overconfidence either at the individual or aggregate level.

We can move this prediction one step ahead with an additional assumption. If subjects hold correct initial beliefs and it is common knowledge, there should be no downgrading or upgrading of subjective priors in *Other* rounds.

ii. Subjects update their beliefs according to the Bayes' rule, given below.

$$logit\ (\mu^2) = logit\ (\mu^1) + \ I\ ln\left(\frac{\alpha}{\beta}\right) + \ (1 - I)\ ln\left(\frac{1-\alpha}{1-\beta}\right) \tag{1}$$

$$where$$

$$logit\ (x) = ln\ (\frac{x}{1-x}) \tag{2}$$

$$\alpha\text{: } probability\ of\ a\ Top\ type\ agent\ getting\ a\ Top\ signal$$

$$\beta\text{: } probability\ of\ a\ Bottom\ type\ agent\ getting\ a\ Top\ signal$$

$$I\text{: } 1\ if\ the\ signal\ is\ Top, 0\ otherwise.$$

Note that in the current design, while α is 0.75 and 0.9 in consecutive stages, β is either 0.25 or 0.1 depending on the signal order. Also, while μ^1corresponds to prior belief, μ^2denotes its Bayesian posterior. As the Bayes' rule suggests, subjects are expected not to treat positive and negative signals differently.

iii. Given the similar priors and same signal sequence, posterior beliefs submitted under self-relevant and self-irrelevant treatments should be similar, as well.

This hypothesis addresses the main research question of this paper. As underlined before, economic theory predicts that context of information does not affect information processing. With our design, we have the opportunity to observe whether people process information differently when it is ego-relevant or not. Also, by generating same signal sequence in the *Self* round and the *Other* round (only when types are the same) we acquire an effective paired comparison.

iv. There is no undervaluation or overvaluation of information either at the individual or aggregate level.

In order to verify this hypothesis, we first need to present what theory suggests for information valuation under the very conditions of our design. Next part demonstrates how much a rational agent should bid for noisy information.

Value of Information: Unlike previously mentioned papers, we give subjects the opportunity to improve their decision making by purchasing noisy, but more accurate information before reporting their last investment decisions. Therefore, we are able to observe how much subjects value information which has an instrumental value, both in self-relevant and self-irrelevant contexts.

According to the standard economic theory, given the belief elicitation mechanism and signal accuracy, subjects report their max WTP as the value which makes them indifferent between purchasing information and not.

As explained before, under the crossover mechanism a subject gets either R (200 ECU) or r (nothing), depending on her investment decision, performance in quiz stage, and chance. Therefore, information is valuable as long as it helps a subject to increase her probability of earning R. The following figure can be seen as a summary of how the *crossover* mechanism works from an agent's point of view.

Figure 1. Game Tree Representation of the Crossover Mechanism

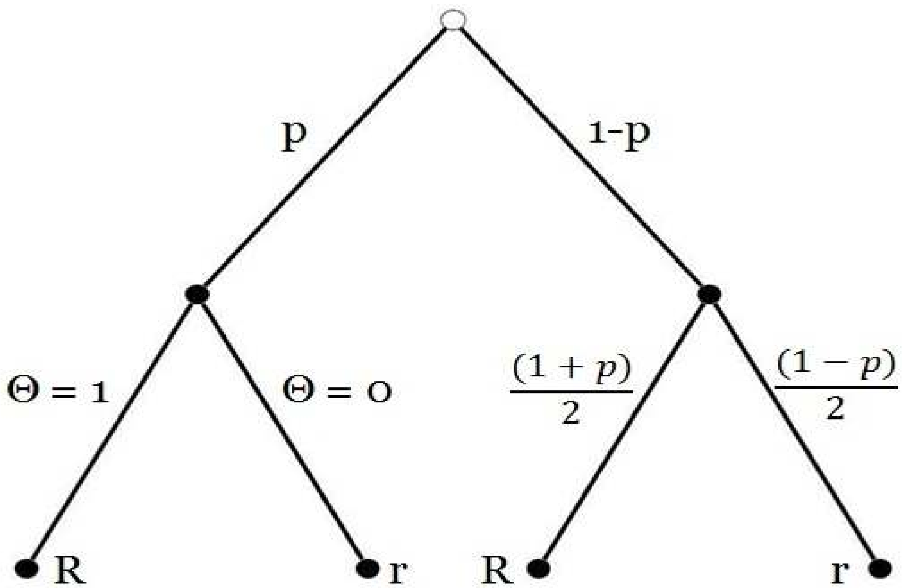

Under this mechanism, the probability of taking R for a rational agent, who believes that she is in the top half with probability p, can be formulated as:[11]

$$P(R) = p.\Theta + (1-p)\frac{(1+p)}{2} \tag{3}$$

Returning to the specific conditions of our design, we have already underlined that if an agent does not purchase information, her second report is automatically assigned as her final decision. Therefore, this agent confronts with bundle (200, 0) with probabilities (P_N, 1-P_N). On the other hand, an agent who gets the final signal ends up with bundle (200-x, -x) with probabilities (P_I, 1-P_I). Hence, the condition $P_I \geq P_N$ must be satisfied in order to make a rational agent bid a non-negative amount of endowment for information. Lastly, we denote the increase in the probability of earning R with ΔP.[12]

The calculation of P_N is quite trivial, since all we do is to replace Θ with its expected value p, in the equality above. Therefore, the probability of taking R for a subject who does not purchase information is:

$$P_N = p^2 + \frac{(1-p)(1+p)}{2} \tag{4}$$

On the other hand, the probability of earning R for a subject who considers buying the noisy signal is a little bit complicated. This subject needs to calculate how much the probability of taking R increases in order to make a rational bid. According to the Bayes' rule, a rational agent with belief p will update her belief in the final investment decision, and report p_G or p_B depending on the type of signal bought.[13] Bayes' rule and 90% signal accuracy rate impose the following equalities for p_G and p_B.[14]

$$p_G = \frac{9p}{(8p+1)} \tag{5}$$

$$p_B = \frac{p}{(9-8p)} \tag{6}$$

[11] Θ denotes the type of agent, either 1 (Top) or 0 (Bottom).

[12] Appendix 3 shows the non-negativity of ΔP for each confidence level.

[13] p_G and p_B denote Bayesian benchmarks for a subject with confidence level p, after a positive and negative feedback, respectively.

[14] See Appendix 2 for derivations of p_G and p_B.

Considering the realization probabilities of each signal and each type from the agent's point of view, we calculate the increase in the probability of earning R as:[15]

$$\Delta P = \frac{p^2(288 + 1472p - 5856p^2 + 6144p^3 - 2048p^4)}{(9 + 64p - 64p^2)^2} \tag{7}$$

Therefore, we are able to calculate P_I using the expression $\Delta P = (P_I - P_N)$ as:

$$P_I = \frac{40.5 + 576p + 1800.5p^2 - 2048p^3 - 2336p^4 + 2048p^5}{(9 + 64p - 64p^2)^2} \tag{8}$$

Having calculated the probabilities for each state under both information and no information scenarios, we are able to determine the value of information for any given confidence level using the following equality.

$$P_I \,.\, U(200 - x) + (1 - P_I\,).\, U(-x) = P_N \,.\, U(200) + (1 - P_N\,).\, U(0) \tag{9}$$

As the above equality suggests, value of information depends on the utility function of subjects, too. Therefore, we prefer to continue further calculations for both risk-neutrality and risk aversion assumptions. As we mentioned before, our design includes a Holt-Laury task at the end, which enables us to estimate risk parameters of subjects for a specific utility function. We prefer to use CARA (constant absolute risk aversion) utility function for calculations under the risk-aversion assumption.

Risk-Neutrality:

Assuming that subjects have risk-neutral preferences leads us to the following value of information.

Inserting U(c) = c into (9) we get:

$$P_I.\,(200 - x) + (1 - P_I\,).\,(-x) = P_N\,.\,200 \tag{10}$$

which is solved for x as:

$$x = \Delta P.\,200 \tag{11}$$

Thus, for a risk-neutral subject, value of information solely depends on her confidence level. As shown in Figure 2, value of information raises proportionately as the agent's uncertainty about her type increases. However, since the range of values that agents can bid for

[15] See Appendix 3 for derivation of ΔP.

information is not a continuum – rather a set of integers from 0 to 50 – they are expected to bid the largest integer that does not exceed the value presented in the formula above. In short, a risk-neutral Bayesian subject is expected to offer the value given below, as her max WTP for information.

$$x = \lfloor \Delta P.200 \rfloor \tag{12}$$

Following figure presents how value of information varies for different confidence levels assuming that subjects are risk-neutral Bayesian belief updaters.

Figure 2. Demand for Information vs. Most Recent Beliefs under Risk-Neutrality Assumption

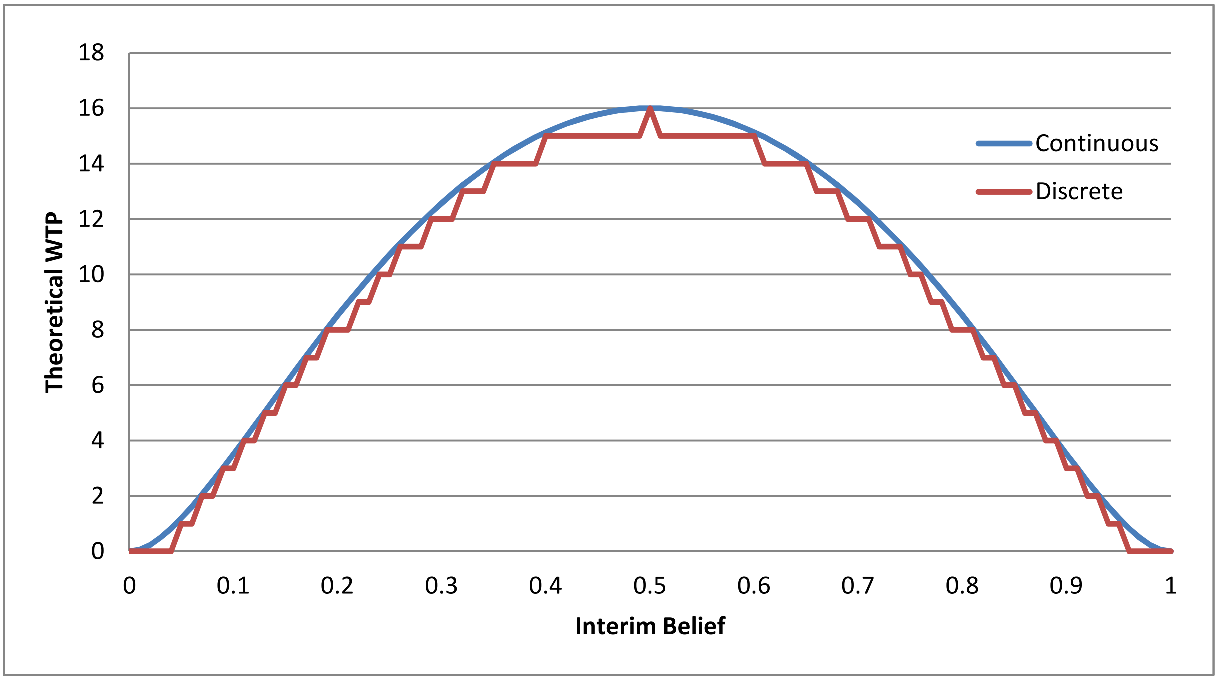

Risk Aversion:

Assuming that subjects have CARA utility function leads us to the following series of calculations for value of information.

Inserting $U(c) = (-1/r).e^{-rc}$ into (9), and eliminating the common denominator we get:

$$P_I\,.\,e^{-r(200-x)} + (1 - P_I)e^{rx} = P_N\,.\,e^{-200r} + (1 - P_N) \tag{13}$$

or more compactly,

$$P_I\,.\,e^{rx}(e^{-200r} - 1) + e^{rx} = P_N\,(e^{-200r} - 1) + 1 \tag{14}$$

which simplifies to:

$$e^{rx}[P_I\,(e^{-200r}-1)+1] = P_N\,(e^{-200r}-1)+1 \tag{15}$$

taking natural logarithms of both sides, value of information is solved as:

$$x = \frac{ln\,(P_N\,(e^{-200r}-1)+1) - ln\,(P_I\,(e^{-200r}-1)+1)}{r} \tag{16}$$

In the formula given above, r denotes the risk parameter of a subject. This parameter is calculated according to a subject's number of safe choices (nosc) in the Holt-Laury task.[16] If we insert equations (4) and (8) into the expression above, (16) reduces to:

$$x = \frac{ln\left[\dfrac{1+\left(\frac{1}{2}(1+p^2)\right)[-1+e^{-200r}]}{1+\dfrac{40.5+576p+1800.5p^2-2048p^3-2336p^4+2048p^5}{(9+64p-64p^2)^2}[-1+e^{-200r}]}\right]}{r} \tag{17}$$

Thus, for a subject with CARA utility function, value of information depends on both her confidence level and her risk parameter. As shown in Figure 3, value of information does not follow a common pattern for all risk groups. However, it is quite clear that the confidence level where value of information reaches its maximum level is under 50% for risk-loving subjects (nosc<4), and over 50% for risk-averse subjects (nosc>4). In addition, these points drift apart from the mid-point as the level of risk-loving/risk aversion increases.

As mentioned in the risk-neutrality case, subjects are expected to offer the largest integer that does not exceed (17). Therefore, the expression for max WTP of a Bayesian agent with CARA utility function, risk parameter r, and confidence level p is:

$$x = \left\lfloor \frac{ln\left[\dfrac{1+\left(\frac{1}{2}(1+p^2)\right)[-1+e^{-200r}]}{1+\dfrac{40.5+576p+1800.5p^2-2048p^3-2336p^4+2048p^5}{(9+64p-64p^2)^2}[-1+e^{-200r}]}\right]}{r} \right\rfloor \tag{18}$$

Figure 4 shows how value of information varies by confidence level and risk preferences, using the equality presented above.

---

[16] See Appendix 1 to behold how risk parameters are obtained.

Figure 3. Demand for Information vs. Most Recent Beliefs under CARA Utility Assumption (Continuous Case)

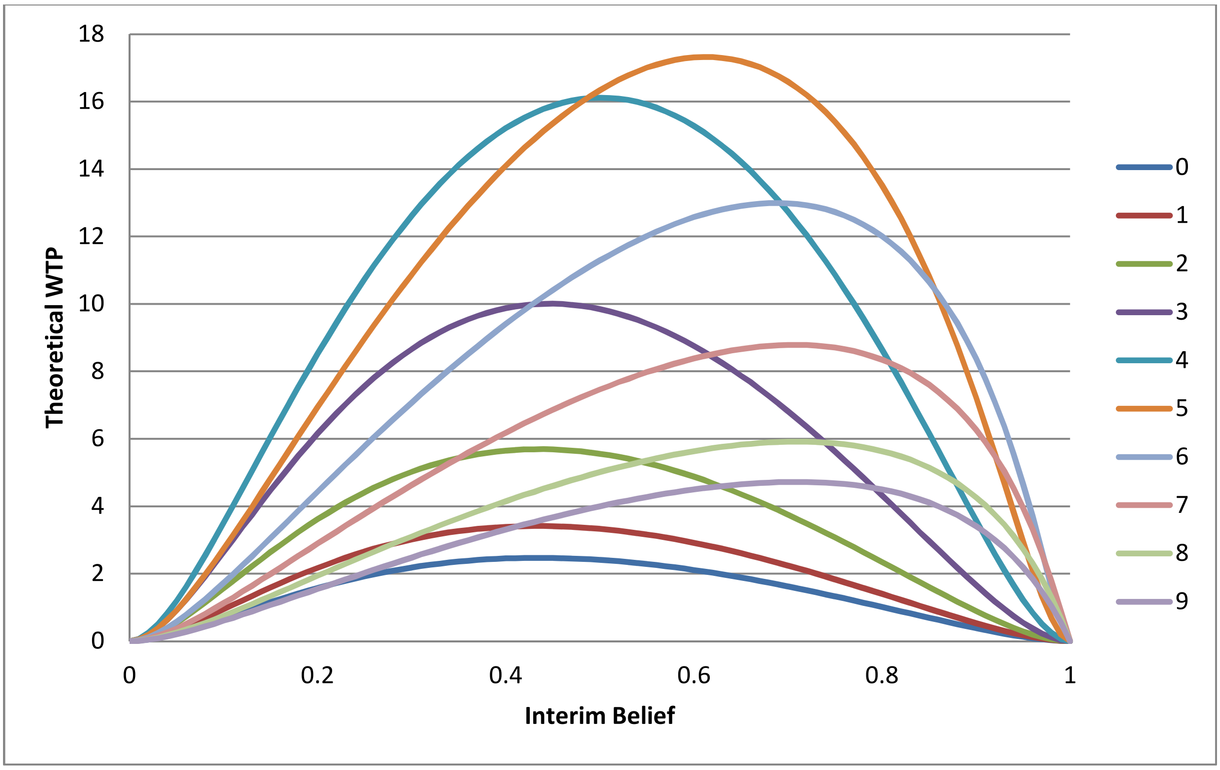

Figure 4. Demand for Information vs. Most Recent Beliefs under CARA Utility Assumption (Discrete Case)

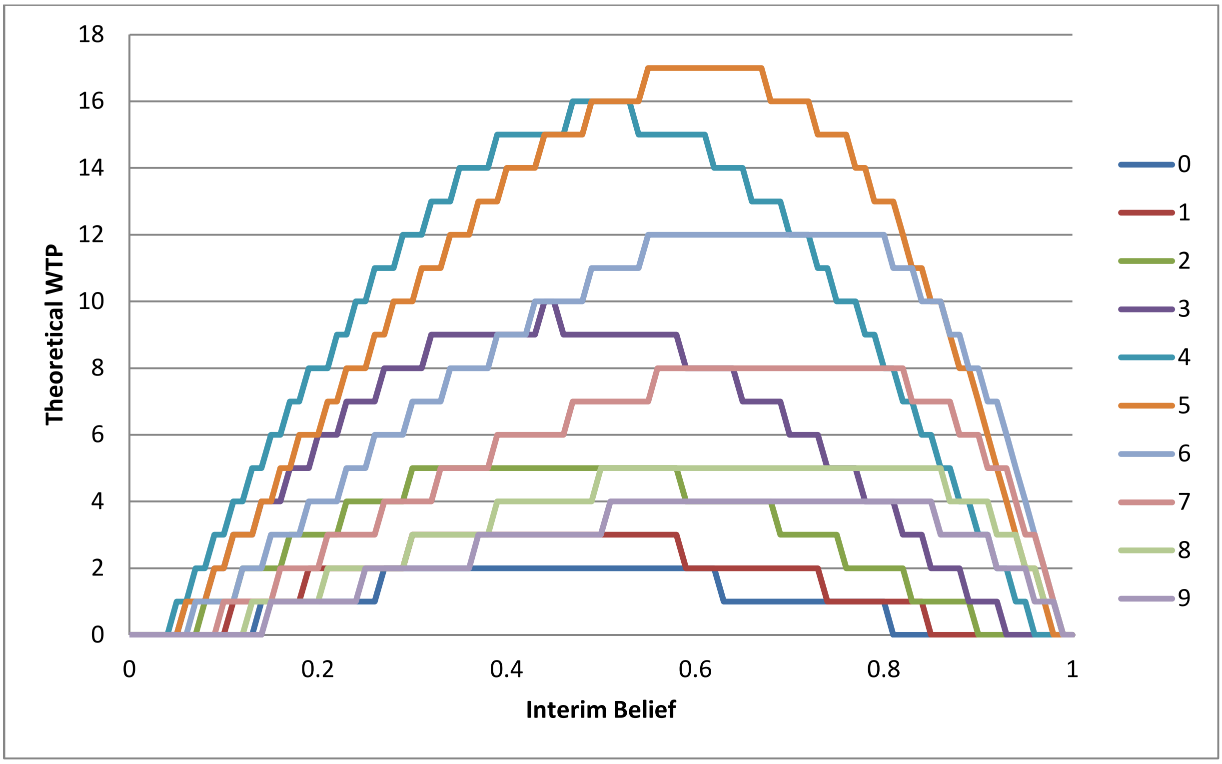

As noted before, the design enables us to check whether subjects' bids match with theoretical predictions. If not, we intend to find out whether they overbid or underbid at both individual and aggregate levels, and if this behavior varies according to their existing beliefs, signal type, etc.

IV. Data and Results

The experiment was programmed and conducted with the software z-Tree (Fischbacher (2007)). While a pilot session and four real sessions were run at the end of Spring 2011, the final session was conducted in Spring 2012, at the computer laboratories of Koç University.[17]

Table 1. Summary of Sessions

Session ID[18]	Date	# of Subjects
0	May 3rd, 2011	14
1	May 11th, 2011	32
2	May 14th, 2011	30
3	May 17th, 2011	38
4	May 25th, 2011	12
5	April 21st, 2012	26

The pilot session was scheduled to check if the program works as it is supposed to do, and the instructions are easily comprehensible. Although our observations from the pilot session were very promising, results from real experiments showed that some subjects did not care or understand the instructions at all. During our analyses, we detected that 35 subjects out of 138 made at least one belief updating in the wrong direction. That is, these subjects got a positive signal and decreased their belief for being in top half, or vice versa, at least once.[19] In addition, two subjects are observed to choose Option A over Option B in the final decision of Holt-Laury task, which conflicts with monotonicity of preferences. As a result, we decided to restrict our sample to subjects who did not make any update in the wrong direction and hold monotonic preferences. Hence, main results in this paper are based on the restricted sample. However, some results on the full sample are also used when necessary.

The majority of the subject pool was undergraduate students enrolled in a diverse set of departments, however there were some graduate students, as well.[20] The restricted sample consists of 39 females and 63 males**.** Since we are left with 102 subjects, the dataset contains 102 observations for the *Self* round, and 204 observations for *Other* rounds. On average, subjects attempted to answer 11.8 questions, from which they managed to answer 6.07 questions correctly, and the average score happened to be 5.04. Post-experiment survey shows

[17] We promoted the experiment mainly via electronic newsletters.

[18] The pilot session is listed as Session #0.

[19] Given that there are 2 updating processes in each round, and 3 different rounds, each subject updates her beliefs 6 different times. A closer look on these subjects reveal that 21 made a unique mistake, 9 made two mistakes, and 5 made three mistakes in total.

[20] Out of 138 subjects 16 were graduate students.

that subjects assess the difficulty of the IQ test as moderate.[21] Additionally, subjects seem to care about solely being in the top half much more than the monetary payoffs of the experiment.[22] Lastly, earnings in all sessions ranged between ₺1.9 and ₺32.7, while the average payment was ₺18.6. Earnings were paid privately in cash at the end of each session.

Self Assessments:

As mentioned earlier, standard economic theory suggests that subjects have correct initial beliefs about themselves. Therefore, we are questioning if there is any confidence bias either at the individual or aggregate level.

To begin with, we utilize a very conservative, but a simple way of analyzing confidence levels. In this method, we only consider which half a subject assigns more probability, and disregard the strength of beliefs. In other words, states perceived as most likely by subjects are assumed to be their choices.

Table 2. Distribution of State Choices by Subject Pools

Sample/Choice	*Bottom*	*Bottom/Top*	*Top*	**Total**
Restricted	13 (12.7%)	7 (6.9%)	82 (80.4%)	102
Unrestricted	30 (21.7%)	7 (5.1%)	101 (73.2%)	138

For the unrestricted sample, we expect to see no significant difference in the amount of people who chose bottom and top states, since they are equally distributed. However, as the above table shows 73.2% of subjects believe that they are more likely to be in the top group. On the other hand, only 21.7% of them pick bottom half as their most probable state. Returning to the restricted sample, on which we intend to construct the main results, it should be noted that top and bottom states are not equally likely. From 102 subjects, the amount of people who are actually in top half are happen to be 58, which corresponds to 56.9% of the subject pool. However, as seen in Table 2, 80.4% of the subjects believe that they are more likely to be in the top half. As a result, this simple methodology seems to be notifying us about an aggregate level overconfidence both at the unrestricted and restricted samples.

Similarly, we apply the same approach to our analysis on overconfidence at the individual level. In order to do that, we only need to pit subjects' chosen states against their actual states.

[21] Mean assessment is 4.69 over a scale of 10.

[22] Mean scores reported for these two survey questions are 6.27 and 3.32, respectively.

The following table presents distributions by chosen and actual states for the restricted sample.

Table 3. Distribution of State Choices by Actual States

Choice/Actual	*Bottom*	*Top*	Total
Bottom	9	4	13
Bottom/Top	4	3	7
Top	31	51	82
Total	44	58	102

Table 3 demonstrates that only 60 subjects (58.8% of the sample) are correct in their choices. On the other side, while 7 people (6.9% of the sample) are underconfident, remaining 35 subjects (34.3% of the sample) seem to have made overconfident initial choices. These results can be used to argue that there is evidence of both underconfidence and overconfidence at the individual level, while noting that the second one is much stronger. On the other hand, as we already noted, the methodology used above is a conservative one, and definitely needs to be supported by additional analyses. Therefore, we intend to present supplementary statistical evidences to prove the existence of overconfidence in initial beliefs.

Figure 5. Distribution of Initial Beliefs in the Self Round

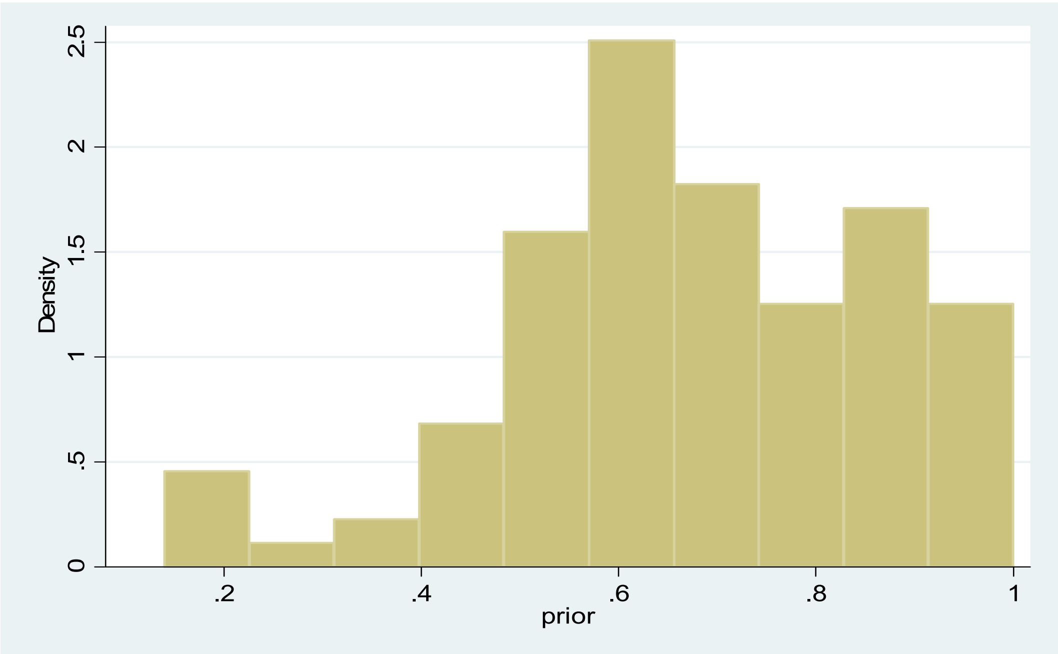

Figure 5 shows how prior beliefs are distributed in the *Self* round. As it can be easily seen, most of the subjects tend to think that they are in the top half. A paired t-test confirms that prior beliefs are significantly higher than actual states of agents ($t = 2.098$, $p = 0.0192$). Also, according to a Wilcoxon signed-rank test, the difference between prior beliefs and actual states are highly statistically significant (z = 2.202, p = 0.0277). Next, we investigate whether there is any gender difference in terms of overconfidence, or not.

Figure 6. Means of Actual States and Initial Beliefs by Gender

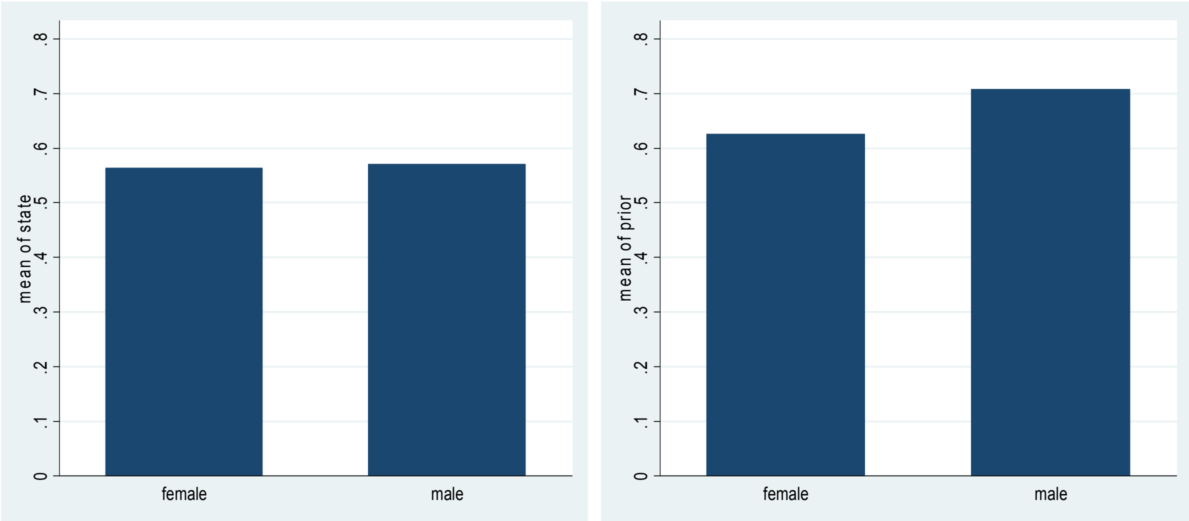

Figure 6 shows how average values of states and priors vary by gender. As a two-sample test of proportion supports, there is no significant difference in the likelihood of a male and a female subject being in the top group (z = -0.0726, p = 0.9421). However, according to a two-sample t-test, the difference between prior beliefs of females and males is significantly less than zero (t = -2.0747, p = 0.0203). Same result holds for a two-sample Mann-Whitney test, as well (z = -2.086, p = 0.037). These results tell us that self-confidence is much stronger in men, but do not say anything about if men are overconfident, or women are underconfident. To resolve this issue, we compare each subject's prior belief with her actual state. Both a paired t-test (t = 0.7422, p = 0.4625) and a Wilcoxon signed rank test (z = 0.81, p = 0.4179) show that women are accurate in their predictions about their actual performance. On the other hand, results of same statistical tests reveal that men hold overconfident initial beliefs about themselves (t = 2.0749, p = 0.0211 and z = 2.01, p = 0.044). Therefore, we conclude that the overconfidence result in the whole subject pool comes mainly from men.

Assessments about Others:

As mentioned before, at the start of each *Other* round, a subject only knows the prior belief reported by her SOP in the *Self* round. Therefore, her first investment decision is expected to depend solely on this information. Assuming that accurate beliefs are common knowledge, individuals should be expected to trust each other's self-assessments. However, our analyses do not seem to be supporting this prediction. Firstly, according to a paired t-test, subjects seem to be significantly downgrading beliefs reported by their SOPs (t = -5.5174, p = 0). A Wilcoxon-signed rank test supports the same finding, too (z = -5.312, p = 0). Besides, a paired t-test (t = 3.632, p = 0.0002) and a Wilcoxon signed-rank test (z = 3.996, p = 0.0001) suggest that in spite of significant downgrading, differences between priors and actual states in *Other* rounds are still significantly positive. Thus, we conclude that even though they underestimate the amount, people are aware of the overconfidence existing in others; however they are not able to recognize that they possess the same bias. This finding can be exhibited as a nice example of the phenomenon called *blind spot bias* in psychology (Pronin and Kugler (2007)).

Additionally, we examine whether trusting other people's self-assessments varies by gender, or not. According to a two-sample Mann-Whitney test, there is some evidence that women downgrade the beliefs reported by their SOPs more than men do (z = -1.899, p = 0.0575).

Finally, we compare prior beliefs submitted in *Other* rounds with beliefs reported by the same subjects in the *Self* round. Results show that while average prior beliefs are 67.7% in the *Self* round, it decreases to 60.8% in *Other* rounds. A Wilcoxon signed-rank test supports that most of the subjects believe that their SOPs are less likely to be in the top group in comparison to themselves (z = -3.908, p = 0.0001). However, since the SOP assignments in the final round are completely random, one can argue that the SOP subject pool may be really worse than the restricted sample, in terms of performance.[23] Thus, we decide to restrict this analysis only with the first *Other* round, where an assigned SOP is the most similar subject to the relevant subject, in terms of both actual states and confidence levels. However, according to a Wilcoxon signed-rank test, similar results still hold, that is subjects tend to think that their SOPs are worse than themselves in terms of actual performance (z = -3.805, p = 0.0001). More interestingly, out of 34 instances where subjects hold the same initial confidence levels with their SOPs, 21 cases end up with a downgrading while 9 subjects choose to stay neutral.[24] Therefore, we conclude that most of the subjects think that they are better than their SOPs, even when they have the same confidence levels.

[23] This is actually the case. While the average state in the *Self*round is 56.86%, it decreases to 48.04% in *Other* rounds.

[24] These 34 observations come from 33 different subjects. That is, there is only one subject whose SOPs in both rounds hold the same intial beliefs with her.

Information Processing:

One of the main goals of this paper is to answer if people have any biases in information processing, and if these biases vary by self-relevance of information. As mentioned earlier, economic theory expects rational agents to update their beliefs according to Bayes' rule. Hence, we investigate how much they stick to the theory and at which points they depart from it. To begin with, we would like to mention possible biases which we may encounter in the experimental data. One well-known bias commonly observed in information processing is conservatism. It is known as the phenomenon that subjects updating their beliefs significantly less than what Bayes' rule predicts in either way. Another well documented bias, asymmetry is subjects weighing one signal significantly more than the other, contrary to what theory suggests.

In order to understand if either of these biases is present in our subject pool, we start with an analysis which is also used in Möbius et al. (2011). Remember that the theory predicts that Bayesian agents should update their beliefs according to (1). Therefore, we run the following regressions for the first and second signals, respectively, to determine how much subjects follow the Bayes' rule.[25]

$$logit\ (\mu^2) = \beta_{1,0}\ .\ logit\ (\mu^1) + \beta_{1,T}\ .\ I\ ln(3)\ +\ \beta_{1,B}\ .\ (1 - I)\ ln\left(\frac{1}{3}\right) \tag{19}$$

$$logit\ (\mu^3) = \beta_{2,0}\ .\ logit\ (\mu^2) + \beta_{2,T}\ .\ I\ ln(9)\ +\ \beta_{2,B}\ .\ (1 - I)\ ln\left(\frac{1}{9}\right) \tag{20}$$

Since our reference equation is (1), we expect all coefficient estimates above to be equal to one. Significant deviations from this prediction may be evidence for any of the biases mentioned above. For instance, while ($\beta_{i,T} < 1$ and $\beta_{i,B} < 1$) corresponds to conservatism in belief updating after signal i, $| \beta_{i,T} - \beta_{i,B} | \neq 0$ is a possible evidence of asymmetric updating following signal i.

Table 4 and Table 5 present regression results for the first and second updating tasks, respectively. According to the output, there is clear evidence of conservatism at the aggregate level for each signal type. According to a Wald test, both $\beta_{1,T}$ and $\beta_{1,B}$ are significantly less

[25] We omit observations where subjects report certain beliefs about their actual states, since the logit function is not defined for such cases. In other words, we only consider data where both priors and their corresponding posteriors are in (0, 1).

than unity in the first updating process not considering the type of round ($F = 241.63$, $p = 0$ and $F = 160.45$, $p = 0$). Similar results hold for the second updating task, as well ($F = 186.95$, $p = 0$ and $F = 77.08$, $p = 0$). Thus, we conclude that disregarding self-relevance of information; conservatism seems to be present for each different signal type. In the next steps, we question the existence of conservatism specifically for each round and signal type.

Table 4. OLS Regressions of the 1st Signal Bayesian Updating (respectively for Self, Other, and Aggregate)

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{1,0}$	.9182623	.0502403	18.28	0.000	.8184203	1.018104
$\beta_{1,T}$	.7038584	.0758053	9.29	0.000	.5532113	.8545056
$\beta_{1,B}$	.4758022	.0683300	6.96	0.000	.3400106	.6115938
Number of Observations: 91					*$R^2 = .8919$*	

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{1,0}$	.6484906	.0495334	13.09	0.000	.5507385	.7462427
$\beta_{1,T}$	.8815994	.0658040	13.40	0.000	.7517380	1.011461
$\beta_{1,B}$	.6634548	.0596279	11.13	0.000	.5457816	.7811279
Number of Observations: 180					*$R^2 = .7687$*	

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{1,0}$	.7560884	.0380542	19.87	0.000	.6811651	.8310117
$\beta_{1,T}$	.8245410	.0530438	15.54	0.000	.7201054	.9289767
$\beta_{1,B}$	.6074525	.0479561	12.67	0.000	.5130338	.7018711
Number of Observations: 271					*$R^2 = .7992$*	

Initially, we would like to present our findings for the first updating process. According to a series of Wald tests, the null hypothesis $\beta_{1,T}$ and $\beta_{1,B}$ are both equal to unity is rejected at both *Self* ($F = 86.21$, $p = 0$ and $F = 48.49$, $p = 0$) and *Other* ($F = 179.49$, $p = 0$ and $F = 123.8$, $p = 0$) rounds. That is, conservatism is observable following the first signal, independent of self-relevance of information.

Running the same series of statistical tests on data for the second updating process give similar results. As in the 75% signal accuracy case, we reject that $\beta_{2,T}$ and $\beta_{2,B}$ are both equal to unity at both *Self* ($F = 59.07$, $p = 0$ and $F = 32.41$, $p = 0$) and *Other* ($F = 125.36$, $p = 0$ and $F = 45.36$, $p = 0$) rounds. In short, we find strong evidence for the existence of conservative belief updating in both self-relevant and self-irrelevant contexts, following the second signal.

Again using the regression results presented in Table 4 and Table 5, we find evidence for asymmetric updating at the aggregate level both after the first and second signals. According to a Wald test, $\beta_{1,T}$ and $\beta_{1,B}$ are significantly different from each other in the first updating process at the aggregate level (F = 7.85, p = 0.0055). Similar results also hold for the second updating task, as well (F = 16.75, p = 0.0001). Hence, we conclude that disregarding self-relevance of information; asymmetry seems to be present for each different signal type.

Table 5. OLS Regressions of the 2nd Signal Bayesian Updating (respectively for Self, Other, and Aggregate)

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{2,0}$	.7696795	.1144194	6.73	0.000	.5397452	.9996139
$\beta_{2,T}$	.6314257	.0821549	7.69	0.000	.4663293	.7965222
$\beta_{2,B}$	.4436052	.0779223	5.69	0.000	.2870145	.6001958
Number of Observations: 52					*R^2 = .8457*	

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{2,0}$	.6073992	.1111092	5.47	0.000	.3852950	.8295034
$\beta_{2,T}$	.8066750	.0720462	11.20	0.000	.6626566	.9506934
$\beta_{2,B}$	.4357222	.0646973	6.73	0.000	.3063941	.5650502
Number of Observations: 65					*R^2 = .8279*	

Regressor	Coefficient Estimate	Standard Error	t	P > |t|	[95% Confidence	Interval]
$\beta_{2,0}$	.6713889	.0776837	8.64	0.000	.5174980	.8252797
$\beta_{2,T}$	.7341444	.0536927	13.67	0.000	.6277795	.8405092
$\beta_{2,B}$	.4296457	.0489368	8.78	0.000	.3327024	.5265891
Number of Observations: 117					*R^2 = .8321*	

We start round specific analyses for the first signal with the *Self* round. A Wald test rejects that $\beta_{1,T}$ and $\beta_{1,B}$ are equal to each other under self-relevance of information (F = 4.15, p = 0.0445). Similarly, in *Other* rounds we find that $\beta_{1,T}$ and $\beta_{1,B}$ are significantly different from each other (F = 5.21, p = 0.0237). In other words, asymmetry is observable following the first signal, independent of the varying context of information.

When we apply the same procedures to the data for the second updating task, we get the following results. First, we cannot reject that $\beta_{2,T}$ and $\beta_{2,B}$ are equal to each other in the *Self* round (F = 2.35, p = 0.132). We believe that, we are not able to obtain statistical significance in the *Self* round despite estimated coefficients for $\beta_{2,T}$ (0.631) and $\beta_{2,B}$ (0.444) seem remarkably different, due to limited number of observations for this specific case. On the other hand, according to a Wald test, there is significant evidence for asymmetry in *Other* rounds (F = 15.03, p = 0.0003).

We continue the results on information processing with an analysis of whether deviations from Bayesian benchmarks vary by signal type. According to a Mann-Whitney test, the amount of absolute bias observed in the *Self* round after the first signal is significantly higher when the feedback is negative (z = 4.691, p = 0).[26] The same result holds for *Other* rounds, too (z = 2.667, p = 0.0076). On the other hand, we find that the same finding is valid after the second signal only in *Other* rounds (z = 2.023, p = 0.0431), but not in the *Self* round (z = 1.291, p = 0.1967). Thus, we conclude that getting bad news mostly makes people deviate more from Bayesian benchmarks in comparison to a good news case. However, we are more interested in the direction of the bias. Therefore, we intend to test how the actual bias varies following a positive/negative first signal across different rounds.[27] Results of two-sample t-tests show that, while the actual amount of bias after a negative initial feedback is significantly more in the *Self* round in comparison to *Other* rounds (t = 2.7409, p = 0.0068), the round type does not generate any remarkable difference in the similar case when the feedback is positive (t = -1.0093, p = 0.3147). Also, Figure 7 shows how actual bias varies by differing types of feedback and round. Additionally, Table 6 presents the frequencies of biases by their directions following a negative initial signal. As one can easily see, while the frequency of optimistic updates is 81.5% when the information is self-relevant, it decreases to 56% in self-irrelevant contexts. Also, a Pearson's chi-squared test rejects that the frequency distribution is similar in both contexts (χ^2 = 10.5172, p = 0.005). Hence, we conclude that people are more likely to interpret negative feedback in a self-serving way when the context of information is ego-relevant.

Table 6. Frequencies of Bias Types by Context of Information after a Negative Initial Feedback

Bias Type	*Self*	*Other*	**Total**
Negative	9 (16.67%)	48 (41.38%)	57 (33.53%)
Neutral	1 (1.85%)	3 (2.59%)	4 (2.35%)
Positive	44 (81.48%)	65 (56.03%)	109 (64.12%)
Total	54	116	170

[26] We define 'bias' as the actual difference between a subject's reported belief and its corresponding Bayesian benchmark. On the other hand, 'absolute bias' corresponds to the absolute value of this bias.

[27] We restrict this analysis with the initial feedback, because given that the second signal is bought by a limited number of subjects, we end up with diluted subsamples for different feedback and round types.

Moreover, we run a within-subject analysis to understand if the amount of absolute error made by a subject varies as signal accuracy rates change. According to a couple of Wilcoxon signed-rank tests, while there is significant evidence for the existence of this situation in *Other* rounds (z = -2.53, p = 0.0114), we do not reject that these errors are similar in the *Self* round (z = -1.355, p = 0.1754).

Figure 7. Mean of Actual Bias after the 1st Signal by Round and Feedback Types

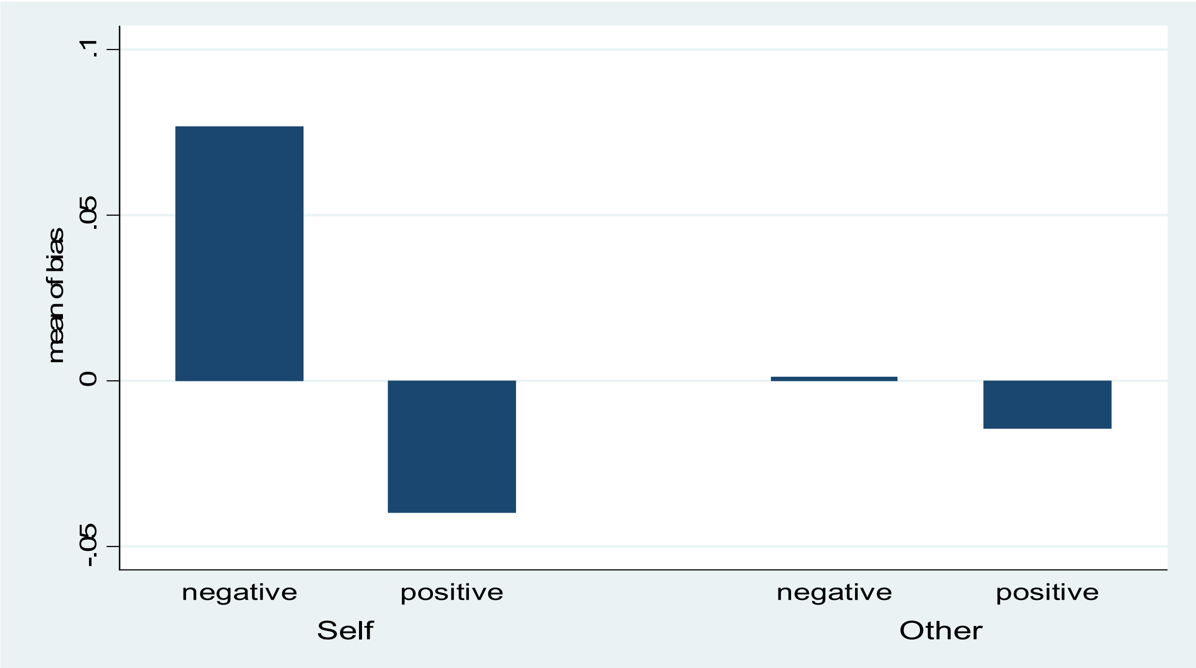

Lastly, we question if there is any relationship with a subject's performance in the IQ test, and her accuracy in belief updating. According to a couple of Mann-Whitney tests, people who happen to be in the bottom half make significantly more absolute errors under self-relevance both after the first and second signals (z = 2.724, p = 0.0065 and z = 2.976, p = 0.0029). On the other hand, there is some evidence for the same result in the *Other* rounds following the first signal, but not the second one (z = 1.785, p = 0.0743 and z = 1.16, p = 0.2461).

<u>*Demand for Information:*</u>

Our initial observation on information demand is the existence of severe overbidding in all rounds. According to a Wilcoxon signed-rank test, amounts offered for information are significantly higher than theoretical predictions (z = 12.625, p = 0).[28] Results show that out of 306 cases, subjects offered more than theoretical predictions 236 times, while they underbid

[28] We use the theoretical values calculated under the assumption of CARA utility function.

only 51 times. Additionally, this situation does not seem to arise from a specific round, but clearly observable in both rounds, as suggested by Figure 8.

Figure 8. Actual Bids vs. Theoretical Bids by Round Type

Our second finding is that negative signals in the *Self* round exacerbate overbidding. According to Mann-Whitney tests, while there is no significant difference in theoretical values after different signals ($z = 1.149$, $p = 0.2507$), people seem to offer somewhat more when they get negative feedback ($z = 1.801$, $p = 0.0717$). Additionally, testing directly for deviations from theoretical predictions after different signals via a two-sample t-test supports presence of some difference ($t = 1.6922$, $p = 0.0937$). On the other hand, in *Other* rounds, there seems no significant difference in amounts offered for information after different signals according to a Mann-Whitney test ($z = 1.019$, $p = 0.3084$). Also, an unpaired t-test on deviations from theoretical predictions supports that signal type does not affect overbidding in *Other* rounds ($t = 1.3207$, $p = 0.1881$). As a result, we conclude that even though overbidding is present in both rounds, it gets even worse when subjects get negative feedback about their own performance.

Having found that the bias in demand for information (overbidding) increases after a negative feedback only when the relevant performance is self-related, make us think that people may have more difficulty in conceding a bad outcome when it is about self-esteem in comparison to a case with no ego-relevance. Therefore, we suspect that this bias might be more prevalent in subjects with high confidence levels. Statistical tests on 37 subjects, whose initial beliefs are at least 75%, show that while the average offer after a positive signal is around 11.5 ECU, it jumps to 24.5 ECU after negative feedback. A slight increase would be understandable, since theory also predicts an increase for these subjects after a disconfirming signal. However, the increase is quite high, and creates some difference in terms of deviations from the theoretical predictions in comparison to the confirming signal case. A Mann-Whitney test seems to present statistical evidences in favor of the above finding ($z = 1.906$, $p = 0.0566$). Also note that, similar tests on other quartiles show no difference in deviations from theoretical values depending on the signal type.

Lastly, we try to figure out if delegation leads to any change in demand for information. In order to understand that, we run a couple of non-parametric paired tests considering subjects in *Other* rounds and their corresponding SOPs in *the Self* round, of course given that they are fed with the same type of signal. Firstly, a Wilcoxon signed-rank test reveals that there is no significant difference regarding theoretical predictions for the two cases ($z = -1.268$, $p = 0.2049$). However, according to the same test, people in self-irrelevant rounds bid significantly less than their corresponding SOPs ($z = -2.007$, $p = 0.0447$). Having said that there is severe overbidding in all rounds, we conclude that delegating information collection might decrease the intensity of overbidding.

Auxiliary Gender Differences:

Beside analyses on gender differences mentioned so far, we want to present some additional results.

Initially, as shown in the following table, there is no clue of a remarkable difference in performance between different genders. Also, series of Mann-Whitney tests show that there is no significant difference in questions attempted to solve ($z = -0.551$, $p = 0.5814$), answered correctly ($z = -0.55$, $p = 0.5823$), and wrongly ($z = -0.332$, $p = 0.7402$) across men and women. As a natural consequence, according to a Mann-Whitney test there is no significant difference in scores of opposite sexes, as well ($z = -0.09$, $p = 0.9284$).

Table 7. Performance Results Broken by Genders

gender	*attempted*			*corrects*			*mistakes*			*score*		
	min	*avg*	*max*	*min*	*avg*	*max*	*min*	*avg*	*max*	*min*	*avg*	*max*
female	6	11.51	18	2	5.97	10	0	1.87	6	0.5	5.04	9.5
male	6	11.98	30	3	6.13	10	0	2.16	11	0.5	5.05	9

Secondly, we find evidence for women being more risk-averse than men. According to the Holt-Laury task results, while women seem to switch to the risky lottery after 6.21 safe choices on average, men do the same after 5.57 safe choices. A Mann-Whitney test supports that women are a bit more risk-averse, too ($z = 1.942$, $p = 0.0522$).

Thirdly, we test if there is a gender difference in willingness to learn the exact state at the end of the experiment without any payment. While 17.9% of women refuse to learn their actual states, only 11.1% of men do the same. However, according to a Mann-Whitney test this difference is not enough to reject the null hypothesis that there is no difference across genders in willingness to learn the actual performance ($z = -0.97$, $p = 0.3318$). Additionally, having found that most of the people are into learning their actual states, we conclude that unlike Möbius et al (2011), there is no remarkable evidence of information aversion in our subject pool. However, we should also note that Möbius et al. (2011) argues more confident agents are less information-averse. Having already shown that there is significant overconfidence in our subject pool; absence of significant information aversion seems to be in line with Möbius et al. (2011).

Lastly, we analyze if any of biases in information processing and valuation is more prevalent in a specific gender. However, statistical tests do not present any evidence of such a difference. For instance, a Mann-Whitney test on absolute amount of deviations from Bayesian benchmarks in the first belief updating task of the *Self* round show that there is no significant difference between genders ($z = 1.504$, $p = 0.1327$). Moreover, according to a Mann-Whitney test there is no significant difference in amounts offered for information in *Self* round across genders ($z = 0.176$, $p = 0.8683$).

V. Conclusion

Decision making under uncertainty constitutes considerable proportion of the decisions in real life. Many of these decisions are relevant to self-esteem since the underlying quantities are mostly personal features like intelligence, beauty, creativity, etc. Therefore, it is important to know how people form their beliefs in the light of new piece of information when the context is self-relevant. What we wonder is if they hold correct initial beliefs about themselves and correctly process new information. If they do both, then there is no problem at the point of decision making. However, if there is a bias in any of these channels, then we may need to consider delegating information collection and processing to disinterested parties in some situations, in order to improve social efficiency.

Our design tries to find out if these two channels – prior belief formation and information processing – work just like the standard economic theory predicts. Moreover, it questions if people value information, which has an instrumental value as they ought to do. We are able to calculate theoretical value of information for both risk-neutrality and risk aversion assumptions. The Holt-Laury task enables us to obtain theoretical value of information for different risk segments.

Using the experimental procedures detailed in this study, we find that people hold overconfident initial beliefs, although they are aware of the existence of the same bias in other people. Additionally, we find evidence for the presence of both conservatism and asymmetry in processing new piece of information both when the information is self-relevant and self-irrelevant. Moreover, people seem to exhibit remarkable overbidding for information in either context. Furthermore, we detect that when a subject gets a negative signal in a self-relevant context, she is more likely to incorporate new information into existing beliefs in a self-serving way, and overvalue a new signal significantly more.

One last thing to note is that our research does consider only monetary payoffs, but not *the anticipatory utility* phenomenon. That is, as analyzing information processing and demand for information, we only focus on how the success of investment decision changes. However, as noted in Möbius et al. (2011), even if the optimistic asymmetry in belief updating, and incorrect valuation in information purchasing are harmful in terms of future monetary payoffs, they might be increasing the anticipatory utilities of subjects. Hence, in a future research, we might try to integrate anticipatory utility in our design to answer some questions more accurately.

Appendices

Appendix 1. Details on Holt-Laury Task

Table 8. Payoffs at Holt-Laury Task

Decision	Option A (Safe)	Option B (Risky)
1	1/10 of 40 ECU, 9/10 of 32 ECU	1/10 of 77 ECU, 9/10 of 2 ECU
2	2/10 of 40 ECU, 8/10 of 32 ECU	2/10 of 77 ECU, 8/10 of 2 ECU
3	3/10 of 40 ECU, 7/10 of 32 ECU	3/10 of 77 ECU, 7/10 of 2 ECU
4	4/10 of 40 ECU, 6/10 of 32 ECU	4/10 of 77 ECU, 6/10 of 2 ECU
5	5/10 of 40 ECU, 5/10 of 32 ECU	5/10 of 77 ECU, 5/10 of 2 ECU
6	6/10 of 40 ECU, 4/10 of 32 ECU	6/10 of 77 ECU, 4/10 of 2 ECU
7	7/10 of 40 ECU, 3/10 of 32 ECU	7/10 of 77 ECU, 3/10 of 2 ECU
8	8/10 of 40 ECU, 2/10 of 32 ECU	8/10 of 77 ECU, 2/10 of 2 ECU
9	9/10 of 40 ECU, 1/10 of 32 ECU	9/10 of 77 ECU, 1/10 of 2 ECU
10	10/10 of 40 ECU, 0/10 of 32 ECU	10/10 of 77 ECU, 0/10 of 2 ECU

In Holt-Laury task, subjects choose one of two options in each decision. Then, a random decision among ten is selected by computer to determine which round will be subject to payment. Eventually, nature decides if subject takes the good or bad outcome from the lottery chosen by the subject at that decision.

As suggested in Holt and Laury (2002) we use the crossover point to risky lottery to estimate risk aversion degrees. Having assumed that subject have CARA type utility functions, we get the following equality to show when a subject is indifferent between two lotteries:[29]

$$(1 - P_i).e^{-32r} + P_i.e^{-40r} = (1 - P_i).e^{-2r} + P_i.e^{-77r} \tag{A1}$$

which simplifies to:

$$e^{-77r}(-e^{45r}(-1 + P_i) + e^{75r}(-1 + P_i) - P_i + e^{37r}.P_i) = 0 \tag{A2}$$

is solved for P_i as:

$$P_i = \frac{e^{45r}(-1 + e^{30r})}{-1 + e^{37r} - e^{45r} + e^{75r}} \tag{A3}$$

[29] P_i denotes the probability of taking the higher reward in any lottery.

Therefore, if a subject prefers lottery A at decision n, but switch to lottery B at decision (n+1) there is no question at assuming,[30]

$$\frac{n}{10} \leq P_i \leq \frac{n+1}{10} \tag{A4}$$

However, we prefer to assign a single risk parameter to all subjects, who make n safe choices. Thus, we choose the midpoint as the probability, which makes the subject indifferent between safe and risky lotteries.

$$P_i = \frac{n}{10} + 0.05 \tag{A5}$$

Using the data from Holt-Laury task we are able to calculate P_i for each subject. Then solving (A3) for each P_i we get corresponding risk parameters.

Table 9. Holt-Laury Task Decisions vs. Corresponding Risk Parameters

# of safe choices	r
0	-0.05000
1	-0.03614
2	-0.02121
3	-0.00982
4	0.00022
5	0.00994
6	0.02016
7	0.03200
8	0.04789
9	0.06000

Appendix 2. Derivations of p_G and p_B

If a Bayesian agent with confidence level p gets a positive signal with 90% accuracy, she is expected to update her belief to p_G as in following calculations.

$$logit\ (p_G) = logit\ (p) + \ ln(9) \tag{A6}$$

using (2) we get,

[30] This assumption is made only for 0<n<9. For cases where n=0 or n=9, we choose a reasonable amount to satisfy the corresponding one-sided inequality.

$$ln\,(\frac{p_G}{1-p_G}) = ln\,(\frac{p}{1-p}) + ln(9) \tag{A7}$$

or more compactly,

$$ln\,(\frac{p_G}{1-p_G}) = ln\,(\frac{9p}{1-p}) \tag{A8}$$

which is solved for p_G as:

$$p_G = \frac{9p}{8p+1} \tag{A9}$$

On the other hand, if a Bayesian agent with confidence level p gets a negative signal with 90% accuracy, she is expected to update her belief to p_B as in following calculations.

$$logit\,(p_B) = logit\,(p) - ln(9) \tag{A10}$$

using (2) we get,

$$ln\,(\frac{p_B}{1-p_B}) = ln\,(\frac{p}{1-p}) - ln(9) \tag{A11}$$

or more compactly,

$$ln\,(\frac{p_B}{1-p_B}) = ln\,(\frac{p}{9-9p}) \tag{A12}$$

which is solved for p_B as:

$$p_B = \frac{p}{9-8p} \tag{A13}$$

<u>Appendix 3.</u> Derivations and Range of ΔP

Regarding that the subject believes she is in top half with probability p, and signal structure imposes 90% accuracy, she must calculate the change in probability of taking R as:

$$\Delta P = (1-p).\left[0.9.\left(P(p_B|\theta=0)-P(p|\theta=0)\right)+0.1.\left(P(p_G|\theta=0)-P(p|\theta=0)\right)\right]$$
$$+(p).\left[0.1.\left(P(p_B|\theta=1)-P(p|\theta=1)\right)+0.9.\left(P(p_G|\theta=1)-P(p|\theta=1)\right)\right] \tag{A14}$$

inserting (4) into (A14) leads to:

$$\begin{aligned}\Delta P = (1-p).&\left[0.9.\left(\frac{(1-p_B)(1+p_B)}{2}-\frac{(1-p)(1+p)}{2}\right)\right.\\ &\left.+0.1.\left(\frac{(1-p_G)(1+p_G)}{2}-\frac{(1-p)(1+p)}{2}\right)\right]\\ &+(p).\left[0.1.\left(p_B+\frac{(1-p_B)(1+p_B)}{2}-p-\frac{(1-p)(1+p)}{2}\right)\right.\\ &\left.+0.9.\left(p_G+\frac{(1-p_G)(1+p_G)}{2}-p-\frac{(1-p)(1+p)}{2}\right)\right]\end{aligned} \tag{A15}$$

or more compactly,

$$\Delta P = (1-p).\left[0.9.\left(\frac{p^2-p_B^2}{2}\right)+0.1.\left(\frac{p^2-p_G^2}{2}\right)\right]$$
$$+(p).\left[0.1.\left(p_B-p+\left(\frac{p^2-p_B^2}{2}\right)\right)+0.9.\left(p_G-p+\left(\frac{p^2-p_G^2}{2}\right)\right)\right] \tag{A16}$$

inserting (A9) and (A13) into (A16) leads to,

$$\Delta P = (1-p)\left[0.45.\left(p^2-\frac{p^2}{(9-8p)^2}\right)+0.05.\left(p^2-\frac{81p^2}{(1+8p)^2}\right)\right]$$
$$+(p)\left[0.1.(-p+\frac{p}{9-8p}+\frac{1}{2}(p^2-\frac{p^2}{(9-8p)^2}))+0.9.(-p+\frac{9p}{1+8p}+\frac{1}{2}(p^2-\frac{81p^2}{(1+8p)^2}))\right] \tag{A17}$$

which is solved as:

$$\Delta P = \frac{p^2(288 + 1472p - 5856p^2 + 6144p^3 - 2048p^4)}{(9 + 64p - 64p^2)^2} \quad \text{(A18)}$$

Table 10. Beliefs and Corresponding Changes in Probability of Earning the Reward in Case of Purchasing Information

Belief	ΔP	Belief	ΔP	Belief	ΔP	Belief	ΔP
0.00	0.00000	-	-	-	-	-	-
0.01	0.00033	0.26	0.05558	0.51	0.07996	0.76	0.05150
0.02	0.00120	0.27	0.05751	0.52	0.07983	0.77	0.04936
0.03	0.00249	0.28	0.05938	0.53	0.07961	0.78	0.04716
0.04	0.00412	0.29	0.06117	0.54	0.07930	0.79	0.04489
0.05	0.00600	0.30	0.06289	0.55	0.07891	0.80	0.04258
0.06	0.00807	0.31	0.06453	0.56	0.07844	0.81	0.04021
0.07	0.01030	0.32	0.06609	0.57	0.07787	0.82	0.03779
0.08	0.01264	0.33	0.06757	0.58	0.07722	0.83	0.03533
0.09	0.01507	0.34	0.06898	0.59	0.07649	0.84	0.03284
0.10	0.01756	0.35	0.07030	0.60	0.07567	0.85	0.03031
0.11	0.02009	0.36	0.07154	0.61	0.07476	0.86	0.02777
0.12	0.02264	0.37	0.07269	0.62	0.07377	0.87	0.02521
0.13	0.02521	0.38	0.07377	0.63	0.07269	0.88	0.02264
0.14	0.02777	0.39	0.07476	0.64	0.07154	0.89	0.02009
0.15	0.03031	0.40	0.07567	0.65	0.07030	0.90	0.01756
0.16	0.03284	0.41	0.07649	0.66	0.06898	0.91	0.01507
0.17	0.03533	0.42	0.07722	0.67	0.06757	0.92	0.01264
0.18	0.03779	0.43	0.07787	0.68	0.06609	0.93	0.01030
0.19	0.04021	0.44	0.07844	0.69	0.06453	0.94	0.00807
0.20	0.04258	0.45	0.07891	0.70	0.06289	0.95	0.00600
0.21	0.04489	0.46	0.07930	0.71	0.06117	0.96	0.00412
0.22	0.04716	0.47	0.07961	0.72	0.05938	0.97	0.00249
0.23	0.04936	0.48	0.07983	0.73	0.05751	0.98	0.00120
0.24	0.05150	0.49	0.07996	0.74	0.05558	0.99	0.00033
0.25	0.05357	0.50	0.08000	0.75	0.05357	1.00	0.00000

Appendix 4. Instructions

[Original instructions were in Turkish and they are available upon request.]

Welcome. Thank you for participating in our experiment. The experiment is on individual decision-making. During the study you are expected to solve some questions in a test and make decisions regarding your performance on that test. Your earnings in the experiment will depend on your performance, accuracy of your decisions, and chance. There is not any kind of misleading or deception in this study. The rules that will be explained soon are totally correct and your payments will be determined according to them.

Participation in this study is entirely voluntary. Following the explanation of rules, you are free to quit until the experiment starts. Any of your decisions during the experiment will not be matched with your real identity, but rather will be recorded under an anonymous subject number. During the experiment, all monetary activities will be done in ECU (experimental currency unit) where 1 ECU equals to ₺0.03. Payments will be made privately in cash at the end of the experiment, and nobody will be informed about your earnings, or your performance.

The experiment consists of five main parts. However, before they start, you will go through a quiz controlling your understanding of the procedures, and a practice round to get you familiar with the program. Shortly after, we will go through a detailed explanation of the first part. Every part will be conducted just after the instructions on it are given. Please do not hesitate to ask questions if you do not understand a particular detail. Also keep in mind that any kind of communication with other participants is prohibited.

Part 1:

In this part, you will perform in an IQ test, which consists of 30 questions and you are expected to answer correctly as many as you can in 6 minutes. Only one question will be displayed on the screen at a time, and you will see the following one as soon as you report your answer or choose to skip the current question. While some of the questions in the test are open-ended, most of them are multiple-choice questions. Wrong answers reported in multiple-choice questions will be subject to a penalty; however incorrect answers in open-ended questions will not affect your total score in any way. In short, your total score in the IQ test will be calculated as follows:

$$Score = \text{\# of correct answers in the whole test} - \frac{\text{\# of incorrect answers in multiple choice questions}}{2}$$

Your payoff in the IQ test will be your score multiplied by 20 ECU. At the end of the IQ test, computer will rank all subjects according to their total scores, and ties will be broken at random. While subjects being among top 50% in rankings will be labeled as a TOP type, the others will be BOTTOM types.

Part 2:

In this part, you are expected to make 3 investment decisions, from which only one will be selected randomly to be actually paid. That requires you to evaluate each decision independently. The essence of each investment decision is the same that is guessing your state accurately increases your chance to take the reward. At each decision success of your investment will depend on one of the following lotteries:

Lottery 1: You take 200 ECU if you are in TOP group and nothing if you are in BOTTOM group.

Lottery 2: You take 200 ECU with probability X and nothing with probability (1-X). The amount of X will be determined by the computer in each round randomly between 0 and 1. X has a uniform distribution that is each number in the interval [0, 1] will be equally likely.

The only thing you will decide is to determine the minimum amount of X to choose the second lottery. Just after your decision, the computer will decide the amount of X. If X turns out to be less than your report, success of your investment will depend on Lottery 1 (or your performance in the IQ test). Otherwise, the success of your investment decision will depend on Lottery 2 (or chance).

You are going to make your first investment decision as soon as this part starts. Following your first decision, a signal stating which half you are actually in will be sent to you. However, this signal has a 75% accuracy that is the transmitted feedback will be telling the truth with 75% probability. On the other hand, the feedback will be wrong with 25% probability. Following this feedback you will be asked to make your second investment decision. After the second investment decision, you will have a chance to get a new signal which is accurate with 90% and will be helpful for the last investment decision. However, this signal will be subject to a fee, which will be randomly determined by the computer in [0, 50] ECU interval. Before the price is determined, you will be asked to report your maximum willingness to pay for this particular information. If the amount reported by you happens to be smaller than the price determined by the computer, you will lose your chance to get further information and not be able to update your investment decision. In other words, your report in the second investment decision will be automatically accepted as your final investment decision, as well. However, if the price of information does not exceed your maximum offer for the information, you will get the information at the exact price determined by the computer. Following that signal you will be asked to make your final investment decision. Moreover, you will be given 50 ECU at the start of this part to compensate your possible expenditures for information purchasing. However, you do not have to spend the whole amount, because you are free to save any amount you want, instead of spending. Nevertheless, note that any saving/spending decision will be relevant only if the final investment decision is chosen to be paid. Otherwise, we will take 50 ECU back from your account, and return the fee of information if you happen to bought it. At the end of this part, the computer will choose one of three investment decisions randomly, and depending on the conditions of that particular investment decision, you will get 200 ECU or nothing. Even though this earning will be directly transferred to your account, you will not be able to learn any monetary payoff until the end of the experiment.

Part 3-4:

In this part, another participant in this session will be assigned to you, and s/he will be called your SOP from now on. At the very beginning, a message declaring the first investment decision of your SOP in Part 2 will be sent to you. Every procedure will be the same as in Part 2, only this time the relevant performance will be your SOP's rather than your own performance. That is to say, Lottery 1 will depend on which half your SOP belongs to, and you will get signals about whether your SOP is in the TOP half or BOTTOM half.

Part 5:

In this part a menu of lottery choices are listed on your screen and you are asked to make 10 decisions in total. At each decision, you will be making a choice between paired lotteries A and B, where Lottery B is riskier in the sense that variation between the good outcome and bad outcome is higher in comparison to Lottery A. As soon as you report your decisions the computer will choose one decision among ten to make an actual payment. Depending on which lottery you have chosen at that decision and the conditions of that particular lottery, you will get either the good outcome or the bad outcome.

The decisions to be made are over. Now you are kindly asked to answer survey questions that will be on your screen soon. Thank you again for your participation.

Appendix 5. Post-Experiment Survey Questions

1. How old are you?

2. What is your major?[31]

3. What is your year/class?[32]

4. What is your gender?[33]

5. What is your CGPA?[34]

6. Aside from monetary payoffs, how important was it for you to be in the TOP half? Please answer on a scale of 0 to 10: 0=not important at all, 10=extremely important.

7. How difficult did you find the IQ test? Please answer on a scale of 0 to 10: 0=extremely easy, 10=extremely hard.

8. How important was monetary payoffs for you in the experiment? Please answer on a scale of 0 to 10: 0=not important at all, 10=extremely important.

Table 11. Summary of Answers to the Post-Experiment Survey

Question	*unrestricted sample*			*restricted sample*		
	min	*avg*	*max*	*min*	*avg*	*max*
1	17	21.74	30	17	21.72	30
2	0	0.36	1	0	0.43	1
3	0	3.02	6	0	3.10	6
4	0	0.59	1	0	0.62	1
5	1	2.62	5	1	2.45	5
6	0	5.99	10	0	6.27	10
7	0	4.46	10	0	4.69	10
8	0	2.87	10	0	3.32	10

[31] Subjects enter exact names of their departments in a text box. However, in Table 11 they are divided into two main groups, where students from social sciences, law, etc. are denoted by O, and students from natural sciences, medicine, engineering, etc. are denoted by 1.

[32] 0 corresponds to preparatory, 1-4 stands for corresponding years in undergraduate study, 5 corresponds to being a master's student, and 6 is for Ph.D. students.

[33] While females are denoted with 0, males are denoted with 1.

[34] We categorize subjects according to their CGPAs as follows: {[3.5, 4], [3, 3.5), [2.5, 3), [2, 2.5), [0, 2)} → {1, 2, 3, 4, 5}

References

Becker, Gordon, Morris DeGroot, and Jacob Marschak. (1964). "Measuring Utility by a Single-response Sequential Method." *Behavioral Science*, 9(3), 226–232.

Eil, David, and Justin M. Rao. (2011). "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself." *American Economic Journal: Microeconomics*, 3(2), 114–138.

Ertac, Seda. (2011). "Does Self-Relevance Affect Information Processing? Experimental Evidence on the Response to Performance and Non-Performance Feedback." *Journal of Economic Behavior and Organization,* 80(3), 532-545.

Fischbacher, Urs. (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2), 171–178.

Grether, David M. (1980). "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." *Quarterly Journal of Economics*, 95, 537–557.

Grossman, Zachary and David Owens. (2010). "An Unlucky Feeling: Overconfidence and Noisy Feedback." working paper, UC Santa Barbara.

Holt, Charles A. and Susan K. Laury. (2002). "Risk Aversion and Incentive Effects." *American Economic Review,* 92(5), 1644-1655.

Karni, Edi. (2009). "A Mechanism for Eliciting Probabilities." *Econometrica,* 77(2), 603-606.

Köszegi, Botond. (2006). "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association*, 4(4), 673-707.

Möbius, Markus, Muriel Niederle, Paul Niehaus, and Tanya Rosenblatt. (2011). "Managing Self-Confidence: Theory and Experimental Evidence." *Working paper.*

Pronin, Emily and Matthew B. Kugler. (2007). "Valuing Thoughts, Ignoring Behavior: The Introspection Illusion as a Source of the Bias Blind Spot." *Journal of Experimental Social Psychology,* 43(4), 565-578.

Tversky, Amos and Daniel Kahneman. (1974). "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185, 1124–1131.