



HAL
open science

Updating beliefs with imperfect signals: experimental evidence

François Poinas, Julie Rosaz, Béatrice Roussillon

► **To cite this version:**

François Poinas, Julie Rosaz, Béatrice Roussillon. Updating beliefs with imperfect signals: experimental evidence. 2010. halshs-00550350

HAL Id: halshs-00550350

<https://shs.hal.science/halshs-00550350>

Submitted on 27 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

W P 1033

Updating beliefs with imperfect signals:
experimental evidence

François Poinas, Julie Rosaz, Béatrice Roussillon

Décembre 2010

GATE Groupe d'Analyse et de Théorie Économique Lyon-St Étienne

93, chemin des Mouilles 69130 Ecully – France

Tel. +33 (0)4 72 86 60 60

Fax +33 (0)4 72 86 60 90

6, rue Basse des Rives 42023 Saint-Etienne cedex 02 – France

Tel. +33 (0)4 77 42 19 60

Fax. +33 (0)4 77 42 19 50

Messagerie électronique / Email : gate@gate.cnrs.fr

Téléchargement / Download : <http://www.gate.cnrs.fr> – Publications / Working Papers

Updating beliefs with imperfect signals: experimental evidence*

François Poinas [†] Julie Rosaz [‡] Béatrice Roussillon [§]

November 22, 2010

Abstract

This article analyses belief updating when agents receive a signal that restricts the number of possible states of the world. We create an experiment on individual choice under uncertainty. In this experiment, the subject observes an urn, containing yellow and blue balls, whose composition is partially revealed. The subject has to assess the composition of the urn and form an initial belief. Then, he receives a signal that restricts the set of the possible urns from which the initial observed sample is drawn. Once again, he has to estimate the composition of the urn. Our results show that, on the whole, this type of signal increases the frequency of correct assessment. However, differences appear between validating and invalidating signals (i.e. signals that either confirm or disprove the initial belief). The later significantly increase the probability to make a correct assessment whereas validating signals reduce the frequency of correct estimations. We find evidences of lack of persistence in choice under uncertainty. The literature shows that people may persist with their choice even when they are wrong. We show that they may also change even if they are right.

JEL Classification: D83, C91

Keywords: Beliefs, Imperfect Information, Experiment

*We are grateful to Marie-Claire Villeval, Stéphane Robin and Anne Rozan for their precious comments. We are also grateful to participants at the GATE research seminars, the ESA meetings in Washington. Financial support from the “Chaire Universitaire en Economie Expérimentale : Finance industrielle et design de réseaux” is gratefully acknowledged.

[†]Toulouse School of Economics (University Toulouse 1 - GREMAQ) - Toulouse F-31000, France - Email: francois.poinas@tse-fr.eu

[‡]Université de Lyon, Université Lyon 2, F - 69007, Lyon, France. CNRS, GATE Lyon-St Etienne, UMR n° 5824, 69130 Ecully, France. Email: rosaz@gate.cnrs.fr

[§]Université Grenoble 2, UMR 1215 GAEL, F-38000 Grenoble, France. INRA, UMR 1215 GAEL, F-38000 Grenoble, France.

I Introduction

We create an assessment problem where there are several (possible) states of the world, and agents have to determine which state they are facing. An agent receives the following two items of information successively: first, a sample of observations from the state of the nature, and second, a signal that restricts the set of the possible states of the nature. With the benefit of the sample of observations from the state of the nature, the agent forms an initial belief and afterward receives the signal. The agent's initial belief can be included or not in the new subset of possible states of the nature, after receiving the signal. The objective of the paper is to analyze how this type of signal is used by the agents to update their belief about the state of the nature.

In a context of incomplete information, the optimal updating behavior is determined by the application of Bayes' theorem. Nevertheless, experiments in both psychology and economics have shown that this theorem is rarely perfectly implemented in decisions taken under uncertainty. Kahneman, Slovic, and Tversky (1982) describe heuristics explaining consistent observed behaviors in probability updating. In uncertain situations, information is key for any economic agent, and may take several forms. It may be the probability of an event, an indication on the move of a player or a restriction of the possible states of the nature.

Signals that reduce the possible states of the nature are found in any situation where an object or a person undertakes a pass/fail test, such as a label, a diploma or an audit notation. This type of test is supposed to improve the evaluation of an object or a person. Many institutions, such as scientific journals, human resource consulting, even state medical licensing, adopt minimum certification standards¹ or pass-fail strategies. These signals do not allow precise quality assessment of an object or a person but indicate whether the object is just above the minimum standard required.

For example, we can consider this problem in the perspective of an employer/employee relationship. The employer, by observing the employee working in his firm, forms an initial belief about the employee's productivity. The employee's productivity may be difficult to assess due to a bad observation or multi-tasking. Thus, an employee's performance appraisal may be undertaken as a signal on the productivity of the worker sent to the employer. For a seller, the performance appraisal can take the form of a sales' objective. What information about the employee's productivity does the employer extract from the observation that the employee succeeds or fails to achieve his sales objectives? This observation does not permit to determine the precise level of the employee's productivity but reduces the range of possible

¹A standard is defined as a minimum level to achieve.

productivity levels.

We examine how this type of signal is used by the agents to update their belief by studying a two-stage game in a laboratory setting. Initially, the subjects observe a partially revealed urn composed of 20 balls, yellow and blue. They form an initial belief about the urn composition and have to estimate the total number of yellow balls contained in the urn. Then, they receive a true but imperfect signal about the urn's composition and have to estimate again the number of yellow balls contained in the urn.

We find that sending these signals improve the subject predictions and especially signals that disprove their initial belief. However, signals that confirm the initial belief may give rise to non optimal updating. In 49% of the cases subjects change their belief when the signal validates their initial belief, which is in contradiction with theoretical predictions. This is particularly true when the initial uncertainty level about the state of the nature is high and the signal does not reduce this uncertainty level. Thus, our result suggest hesitation and choice changes whereas the literature seems to show persistence in choice under uncertainty, in particular when they are made voluntarily. Investigating this behavior, we find that the absence of initial belief reinforcement by a validating signal seems to trigger belief updating.

This paper is organized as follows. Section II studies the literature on updating beliefs. Section III describes the experimental design and delivers predictions. The data are analyzed in Section IV. In Section V, we discuss our results and conclude.

II Related Literature

When facing uncertainty, economic agent use the available information to form their beliefs. In theory, they should use Baye's rule to update their belief when they receive new information. However, many experiments in psychology and in economics have shown that people are not perfect Bayesian thinkers and tend to use heuristics to deal with uncertain events. Kahneman, Slovic, and Tversky (1982) analyze the three heuristics used in the updating process: representativeness, availability, and anchoring. We will focus our attention on the representativeness, which is the most related to our work. The intuition behind representativeness is that people tend to associate the observed sample to the population it looks like the most. When the representativeness and Bayes reasoning are not in line, the former may drive the updating process to non-maximizing behavior. The representativeness leads to several errors. Besides, Kahneman, Slovic, and Tversky (1982) underline insensitivity to prior probability of outcomes and the insensitivity of the sample size.

To study the insensitivity to prior probability, Kahneman and Tversky (1973) conduct an

experiment where two groups of subjects face different descriptions of people. The first group knows that a population of 100 people is composed by 30 engineers and 70 lawyers. The second group knows that the 100 people population is composed by 70 engineers and 30 lawyers. The participants in both groups have to give the probability that a certain description corresponds to an engineer. When the description is uninformative, the participants in the first group should declare a probability of 0.3 and the participants in the second group should declare a probability of 0.7 according to the Bayes' rule. However, the results of the experiment are really different since the participants in both groups declare a probability of 0.5. The prior probability stated for the first(second) group that there are 30%(70%) of engineers in the population is ignored by the participants. Grether (1980) confirms with an experiment conforming economics rules that individuals do not use the base rate when updating belief.

When using the representativeness, people tend to be insensitive to the sample size to judge the probability of the sample to be a good representation of the urn (Tversky and Kahneman (1971)). A high proportion of one color in the sample, even if the sample is small, induces the subjects to believe that it comes from the urn that has a large amount of balls of that color, and they believe it more than if they have observed a bigger sample but with less extreme proportion². Thus, the composition of the sample seems to matter as much as the size of the sample. In our experiment, composition and size of the sample vary to identify for this possible source of error.

Finally, the underestimation of the impact of evidence has also been observed repeatedly and is labeled "conservatism". Individuals are underweighting the likelihood of information (Grether (1980)), they update but less than that they should do. Indeed, conservatism is the underconfidence that results when people underemphasize the big size of a sample of weak proofs. Without specifically testing for the conservatism, our experiment wishes to investigate the importance of the signal in comparison with the agent's initial belief in the updating process. The question of what makes people modify their behavior or change their belief is crucial in economics. With the concept of conservatism, we understand that people may unreasonably keep with the same belief or continue with the same behavior regardless the new information they receive.

²For instance, consider a urn composed of balls of two colors: $2/3$ are red and $1/3$ are white. A first subject draws 6 balls (5 are red and 1 is white), and a second subject draws 20 balls (12 are red and 8 are white). Asking individuals who of the two subjects should feel more confident that the urn where the sample has been drawn contains $2/3$ of red balls, they answer the first one.

Charness and Levin (2005) and Friedman (1998) have also shown the persistence in subject choices, and thus in belief, even when Bayesian updating requires changes. At the opposite, our results suggest changes in subject choices, even when Bayesian updating requires persistence.

Our design shares similarity with Charness and Levin (2005) as we also ask for repeated assessment of a state of the nature represented by a urn and balls. However, they examine how the difference between clashes and reinforcement affects the propensity to use the Bayes' rule whereas we look at the importance of another type of signal. The reinforcement supposes that people are more likely to choose actions associated with successful past outcomes. In their experiment, the subjects face two equally likely states of the world, Up and Down, and two urns, Left and Right, containing black or white balls. The Up state of the world is characterized by more black balls in both urns. Moreover, the Left urn has a mixed composition whereas the Right urn contains balls of only one color (only black balls in Up and only white balls in Down) indicating the state of the world for sure. Two draws are done and the subjects choose whether to continue drawing in the same urn or to switch. Drawing a black ball is considered as a success. All the subjects start by drawing a ball in the Left urn. The participants should draw a second time in the Left urn after a white ball: failure, and switch to the Right urn: success. They find that when Bayesian updating and reinforcement are not aligned, around 50% of the decisions are inconsistent with bayesian updating. They also find that when the urn used for the first draw is a choice, it is more likely to be chosen again than when the urn is imposed to the participant. This experiment points out the role of the past success and the degree of freedom in the decision of people to change or not an initial estimation. In our experiment, we do not play on these two aspects in order to focus on the signal impact and the initial belief precision.

The consistency of a subject choice has also being studied by Friedman (1998) in his analysis of the three-door paradox. The subject faces three face-down cards. The experimentalist turns a non-winning card that the subject did not choose. Then the subject has the opportunity to switch between the two remaining cards. Finally the two cards are turned over. The participants play 10 trials. The Bayes' rule predicts that the subjects should switch. However, only 6/104 subjects switch more than half the time. Friedman (1998) explains this recurrent behavior by the illusion of control (Camerer (1995)), i.e. the subjects think that their initial choice is the most likely, or by the non rational escalation of commitment (Bazerman (1990)), i.e. the persistence in choices is viewed as a virtue and flip flopping as a vice. Friedman (1998) proposes as an exploration of the three-door anomaly to manipulate the subject's initial choice by increasing the number of doors choice. In somehow, our experiment is close to that purpose

as the number of possible urn is higher than the number of card. Yet, the subject's choice is about the same: choosing the correct urn/card in between all possible ones in two steps, where the first step represents the formation of an initial belief. The main difference in between the two experiments is the fact that the signal may invalidate the initial belief of the subject. Charness and Levin (2005) and Friedman (1998) experiments are of particular interest for ours as they underline the difficulty to predict the subject choice. They both find that people may persist with their choice even if they are wrong, while we show that they may also change even when they are right.

Finally, Ashton and Ashton (1988) study the sequential belief revision in five experiments and how subject's beliefs are influenced by the order and the presentation of new evidence. They study the impact of the validating/invalidating new information and the degree that it confirms the initial belief. The participants are auditors. They have to investigate the payroll records of a hypothetical client. They observe the result of a preliminary investigation that indicates the likelihood the controls would prevent errors was either 0.20, 0.50, or 0.80. Then they receive positive or negative evidences that can be either strong or weak evidences. They find that subjects are more willing to revise their belief after receiving new information in contrast with the theoretical behavioral literature. Moreover, the subjects revise their belief in a greater extent after receiving new evidence invalidating their initial belief contrary to the literature that suggests that people are more influenced by validating information. Ouwersloot, Nijkamp, and Rietveld (1998) study the impact of the characteristics of the messages on the size of the errors in updating belief. The subject is asked to assess a quantity. He is given a prior distribution of the unknown value. Then, he receives information and assesses the quantity. They confirm that people do not apply the Bayes' rule. They find that the deviations from the Bayesian behavior are impacted by the characteristics of the message such as its precision, reliability, relevance and timeliness. However, the precision of the message has a small impact on the error of probability updating. The main limit of their experiment is that no proper incentives are given to subjects. They receive a fixed reward that does not depend on the correctness of the answers. Moreover the subjects face different hypothetical stories. In our experiment people have to assess an uncertain state of the nature incorporating the information given by a pass/failed test and are incentivized such as to give the correct answer.

The updating process is still to be investigated and understood. People do not seem to be exclusively Bayesian thinkers and there is a need for understanding the different processes used in their decision making. Our experiment contributes to this goal by analyzing specifically one kind of information: a signal that restricts the possible state of the world, or more precisely

the use of pass/failed test to assess the characteristics of an object.

III Experimental Design

III.1 The game

Subjects have to determine the number of yellow balls contained in an urn. The urn is composed of 20 balls that can be either yellow or blue. The number of yellow balls (between 0 and 20) is randomly drawn. Therefore, each of the 21 possible urn compositions (representing a particular state of the nature) has the same probability of being drawn. This is common information.

The subjects observe only a random sample of this urn, whose size is either 6, 10 or 14 balls. The size of the sample is randomly drawn³.

Each period consists of two parts. In the first part, subjects observe the sample drawn from the urn. They have to give an estimation of the number of yellow balls in the urn only observing that sample. We call this first estimation the *initial belief*. In the second part, the same sample is kept visible and the subjects receive a signal about the urn composition. The signal indicates that the urn is composed either of at least or less than X yellow balls. The value on which the test is based is randomly chosen among 6, 10, 14⁴ and this is common information.

Therefore one of the following signals is sent to the subjects:

- < 6 : There are LESS THAN 6 yellow balls in the urn
- ≥ 6 : There are AT LEAST 6 yellow balls in the urn
- < 10 : There are LESS THAN 10 yellow balls in the urn
- ≥ 10 : There are AT LEAST 10 yellow balls in the urn
- < 14 : There are LESS THAN 14 yellow balls in the urn
- ≥ 14 : There are AT LEAST 14 yellow balls in the urn.

They have, once again, to estimate the number of yellow balls in the urn. We call this second estimation the *posterior belief*. Each subject plays 50 periods of this game. At each period, the urn composition, the observed sample drawn and the signal change.

Subjects were paid in cash in a separate room. They received a show-up fee of €10. Three

³As an illustration, see the instructions in the appendix VI.

⁴Given the urn and the sample composition are random, the numbers 6, 10 and 14 are chosen so as to insure enough observations in each case.

periods out of the 50 were randomly drawn randomly at the end of the session. Subjects were paid according to their estimation in these periods. They tossed up three times to determine which part is paid from each selected period when they entered the payment room. If one of the estimations over the 3 selected is correct, subjects received €10, if two, they received €15 and if three €20. This payoff rule induces players to give an estimation of the number of yellow balls associated to the mode of their subjective probability distribution over the possible urns they are facing. In what follows, we derive theoretical predictions.

III.2 Theoretical predictions

In our set up, subjects need to infer the number of yellow balls in the urn thanks to the available information. The way how the probability distribution is computed is provided in Appendix VII.

In the first part, the available information is the composition of the observed sample, whereas it is completed by the signal in the second part. Given this available information, subjects compute their (subjective) probability distribution over the set of potential urns whose composition is consistent with available information. In order to maximize their payoff, subjects have to give their initial belief of the number of yellow balls contained in the urn associated to the mode of the initial probability distribution, that we call the modal urn. The payoff rule gives to the subjects the incentive to declare as a belief the number of balls contained in the urn associated with the mode of their probability distribution rather than the expected number of yellow balls in the urn. This rule has been chosen for two reasons. First, we think that it is easier to calculate the mode than the expected number of yellow balls. Second, the mode of the distribution changes between parts 1 and 2 only when the signal removes the urn associated to the initial mode from the set of possible urns in part 2. Therefore it is easy to identify what information is taken into account by the subjects. If subjects are rational, i.e. if they use Baye's rule, their subjective probability distribution coincides with the objective one. Since the observed sample is a random draw from the unobserved urn without repetitions, the mode of the objective probability distribution in the first part corresponds to the urn in which the proportion of yellow balls is equal to the proportion of yellow balls observed in the sample. For instance, if subjects see 2 yellow balls out of 10 in the observed sample, the most likely number of yellow balls contained in the urn is 4. Moreover, the larger the observed sample, the smaller the set of possible states of the nature, so the higher the mode of the initial probability distribution probability. Hence, the larger the observed sample, the higher the probability of the optimal answer to be right. Therefore, in our setup, Bayesian

revision and representativeness are in line. The optimal strategy is to play accordingly to the representativeness heuristic.

Our first assumption is:

A_1 Subjects estimate the number of yellow balls in part 1 according to the representativeness heuristic.

The determination of the optimal belief in the second part follows the same logic: subjects update their initial belief using the information provided by the signal they receive. The type of signal used in this experiment consists in reducing the number of possible urns from which the sample may have been drawn⁵. Therefore, the posterior probability distribution is computed over a smaller number set of potential urns. Given the signal is randomly selected, the urns that remain in the set of possible urns all have an equal or higher probability of occurrence as they had initially. Therefore, the mode of the objective posterior probability distribution remains the same as that of the initial distribution when the initial modal urn is still included in the posterior set of possible urns. When it is not, the mode of the posterior distribution corresponds to the urn containing the number of yellow balls which is the closest to the one in the initial modal urn. As an illustration, figures 1 and 2 displays the probability distribution that different estimations are correct for an urn where the observed sample is composed by 6 yellow balls and 4 blue balls.




Figure 1: Theoretical predictions for a validating signals.

⁵Some signals do not reduce the number of possible urns. This is the case, for instance, when the signal indicates that there are at least 6 yellow balls in the urn whereas there are already 10 yellow balls in the observed sample. In this case, the set of possible urns and the probability distribution remain the same.




Figure 2: Theoretical predictions for an invalidating signals.

From the subjects’ point of view, receiving a signal that validates the initial belief may increase the probability that the first estimation is correct. When the signal invalidates the first estimation, it proves it to be wrong and the optimal answer changes. Taking back our previous example, the optimal answer is 4 yellow balls before the signal is sent. After receiving the signal, if it states “there are less than 10 yellow balls” in the urn, the optimal answer remains the same since the signal confirms the initial optimal answer. At the opposite, if the signal received states that “there are at least 10 yellow balls”, the optimal answer becomes 10.

Two assumptions can be derived, regarding the subjects’ behaviors:

A_{2a} : Subjects change their initial belief when they observe an invalidating signal.

A_{2b} : Subjects do not change their initial belief when the signal validates it.

III.3 Procedures

The experiment was conducted at the GATE laboratory, Lyon, France using the Regate software (Zeiliger (2000)). Using the ORSEE software (Greiner (2004)), we recruited 81 under-graduate students from local business and engineering schools.

At the beginning of the session, each subject was randomly assigned to a computer. The subjects played the 50 periods of the game. At the end of the session, they filled out a demographic questionnaire and are asked which strategy do they adopt.

A session lasted 60 minutes on average. On average subjects earned €15.

IV Results

First, we analyze if the signal helps the subjects in assessing the exact number of yellow balls in the urn. Then, we study if the signals help the subjects to behave optimally in forming their belief. Finally, we explore the reasons why subjects change their initial belief even when the signal validates their initial belief.

IV.1 Consequences of belief updating on estimation accuracy

On average, signals sent to the agents help them to correctly assess the number of yellow balls contained in the urn.




Figure 3: Density of the proportion of correct estimations.

For each subject, we compute the proportion of correct estimations as the number of estimations equal to the number of yellow balls contained in the urn divided by the number of decisions taken. Figure 3 plots the density of the proportion of correct estimation in the whole sample. Figure 3.A (resp. Figure 3.B) shows this density for estimations realized in the first (resp. second) part of the game. A Wilcoxon matched-pairs signed-ranks test⁶ shows that the two distributions of the proportion of correct estimation in part 1 and in part 2 are significantly different (p-value = 0.0000) and that the median of the distribution of correct estimation in part 2 is higher than in part 1 (p-value = 0.0000). However, this pattern hides some heterogeneity according to the type of signal received.

⁶Test of equality of distributions on matched data.




Figure 4: Density of the proportion of correct estimations in part 2, distinguishing between validating and invalidating signals.

Figure 4 represents the distribution of the proportion of correct estimations in part 2, distinguishing between validating and invalidating signals. The left panel represents the density when the signal validates the initial belief and the right panel represents the density when the signal invalidates the initial belief. It has to be noted that there are 85% of validating signals, i.e. signals that confirm the subject's initial belief, and therefore 15% of invalidating signals. The low number of observations corresponding to invalidating signal explains the particular form of the density of Figure 4.B. A Wilcoxon signed-ranks test shows that the two distributions of the proportion of correct estimation in part 2 with validating and invalidating signals are significantly different ($p\text{-value} = 0.0000$) and that the median of the distribution of correct estimation with invalidating signals is higher than the one after validating signals ($p\text{-value} = 0.0000$). Thus, invalidating signals help the agents to correctly estimate the number of yellow balls. Invalidating signals occur in two situations: when the observed sample provides a bad representation of the urn composition and/or when the subjects give in part 1 a number that is far from the optimal estimation. Therefore, invalidating signals disapprove wrong beliefs, and allow the subjects to improve their estimation of the urn composition.

If invalidating signals help the subjects in their urn assessment, at the opposite, validating signals may have perverse effects. According to the theoretical predictions, when receiving a validating signals, the subjects should not change their belief.




Figure 5: Density of the proportion of correct estimations when the signal validates the initial belief and the estimation does not change between parts 1 and 2.

Figure 5 displays the density of the proportion of correct estimations when the signal validates the initial belief and the estimation does not change between parts 1 and 2. A Wilcoxon signed-ranks test shows that the distributions of the proportion of correct estimation in part 2 with invalidating signals (plotted in Figure 4.B) and the proportion of correct estimation in part 2 after validating signals, with no change of estimation between the parts 1 and 2, are significantly different ($p\text{-value} = 0.0000$). Moreover, the median of the distribution of correct estimation with invalidating signals is higher than the one after validating signals and no change of estimation between part 1 and part 2 ($p\text{-value} = 0.0068$).

The subjects may also decide to change their estimation between part 1 and part 2 when they receive a validating signal whereas they are not supposed to do so according to the optimal strategy.




Figure 6: Density of the proportion of correct estimations in part 1 and in part 2 when the signal validates the initial belief and the estimation changes between parts 1 and 2.

Figure 6 displays the density of the proportion of correct estimations in each of the two parts when the signal validates the initial belief and the estimation changes between parts 1 and 2. A Wilcoxon matched-pairs signed-ranks test shows that the distributions of the proportion of correct estimations in parts 1 and 2 are significantly different (p -value = 0.0000) and that the median of the distribution of correct estimations in part 1 is higher than the one in part 2 (p -value = 0.0000). Therefore, when subjects change their initial belief whereas they are not supposed to do so, the proportion of correct estimations decreases between the two parts.

IV.2 Do subjects play optimally?

In this section, we explore how the estimations made by the subjects differ from the estimations predicted by Bayesian updating. On average, the optimal number is given in 36% of the situations⁷. In 60% of the cases, their estimation error is only of one ball with respect to the optimal number. Then, A_1 is only partially verified. The literature on Bayesian updating shows that, when facing an uncertain question, agents use the representativeness heuristics whereas the optimal strategy is to use the Bayesian updating rule. Even if, in our setup, the Bayesian updating rule and the representativeness heuristic are in line, our results suggest that the subjects do not always follow the optimal strategy. On average, subjects earned €15.16. If they had behaved rationally, they should have earned €17.70 on average (see calculus in Appendix VIII).

We observe that 20% of the subjects play optimally in more than 50% of the cases. This observation suggests that a limited number of subjects keep playing optimally in more than

⁷In a post experimental questionnaire, subjects are asked to explain the strategy they used during the game. 28% of the subjects' answers correspond to a description of the representativeness strategy.

a half of the periods. The data do not reveal any other particular recurrent strategy.

When the subjects receive a signal that invalidates their initial belief, they modify their estimation in 98% of cases, which is consistent with assumption A_{2a} . However, subjects choose the number of yellow balls which is the closest to their first estimation in the new set of possible urns in 30% of the cases.

When the subjects receive a signal that validates their initial belief, their updating behaviors are quite different than what the theory predicts. In such situations, they change their estimation in 50% of the cases, which is in opposition with our assumption A_{2b} .

IV.3 Explaining updating behavior when the signal validates the initial belief

On top of its contradiction with the rational behavior prediction, the fact that subjects may change their initial belief when the signal validates it, deserves a particular interest. Indeed, it is of particular interest to identify when and why the signal makes subjects modify their initial belief. Indeed, in the previous section, we showed that subjects revise their belief, as predicted by the theory, when the signal invalidates their initial belief. The fact that they revise their belief also when the signal validates their initial belief, which is not consistent with the optimal behavior, suggests that there may be some situations on which the particular information revealed to the subjects makes them more likely to update their belief.

Each subject plays the same game during a subsequent number of periods. The post-experimental questionnaire shows an important heterogeneity in the playing strategy as self-reported by the subjects. This heterogeneity also appears in the data, as illustrated, for instance, in Figure 3 that shows the distribution of the frequency of correct estimations across periods in the population. Therefore, there might be some underlying characteristics (observable or not) that affect the decisions of the players, and, in particular, their updating strategies. This is why we account for the individual characteristics persistent over time by (i) introducing individual observed characteristics (reported by the subjects in the administrative questionnaire) (ii) allowing for the presence of unobserved permanent characteristics by exploiting the panel dimension of our data. More precisely, we estimate a random effect logit model⁸ in which the explained variable is the indicator of changing the estimation between parts 1 and 2. This estimation is performed on observations when the signal confirms the initial belief.

We explore the role of the two factors that may explain why people change their beliefs:

⁸Hausman test does not reject the random effect against the fixed effect specification at 90%.

the evolution of the uncertainty level about the state of the nature between the first and the second part and the (objective) probability that the estimation in part 1 is correct, according to theoretical predictions.

The evolution of the uncertainty level takes into account both the uncertainty level in the first part, induced by the number of visible balls, and the reduction of the uncertain level between parts 1 and 2, induced by the signal. Given our experimental design, the initial uncertainty level may be either high (when the observed sample is composed by 6 visible balls), medium (10 visible balls) or low (14 visible balls). Therefore, a high uncertainty level means that there are 15 possible urns from which the observed sample is drawn. Medium and low uncertainty levels correspond respectively to 11 and 7 possible urns. The signal consists in a reduction of the number of possible urns from which the sample is drawn. In order to obtain the final uncertainty level, we discretize the number of possible urns left after reception of the signal into three categories, high, medium and low, if the number of possible urns is respectively 12 or more, between 8 and 11, and 7 or less. The evolution of the uncertainty level between the two parts falls therefore into six categories: high-to-high, high-to-medium, high-to-low (taken as the reference category because it correspond to the highest decrease in uncertainty between the part 1 and the part 2)), medium-to-medium, medium-to-low, and low-to-low.

The probability of correct estimation in part 1 corresponds to the probability that the number of yellow balls contained in the urn is equal to the estimation given by the subject in part 1, given the signal received. This factor takes into account both the composition of the observed sample and the estimation of the players given in part 1. The higher this probability, the higher the subject should be confident about the estimation given in part 1, and the lower should be the probability that he changes his estimation.

Among the explanatory factors, we also add some observed individual characteristics such as the gender, the age, the field of study at university and the type of high-school diploma passed (either in science, literature or social sciences).

Dependent variable: indicator for belief change

Evolution of the uncertainty level about the state of the nature:		
high → high	.5143**	(.2599)
high → medium	.3812	(.2557)
high → low	ref.	
medium → medium	.3545	(.2371)
medium → low	.0556	(.2606)
low → low	.3268	(.2305)
Probability of correct estimation at period 1	-2.5734***	(.3480)
Constant	1.6821	(5.0563)

Observations	3316
Log-likelihood	-1480.1101
ρ	.5171*** (.0468)

Logit model with random effects (Hausman test doesn't reject the random effect against the fixed effect specification at 90%). Standard errors in parentheses. Levels of significance: * 10%; ** 5%; *** 1%. ρ denotes the share of the individual (within-group) variance in the total variance.

Table 1: Determinants of belief's updating

Table 1 shows the results of this estimation. We observe that only being in a situation of high initial level of uncertainty that is not decreased after receiving the signal increases the propensity to update the belief when the signal received validates the initial belief. This result suggests that the decision of updating the initial belief is neither affected only by the initial uncertainty level nor affected only by the the fact that the signal reduces the number of possible urns. This result rather shows that belief updating is more likely when the initial uncertainty level is high and remains high after receiving the signal. Concerning the impact of the sample composition and the subjects' first estimation, we find that the probability that the estimation in part one is correct decreases the propensity to update the belief when receiving a validating signal. This result suggested that, when the subject is more confident about his initial belief, he is less likely to change it after a validating signal.⁹

Concerning the individual characteristics¹⁰, we notice that women are less likely to update

⁹A closer look at the data makes it apparent that the subjects are the least likely to change their estimation when the observed sample is composed of around 100% or 0% of yellow balls and subjects play the optimal strategy (their estimation is close to 20 or 0 yellow balls respectively).

¹⁰The results are not reported here, but are available upon request.

their belief after receiving validating signal. All the other characteristics do not affect the probability to change the estimation after a validating signal.

V Conclusion

We analyze how individuals update their belief when they receive an imperfect signal. The originality of our paper stands in the nature of the signal sent to the agents, which is a restriction of the possible states of the nature. More precisely, we analyze how subjects update their belief after receiving these signals for different levels of uncertainty about the state of the nature, for signals that confirm or not their initial belief and for different signal accuracies. We examine these questions in a laboratory setting with a two-stage game. The subjects observe a partially revealed urn composed of 20 balls (yellow or blue) and have to assess the total number of yellow balls in the urn twice, once before a signal on the urn's composition is sent, and next after receiving the signal.

We find that, on average, signals significantly help the subjects to estimate the number of yellow balls contained in the urn. In line with the theoretical predictions, a signal that disproves a wrong initial belief, i.e. an invalidating signal, increases the probability to find the exact composition of the urn. However, in 49% of the cases, subjects change their belief when the signal validates their initial belief, though they are not supposed to do so. In this case, the signal has an adverse effect: the subjects are mistaken since their second prediction is, on average, less correct than the first one. We show that, when the signal confirms their initial belief, the subjects are the most likely to change their estimation when the uncertainty level of the state of the nature is high, i.e. the observed sample is small, and is not reduced by the signal.

Our results suggest that discrete signals are useful to improve the assessment of the state of the nature, especially when the first estimation is wrong. However, we find a limitation on the ability of discrete signals to inform agents about the state of the nature and the risk that they may even mistake them in some situations. These pass/fail tests keep agents uncertain about the state of the nature, especially when they confirm what the agent already knows, but are still very inaccurate. Moreover, it seems, from our experiment, that individuals who are more uncertain about the state of the nature are more likely to revise their belief.

These results on agents' belief revision are obtained in a simple environment. In order to complete our knowledge on the power of imperfect signals, it would be of particular interest to extend this framework by introducing a measurement of the degree of confidence of the subjects' estimations. As seen in the literature on updating behavior, several studies find that

people do not want to change their belief or behavior even when they should. At the opposite, our experiment underlines situations where people modify their initial belief when in fact they should not. Our results indicate that the level of uncertainty about the initial state of the world significantly affects the subject's changing behavior. Therefore, in a new experimental design, we could, for example, ask the subjects to give a confidence interval that their answer is true. We expect that the higher the confidence level of the subjects in their initial belief (first estimation), the lower the likelihood to change their initial belief. In this new setting, we would also be willing to investigate if there exists some specific signals or situations that make the agents change their initial belief even if they are initially confident about it.

References

- Ashton, A.H., and R.H. Ashton. 1988. "Sequential Belief Revision in Auditing." *The Accounting Review* 63:pp. 623–641.
- Bazerman, M. 1990. *Judgement in managerial decision making*. New York: Wiley.
- Camerer, C. 1995. *The handbook of experimental economics*, Princeton University Press, chap. Individual Decision Making. pp. 587–703.
- Charness, G., and D. Levin. 2005. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect." *American Economic Review* 95:1300–1309.
- Friedman, D. 1998. "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly." *American Economic Review* 88 (4):933–946.
- Greiner, B. 2004. "An Online Recruitment System for Economic Experiments." *Kurt Kremer, Voler Macho (eds). Forschung und wissenschaftliches Rechnen, GWDG Bericht 63, Gttingen: Ges. fr Wiss. Datenverarbeitung*, pp. 79–93.
- Grether, D. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." *Quarterly journal of Economics* 95:537–557.
- Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Kahneman, D., and A. Tversky. 1973. "On the psychology of prediction;" *Psychological Review* 80:237–251.
- Ouwensloot, H., P. Nijkamp, and P. Rietveld. 1998. "Errors in probability updating behaviour : Measurement and impact analysis." *Journal of Economic Psychology* 19:535–563.
- Tversky, A., and D. Kahneman. 1971. "Belief in the law of small numbers." *Psychological Bulletin* 76:105–110.
- Zeiliger, R. 2000. "A Presentation of Regate, Internet Based Software for Experimental Economics."

VI Appendix: Instructions (original in French)

You are about to participate in an experiment during which you can earn money. Your earnings will depend on your decisions during the experiment. Your earnings will be paid to you in cash in private and in a separate room.


There are 50 periods in the experiment. Each period has two parts. In each period, you see a partially revealed urn containing 20 balls. The balls may be yellow or blue.

In each part, you have to estimate the total number of yellow balls in the urn. Your earning depends on your estimation of the total number of yellow balls. You increase your earning if you find the exact total number of yellow balls in the urn.


Period description An urn is composed by **20 balls**: yellow or blue. The number of yellow balls is randomly determined between 0 and 20. There are equally chances that the urn is composed by 0 yellow ball, or 1 yellow ball, ... or 20 yellow balls. The urn stays the same during the period. You can observe only 6, 10 or 14 balls in the urn. The visible balls are randomly drawn from the urn.

Each period has 2 parts. The visible balls are identical for the 2 parts.

PART 1: you have to estimate the total number of yellow balls in the urn. Enter your answer with the scroll bar and validate it.



PART 2: The urn stays the same as in the 1st part. You receive complementary information on the urn's composition. This information indicates if the urn is composed with at most, or at least 6, 10 or 14 yellow balls. Then, you have to estimate a second time the total number of yellow balls in the urn. Enter your answer with the scroll bars, even if your answer does not differ from your estimation in the 1st part, and validate it.



Description of the added information on the urn's composition You can receive one of the following information:

- ≥ 6 , there are AT LEAST 6 yellow balls in the urn.
- < 6 , there are LESS THAN 6 yellow balls in the urn.
- ≥ 10 , there are AT LEAST 10 yellow balls in the urn.
- < 10 , there are LESS THAN 10 yellow balls in the urn.
- ≥ 14 , there are AT LEAST 14 yellow balls in the urn.
- < 14 , there are LESS THAN 14 yellow balls in the urn.

All the informations are not possible depending on the composition of the urn. For example, when an urn has a total number of 8 yellow balls, only three informations are possible:

- ≥ 6 , there are AT LEAST 6 yellow balls in the urn.
- < 10 , there are LESS THAN 10 yellow balls in the urn.
- < 14 , there are LESS THAN 14 yellow balls in the urn.

The computer draws randomly and shows one of the possible information.

At the end of this 2nd part, you observe the total urn and your estimations in part 1 and in part 2.

A new period starts automatically when you validate the last screen. Periods are identical except:

- the urn's composition
- the number of visible balls
- the information on the maximum or minimum total number of yellow balls in the urn.

Payment Three periods over the 50 played are selected for your payoff. The three periods are randomly drawn by the computer at the end of the experiment. They can be different among the participants. Thus, you do not know in advance which periods will be selected. Each period has the same chance to be selected.

For each of the three selected periods, only one part is chosen for your payoff. When you enter the payment room, you determine the part selected for the payoff by tossing a coin:

- if you toss head, the 1st part is selected for your payoff.
- if you toss face, the 2nd part is selected for your payoff.

To summarize, three periods are selected randomly by the computer and, for each of the selected periods, you determine which part will be paid by tossing a coin. Thus, three parts are selected for your payoff. Your payoff depends on the number of exact estimations you have done for these three parts:

- If none of your three estimations is exact: your payoff is €0.
- If one of your three estimations is exact: your payoff is €10.
- If two of your three estimations are exact: your payoff is €15.
- If all your three estimations are exact: your payoff is €20.

In all the cases, you will receive €10 more for your participation in this experiment.

It is forbidden to communicate with other subjects during the experiment. If you have any question regarding these instructions, please raise your hand.

VII Appendix: Predicted Probabilities

Predicted Probabilities in Part 1

An urn is composed by 20 balls and the number of yellow balls contained in the urn is chosen randomly (from 0 to 20). Therefore, there are 21 possible urns, which one having the same probability (1/21) of being drawn. Then, the probability to observe sample S , given the urn from which the sample is drawn is U_i is:

$$P(S|U_i) = \frac{C_{y_u}^{y_s} * C_{n_u - y_u}^{n_s - y_s}}{C_{n_u}^{n_s}}$$

where n_u (resp. n_s) is the number of balls contained in urn U_i , $n_u = 20$, (resp. in sample S) and y_u (resp. y_s) represents the number of yellow balls contained in urn U_i (resp. in sample S). Then, according to Bayes' rule and given an observed sample S , the probability that this sample is drawn from urn U_i is equal to:

$$P(U_i|S) = \frac{P(S|U_i) \cdot P(U_i)}{\sum_{x=1}^{21} P(S|U_x) \cdot P(U_x)}$$

Given an observed sample S , the distribution of probabilities over the 21 possible urns is obtained using this formula. The optimal choice for the subject is thus to give as an estimation the number of yellow balls contained in the urn having the highest probability.

Predicted Probabilities in Part 2

Now, we have to compute the probability that a particular observed sample is drawn from each of the 21 possible urns, after a signal is received. According to Baye's rule, the probability that sample S is drawn from urn U_i given signal I (I stand for information) is:

$$P(U_i|S, I) = \frac{P(S, I|U_i) \cdot P(U_i)}{\sum_{x=1}^{21} P(S, I|U_x) \cdot P(U_x)}$$

The signal sent to the agent is only based on the composition of the urn. In particular, it is independent from the composition of the observed sample drawn from the urn in part 1. Therefore, $P(S, I|U_i) = P(S|U_i) \cdot P(I|U_i)$. $P(I|U_i)$ is computed according to the following logic. For a particular urn, among the six existing signals, only three can be sent to the subject. Indeed, the number of yellow balls contained in an urn cannot be both inferior and superior to the same number. Since the signal is randomly chosen, the probability that a signal is sent is equal to one third if the signal is a possible one, and 0 if the signal is inconsistent with the composition of the urn. This leads to the following:

$$P(U_i|S, I) = \frac{P(S|U_i) \cdot P(I|U_i) \cdot P(U_i)}{\sum_{x=1}^{21} P(S|U_x) \cdot P(I|U_x) \cdot P(U_x)}$$

with

$$P(I|U_i) = \begin{cases} 1/3 & \text{if signal } I \text{ is consistent with the composition of urn } U_i \\ 0 & \text{if not} \end{cases}$$

As is the case in the first part, the optimal choice for the subject is to give as an estimation the number of yellow balls contained in the urn having the highest probability, given the sample and the signal.

VIII Appendix: Expected revenue

The expected gain for one subject depends both on the number of correct estimations he gives and on the likelihood that the period(s) at which he makes correct estimations is(are) drawn at the end of the game. Three periods are drawn. The probability that, among those three periods, a subject makes i correct estimations (and therefore makes $3 - i$ wrong estimations) is given by:

$$P(i) = \frac{C_n^i * C_{n-c}^{3-i}}{C_n^3}, \text{ for } i = 0, \dots, 3$$

where n represents the total number of periods played and c the total number of correct estimations.

The expected gain is the average of the gains associated to 0, 1, 2 and 3 correct answers drawn, weighted by the probability that the corresponding number of correct answer is drawn:

$$ER(n, c) = \frac{C_c^0 * C_{n-c}^3}{C_n^3} * 10 + \frac{C_c^1 * C_{n-c}^2}{C_n^3} * 20 + \frac{C_c^2 * C_{n-c}^1}{C_n^3} * 25 + \frac{C_c^3 * C_{n-c}^0}{C_n^3} * 30$$

Given that, for the whole experience, the subjects made 7900 estimations, among which 1492 were correct ones, the total expected revenue is $ER(7900, 1492) = 15.16\text{€}$. If every player had played with the optimal strategy (consisting in giving as an estimation the urn associated with the highest probability), they would have made 2340 correct estimations. In this case, the expected revenue would have been $ER(7900, 2340) = 17.70\text{€}$.