



HAL
open science

Confidence Biases and Learning among Intuitive Bayesians

Louis Lévy-Garboua, Muniza Askari, Marco Gazel

► **To cite this version:**

Louis Lévy-Garboua, Muniza Askari, Marco Gazel. Confidence Biases and Learning among Intuitive Bayesians. 2015. halshs-01243584


HAL Id: halshs-01243584

<https://shs.hal.science/halshs-01243584>

Submitted on 15 Dec 2015

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.



**Confidence Biases and Learning
among Intuitive Bayesians**

Louis LEVY-GARBOUA, Muniza ASKARI, Marco GAZEL

2015.80



Confidence Biases and Learning among Intuitive Bayesians

Louis Lévy-Garboua*, Muniza Askari†, Marco Gazel‡

October 2015

Abstract

We design a double-or-quits game to compare the speed of learning one's specific ability with the speed of rising confidence as the task gets increasingly difficult. We find that people on average learn to be overconfident faster than they learn their true ability and we present a simple Bayesian model of confidence which integrates these facts. We show that limited discrimination of objective differences, myopia, and uncertainty about one's true ability to perform a task in isolation can be responsible for large and robust confidence biases, namely the hard-easy effect, the Dunning-Kruger effect, conservative learning from experience and the overprecision phenomenon (without underprecision) if subjects act as Bayesian learners. Moreover, these biases are likely to persist since the Bayesian aggregation of past information consolidates the accumulation of errors and the perception of contrarian illusory signals generates conservatism and under-reaction to events. Taken together, these two features may explain why intuitive Bayesians make systematically wrong predictions of their own performance.

Keywords: Confidence biases, Bayesian learning, double or quits experimental game, hard-easy effect, Dunning-Kruger effect, illusory signals.

JEL classification: C11, C91, D03

*Corresponding Author: louis.levy-garboua@univ-paris1.fr, Phone:+33 6 85 11 18 68. Paris School of Economics, Université Paris 1 Pantheon-Sorbonne, and Centre d'Economie de la Sorbonne, 106-112 Bd de l'Hôpital 75013, Paris, France; and CIRANO Montreal

†Centre d'Economie de la Sorbonne

‡Paris School of Economics, Université Paris 1 Panthéon Sorbonne, Centre d'Economie de la Sorbonne, 106-112 Bd de l'Hôpital - 75013 Paris, marco-antonio.gazel@univ-paris1.fr

1 Introduction

In many circumstances, people appear to be "overconfident" in their own abilities and good fortune. This may occur when they compare themselves with others, massively finding themselves "better-than-average" in familiar domains (eg., Svenson 1981, Kruger 1999), when they overestimate their own absolute ability to perform a task (eg., Lichtenstein and Fischhoff 1977, Lichtenstein, Fischhoff and Phillips 1982), or when they overestimate the precision of their estimates and forecasts (eg., Oskamp 1965). Moore and Healy (2008) designate these three forms of overconfidence respectively as overplacement, overestimation, and overprecision. We shall here be concerned with how people overestimate, or sometimes underestimate, their own absolute ability to perform a task in isolation. Going deeper into this bias, we shall decompose it in two parts, imperfect discrimination and illusion, and we shall see that our explanation of the latter bias predicts the overprecision phenomenon as well.

The estimation bias refers to the discrepancy between ex post objective performance (measured by frequency of success in a task) with ex ante subjectively held confidence (Lichtenstein, Fischhoff and Phillips, 1982). It has first been interpreted as a cognitive bias caused by the difficulty of the task (eg., Griffin and Tversky 1992). It is the so called "hard-easy effect" (Lichtenstein and Fischhoff, 1977): people underestimate their ability to perform an easy task and overestimate their ability to perform a difficult task. However, a recent literature has challenged this interpretation by seeking to explain the apparent over/underconfidence by the rational-Bayesian calculus of individuals discovering their own ability through experience (Moore and Healy, 2008, Grieco and Hogarth, 2009, Benoît and Dubra, 2011, Van den Steen, 2011). While the cognitive bias view describes self-confidence as a stable trait, the Bayesian perspective points at the experiences leading to over- or under-confidence.

We design a real-effort experiment which enables us to test the respective strengths of estimation biases and learning. People enter a game in which the task becomes increasingly difficult -*i.e.* risky- over time. By comparing, for three levels of difficulty, the subjective

probability of success (confidence) with the objective frequency at three moments before and during the task, we examine the speed of learning one's ability for this task and the persistence of overconfidence with experience. We conjecture that subjects will be first underconfident when the task is easy and become overconfident when the task is getting difficult. However, "difficulty" is a relative notion and a task that a low-ability individual finds difficult may look easy to a high-ability person. Thus, we should observe that overconfidence declines with ability and rises with difficulty. The question raised here is the following: if people have initially an imperfect knowledge of their ability and miscalibrate their estimates, will their rising overconfidence as the task becomes increasingly difficult be offset by learning, and will they learn their true ability fast enough to stop the game before it is too late?

The popular game "*double or quits*" fits the previous description and will thus inspire the following experiment. A modern version of this game is the world-famous TV show "*who wants to be a millionaire*". In the games of "*double or quits*" and "*who wants to be a millionaire*", players are first given a number of easy questions to answer so that most of them win a small prize. At this point, they have an option to quit with their prize or double by pursuing the game and answering a few more questions of increasing difficulty. The same sort of double or quits decision may be repeated several times in order to allow enormous gains in case of repeated success. However, if the player fails to answer one question, she must step out of the game with a consolation prize of lower value than the prize that she had previously declined.

Double-or-quits-type behavior can be found in many important decisions like addictive gambling (Goodie, 2005), military conquests (Johnson, 2004), business expansion (Malmendier and Tate, 2005), speculative behavior (Shiller, 2000), educational choices (Breen, 2001), etc. Overconfident players, chiefs, entrepreneurs, traders, or students are inclined to take excessive risks; they are unable to stop at the right time and eventually fail more than well-calibrated persons¹ (e.g., Barber and Odean 2001, Camerer and Lovo 1999).

¹However, overconfidence may pay off when there is uncertainty about opponents' real strengths, and when the benefits of the prize at stake are sufficiently larger than the costs (e.g., Johnson and Fowler 2011,

In contrast, under-confident individuals won't take enough risks and stay permanently out of successful endeavors. Our experimental data reproduces the double or quits game. We observe that subjects are under-confident in front of a novel but easy task, whereas they feel overconfident and willing to engage in tasks of increasing difficulty to the point of failing.

We propose a new model of "intuitive Bayesian learning" to interpret the data. Like Erev, Wallsten and Budescu (1994), our model retains the view that confidence judgments are subject to errors but it goes beyond by eliciting systematic components of such errors. Like Moore and Healy (2008), our model is Bayesian but it is more general in that it incorporates confidence judgments prior to experience of the task and learning from experience. In contrast with Gervais and Odean (2001), we don't need to assume a self-attribution bias (Langer and Roth, 1975, Miller and Ross, 1975) combined with Bayesian learning to produce overconfidence². Signals of future success and failure are treated symmetrically³. Finally, unlike models of confidence management (e.g. Brunnermeier and Parker 2005, Köszegi 2006, Möbius et al. 2014), we don't have to postulate that individuals manipulate their beliefs and derive direct utility from optimistic beliefs about themselves. Subjects in our view derive their beliefs exclusively from their prior and the informative signals that they receive. In our model, "intuitive Bayesians" fail to predict their own performance even with prior experience of the task because they are uncertain of their prior belief and update it systematically both by cues on their current performance and by *a contrarian illusory signal* triggered by their doubt.

Section 2 lays down the structure of the experiment and incentives, and provides the basic descriptive statistics. Three major facts about confidence are reported in section 3 from our data: *(i)* limited discrimination among different tasks; *(ii)* miscalibration of subjective prob-

Anderson et al. 2012).

²Using German survey data about stock market forecasters, Deaves, Lüders and Schröder (2010) does not confirm that success has a greater impact than failure on self-confidence, which casts doubt on the self-attribution bias explanation.

³In studies where subjects are free to stay or to leave after a negative feedback, subjects who update most their confidence in their future success to a negative feedback are selectively sorted out of the sample. This creates an asymmetry in measured responses to positive and negative feedback. Such spurious asymmetry does not exist in the present experiment, because subjects who fail to reach one level must drop out of the game.

abilities of success elicited by the "hard-easy effect"; and (iii) differential, ability-dependent, calibration biases known as the Dunning-Kruger (or ability) effect. Subjects have biased estimates of their unknown ability partly because they have a limited power of discrimination. They do not perceive differences of difficulty between tasks unless such differences are extremely salient; and they are not forward-looking, in the sense that they are unable to anticipate the increased likelihood of their success at the high level conditional on passing the middle level. However, confidence biases are not just a matter of mistaken measurement. Bayesian people also miscalibrate their own probability of success after making a correct measure of their own ability without knowing that it is accurate. This calibration problem is responsible for the hard-easy effect and for the 'ability' or Dunning-Kruger effect (Kruger, 1999), of which we provide strong evidence. It is then natural to ask if these biases can be eliminated by learning. We find that individuals behave rationally in a local or intuitive sense. Overconfidence is somehow limited by experience but by far not eliminated. A theoretical framework for understanding why local learning coexists with miscalibration is proposed in section 4, and results are presented in section 5. An 'intuitive Bayesian' model that predicts reported self-confidence as a weighted average of a prior and cues received during the game is consistent with local learning. However, we claim that an intuitive Bayesian will also perceive 'available illusory signals' generated by the uncertainty surrounding his estimated probability of success and be vulnerable to illusions. The type of illusory signals induced by doubt predicts the hard-easy effect, the ability or Dunning-Kruger effect, and conservatism, that is, under-reaction to received information. It also predicts overprecision, another facet of overconfidence, that is, the tendency to attribute excessive precision to own ability estimates. We show in section 6 that the cues upon which subjects construct their own estimate of success, *i.e.* confidence, widely differ from the genuine predictors of success. A brief conclusion follows in section 7.

2 The experiment

2.1 Task and treatments

Participants perform a real-effort, rather long and difficult, task for which they get paid according to their degree of success. The task consists in solving anagrams ranked in three levels of increasing difficulty. It is performed during a maximum of 15 rounds lasting no more than 8 minutes each. These 15 rounds are structured in three successive levels of increasing difficulty, designated respectively as the training level, the middle level, and the high level.

Participants are successful at one level when they manage to decode $2/3$ of the anagrams at this level. The training level consists of 9 rounds of low difficulty (*i.e.* 6 anagrams per round to be solved in no more than eight minutes). It is long enough to let participants feel that a large effort and ability is required of them to succeed at the optional upper levels. It does also let them ample time to learn the task. The middle and high levels, which come next, comprise 3 rounds each.

The gradient of task difficulty was manipulated after completion of the training level and two conditions are available: (*i*) in the '*wall*' condition, the difficulty jumps sharply at middle level, but remains constant at high level; (*ii*) in the '*hill*' condition, the difficulty always rises from one level to the next, slowly first at middle level, then sharply at high level.

By the end of the experiment, the required number of anagrams is the same for the '*wall*' and '*hill*' conditions. However, the distribution of anagrams to be decoded differs for these two conditions. In the wall condition, ten anagrams per round are proposed at the middle and high levels, of which 20 anagrams at least must be decoded per level. In the hill condition, eight anagrams per round are proposed at middle level, and this rises to twelve anagrams at high level. Decoding sixteen anagrams in three rounds is required for middle level; and decoding twenty-four anagrams in three rounds is required for high level. This design can be visualized in figure 1.

The manipulation of the '*wall*' and '*hill*' conditions gave rise to three treatments:

- Wall treatment (*wall*): the wall condition is imposed to participants who passed the




Figure 1 – Decision problem perceived by participants at the start of level 2 of the choice treatment.

training level;

- Hill treatment (*hill*): the hill condition is imposed to participants who passed the training level;
- Choice treatment (*choice*): a choice among the two conditions (*wall* or *hill*) is proposed to participants who passed the training level.

The double or quits game is played under these three treatments. All subjects first go through the training level. Those who were successful -i.e., those who solved at least 36 anagrams during the training level- will then be asked to double or quits:

- **Double**: Continue to the next level to win a substantial increase in earnings;
- **Quits**: Stop the experiment and take your earnings.

Participants who decide to go to middle level get a consolation prize that is lower than the foregone earnings if they fail or drop out before the third round. If they succeed middle level, they will be asked again to double or quits. The same rules apply for high level at rising levels of earnings. The potential gains (in Euros) were (10, 2) at the training level,

that is, 10€ for successful quitters and 2€ for failures, (14, 4) at middle level, and (26, 11) at high level.

2.2 Experimental sessions

We ran 24 sessions for a total of 410 participants, half for the choice treatment and the other half equally split between the *'wall'* and *'hill'* treatments. Eight sessions were run in the BULCIRANO lab (Center for Interuniversity Research and Analysis on Organizations), Montreal (Canada), and the same number of sessions were conducted at the LEEP (*Laboratoire d'Economie Expérimentale de Paris*), Pantheon-Sorbonne University. The difference between Paris and Montreal was observed to be insignificant. Thus, eight additional sessions were conducted at LEEP in order to acquire robust results. A show-up fee of 5€ in Paris and Can\$ 5 in Montreal was paid to the participants (from now on, all money amounts will be given in Euros). About 80% of the participants were students.

At the start, instructions were read out and a hard copy of it was also provided individually. Participants answered six questions to test their full comprehension of the experiment. Information on gender, age, educational level and labor market status was required. The last question was a hypothetical choice between 5€ for sure and an ambiguous urn containing 100 balls of two colors (white and black) in unknown proportions. Ten Euros (10€) were to be earned if a black ball was drawn. Choice of the sure gain provided a rough but simple measure of risk aversion in the uncertainty context of the experiment.

2.3 Descriptive statistics

The main descriptive statistics for the three treatments are reported in Table 1:

The results of tests show that the three samples are homogeneous. No significant difference is observed among the samples' means for individual characteristics. As expected, the *'wall'* and *'hill'* treatments had a substantial impact on the decision to double upon reaching

⁴Ability is measured by the number of anagrams solved per minute in the first 4 rounds. It lies in the interval [0,6].

Table 1 – Descriptive statistics for the three treatments

Variables	Treatments		
	Wall	Hill	Choice
Male	56%	48%	49%
Age	24.5	25.8	25.1
Risk Averse	54%	59%	51%
Payments	9.1	8.9	7.8
Total anagrams solved	55.6	53.7	54.3
Ability ⁴	2.8	2.7	2.6
Number of observations	101	106	203
<i>Decision to double conditional on success at previous level:</i>			
Middle level	78% (91)	76% (90)	77% (176)
High level	95% (22)	72% (29)	82% (34)

Notes: Decision to double to High level: difference between the "wall" and "hill" treatments is significant at 10%; all other differences are not significant at 10% level. (Bonferroni multiple comparison tests). Number of participants successfully clearing the previous level is in parentheses.

the middle level (10% significance by Bonferroni multiple comparison tests; F -value= 2.33; p -value=0.100). Almost everybody doubles in the 'wall' treatment on reaching middle level because the high level is no more difficult than the middle level. In contrast, only 72% enter the high level in the 'hill' treatment as the difficulty gradient is very steep. In spite of these differences, the number of anagrams solved and payments may be considered equal among treatments at the usual level of significance.

Subjects can also be grouped in three different levels of ability, according to the number of anagrams solved per minute in the first 4 rounds: high ability (first tercile), medium ability (second tercile) and low ability (last tercile). Some descriptive statistics for the three treatments are reported on Table 2. The three groups are homogeneous in terms of gender and risk aversion but a slightly greater proportion of low-ability subjects can be found among older, probably non-student, participants.

Table 2 shows that "ability" strongly discriminates among participants in terms of performance (total anagrams solved, payments) and quits before the middle level. However, the training level was meant to be easy enough that three-quarters (102:137) of low-ability subjects would pass it.

Table 2 – Descriptive statistics by ability level

Variables	Level of ability			Difference		
	High	Medium	Low	M-H	L-M	L-H
Male	47%	54%	50%	<i>ns</i>	<i>ns</i>	<i>ns</i>
Age	23.6	24.5	27.2	<i>ns</i>	***	***
Risk Aversion	53%	50%	59%	<i>ns</i>	<i>ns</i>	<i>ns</i>
Payments	11.7	7.7	6.0	***	<i>ns</i>	***
Number anagrams solved	67.7	53.8	42.6	***	***	***
Ability	4.5	2.4	1.1	***	***	***
Number of observations	131	142	137			
<i>Decision to double conditional on success at previous level:</i>						
Middle level	91% (128)	81% (127)	54% (102)	<i>ns</i>	***	***
High level	87% (55)	72% (25)	80% (5)	<i>ns</i>	<i>ns</i>	<i>ns</i>

Notes: Significance level: * 10%; ** 5%; ***1%; ns: not significant at 10% level (Bonferroni multiple comparison tests). Number of participants successfully clearing the previous level is in parentheses

2.4 Confidence judgments

Participants were asked to state their subjective probability of success for the three levels and at three moments: before, during, and after the training level. Before beginning the game, they were shown a demonstration slide which lasted one minute. Anagrams of the kind they would have to solve appeared on the screen with their solution. Then, they were asked to assess their chances of success on a scale of 0 to 100 (Adams, 1957), and the game started for real. After four rounds of decoding anagrams, players were asked again to rate their confidence. Lastly, players who had passed the training level and decided to double re-estimated their chances of success for the middle and high levels.

The Adams’s (1957) scale that we used is convenient for quantitative analysis because it converts confidence into (almost) continuous subjective probabilities. It was required for consistency that the reported chances of success do not increase as the difficulty level increased. Answers could not be validated as long as they remained inconsistent. Subjects actually used the whole scale but, before the experiment, 14% expressed absolute certainty that they would succeed the first level and only 0.24% was sure that they would fail.

We did not directly incentivize beliefs because our primary aim was not to force subjects to make optimal forecasts of their chances of success but to have them report sincerely their true beliefs in their attempt to maximize their subjective expected utility, and to observe

the variation of such beliefs with experience. The true beliefs are those which dictate actual behavior following such prediction, and the latter was incentivized by the money gains based on subjects' decisions to double or quits and performance in the task. Armantier and Treich (2013) have recently generalized previous work on proper scoring rules (see their extensive bibliography). They show that, when subjects have a financial stake in the events they are predicting and can hedge their predictions by taking additional action after reporting their beliefs, use of any proper scoring rule generates complex distortions in the predictions and further behavior since these are not independent and are in general different from what they would have been if each had been decided separately. In the present context, final performance yields income and does not immediately follow the forecast. Hence, incentivizing forecasts might force subjects to try and adjust gradually their behavior to their forecast and, therefore, unduly condition their behavior. A further difficulty encountered in this experiment was that, by incentivizing beliefs on three successive occasions, we induced risk-averse subjects to diversify their reported estimates as a hedge against the risk of prediction error. Self-report methods have been widely used and validated by psychologists and neuroscientists; and recent careful comparisons of this method with the quadratic scoring rule⁵ found that it performed as well (Clark and Friesen, 2009) or better (Hollard, Massoni and Vergnaud, 2015) than the quadratic scoring rule⁶. Considering that self-reports perform nicely while being much simpler and faster than incentive-compatible rules, use of the self-report seemed appropriate in this experiment.

⁵After the subject has reported a probability p , the quadratic scoring rule imposes a cost that is proportional to $(1 - p)^2$ in case of success and to $(0 - p)^2$ in case of failure. The score takes the general form: $S = a - b$. Cost, with $a, b > 0$.

⁶The second study also included the lottery rule in the comparison and found that the latter slightly outperformed self-report. The lottery rule rests on the following mechanism: after the subject has reported a probability p , a random number q is drawn. If q is smaller than p , the subject is paid according to the task. If q is greater than p , the subject is paid according to a risky bet that provides the same reward with probability q .

3 Confidence biases

3.1 Limited Discrimination

About half of our subjects were selected randomly into the 'wall' and 'hill' treatments and could not choose between the two. Those selected in one path were informed of the characteristics of their own path but had no knowledge whatsoever of the characteristics, nor even the existence, of the other path.

Result 1 (Limited discrimination): *Subjects do not perceive differences of difficulty between two different tasks in the future unless such differences are particularly salient. Moreover, they are not forward-looking, in the sense that they are unable to anticipate the increased likelihood of their success at the high level conditional on passing the middle level.*

However, they can be sophisticated when it is time for them to choose.

Support of result 1: Table 3 compares confidence judgments regarding the three levels of difficulty among the 'wall' and the 'hill' subjects before, during, and after the training period. Although the 'wall' and 'hill' were designed to be quite different at the middle and high levels, the subjective estimates of success exhibit almost no significant difference at any level. The single exception concerns the early estimate (before round 1) regarding the high level for which the difference of gradient between the two paths is particularly salient. However, the difference ceases to be significant as subjects acquire experience of the task. This striking observation suggests that individuals are unable to discriminate distinctive characteristics of the task unless the latter are particularly salient.

Perhaps even more disturbing is the fact that, in table 3, subjects discount their confidence level from the middle to the high level as much in the Wall as in the Hill treatment. For instance, just before the middle level, the ratio of confidence in passing the high level to confidence in passing the middle level was close to 0.70 in both treatments. However, a perfectly rational agent should realize that the high level is no more difficult than the middle level in the Wall treatment whereas it is much more difficult in the Hill treatment. Thus, she should report almost the same confidence at both levels in the Wall treatment, and a

considerably lower confidence at the high level in the Hill treatment. The latter observation suggests that most individuals are unable to compute conditional probabilities accurately even when the latter is equal to one as in the Wall treatment. They don't anticipate that, if they demonstrate the ability to solve 20 anagrams or more at middle level, they should be almost sure to solve 20 or more at the high level. However, subjects do make the right inference when it is time for them to make the decision since 95% of subjects who passed the middle level in the Wall treatment decided to continue (Table 1). And, if they have a choice between Wall and Hill, they do make a difference between these two tracks: 71.4% of doublers then prefer the Wall track although they would have greater chances of success at the middle level if they chose Hill. This observation suggests that subjects did not maximize their immediate probability of success but made a sophisticated comparison of the expected utility of both tracks, taking the option value of Wall in consideration before making an irreversible choice of track spanning over two periods⁷.

Table 3 – A comparison of confidence for the wall and hill treatments shown separately

Subjective confidence	No-choice treatment			
		Wall (%)	Hill (%)	Difference
Before round 1:	<i>Level 1</i>	80	77	<i>ns</i>
	<i>Level 2</i>	62	58	<i>ns</i>
	<i>Level 3</i>	47	40	**
Before round 5:	<i>Level 1</i>	71	71	<i>ns</i>
	<i>Level 2</i>	53	52	<i>ns</i>
	<i>Level 3</i>	40	36	<i>ns</i>
Before round 10:	<i>Level 2</i>	60	56	<i>ns</i>
	<i>Level 3</i>	43	39	<i>ns</i>

Number of observations: Before round 1 and 5 (before round 10): 101 (71) for wall and 106 (68) for hill. Significance Level: * $p < .10$, ** $p < .05$, *** $p < .01$, ns: not significant at 10% level.

3.2 Miscalibration

Result 2 (The hard-easy effect): *In comparison with actual performance, confidence in one's ability to reach a given level is underestimated for a novel but relatively easy task (the*

⁷We are grateful to Luis Santos-Pinto for making the last point clear in early discussions.




Figure 2 – Hard-easy effect for middle and high levels

Notes:

1. Observations: before training level (N: 410); after training level (N: 275 - analysis restricted to doublers).
2. Differences between frequency of success and confidence (before and after) are significant at 1% at all levels (Training, Middle and High). (paired *t-test*)

training level); and it is overestimated for the subsequent more difficult tasks (the middle and high levels). Overconfidence increases in relative terms with the difficulty of the task. Conditional on an initial success (training level) and on the decision to continue, confidence in one's ability to reach higher levels is still overestimated. Thus, initially successful subjects remain too optimistic about their future.

Support for result 2: Figure 2 compares the measured frequency of success with the reported subjective confidence in the three successive levels of increasing difficulty. For the middle and high levels, we also indicate these two probabilities as they appear before the training period and after it conditional on doubling. The Choice and No-choice conditions have been aggregated on this figure because no significant difference was found in the result of tests.

The task required at the training level was relatively easy for our subjects since 87% passed this level. However, subjects started it without knowing what it would be like and, even after four rounds of training, they underestimated their own ability to a low 77% probability of success. The difference among the two percentages is significant ($t=5.77$,

$p=0.000$; paired t-test). Hence, individuals are under-confident on the novel but relatively easy task.

In contrast, subjects appear to be overconfident as the task gets increasingly difficult. They consistently diminish their estimated probabilities of success but do not adjust their estimates in proportion to the difficulty of the task. Thus, individuals tend to overestimate their own chances for the advanced levels. The difference between the frequency of success and confidence before the task is always significant, both at the middle level ($t=18.3$, $p=0.000$) and at the high level ($t=17.1$, $p=0.000$).

The same conclusions hold conditional on passing the training level and choosing to double. Subjects remain overconfident in their future chances of success. However, their confidence does not rise after their initial success in proportion to their chances of further success.

3.3 The ability effect

Result 3 (The ability effect): *Overcalibration diminishes with task-specific ability.*

Support for result 3: The hard-easy effect is reproduced on figures 3a, 3b, 3c for the three ability terciles⁸. Low-ability subjects are obviously more overconfident at middle and high levels relative to high and medium-ability individuals. This result confirms earlier observations of Kruger (1999) among others (see Ryvkin, Krajč and Ortmann (2012) for a recent overview and incentivized experiments). The so-called Dunning-Kruger effect has been attributed to a metacognitive inability of the unskilled to recognize their mistakes⁹. We give here another, and in our opinion, simpler explanation¹⁰. The ability (or Dunning-Kruger) effect may be seen as a corollary of the hard-easy effect because "difficulty" is a

⁸Difference between confidence and frequency of success is significant at 1% for all ability levels. For these figures, we selected confidence reported after 4th round (during training level) in order to minimize the impact of mismeasurement.

⁹The Dunning-Kruger effect initially addressed general knowledge questions whereas we consider self-assessments of own performance in a real-effort task.

¹⁰Our explanation may also be better than the initial explanation such that the unskilled are unaware of their lower abilities. Miller and Geraci (2011) found that students with poor abilities showed greater overconfidence than high-performing students, but they also reported lower confidence in these predictions.




Figure 3a – Under-confidence at the training level, by ability.




Figure 3b – Overconfidence at middle level, by ability.

relative notion and a task that a low-ability individual finds difficult certainly looks easier to a high-ability person. Thus, if overconfidence rises with the difficulty of a task, it is natural to observe that it declines on a given task with the ability of performers.

4 Learning to be over (under-) confident: Theory

We present in this section a simple Bayesian model that integrates the discrimination and calibration biases, explains why people apparently don't learn completely, why they are conservative by under-reacting to experience, and predicts the hard-easy effect and the ability

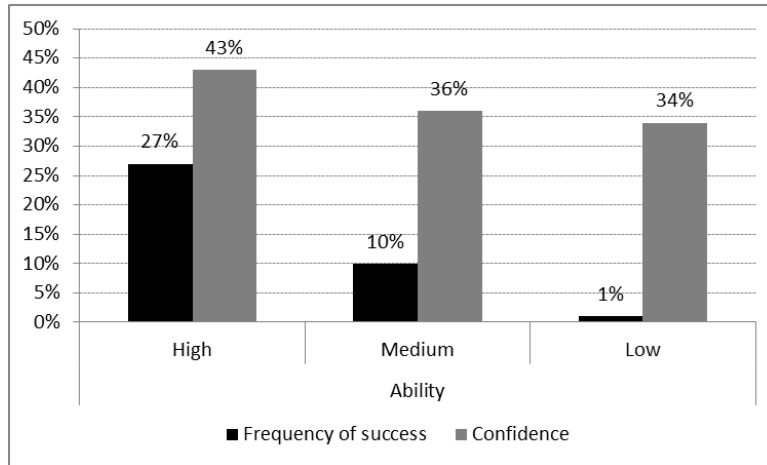


Figure 3c – Overconfidence at high level, by ability.

effect. It describes the behavior of subjects who are myopic, poorly discriminate, and thus make measurement errors; and who are uncertain of their true probability of success and may thus be victims of cognitive illusions. We designate these subjects as "intuitive Bayesians". It turns out, unexpectedly, that the same model also predicts the overprecision bias of confidence, which we consider as a further confirmation of its validity.

4.1 Doubt, miscalibration and overprecision

Intuitive Bayesians may miscalibrate their own probability of success even if they have an unbiased estimate of their own ability to succeed. This can occur if they are uncertain of the true probability of success because they can be misled by "available" illusory signals triggered by their doubt. The direction of doubt is entirely different depending on whether their prior estimate led them to believe that they would fail or that they would succeed. We thus distinguish miscalibration among those individuals who should normally believe that they will fail and those who should normally believe that they will succeed.

Let us consider first a subject who is almost sure to succeed a task, either because the task is easy or because the subject has high-ability (H). However, the "availability" of a possible failure often leads to overweighting this possibility (Tversky and Kahneman, 1973),

and underweighting her subjective probability of success¹¹ Ep_H , *i.e.* underconfidence:

$$q_H = \mu_H Ep_H + (1 - \mu_H)0 = \mu_H Ep_H \leq Ep_H, \quad (1)$$

with $0 < \mu_H \leq 1$

Even though high-ability agents are almost sure of succeeding the training level, their confidence is way below 1, confirming the Dunning-Kruger effect where high-ability subjects underestimate their abilities. An estimate of this undercalibration bias is derived from Figure 3aa:

$$\mu_H = \frac{0.79}{0.98} = 0.806 \cong q_H \text{ (training level)}$$

The undercalibration bias is: $1 - 0.806 = 0.194$.

At the other end of the spectrum, consider now a subject who is almost sure of failing, either because the task is very difficult or because the subject has low-ability (L). However, the "availability" of a possible success often leads to overweighting her subjective probability of success Ep_L *i.e.* overconfidence:

$$q_L = \mu_L Ep_L + (1 - \mu_L)1 \geq Ep_L, \quad (2)$$

with $0 < \mu_L \leq 1$

Thus, even though low ability agents should give up a task, they are overconfident and are thus tempted by the returns to success¹². In the limit, confidence remains positive if one is almost certain to fail. This means that low-ability individuals always exhibit a positive bottom confidence, which is in line with the Dunning-Kruger effect (they overestimate their abilities). An estimate of this overcalibration bias for the high level is derived from Figure 3c:

$$1 - \mu_L = \frac{0.34 - 0.01}{1 - 0.01} = 0.333 \cong q_L \text{ (high level)}$$

¹¹The time $t = (1, 5, 10)$ when confidence is reported is omitted in this sub-section to alleviate notations.

¹²This should not be confounded with motivated inference as it applies symmetrically to undesirable and desirable outcomes.

Similarly, the overcalibration bias for the middle level is derived from Figure 3b:

$$1 - \mu'_L = \frac{0.45 - 0.04}{1 - 0.04} = 0.427 \cong q_L \text{ (middle level)}$$

To sum up, we explain both the hard-easy effect and the ability effect by an availability bias triggered by the doubt about one's possibility to fail a relatively easy task (underconfidence) or to succeed a relatively difficult task (overconfidence). If probabilities are updated in a Bayesian fashion, the calibration bias is the relative precision of the illusory signal. The latter is inversely related with the absolute precision of the prior estimate and positively related with the absolute precision of the illusory signal. Thus, we mustn't be surprised to find that our estimate of the calibration bias is lower for the training level (19.4%) than for upper levels (42.7% and 33.3% respectively) because experience in the first rounds of the training level must be more relevant for predicting the probability of success in the training level than in subsequent levels. And, when comparing upper levels, the illusion of success should be more credible for the near future (middle level) than for the more distant future (high level).

This explanation is also consistent with the other measures displayed by figures 3a, 3b, 3c, given the fact that they aggregate overconfident subjects who should not undertake the task with underconfident subjects who should undertake it¹³. If λ_L is the proportion who should stop and λ_H the proportion who should continue ($\lambda_L + \lambda_H \equiv 1$), the average confidence is: $\lambda_L(\mu Ep_L + 1 - \mu) + \lambda_H \mu Ep_H = \mu \overline{Ep} + (1 - \mu)\lambda_L$. Confidence is overcalibrated on average iff $\lambda_L > \overline{Ep}$ and undercalibrated iff the reverse condition holds. The apparent overcalibration of confidence for a difficult task takes less extreme values when the average measured ability of the group rises. For instance, the results displayed by figure 3c are consistent with our estimate for the overcalibration bias if the proportion of successful middle-ability subjects is

¹³The rational decision to undertake a non-trivial task of level l (with a possibility to fail and regret) is subjective. The economic criterion for making this decision rests on the comparison of the expected utilities of all options conditional on the estimated probabilities of success at the time of decision. A rational subject should refuse the task if the expected utility of continuing to level l or above is no higher than the expected utility of stopping before level 1. We make use of this criterion for writing equations 6 and 7 in the next sub-section (4.2).

12% and that of successful high-ability subjects is 25%, since these two predicted values are close to the observed frequency of success in these groups, respectively 10% and 27%.

A further implication of Bayesian updating is that, in the subject's mind, the precision of the posterior estimate for probabilities of success, *i.e.* confidence in his estimate, is increased by reception of the illusory signal, whatever the latter may be¹⁴. Therefore, our theory of confidence predicts the overprecision phenomenon! In contrast with the other distortions of confidence, underprecision will never be observed, a prediction which is corroborated by Moore and Healy (2008) who do not quote any study in their discussion of "underprecision". The overestimation of the precision of acquired knowledge is an additional manifestation of the hidden search undertaken by intuitive Bayesians. Our analysis of overprecision is congruent with the observation that greater overconfidence of this kind was found for tasks in which subjects considered they were more competent (Heath and Tversky, 1991).

4.2 Confidence updating by intuitive Bayesians

In our experiment, confidence is reported prior to the task E_1p , after four rounds E_5p , and after nine rounds (only for doublers) $E_{10}p$.

After going through four rounds of anagrams, a number of cues on the task have been received and processed. Participants may recall how many anagrams they solved in each round and in the aggregate, whether they would have passed the test in each round or on the whole at this stage of the task, whether their performance improved or declined from one round to the next, how fast they could solve anagrams, and so forth. For the purpose of decision-making, each cue is converted into a discrete variable reflecting the predicted performance (subjective success) of individuals for each level l ($l = 1, 2, 3$): x_{tl} takes value 1 if an individual is expected to reach his goal for level l in round t and 0 otherwise. Individuals observe the *i.i.d.* sequence $(x_{1l}, x_{2l}, x_{3l}, x_{4l})$ of a Bernoulli variable of performance at level l during the first four rounds. The single parameter of the random variable is its mean which

¹⁴If ν_i denotes the prior precision of subject I 's estimate of her future success (omitting level l for simplicity) $\nu_i + 1 \equiv \Phi_i$ will be the posterior precision after reception of a signal. Thus, $\Phi_i > \nu_i$. Notice that $\mu_i = \frac{\nu_i}{\nu_i + 1}$.

defines the expected likelihood of success. However, this mean is essentially unknown to that individual. Thus, let it be denoted by \tilde{p} which is randomly distributed within the interval $[0, 1]$. Assume that the prior distribution of \tilde{p} is a Beta-distribution with a reported mean E_1p .

Behaving like intuitive Bayesians, participants update their prior expectation of success at level l ($l = 1, 2, 3$) before the training session E_1p_l in the following manner (see DeGroot 1970, Chapter 9):

$$E_5p_l = \frac{\nu_l}{\nu_l + \tau_{4l}} E_1p_l + \frac{1}{\nu_l + \tau_{4l}} X_{1-4,l} \quad (3)$$

with $\nu_l, \tau_{4l} > 0$ designating the respective precisions of the prior expectation and all of the independent cues perceived during the first four rounds, $\nu_l + \tau_{4l}$ being the precision of the posterior expectation E_5p_l , and $X_{1-4,l} = x_{1l} + x_{2l} + x_{3l} + x_{4l}$ defining the number of independent cues predicting future success at level l at this stage of the task. Equation (3) cannot be directly estimated on the data because the estimated probabilities E_1p_l and E_5p_l are unobservable. However, it may be rewritten in terms of reported confidence $q_1(l)$ and $q_5(l)$ with the help of the miscalibration equations (1) and (2). Let us express generally the Bayesian transformation of the probability estimates into confidence as:

$$q_5(l) = \mu_{5l} E_5p_l + (1 - \mu_{5l}) D_{5,l}, \quad l = (1, 2, 3) \quad (4)$$

with $\mu_{5l} = \frac{\nu_l + \tau_{4l}}{\nu_l + \tau_{4l} + 1}$ and¹⁵

$$D_{(5,l)} = 1 \text{ if } \operatorname{argmax} EU(l \mid E_5p_l, l = (0, \dots, l-1)) \leq \operatorname{argmax} EU(l' \mid E_5p_{l'}, l' = (l, \dots, 3)) \\ = 0 \text{ otherwise.}$$

And

$$q_1(l) = \mu_{1l} E_1p_l + (1 - \mu_{1l}) D_{1,l} \quad (5)$$

with $\mu_{1l} = \frac{\nu_l}{\nu_l + 1}$ and

$$D_{(1,l)} = 1 \text{ if } \operatorname{argmax} EU(l \mid E_1p_l, l = (0, \dots, l-1)) \leq \operatorname{argmax} EU(l' \mid E_1p_{l'}, l' = (l, \dots, 3))$$

¹⁵In order to have an unambiguous definition of $D_{(5,l)}$ and $D_{(1,l)}$ below, we use the expected utility (EU) criterion, as explained in note 13.

= 0 otherwise.

Combining (3), (4) and (5), we get:

$$q_5(l) = \frac{\nu_l + 1}{\nu_l + \tau_{4l} + 1} q_1(l) + \frac{1}{\nu_l + \tau_{4l} + 1} X_{1-4,l} + \frac{1}{\nu_l + \tau_{4l} + 1} (D_{5,l} - D_{1,l}) \quad (6)$$

By the same reasoning, we can express the confidence of doublers for upper levels $l = (2, 3)$ as:

$$q_{10}(l) = \frac{\nu_l + \tau_{4l} + 1}{\nu_l + \tau_{9l} + 1} q_5(l) + \frac{1}{\nu_l + \tau_{9l} + 1} X_{5-9,l} + \frac{1}{\nu_l + \tau_{9l} + 1} (D_{9,l} - D_{5,l}) \quad (7)$$

with $\tau_{9l} \geq \tau_{4l}$ designating the precision of all of the independent cues perceived during the training level (9 rounds), $\nu_l + \tau_{9l}$ the precision of the posterior expectation $E_{10}p_l$, and $X_{(5-9,l)} = x_{5l} + \dots + x_{9l}$ defining the number of independent cues predicting future success at level l at this stage of the task.

Equations (6) and (7) are essentially the same with a moving prior of increasing precision. In the absence of miscalibration, confidence reported before round t ($t = (5, 10)$) would be a weighted average of prior confidence and the mean frequency of cues predicting future success at level l since the last time confidence was reported. With miscalibration, another term is added which can only take three values, reflecting the occurrence and direction of change in subjects' estimated ability with experience. If experience confirms the prior intention to stop or continue to level l , this additional term takes value 0 and confidence is predicted by the rational-Bayesian model (with perfect calibration). However, if experience disconfirms the prior intention to stop or continue to level l , confidence rises above this reference value with disappointing experience and declines symmetrically below this reference value with encouraging experience. *Thus, our model predicts that intuitive Bayesians be conservative and under-react symmetrically to negative experience (by diminishing their confidence less than they should) and to positive experience (by raising their confidence less than they should).* Below, we report indeed rather small variations of confidence in our experiment in the form of local, but not global, learning.

5 Learning to be over (under-) confident: Results

5.1 Learning

Result 4 (Learning is local, not global): *Confidence and performance co-vary during the task. Subjects learned **locally** upon experiencing variations in their performance. However, they didn't learn **globally** in our experiment, since doublers remained as confident as before after completing the training level irrespective of their true ability level.*

Support for result 4: Figures 4 and 5 describe confidence by ability group before, during, and after the training period¹⁶¹⁷ for the middle and high level respectively whereas figure 6 describes the variation of performance of the same groups within the same period. These graphs, taken together, show a decline in both (ability-adjusted) confidence and performance during the first four rounds, followed by a concomitant rise of confidence and performance in the following rounds¹⁸. The observed decline of confidence at the beginning of the training period can be related on Figure 6 to the fact that participants solved less and less anagrams per period during the first four periods: 5.51 on average in period 1, 5.18 in period 2, 4.60 in period 3, and 4.17 in period 4¹⁹. Subjects kept solving at least two-thirds of the anagrams available during the training session but probably lost part of their motivation on repeating the task. On sequentially observing their declining performance, they revised their initial estimate of future success downward. However, on being asked to report their confidence after four rounds, they became conscious of their performance decline and responded to this information feedback. Performance rose sharply but momentarily during the next two

¹⁶No significant difference was found between the Choice and No-choice conditions, suggesting that the option to choose the preferred path does not trigger an illusion of control.

¹⁷The participants who reported confidence after the training period were more able than average since they had passed this level and decided to double. Thus, we compare ability-adjusted confidence "before" and "during" with the reported confidence "after". The ability-adjusted confidence "before" and "during" are obtained by running a simple linear regression of confidence "before" and "during" on ability, measured by the average number of anagrams solved per minute in the first 4 rounds of the training level. The estimated effect of superior ability of "doublers" was added to confidence "during" or "before" to get the ability-adjusted confidence which directly compares with the observed confidence "after".

¹⁸With a single exception, confidence variations are statistically significant at 1% at the middle and high levels.

¹⁹There was no significant difference between treatments.




Figure 4 – Variation of confidence with experience, by level of ability: middle level

Notes²⁰: Differences between the three levels of ability are significant at 1% "Before" and "During". Differences After are not significant at 10% "After". (Bonferroni multiple comparison tests). Sample size: 410 individuals for "Before" and "During"; 275 for "After" (only doublers).

rounds. The average performance first rose to 4.37 in period 5 and 5.05 in period 6 then sharply declined to 4.39 in period 7, 4.06 in period 8 and 3.48 in period 9. As soon as subjects became (almost) sure of passing the training level, they diminished their effort. During the experiment it was also observed that individuals stopped decoding further anagrams as soon as the minimum requirement to clear a level was fulfilled.

Subjects experiencing low (medium) performance in the first rounds seem to learn locally that they have a low (medium) ability since the confidence gap widens during the first four periods. However, this learning effect is short-lived since the confidence gap shrinks back to its initial size after low (medium)-ability subjects strove to succeed, increasing their performance (as reported on Figure 6) and regaining confidence. Eventually, experienced "doubblers" are as confident to succeed at higher levels as they were before the task, irre-

²⁰Differences on confidence: High ability: During-Before: ***; After-During: ns; After-Before: ns. Medium ability: During-Before: ***; After-During: ***; After-Before: **. Low ability: During-Before: ***; After-During: *** ; After-Before: ns. Significance level: *** 1%; ** 5%; * 10%; ns: not significant at 10% level. (T-test).

²¹Differences on confidence: High ability: During-Before: ns; After-During: ns; After-Before: ns. Medium ability: During-Before: ***; After-During: **; After-Before: ns. Low ability: During-Before: ***; After-During: *** ; After-Before: ns. Significance level: *** 1%; ** 5%; * 10%; ns: not significant at 10% level. (T-test).

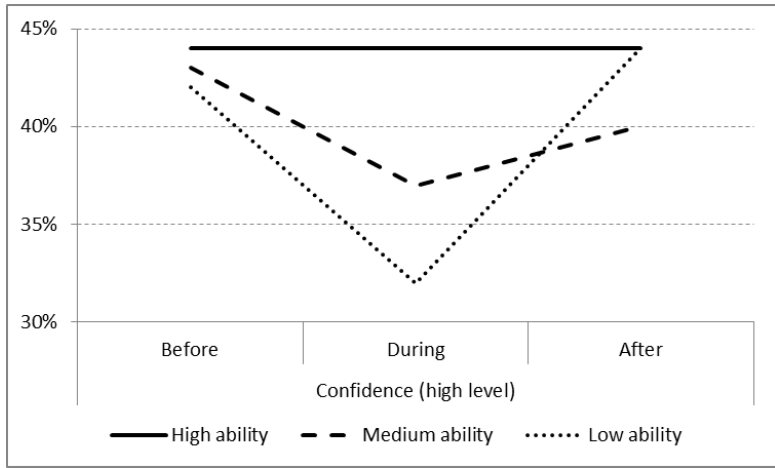


Figure 5 – Variation of confidence with experience, by level of ability: high level

Notes²¹: Differences between the three levels of ability are significant at 1% "Before" and "During". Differences After are not significant at 10% "After". (Bonferroni multiple comparison tests). Sample size: 410 individuals for "Before" and "During"; 275 for "After" (only doublers).




Figure 6 – Number of anagrams solved per round by level of ability

spective of their ability level: there is no global learning effect. We share the conclusion of Merkle and Weber (2011) that the persistence of prior beliefs is inconsistent with fully rational-Bayesian behavior(see also Benoît, Dubra and Moore 2015).

5.2 Regression analysis

The models of Bayesian estimation of confidence described by equations (6) and (7) are tested by an OLS in Tables 4 and 5 ²² respectively. Reported confidence in participant i 's ability to reach one level of the double-or-quits game is regressed in table 4 (table 5) on the confidence that he reported before the first (fifth) round and on a vector Z_{li} of level-specific cues observable in the first four (last five) rounds, assuming that $X_{1-4,li}(X_{5-9,li}) = \beta_l Z_{li} + \epsilon_{li}$ where β_l is a vector of coefficients and ϵ_{li} an error term of zero mean. Two dummy variables for the hill and choice treatments (wall as reference) have been added to the regression.

Table 4 – OLS estimation of the Bayesian model of confidence before round 5

	Training Level	Middle Level	High Level
Confidence before training session	0.79***	0.86***	0.90***
Freq. of rounds with 4 anagrams solved	0.14***	0.06 ^{ns}	0.01 ^{ns}
Freq. of rounds with 5- 6 anagrams solved	0.29***	0.19***	0.13***
Freq. of rounds with non-declining performance	0.12***	0.10***	0.09***
Anagrams solved per minute on rounds 1-4	0.01***	0.01***	0.01***
Hill	0.03*	0.04**	0.03*
Choice	0.01 ^{ns}	0.02 ^{ns}	0.00 ^{ns}
Constant	-0.25***	-0.25***	-0.20***
R^2	67%	70%	76%

N=410

Significance Level: * $p < .10$, ** $p < .05$, *** $p < .01$, ns: not significant at 10% level.

Frequency of rounds with non-declining performance: percentage of rounds (in rounds 2-4) in which number of anagrams solved was equal or higher than in the previous round. Takes four values (0,.33,.67,1).

Hill and Choice: Wall as reference.

The regressions confirm the existence of local learning. Subjects did revise their expectations with experience of the task as several cues have highly significant coefficients (at 1% level) with the right sign. Moreover, they analyze their own performance correctly by setting stronger pre-requisites for themselves when the task gets more difficult. For example, their ability to solve just four anagrams per round in the training period increases their confidence for this period only because, if such performance is enough to ensure success in this period, it is no longer sufficient when the task becomes more difficult. Another interesting result in table 5 consistent with the miscalibration term in equation (7) concerns low achievers who

²²The discrete value of confidence between 0 and 100 can be safely treated as continuous.

Table 5 – OLS estimation of the Bayesian model of confidence for doublers before the middle level

	Middle Level	High Level
Confidence after round 4	0.772***	0.872***
Freq. of rounds with 4 anagrams solved (5-9)	0.017 ^{ns}	-0.024 ^{ns}
Freq. of rounds with 5- 6 anagrams solved (5-9)	0.120***	0.073*
Freq. of rounds with non-declining performance (5-9)	0.034 ^{ns}	0.088**
Number of rounds used to solve 36 anagrams	0.027***	0.021**
Anagrams solved per minute on rounds 5-9	0.003 ^{ns}	-0.003 ^{ns}
Hill	-0.047***	-0.022 ^{ns}
Choice	-0.017 ^{ns}	0.022 ^{ns}
Constant	-0.136 ^{ns}	-0.186**
R^2	74%	81%

N=275

Significance Level: * $p < .10$, ** $p < .05$, *** $p < .01$, ns: not significant at 10% level.

Frequency of rounds with non-declining performance: percentage of rounds (in rounds 5-9) in which number of anagrams solved was equal or higher than in the previous round.

Hill and Choice: Wall as reference.

(5-9): measures between rounds 5 and 9

Number of rounds used to solve 36 anagrams (between rounds 6 and 9)

double. The later they ended up solving the required number of anagrams in the training period, the more abruptly their confidence rose. It is indeed an implication of subjects' vulnerability to illusory signals that low-ability doublers find themselves almost as confident as high-ability doublers in spite of widely different performances. This result appears too on figures 4 and 5, where the confidence of low-ability doublers jumps from bottom to top during the second stage of the training period.

A major testable implication of the Bayesian model lies in the coefficient of the prior confidence, which must be interpreted as the precision of prior information relative to the information collected by experience of the task during the training period. This coefficient is always high in tables 4 and 5 with a minimum value of 0.77. Observing such high weights for the prior favors the hypothesis of rational-Bayesian updating over adaptive expectations as the latter would considerably underweight the prior relative to the evidence accumulated in the first four rounds. Successful experience of the easier task in the early rounds is expected to be more predictive of final success on the same task than in future tasks of greater difficulty. Thus, the relative weight of experience should diminish in the confidence

equation at increasing levels or, equivalently, the relative weight of prior confidence should rise. Indeed, the coefficient of prior confidence increases continuously with the level. It rises from 0.79 to 0.86 and 0.90 in table 4; and, from 0.77 to 0.87 in table 5. In parallel, the coefficients of cues signaling a successful experience continuously diminish when the level rises. We can use the mathematical expressions of the two coefficients of prior confidence derived from equations (6) and (7) to calculate the precision of early experience relative to prior confidence (before the task) $\frac{\tau_{4l}}{\nu_l}$ ($l = 1, 2, 3$). With the data of table 4, we get 0.266 for the training level, 0.163 for the middle level, and 0.111 for the high level. Similarly, we compute the precision of late experience relative to prior confidence (before the task) $\frac{\tau_{9l}}{\nu_l}$ ($l = 2, 3$). With the data of table 5, we get 0.506 for middle level and 0.274 for high level. The impact of learning from experience appears to be substantial and with increasing returns. By elimination of ν_l , we finally calculate the precision of early experience relative to total experience during the training period $\frac{\tau_{4l}}{\tau_{9l}}$ ($l = 2, 3$). We obtain 0.322 for middle level and 0.405 for high level. The rate of increase of precision resulting from longer experience (from 4 to 9 rounds) $\frac{\tau_{9l} - \tau_{4l}}{\tau_{4l}}$ reaches a considerable 211% at middle level and 147% at high level, which forms indirect evidence of the overprecision phenomenon.

6 Why do intuitive Bayesians make wrong predictions of performance?

The answer to this important question is contained in table 6, which uses the same set of potential predictors to forecast confidence in succeeding the middle level after doubling and ex post chances of success²³: prior confidence, ability, and performance cues observed subsequently (during rounds 5 to 9). The mere comparison of coefficients between the two columns of table 6 demonstrates that posterior confidence is based on both objective performance cues and subjective variables, whereas the chances of success are predicted by the objective performance cues and ability only. The latter are the frequencies of rounds with

²³We used an OLS to predict probabilities of success so as to make the comparison with confidence transparent. Estimating an OLS instead of a Probit in columns 3 and 4 didn't affect the qualitative conclusions.

4 and with 5-6 anagrams solved respectively (effort) and the speed of anagram resolution (ability); and the subjective variables are essentially the prior confidence and the illusory signal given to low achievers by their (lucky) initial success. Remarkably, the number of rounds needed for solving 36 anagrams (varying from 6 to 9), which indicates low achievement and recommends quitting the game at an early stage, acts as an illusory signal with a significantly positive effect on confidence in column 1; but the same variable acts as a correlate of low ability in column 2 with a strong negative effect on the chances of success at middle level. Indeed, the subjective predictors of posterior confidence do not predict success when the objective performance cues are held constant. Prior confidence predicts the posterior confidence that conditions the decision to double²⁴ but fails to predict success because it is based on an intuitive reasoning which suffers from systematic biases. Past errors convey to the prior through the aggregation procedure of Bayesian calculus and may add up with further errors caused by the perception of illusory signals.

²⁴Conditional on initial success, prior confidence is a good predictor of the future decision to double (regression not shown). This is good news for the quality of confidence reports; and it confirms that subjects behave as intuitive Bayesians who rely on their own subjective estimates of success to make the choice of doubling.

Table 6 – Estimation of posterior confidence (after doubling) and ex post chances of success at the middle level

	Level 2	
	Confidence After	Chances of success
Confidence after round 4	0.778 ^{***}	0.034 ^{ns}
Freq. of rounds with 4 anagrams solved (5-9)	0.014 ^{ns}	0.276 [*]
Freq. of rounds with 5- 6 anagrams solved (5-9)	0.107 ^{***}	0.348 ^{**}
Freq. of rounds with non-declining performance (5-9)	0.043 ^{ns}	-0.036 ^{ns}
Number of rounds used to solve 36 anagrams	0.024 ^{***}	-0.115 ^{***}
Anagrams solved per minute on rounds 5-9	0.009 ^{ns}	0.070 ^{***}
Ability	-0.007 ^{ns}	0.062 ^{***}
Hill	-0.046 ^{***}	0.097 ^{ns}
Choice	-0.018 ^{ns}	-0.100 [*]
Constant	-0.106 ^{ns}	0.598 [*]
R^2	74%	30%

N=275 (only those succeeded first level and decided to double to second level)

Significance Level: * $p < .10$, ** $p < .05$, *** $p < .01$, ns: not significant at 10% level.

Frequency of rounds with non-declining performance: percentage of rounds (in rounds 5-9) in which number of anagrams solved was equal or higher than in the previous round.

Hill and Choice: Wall as reference.

(5-9): measures between rounds 5 and 9

Number of rounds used to solve 36 anagrams (between rounds 6 and 9)

7 Conclusion

We designed an experimental analog to the popular double-or-quits game to compare the speed of learning one's ability to perform a task in isolation with the speed of rising confidence as the task gets increasingly difficult. In simple words, we found that people on average learn to be overconfident faster than they learn their true ability. We present a new intuitive-Bayesian model of confidence which integrates these facts and makes further predictions. The distinctive feature of our model of self-confidence is that it rests solely on a Bayesian representation of the cognitive process: people predict their own probability of performing a task on the basis of incomplete, sequentially collected, information related to the task. Confidence biases arise, not from an irrationality of the treatment of information, but from the poor quality and subjectivity of the information being treated. For instance, we rule out self-attribution biases, motivated cognition, self-image concerns and manipulation of beliefs but we describe people as being fundamentally uncertain of their future performance

and taking all the information they can get with limited discrimination, including cognitive illusions. Above all, a persistent doubt about their true ability is responsible for their perception of contrarian illusory signals that make them believe, either in their failure if they should succeed or in their success if they should fail.

Our analysis of overconfidence is restricted to the overestimation bias. In our opinion, the latter should be carefully distinguished from the overplacement bias since the hard-easy effect that we observed here with absolute confidence has often been reversed when observing relative confidence: overplacement for an easy task (like driving one's car) and underplacement for a novel or difficult task. A remarkable feature of our intuitive-Bayesian theory of estimation, though, is that the latter predicts the overprecision phenomenon and even rules out underprecision. This demonstrates that overestimation and overprecision are related but different biases. Going one step further, we think that the literature is sometimes misleading when it emphasizes "overconfidence" as a general phenomenon. Indeed, the hard-easy effect and the ability or Dunning-Kruger effect are more descriptive of the data on estimation biases. Pure overconfidence is truly limited to the overprecision phenomenon.

We showed that limited discrimination of objective differences and myopia can be responsible for large prediction errors which learning should help to reduce. However, it is mainly the fundamental uncertainty about one's true ability which causes systematic and robust confidence biases, namely the hard-easy effect, the Dunning-Kruger effect, conservative learning from experience and the overprecision phenomenon (without underprecision) if subjects act as Bayesian learners. Moreover, these biases are likely to persist since the Bayesian aggregation of past information consolidates the accumulation of errors, and the perception of illusory signals generates conservatism and under-reaction to events. Taken together, these two features may explain why intuitive Bayesians make systematically wrong predictions of their own performance. Don't we systematically underestimate the time needed to perform a new (difficult) task and never seem to learn?

Acknowledgements

We thank the French *Ministère de la Recherche* (ACI "*Contextes sociaux, contextes institutionnels et rendements des systèmes éducatifs*") for generous support, Claude Montmarquette for offering an opportunity to conduct part of the experimental sessions at CIRANO (Montreal), and Noemi Berlin for numerous discussions. We remain responsible for any error.

References

- Adams, Joe K.** 1957. "A confidence scale defined in terms of expected percentages." *The American journal of psychology*, 432–436.
- Anderson, Cameron, Sebastien Brion, Don A Moore, and Jessica A Kennedy.** 2012. "A status-enhancement account of overconfidence." *Journal of personality and social psychology*, 103(4): 718.
- Armantier, Olivier, and Nicolas Treich.** 2013. "Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging." *European Economic Review*, 62: 17–40.
- Barber, Brad M, and Terrance Odean.** 2001. "Boys will be boys: Gender, overconfidence, and common stock investment." *Quarterly journal of Economics*, 261–292.
- Benoît, Jean-Pierre, and Juan Dubra.** 2011. "Apparent overconfidence." *Econometrica*, 79(5): 1591–1625.
- Benoît, Jean-Pierre, Juan Dubra, and Don A Moore.** 2015. "Does the better-than-average effect show that people are overconfident?: Two experiments." *Journal of the European Economic Association*, 13(2): 293–329.
- Breen, Richard.** 2001. "A rational choice model of educational inequality." *Estudios/Working Papers (Centro de Estudios Avanzados en Ciencias Sociales)*, , (166): 1.
- Brunnermeier, Markus K., and Jonathan A. Parker.** 2005. "Optimal Expectations." *American Economic Review*, 95(4): 1092–1118.
- Camerer, Colin, and Dan Lovallo.** 1999. "Overconfidence and excess entry: An experimental approach." *American economic review*, 306–318.
- Clark, Jeremy, and Lana Friesen.** 2009. "Overconfidence in Forecasts of Own Performance: An Experimental Study*." *The Economic Journal*, 119(534): 229–251.
- Deaves, Richard, Erik Lüders, and Michael Schröder.** 2010. "The dynamics of overconfidence: Evidence from stock market forecasters." *Journal of Economic Behavior & Organization*, 75(3): 402–412.
- DeGroot, Morris.** 1970. *Optimal Statistical Decisions*. McGraw-Hill.
- Erev, Ido, Thomas S Wallsten, and David V Budescu.** 1994. "Simultaneous over-and underconfidence: The role of error in judgment processes." *Psychological review*, 101(3): 519.

- Gervais, Simon, and Terrance Odean.** 2001. "Learning to be overconfident." *Review of Financial studies*, 14(1): 1–27.
- Goodie, Adam S.** 2005. "The role of perceived control and overconfidence in pathological gambling." *Journal of Gambling Studies*, 21(4): 481–502.
- Grieco, Daniela, and Robin M Hogarth.** 2009. "Overconfidence in absolute and relative performance: The regression hypothesis and Bayesian updating." *Journal of Economic Psychology*, 30(5): 756–771.
- Griffin, Dale, and Amos Tversky.** 1992. "The weighing of evidence and the determinants of confidence." *Cognitive psychology*, 24(3): 411–435.
- Heath, Chip, and Amos Tversky.** 1991. "Preference and belief: Ambiguity and competence in choice under uncertainty." *Journal of risk and uncertainty*, 4(1): 5–28.
- Hollard, Guillaume, Sébastien Massoni, and Jean-Christophe Vergnaud.** 2015. "In search of good probability assessors: an experimental comparison of elicitation rules for confidence judgments." *Theory and Decision*, 1–25.
- Johnson, Chalmers.** 2004. "America's Empire of Bases." *TomDispatch. com*, 15.
- Johnson, Dominic DP, and James H Fowler.** 2011. "The evolution of overconfidence." *Nature*, 477(7364): 317–320.
- Kőszegi, Botond.** 2006. "Ego utility, overconfidence, and task choice." *Journal of the European Economic Association*, 4(4): 673–707.
- Kruger, Justin.** 1999. "Lake Wobegon be gone! The " below-average effect" and the egocentric nature of comparative ability judgments." *Journal of personality and social psychology*, 77(2): 221.
- Langer, Ellen J, and Jane Roth.** 1975. "Heads I win tails it's chance: The illusion of control as a function of the sequence outcomes in a purely chance task." *Journal of Personality and Social Psychology*, 32: 951–955.
- Lichtenstein, Sarah, and Baruch Fischhoff.** 1977. "Do those who know more also know more about how much they know?" *Organizational behavior and human performance*, 20(2): 159–183.
- Lichtenstein, S, B Fischhoff, and L Phillips.** 1982. "Calibration of probabilities: The state of the art to 1980. D. Kahneman, P. Slovic, and A. Tverski (Eds.) *Judgement under uncertainty: Heuristics and biases.*"

- Malmendier, Ulrike, and Geoffrey Tate.** 2005. "CEO overconfidence and corporate investment." *The journal of finance*, 60(6): 2661–2700.
- Merkle, Christoph, and Martin Weber.** 2011. "True overconfidence: The inability of rational information processing to account for apparent overconfidence." *Organizational Behavior and Human Decision Processes*, 116(2): 262–271.
- Miller, Dale T, and Michael Ross.** 1975. "Self-serving biases in the attribution of causality: Fact or fiction?" *Psychological bulletin*, 82(2): 213.
- Miller, Tyler M, and Lisa Geraci.** 2011. "Unskilled but aware: reinterpreting overconfidence in low-performing students." *Journal of experimental psychology: learning, memory, and cognition*, 37(2): 502.
- Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat.** 2014. "Managing Self-Confidence." Working Paper.
- Moore, Don A, and Paul J Healy.** 2008. "The trouble with overconfidence." *Psychological review*, 115(2): 502.
- Oskamp, Stuart.** 1965. "Overconfidence in case-study judgments." *Journal of consulting psychology*, 29(3): 261.
- Ryvkin, Dmitry, Marian Krajč, and Andreas Ortmann.** 2012. "Are the unskilled doomed to remain unaware?" *Journal of Economic Psychology*, 33(5): 1012–1031.
- Shiller, Robert J.** 2000. "Measuring bubble expectations and investor confidence." *The Journal of Psychology and Financial Markets*, 1(1): 49–60.
- Svenson, Ola.** 1981. "Are we all less risky and more skillful than our fellow drivers?" *Acta Psychologica*, 47(2): 143–148.
- Tversky, Amos, and Daniel Kahneman.** 1973. "Availability: A heuristic for judging frequency and probability." *Cognitive psychology*, 5(2): 207–232.
- Van den Steen, Eric.** 2011. "Overconfidence by Bayesian-rational agents." *Management Science*, 57(5): 884–896.