

Learning about One's Self

Yves Le Yaouanq, Peter Schwardmann

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Learning about One's Self

Abstract

How can naiveté about present bias persist despite experience? To answer this question, our experiment investigates participants' ability to learn from their own behavior. Participants decide how much to work on a real effort task on two predetermined dates. In the week preceding each work date, they state their commitment preferences and predictions of future effort. While we find that participants are present biased and initially naive about their bias, our methodology enables us to establish that they are Bayesian in how they learn from their experience at the first work date. A treatment in which we vary the nature of the task at the second date further shows that learning is unencumbered by a change in environment. Our results suggest that persistent naiveté cannot be explained by a fundamental inferential bias. At the same time, we find that participants initially underestimate the information that their experience will provide - a bias that may lead to underinvestment in experimentation and a failure to activate self-regulation mechanisms.

JEL-Codes: D830.

Keywords: naiveté, present bias, learning.

Yves Le Yaouanq
Department of Economics
University of Munich (LMU)
Kaulbachstr. 45
Germany – 80539 Munich
yves.leyaouanq@econ.lmu.de

Peter Schwardmann
Department of Economics
University of Munich (LMU)
Ludwigstr. 28
Germany – 80539 Munich
peter.schwardmann@econ.lmu.de

January 2019

This experiment was pre-registered in the AEA RCT Registry (*AEARCTR-0003060*) and approved by the Ethics Commission of the Department of Economics at the University of Munich (ID 2018-03). We report all experimental data and all experimental conditions performed in this research project. We thank Florian Englmaier, Lena Greska, Taisuke Imai, Alex Imas, George Loewenstein, Michael Muehlegger, Takeshi Murooka, Matthew Rabin, Klaus Schmidt, Séverine Toussaert, and various seminar audiences for helpful comments and we gratefully acknowledge financial support from the *Deutsche Forschungsgemeinschaft* through CRC TRR 190.

1 Introduction

When deciding how much to work next week, we tend to set ambitious goals for ourselves. But come next week, we often work less than we originally intended. Such time inconsistencies are especially harmful to someone who is naive about their existence. Whereas a sophisticated individual will commit to her desired course of action and avoid being exploited in markets, a naive individual bears the brunt of her present bias.¹ Naiveté about present bias has been documented in exercising (DellaVigna and Malmendier, 2006), tobacco consumption (Giné et al., 2010), saving behavior (John, 2018) and real-effort experiments (Augenblick and Rabin, 2018b). Yet it is puzzling that naiveté should persist in repeated behaviors, which afford ample opportunity to learn. To resolve this puzzle we require a better understanding of how individuals draw or fail to draw lessons from their experience.

We report the results of a longitudinal experiment in which participants update their beliefs about future effort based on their past effort choice. Our methodology permits a non-parametric investigation of the updating process and the definition of a Bayesian benchmark against which participants' learning biases can be detected. We are able to ascertain not just whether participants learn, but how well they learn, how much they expect to learn, and whether they are able to transport what they learn from one setting to the next.

Over three weeks, 187 subjects participated in five experimental sessions. The first session took place in the experimental laboratory and allowed participants to familiarize themselves with the experimental tasks and the mechanism used to elicit beliefs. The following four sessions took place online and are labeled date 1, 2, 3 and 4. At dates 2 and 4 participants had the opportunity to complete a maximum of 40 computer screens of an unpleasant task, which either involved positioning sliders on pre-specified targets or counting zeroes in tables of ones and zeroes. At dates 1 and 3, on average five days before dates 2 and 4 respectively, we elicited participants' commitment preferences over and predictions of future effort. Commitment preferences were binding with a probability of 5 percent, in which case participants had to complete the stated number of screens at the next date or forgo any payment for the task. Subjects were paid a few days after the experiment, conditionally on participating in all sessions.

Our methodology leverages two pieces of data. The first primitive describes the on-the-spot effort choices at dates 2 and 4. We transform effort at a given date t into a binary variable a_t that equals 1 if the subject completes at least 20 of 40 screens

¹See section 2.1 of DellaVigna (2009) and Laibson (2015) for the link between naiveté and a lack of demand for commitment, and Köszegi (2014) for the exploitation of naive agents in markets.

(*high effort*) and 0 otherwise (*low effort*). The true data-generating process is given by the probability distribution $q(a_2, a_4)$ over the four possible intertemporal events $(a_2, a_4) \in \{0, 1\}^2$ and can be constructed from the empirical frequencies of high and low effort. This primitive characterizes not only the marginal distributions of a_2 and a_4 , which describe participants' probability of high effort at a given date, but also the informativeness of effort at date 2 for effort at date 4. We measure the *true informativeness* of date 2 effort by the likelihood ratios $q(a_2 | a_4 = 1)/q(a_2 | a_4 = 0)$, defined for both $a_2 = 0$ and $a_2 = 1$. These likelihood ratios provide the answer to the question of how much more likely a participant is to observe a_2 , if she is the type to exert high effort at date 4.

Analogously, the second primitive describes participants' beliefs about their future effort. At date 1, we elicit prior belief distributions p_1 , defined over the four possible future events $(a_2, a_4) \in \{0, 1\}^2$. At date 3, after observing participants' realized effort a_2 , we elicit their posterior beliefs p_3 about effort at date 4. The prior belief distribution captures both a participant's perceived probability of exerting high effort at a given date and the *anticipated informativeness* of date 2 effort, given by the subjective likelihood ratios $p_1(a_2 | a_4 = 1)/p_1(a_2 | a_4 = 0)$. Our experiment is purposefully simple and features only two work dates in order to allow for the meaningful elicitation of the complete prior belief distribution.

We find that subjects are present biased at both work dates. The fraction of subjects that commits to high effort at date 1 (date 3) is 16.4 (14.8) percentage points higher than the fraction of subjects that ends up exerting high effort on the spot. Moreover, effort choices at date 2 are highly informative about effort choices at date 4. Of those subjects that exert high effort at date 2, 73.3 percent exert high effort at date 4, whereas of those that exert low effort at date 2, only 9 percent exert high effort at date 4. The fact that subjects exhibit present bias and that their past behavior is highly informative about future behavior implies that there is both something to learn about and something to learn from.

As a first step in the analysis of learning we can compare beliefs about date 4 effort at date 1 (*prior*) and at date 3 (*posterior*). The average prior belief that effort will be high at date 4 is 61.4 percent, the average posterior belief is 53.7 percent, and the actual likelihood of high effort is 43.6 percent. Therefore, participants are initially naive and become less naive after having experienced the task.² However,

²Theoretically, apparent naiveté about self-control could be the result of sophisticated subjects stating high beliefs in an attempt to use the belief elicitation as a soft commitment to exert high effort. We test for this confound by, after beliefs have been elicited, randomly varying whether the belief elicitation is payoff-relevant. Payoff-relevance does not lead to higher effort, suggesting that belief statements are not used as commitment devices.

our goal is to ascertain not only whether subjects learn at all, but whether the average updating that we observe is quantitatively appropriate.

To this end, we use Bayes' rule to construct normative posteriors from participants' average date 1 prior $p_1(a_4)$ and the true informativeness of behavior $q(a_2 | a_4 = 1)/q(a_2 | a_4 = 0)$. We find that participants' elicited posteriors are *not* significantly different from these normative posteriors after both high and low effort at date 2. Therefore, we find no evidence for an inferential bias that would explain a persistent overestimation of future effort. Our analysis of learning asks whether participants are good intuitive statisticians in revising their predictions based on their past behavior, but it does not, by itself, pin down the exact variable that subjects learn about. An additional set of results suggests that subjects are learning about their present bias: Controlling for subjects' commitment preferences, beliefs and on-the-spot effort are much more correlated at the second work date, after learning took place, than at the first work date.

To investigate whether subjects expect to learn from their behavior, we construct anticipated posterior beliefs based on each subjects' unconditional prior $p_1(a_4)$ and anticipated informativeness of behavior $p_1(a_2 | a_4 = 1)/p_1(a_2 | a_4 = 0)$. We find that the anticipated posteriors are not very sensitive to date 2 effort choices. We call this bias the non-belief in the propensity to learn. The severity of the bias in our participants is reflected in the difference between the actual and anticipated improvement in predictions as we move from priors to posteriors. The former captures how much better a participant is able to predict her date 4 effort once she experienced date 2 and the latter her expectation of this improvement at date 1. We find that the actual improvement is 8 times higher than the anticipated improvement.

Finally, we introduce a treatment to investigate whether participants are able to transport what they learn from their behavior in one environment to a different environment. While all participants work on the sliders task at date 2, the treatment varies whether date 4 features the sliders or the counting zeroes task. We find that subjects learn equally well and underestimate their future learning to the same extent in the two conditions.

This paper makes three contributions. First, we find that naiveté is unlikely to be the result of either a fundamental bias in how people learn from their own behavior in simple settings or of an inability to transport what they learn to different environments. Taken together with theoretical work that shows that optimism about self-control should be self-limiting in dynamic environments (Ali, 2011; Hestermann and Le Yaouanq, 2018), our results therefore greatly diminish the space of possible explanations for persistent naiveté. Two remaining candidates are the

misattribution of failures in richer settings and the imperfect or biased recall of past behavior in settings that feature larger time lags.

Second, our paper is the first to not just study retrospective learning, but also the prospective learning implied by Bayes' rule. In doing so, we uncover a hitherto unexplored bias in participants' expectations: The non-belief in the propensity to learn. An agent subject to this bias will underinvest in experimentation. She may also fail to activate self-regulation mechanisms, emphasized by [Ainslie \(1975\)](#), whereby achieving self-control is facilitated by seeing each self-control choice as being correlated with many similar future choices.

Third, we propose a simple methodology for the analysis of learning from naturally occurring signals (e.g., past savings decisions, school grades, health or employment status, etc.). This may allow economists and psychologists to take the old question of whether people are good intuitive Bayesians from the laboratory to the field, where any such intuition would have more plausibly developed. In its simplest form, our methodology imposes only two requirements on the collection of data. First, it requires that the researcher partitions the outcome space into two events and measures their occurrence at two points in time. Second, she needs to measure probabilistic beliefs over all future events before and after the first event has been observed.³

Our paper connects two sizable literatures, one on time discounting and one on updating biases. We follow recent experiments on present bias in studying the intertemporal allocation of effort, a non-fungible carrier of utility for which the resulting utility flow can be dated precisely.⁴ Experiments by [Ariely and Wertenbroch \(2002\)](#), [Augenblick et al. \(2015\)](#), [Kaur et al. \(2015\)](#) and [Bisin and Hyndman \(2018\)](#) all find substantial present bias in the intertemporal allocation of consumption events and some demand for commitment, suggesting that participants are at least partially aware of their present bias.⁵ In the experiment most closely related to ours, [Augenblick and Rabin \(2018b\)](#) compare incentivized predictions and commitment choices with on-the-spot effort choices and estimate the parameters

³Our approach is similar in spirit to that of [Augenblick and Rabin \(2018a\)](#), who detect deviations from Bayesian inference by comparing belief movements with uncertainty reduction in people's beliefs. A key difference is our reliance on measuring the data generating process as well as participants' beliefs over it.

⁴Early studies in the literature on present bias elicit preferences over time-dated monetary payments (see [Frederick et al., 2002](#), for a survey), but several papers have argued that the monetary domain is not appropriate for the measurement of preferences over utility streams due to the fungibility of money ([Cubitt and Read, 2007](#); [Chabris et al., 2008](#); [Augenblick et al., 2015](#); [Cohen et al., 2016](#)).

⁵Relatedly, [Toussaert \(2018\)](#) documents that experimental subjects demand commitment, not only in order to implement their preferred decision, but also to avoid the experience of future temptation and the associated self-control costs.

of a quasi-hyperbolic discounting model (Laibson, 1997) to uncover near-complete naiveté. Possibly owing to our elicitation of probabilistic beliefs rather than point estimates, our measurements suggest that participants are only partially naive.

Our paper goes beyond the current literature on present bias in a number of ways. By collecting data about the actual and perceived correlation of intertemporal decisions, we are able to retrieve and analyze the true and perceived information structure underlying participants' behavior. By using probabilistic predictions and outcomes instead of point estimates, we are able to apply Bayes' rule, which does not constrain the evolution of subjective point estimates. Finally, our use of a non-parametric model allows for the detection and measurement of present bias and naiveté independently of functional form assumptions about preferences.

A series of experimental papers, starting with Phillips and Edwards (1966) and culminating in Eil and Rao (2011) and Möbius et al. (2014), analyze individuals' updating behavior compared to a Bayesian benchmark.⁶ These experiments tend to generate noisy feedback about an objective variable (e.g., the performance in an IQ test) and ask participants to report both their prior and their posterior beliefs. Since the information structure that generates the signal is exogenously given and communicated to subjects, it is easy for the researcher to derive a Bayesian benchmark. At the same time, the updating task that subjects face is artificial and may draw on mathematical ability as much as it draws on intuition. In our setting, the signal structure is initially unknown and we need to infer it from the data. However, this allows us to ask whether people are good intuitive Bayesians when they learn from familiar, naturally-occurring information like their own behavior.

In a similar vein, our work relates to experiments by Eyster and Weizsäcker (2010), Enke and Zimmermann (2018) and Hossain and Okui (2018), who use artificial, exogenous information structures to document systematic misperceptions of correlation in participants' updating behavior and attitudes towards information sources.⁷ To our knowledge, our paper is the first to document an underestimation of the autocorrelation of intertemporal decisions and to draw out its implications. Our participants' non-belief in their propensity to learn is in line with results in Charness et al. (2018), who find that, while experimental subjects are able to learn from complicated information structures, this does not translate into them choosing

⁶Also see Ertac (2011); Buser et al. (2018); Schwardmann and van der Weele (2018); Coutts (2018); Zimmermann (2018) for updating in the domain of ego-relevant information, Gotthard-Real (2017); Coutts (2018); Barron (2016) for updating in the financial domain, and Benjamin (2018) for a recent review.

⁷See DeMarzo et al. (2003), Ortoleva and Snowberg (2015), and Levy and Razin (2015) for applications of correlation neglect and Spiegel (2016) for a general framework for analyzing misperceptions of causal or statistical relationships between decision-relevant variables.

the correct information structure.

The next section introduces the framework on which our analysis of learning is based. Section 3 describes our experimental design and Section 4 our results. Section 5 discusses the implications of our results and suggests avenues for future research.

2 Theoretical framework

Participants in the experiment work on real effort tasks at dates 2 and 4 and make predictions about their future effort choices at dates 1 and 3. Figure 1 depicts this timeline. We denote effort choices at date t by $a_t = 1$ for high and $a_t = 0$ for low effort. Then, $q(a_2, a_4)$ is the probability distribution over the four possible intertemporal events $(a_2, a_4) \in \{0, 1\}^2$, which we construct from the frequency of each combination of effort levels in the data. Participants' prior beliefs, at date 1, are given by the probability distributions $p_1^i(a_2, a_4)$ and their posterior beliefs, at date 3, by the probability distributions $p_3^i(a_4 | a_2^i)$, which are conditioned on realized effort a_2^i . The superscript i marks individual-level variables and is dropped when we refer to average quantities. From the primitive q we can construct the marginal distributions $q(a_t)$ and, whenever $q(a_2) > 0$, the conditional distributions $q(a_4 | a_2)$. We can also construct marginal and conditional belief distributions from the individual prior belief distributions p_1^i . For ease of exposition, we assume that q , p_1^i and p_3^i have full support.

Sophistication is the correct belief about (average) future effort, and naiveté the statistical overestimation of future effort. Naiveté can be detected at the sample level by comparing participants' average beliefs about future effort with the probability with which high effort is actually realized.

Our objective is to make normative statements about the learning of an initially naive population. We analyze separately the reactions to a good signal ($a_2 = 1$) and to a bad signal ($a_2 = 0$). We propose a benchmark that takes the population's average prior belief $p_1(a_4)$ as given and simulates the belief that the average participant would form conditional on a_2 following the Bayesian revision of her prior. To this end, we need to identify the actual information that the signal a_2 contains for predicting a_4 . We measure the *actual informativeness* of a_2 by the likelihood ratio $q(a_2 | a_4 = 1)/q(a_2 | a_4 = 0)$, which pins down the intensity at which a rational




Figure 1 – Timeline.

population equipped with the correct model would update following a_2 .⁸

Definition 1. Given the average prior beliefs p_1 , the *normative posterior beliefs* p_3^N , conditional on a_2 , are defined by

$$\underbrace{\frac{p_3^N(a_4 = 1 | a_2)}{p_3^N(a_4 = 0 | a_2)}}_{\text{normative posterior likelihood ratio}} = \underbrace{\frac{p_1(a_4 = 1)}{p_1(a_4 = 0)}}_{\text{prior likelihood ratio}} \underbrace{\frac{q(a_2 | a_4 = 1)}{q(a_2 | a_4 = 0)}}_{\text{actual informativeness}}.$$

For each $a_2 \in \{0, 1\}$, we compare the average elicited posterior beliefs $p_3(a_4 | a_2)$ in this subgroup with the normative posterior beliefs $p_3^N(a_4 | a_2)$ implied by the signal a_2 . This comparison enables us to decompose posterior naiveté into two sources: prior naiveté and erroneous learning (e.g., an underreaction to information). In particular, an initially naive population who updates in line with the normative benchmark remains naive ex post due to the inflated prior, but does not exhibit an updating bias.

Our dataset also allows us to analyze how participants forecast their future updating from the perspective of date 1. We measure the *anticipated informativeness* of the signal a_2 at the individual level by means of the subjective likelihood ratio $p_1^i(a_2 | a_4 = 1)/p_1^i(a_2 | a_4 = 0)$. Applying this measure to the prior beliefs pins down the posterior beliefs that the participant initially expects to form following the event a_2 .

Definition 2. Given a subject's prior beliefs p_1^i , the *anticipated posterior beliefs*

⁸By Bayes' rule, the correct model satisfies

$$\underbrace{\frac{q(a_4 = 1 | a_2)}{q(a_4 = 0 | a_2)}}_{\text{posterior likelihood ratio}} = \underbrace{\frac{q(a_4 = 1)}{q(a_4 = 0)}}_{\text{prior likelihood ratio}} \underbrace{\frac{q(a_2 | a_4 = 1)}{q(a_2 | a_4 = 0)}}_{\text{actual informativeness}}.$$

The likelihood ratio $q(a_2 | a_4 = 1)/q(a_2 | a_4 = 0)$ is a sufficient statistic for the informativeness of the signal a_2 , as it measures the rate at which the prior likelihood ratio must be multiplied in order to incorporate the information contained in a_2 . Our normative benchmark is based on a similar equation, where we replace the actual prior probabilities of high effort by the population's average prior beliefs in order to account for the population's initial naiveté.

$p_3^{i,A}$, conditional on a_2 , are defined by

$$\underbrace{\frac{p_3^{i,A}(a_4 = 1 | a_2)}{p_3^{i,A}(a_4 = 0 | a_2)}}_{\text{anticipated posterior likelihood ratio}} = \underbrace{\frac{p_1^i(a_4 = 1)}{p_1^i(a_4 = 0)}}_{\text{prior likelihood ratio}} \underbrace{\frac{p_1^i(a_2 | a_4 = 1)}{p_1^i(a_2 | a_4 = 0)}}_{\text{anticipated informativeness}}. \quad (1)$$

The anticipated posterior beliefs simulate the learning process of a participant who updates consistently with her forecast. For each effort level a_2 , we compare the average posterior beliefs $p_3(a_4 | a_2)$ with the average anticipated posterior beliefs $p_3^A(a_4 | a_2)$. This allows us to detect mispredictions of future updating (e.g., the underestimation of future learning).

A rational population that is equipped with the correct model q and uses Bayes' rule learns in such a way that, for each signal a_2 , the average elicited posterior $p_3(a_4 | a_2)$, the normative posterior $p_3^N(a_4 | a_2)$, the average anticipated posterior $p_3^A(a_4 | a_2)$, and the true conditional distribution of high effort $q(a_4 | a_2)$ all coincide. For a Bayesian population whose only misspecified variable is the prior belief, elicited posterior, normative posterior, and anticipated posterior coincide.

Our measures of actual and anticipated informativeness leverage the (actual and anticipated) statistical relationships between the variable subjects learn from (a_2) and the variable they learn about (a_4). As a result, our framework is agnostic about the underlying model that generates the effort decisions. Let us highlight two consequences of this fact.

First, our framework can be applied in any setting in which individuals learn from natural signals. For example, our conceptual apparatus can be used to analyze how students learn from their grade in one semester (a_2) about their likely grade in the next semester (a_4). In this paper, we use our methodology to detect biases in learning about one's effort, as mispredictions of future behavior have been shown to be particularly problematic (O'Donoghue and Rabin, 1999) and pronounced (Augenblick and Rabin, 2018b) in this domain.

Second, our framework does not require assumptions about why a_2 and a_4 are correlated. The standard foundation in learning models is that a_2 and a_4 both depend on a fundamental preference parameter that is initially unknown. But it could also be the case that a_2 and a_4 are affected by a common shock that occurs before date 2, or that a_2 has a causal effect on effort a_4 due to habit formation. Our measures of informativeness are independent of the exact foundation and our detection of biases does not require knowledge of the true data-generating process of effort choices. The flipside of this generality is that, by itself, our methodology cannot pin down what exactly it is that subjects are learning about. In section 4.4

we therefore supplement our analysis and show that participants in the experiment are in fact learning about their self-control.

Finally, we note that our analysis has to be conducted at the sample level. At the individual level, an issue arises if a subject receives private information between dates 1 and 3. For example, consider a subject who is affected by an adverse and durable productivity shock just before date 3. She would legitimately report a posterior belief about a_4 that is more pessimistic than her anticipated posterior and our test would falsely categorize her as an erroneous learner. However, rational learning from sources unknown to the analyst cannot explain a deviation from the Bayesian posterior in the aggregate because the law of iterated expectations precludes any systematic effect of private signals on average posterior beliefs. Therefore, our analysis focuses on average updating, just as naiveté and sophistication are statements made about average beliefs.⁹

3 Experimental design

The experiment was conducted in June 2018 and preregistered. Subjects were recruited via standard ORSEE procedures (Greiner, 2015) and participated in five experimental sessions across three weeks. The initial session took place at the Munich Experimental Laboratory for the Social Sciences and the remaining four sessions took place online. The four online sessions are labeled $t = 1, 2, 3, 4$ and featured either one of two real effort tasks (at dates 2 and 4) or the elicitation of predictions of and preferences over future effort choices (at dates 1 and 3).

At the initial session, we told subjects the dates of all future sessions and the content of the sessions at dates 2 and 4. We also provided information on payment rules and technical requirements for accessing the experimental website. Subjects then had to complete 5 practice screens of each of the real effort tasks before being introduced to the BDM mechanism used to elicit beliefs and answering some comprehension questions about the mechanism. After the first session in the lab, subjects received an email summarizing all relevant information.

⁹This argument requires the assumption that the private signals received by subjects in the course of the experiment are uncorrelated to each other. Assuming that unobserved events are uncorrelated across subjects is not only necessary for the study of learning, but also for the detection of naiveté, in our experiment as well as in existing ones. In particular, the statistical comparison of $p_1(a_2)$ and $q(a_2)$ implicitly assumes that the realization of the uncertainty is independent across subjects. In the experiment, we introduced exogenous variation in the precise days of dates 2 and 4 in order to reduce the likelihood of aggregate shocks like a particularly hot day that would reduce aggregate effort. Nevertheless, this solution is necessarily imperfect and the concern can only be put to rest by the replication of experimental findings in different contexts, as it is not possible to completely eliminate correlated uncertainty in any single study.


On the day of an online session a participant had a 24-hour window, from midnight to midnight, to log on to the experimental website and make her choices. Experimental instructions and decision environments can be found on a “tourist version” of the website we set up for the experiment (<https://www.lsc-experiment.com>).

Subjects who missed a session were excluded from the experiment and did not receive any payment, irrespective of their previous decisions and earnings. Subjects who completed all sessions received a participation fee of 25 Euro and a bonus payment determined by their effort choices in the real effort tasks and their earnings from the belief elicitation. On the day of an online session, we sent three reminder emails. All payment rules were made transparent at the beginning of the experiment and participants were reminded of them every time it was relevant for their decisions. Participants were paid via bank transfer a few days after the last session.

Timeline. Date 1 took place two days after the initial session. The two-day lag between the initial session in the lab and the first elicitation of beliefs is intended to eliminate the possible effect of projection bias (Loewenstein et al., 2003; Kauffmann, 2018), i.e. to avoid that subjects report a lower willingness to engage in the task in the future and pessimistic predictions because they very recently exerted a lot of effort and feel tired. Date 2 took place roughly one week after the initial session. We randomized the exact date at the individual level within a window of three consecutive days in order to reduce the incidence of correlated shocks on subjects’ effort cost (e.g., due to weather). A participant’s date 3 session took place exactly two days after her date 2 session. Date 4 took place in the following week and we again randomized the date between three possible dates, independently of date 2.¹⁰

Experimental tasks. Our real effort tasks are depicted in Figures 2a and 2b. In the *sliders* task subjects saw a screen with 40 slider bars, initially positioned on 50 and associated with a random target between 0 and 100. Completing one screen required positioning all 40 sliders on their target with the mouse or the keyboard. In the *counting zeroes* task subjects saw a screen with 10 matrices of 4 rows and 10 columns of ones and zeroes each. Completing one screen required counting the number of zeroes in each matrix and reporting this number in a text area. A screen submitted with incorrect entries had to be redone. On average, subjects completed one screen in 3 minutes and 26 seconds for the sliders task, and in 3 minutes and

¹⁰These dates were not consecutive because we excluded one day (June 27th, 2018) due to a scheduled football world cup game involving Germany which could have affected subjects’ aggregate effort.



(a) Sliders

Table 1

0	1	0	1	0	0	1	1	0	0
1	1	1	1	1	1	1	0	0	0
0	0	0	1	1	0	1	1	0	1
0	0	1	1	1	1	0	1	0	0

How many zeroes does Table 1 contain ?

(b) Counting zeroes

Figure 2 – Experimental tasks (translated from German).

54 seconds for the counting zeroes task.

We construct a binary variable a_2 (a_4) equal to 1 if the subject completed more than 20 screens at date 2 (4) and 0 otherwise. The threshold of 20 screens was calibrated to obtain balanced subsamples for whom $a_2 = 0$ and $a_2 = 1$ respectively, as we need to analyze learning separately for these two subgroups. The calibration was based on the results of a small pilot that featured only dates 1 and 2 of the experiment.

Participants were paid for every batch of 5 screens they completed successfully, according to the concave payment scheme in Table 1.

Treatments. We implemented two experimental conditions. In the Same Tasks condition, subjects worked on the sliders at dates 2 and 4. In the Different Tasks condition, they worked on sliders at date 2 and on the counting zeroes task at date 4.

Number of screens completed	Payment for this batch of 5 screens	Cumulative payment
5	5 Euros	5 Euros
10	4 Euros	9 Euros
15	3 Euros	12 Euros
20	2 Euros	14 Euros
25	1.5 Euros	15.5 Euros
30	1 Euro	16.5 Euros
35	0.5 Euros	17 Euros
40	0.1 Euros	17.1 Euros

Table 1 – Payment scheme

Date 1: Commitment choices and predictions. First, subjects committed to the number of screens (between 0 and 40) that they wanted to complete at future dates 2 and 4. We call this object the *commitment choice*. We randomly selected whether this decision was binding or not and communicated the outcome of the randomization immediately after subjects made their choice.

If the decision was binding (with probability 5%), then at both dates 2 and 4, subjects had to complete the exact number of screens they committed to in order to receive the bonus for that session. Completing more screens was not permitted. These committed subjects are excluded from the analysis, since we do not observe their effort choices under the same incentive scheme at different dates.

If the decision was not binding (with probability 95%), then subjects were free to choose how much to work on the task at the future date. For this group, we elicited predictions of future effort. We described the four possible future events (a_2, a_4) (called *scenarios*) in a table and asked subjects to report their subjective probability distribution p_1 (see Figure 3). We required subjects to state four non-negative integers that sum up to 100. The order of the events presented in the table was randomized at the individual level.

We incentivized predictions with a BDM mechanism (Karni, 2009) associated with a 3 Euro prize. We selected one of the joint events $(a_2, a_4) \in \{0, 1\}^2$ at random and paid the subject according to the BDM mechanism applied to the stated probability.¹¹

One potential concern about the belief elicitation mechanism is its possible lack of incentive-compatibility due to the endogeneity of effort decisions: Subjects could report a probability of high effort that is more confident than their true belief as a commitment strategy to work longer in future sessions. To solve this issue, the BDM mechanism was implemented with probability 50% and the uncertainty was

¹¹Suppose that the subject states a subjective probability equal to x percent for a given event. The BDM mechanism for this event selects a random integer y between 0 and 100 with uniform probability. If $x < y$, then the subject receives 3 Euro with probability y , otherwise the subject receives 3 Euro if and only if the event occurs.

	Session 2 (18.06.2018)		Session 4 (26.06.2018)	
	Sliders		Counting zeroes	Your prediction
Scenario 1	20 screens or more	AND	20 screens or more	<input type="text"/> %
Scenario 2	20 screens or more	AND	fewer than 20 screens	<input type="text"/> %
Scenario 3	Fewer than 20 screens	AND	20 screens or more	<input type="text"/> %
Scenario 4	Fewer than 20 screens	AND	fewer than 20 screens	<input type="text"/> %

Figure 3 – Elicitation of prior beliefs (translated from German).

realized and communicated immediately after the elicitation of beliefs. We test for the (successful) strategic use of the elicitation mechanism by measuring whether subjects whose beliefs were payoff-relevant exerted higher effort in the corresponding work session. As an additional step we asked subjects to state whether they reported their true beliefs or not in the post-experimental survey.

Date 2: Real effort task. All subjects were required to log on and then decide how long to work on the task, with no minimum effort requirement and a maximum of 40 screens. While this is not crucial for the study of naiveté and learning, we required all subjects to log on even if they did not wish to complete any screen in order to interpret their effort choice as an active decision. Otherwise, we would not be able to tell whether a subject who skipped a session had decided not to work on the task or simply forgot to log on. This would confound the identification of present bias in effort choice (Ericson, 2011, 2017). Each screen was numbered (from 1 to 40) and contained a “Finish” button that terminated the session. Upon submission of a correctly entered screen, the next screen was automatically called up.¹² Taking breaks was not allowed.

Subjects who were committed to the number of screens chosen at date 1 earned nothing before they reached their target and their session was automatically terminated if they reached the target. This information was provided to them on every screen. Subjects who were not committed saw the payment scheme as well as their

¹²Subjects faced a stopping problem, since the decision of terminating a session was made at every point in time. In contrast, Augenblick et al. (2015), Augenblick and Rabin (2018b) and Fedyk (2016) ask subjects to commit to their effort choice at the beginning of the session. These two paradigms are equivalent if time preferences inside the experimental session are dynamically consistent, but not if subjects prefer higher effort when they start working than when they have already completed some screens. We deemed the stopping problem attractive because it better captures most real-life behaviors and because it is arguably the more natural choice for participants to form expectations over.

accumulated bonus on every screen. Upon clicking on “Finish” a dialog box opened and the participant was required to confirm that she wanted to stop working and receive her current earnings.

Date 3: Commitment choices and predictions. Subjects that were not committed to their date 1 effort choice selected a number of screens that they committed to completing at date 4. The decision was binding with probability 5% and its consequence was identical to the commitment choice made at date 1. Subjects whose commitment choice was binding were excluded from the analysis at this stage. Our final sample thus consists only of participants for whom we observe commitment choices at dates 1 and 3 and unconstrained effort choices at dates 2 and 4. These subjects reported their posterior beliefs p_3 over the two events $a_4 \in \{0, 1\}$ (in a random order), in a table similar to that used at date 1. We incentivized the prediction with the BDM mechanism with a 3 Euro prize, and made it payoff-relevant with probability 50%.

Date 4: Real effort task. This session mirrors the session at date 2, except that subjects in the Different Tasks condition worked on the counting zeroes task instead of the sliders task. After they worked on the task, earnings were announced and we asked subjects to fill in a short post-experimental survey (e.g. about their gender, age, parents’ income and feedback about their decisions).

Sample size. In total, 201 subjects came to the initial session, where 11 decided not to take part in the experiment, e.g. because of scheduling conflicts with their randomly assigned sequence of dates for the online sessions. Only 3 subjects started the experiment but subsequently missed an online session. Therefore, attrition was minor. Our sample thus includes 187 participants. 19 of these subjects were committed to their effort decisions made in advance (either at date 1 or 3) and are thus excluded from the analysis. Our final sample consists of 168 individuals, 88 in the Same Tasks condition and 80 in the Different Tasks condition.

4 Results

In what follows we provide evidence for present bias and naiveté about present bias and then study how subjects learn and anticipate to learn from their behavior. Throughout, we pool data from our two conditions, before analyzing treatment differences in section 4.6.

4.1 Present bias and naiveté

Figure 4 compares commitment choice, beliefs about future effort, and effort relative to the 20 screen threshold. Since beliefs are probabilistic over the binary state of the world that a subject completes 20 screens or more, we make commitment and effort choices comparable by coding them as equal to 100 percent if the threshold is passed (high effort) and 0 otherwise (low effort). In Appendix A we also document present bias using the exact number of screens.

At both dates, more participants commit to high effort than end up exerting high effort on the spot (two-sided t-test, dates 1-2: $p < 0.001$; dates 3-4: $p < 0.001$), confirming that participants are present biased. Participants' average beliefs are also significantly higher than their effort choice (two-sided t-test, dates 1-2: $p = 0.005$; dates 3-4: $p < 0.001$), but significantly lower than their commitment choice (two-sided t-test, dates 1-2: $p = 0.002$; dates 3-4: $p = 0.003$), revealing partial naiveté about self-control.

Result 1 (Present bias and naiveté). *Effort choices are time-inconsistent at both dates, and participants are partially naive about their future effort.*

We find that 90 out of 168 subjects complete 20 or more screens at date 2. This variation in effort choices assures that our analysis of updating, where we will need to separately look at participants who exerted high effort (good signal) and low effort (bad signal) at date 2, is well-powered. We also observe that aggregate effort decreases over time, with the proportion of participants exerting high effort being significantly lower at date 4. This vindicates the usefulness of a methodology that allows us to study learning without relying on the stability of aggregate behavior.

4.2 Do subjects learn from their behavior?

Is there scope for learning? While the unconditional probability of high effort at date 4 equals 43.5%, this probability goes up to 73.3% if we condition it on high effort at date 2, and goes down to 9.0% if we condition it on low effort at date 2. A Fischer's exact test confirms that $q(a_4 = 1 | a_2 = 1)$ is larger than $q(a_4 = 1 | a_2 = 0)$ ($p < 0.001$). Therefore, behavior is highly correlated across periods and a_2 is highly predictive of a_4 .

We measure the amount of information contained in a_2 by computing the likelihood ratios $LR^q(a_2) = q(a_2 | a_4 = 1)/q(a_2 | a_4 = 0)$ based on the full sample. We find that $LR^q(a_2 = 1) = 3.57$, and that $LR^q(a_2 = 0) = 0.12$. To give an element of comparison, consider a hypothetical information structure generated by an oracle




Figure 4 – Average commitment choice, beliefs about effort and effort choice.

who observes a_4 and sends a probabilistic, informative message to the participant in the past: The oracle speaks the truth with probability $x\%$ and lies with the complementary probability. In our dataset, the signal $a_2 = 1$ contains as much information as a message for which $x = 78.1\%$, while the signal $a_2 = 0$ contains as much information as a message for which $x = 89.3\%$.


Do subjects learn appropriately? As defined in Section 2, the normative posterior beliefs against which we compare the elicited posterior beliefs are constructed by updating the average prior belief of the population $p_1(a_4 = 1)$ (which equals 61.4%) at the rate prescribed by the likelihood ratio $LR^q(a_2)$, for each a_2 .¹³ The elicited posterior after high effort at $t = 2$ is equal to 83.6%, while the normative posterior is equal to 85.0% (two-sided t-test, $p = 0.615$). The elicited posterior after low effort at $t = 2$ is equal to 19.2%, while the normative posterior is equal to 16.9% (two-sided t-test, $p = 0.526$). Therefore, we can not reject the null hypothesis that the elicited posterior is equal to the benchmark provided by the normative posterior, and that posterior beliefs at date 3 result from the correct updating of naive prior beliefs.¹⁴

Result 2 (Actual learning). *Effort at date 2 is highly informative about effort at*

¹³That is,

$$p_3^N(a_4 = 1) = \frac{p_1(a_4 = 1)LR^q(a_2)}{p_1(a_4 = 1)LR^q(a_2) + p_1(a_4 = 0)}.$$

¹⁴With this construction we only test whether subjects learn well from the information inherent to their binary effort a_2 . However, it is plausible that the precise effort level contains additional information. For instance, completing 40 screens at date 2 is likely to be a stronger signal of a high effort at date 4 than completing only 20 screens. We discuss this point and test for Bayesian updating using a finer information structure in Appendix D.



(a) Beliefs and effort after *low* effort at date 2 (b) Beliefs and effort after *high* effort at date 2

Figure 5 – Priors, average posterior beliefs and actual effort at date 4 after both low (panel a) and high (panel b) effort at date 2.

date 4. Participants incorporate the appropriate amount of information into their posterior beliefs.

4.3 Do subjects anticipate their future learning?

We construct the anticipated posterior beliefs $p_3^{i,A}$ at the individual level as a function of prior beliefs p_1^i and effort a_2^i .¹⁵ We only do this for subjects for whom $p_1^i(a_2^i) > 0$, as Bayes' rule is silent about updating after a zero-probability event. For this section, we therefore exclude 16 subjects from the dataset, 11 who provided low effort and 5 who provided high effort. Their updating behavior is reported in Appendix E.

Figure 5 compares participants' average prior, anticipated posterior, elicited posterior, normative posterior and true likelihood of exerting high effort at date 4.¹⁶ Panel 5a depicts the case of a bad signal, i.e. low effort, at date 2. Panel 5b depicts the case of a good signal, i.e. high effort, at date 2. Participants who exert high effort at date 2 have substantially higher priors, as reflected in the respective first bars in the panels.

For the case of a bad signal, we see that the anticipated posterior reflects less learning than both the elicited and the normative posterior. On the other hand,

¹⁵That is,

$$p_3^{i,A}(a_4 = 1) = \frac{p_1^i(a_2^i, a_4 = 1)}{p_1^i(a_2^i)}.$$

¹⁶For the sake of comparability, we compute these measures based on the sample consisting only of the subjects for whom the anticipated posterior beliefs are defined. This explains why some variables have different values than the ones reported in section 4.2.

<i>Low effort at date 2 (N = 67)</i>					<i>High effort at date 2 (N = 85)</i>						
Variable 1	Variable 2	Diff.	p-value		Variable 1	Variable 2	Diff.	p-value			
Prior	37.3	Ant. post.	29.6	7.7	0.004	Prior	79.4	Ant. post.	80.6	-1.2	0.113
		Elic. post.	17.8	19.5	< 0.001			Elic. post.	83.8	-4.5	0.155
		Nor. post.	13.1	24.2	< 0.001			Nor. post.	85.1	-5.7	0.033
		Effort	7.4	29.9	< 0.001			Effort	75.3	4.1	0.122
Ant. post.	29.6	Elic. post.	17.8	11.8	0.002	Ant. post.	80.6	Elic. post.	83.8	-3.2	0.293
		Nor. post.	13.1	16.5	< 0.001			Nor. post.	85.1	-4.5	0.095
		Effort	7.4	22.2	< 0.001			Effort	75.3	5.3	0.048
Elic. post.	17.8	Nor. post.	13.1	4.7	0.173	Elic. post.	83.8	Nor. post.	85.1	-1.2	0.658
		Effort	7.4	10.3	0.003			Effort	75.3	8.6	0.002

Table 2 – Pairwise comparisons of different posteriors with the normative posterior and the true probability of high effort at date 4. P-values derive from a two-sided t-test under the null hypothesis that the difference between the two variables is equal to zero.

elicited posteriors are not significantly different from the normative posterior. All posteriors lie above the true likelihood of high effort. Therefore, subjects learn more than they anticipated and arrive at posteriors that are naive but consistent with the Bayesian benchmark.

In the case of a good signal, we see that the elicited posterior is consistent with the anticipated posterior and the normative posterior. Owing to the inflated priors, all posteriors once again are significantly higher than the true likelihood of high effort.

The statistical relationships mentioned above are summarized and made precise in Table 2. The main message emerging from this table is that subjects update slightly less than they should following both low and high effort, but not significantly so. However, after a failure they update substantially more than they expect initially. Subjects learn rather well once they are confronted with the information inherent in their effort choice, but they underappreciate this fact ex ante.¹⁷

Result 3 (Non-belief in the propensity to learn). *At date 1, participants underestimate their future learning.*

¹⁷This result cannot be driven by a taste for consistency (Falk and Zimmermann, 2016), as subjects who want to make consistent reports would state posterior beliefs that are aligned with their anticipated posterior beliefs.

4.4 Do subjects learn about their normative preferences or their present bias?

In our analysis of learning thus far we have not attempted to disentangle whether subjects learn about their normative preferences (expressed in advance) or about their self-control, i.e. their ability to complete the number of tasks that they judge desirable *ex ante*. The literature has generally interpreted naiveté as pertaining to the degree of present bias, but it is at least possible that individuals also learn about their normative preferences in situations where they are not deeply familiar with all aspects of the decision-making environment.

Table 3 presents linear regressions of participants' beliefs on both their commitment choice and their effort choice. The regressions indicate whether and by how much subjects' beliefs are reflective of their normative preferences and their ultimate effort choice. The first column features beliefs and commitment at date 1 and effort at date 2, whereas the second column features beliefs and commitment at date 3 and effort at date 4. We see that the coefficient of the commitment choice stays constant as subjects move from date 1 to date 3 and learn. On the other hand, the coefficient of effort increases in size and significance.

Column 3 features a regression that pools data across dates and adds a time dummy and its interactions with the previous explanatory variables. This regression confirms that effort, but not commitment choice, becomes more predictive of beliefs after learning takes place. Therefore, our results suggest that, after the learning opportunity, participants' beliefs reflect a greater awareness of their ultimate effort choice conditional on the normative preference inherent in their commitment choice. This, in turn, is indicative of learning about present bias.

4.5 Anticipated and actual improvement in predictions

Next we analyze anticipated and actual learning from a perspective that explicitly takes the quality of individual predictions into account. We measure the mistake inherent in a belief-effort pair $(p(a_4 = 1), a_4)$ by the absolute distance between prediction and realized effort $|p(a_4 = 1) - a_4|$.

Figure 6a compares the mistakes implied by our different elicited and hypothetical beliefs. Comparing the mistakes implied by participants' priors with the mistakes implied by their elicited posteriors, we see that subjects' predictions become much more accurate between date 1 and 3 (two-sided t-test, $p < 0.001$). We also see that the average mistake implied by the anticipated posterior is only a little smaller than the mistake implied by the prior (two-sided t-test, $p = 0.003$), but

Dep. Variable	(1) Belief ($t = 1$)	(2) Belief ($t = 3$)	(3) Belief ($t = 1$ & $t = 3$)
Commitment choice	0.637*** (0.0428)	0.660*** (0.0346)	0.637*** (0.0428)
Effort	0.0671* (0.0361)	0.198*** (0.0344)	0.0671* (0.0361)
$t = 3$ (d)			-8.911** (3.873)
$t = 3$ * Commitment			0.022 0.061
$t = 3$ * Effort			0.131** 0.056
Constant	15.53*** (3.464)	6.62*** (1.910)	15.53*** (3.464)
Observations	168	168	168
R^2	0.689	0.865	0.793

Table 3 – OLS regressions of the determinants of beliefs with robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

much larger than the mistake implied by the elicited posterior (two-sided t-test, $p < 0.001$).¹⁸

A consequence of the non-belief in the propensity to learn is that participants would be less willing to experiment than they should be. In particular, an individual who does not believe in the informativeness of her own behavior will deem it less worthwhile to first engage in a task before abandoning it or committing to it.

To make these ideas more precise, we calculate participants' anticipated improvement in their predictions as they move from their prior to their posterior. We measure this improvement as the increase in the probability of winning the prize of the BDM belief elicitation from being paid for a date 3 rather than an unconditional date 1 belief.

Figure 6b compares by how much participants anticipate their predictions to improve, and by how much their predictions actually improve. We see that the actual value of information is more than 8 times higher than the improvement implied by participants' perceived informativeness of effort at date 2 (two-sided t-test, $p < 0.001$).

Result 4. *Subjects initially underestimate the value of information.*

¹⁸ It could be the case that participants receive some information about a_4 that is uncorrelated with a_2 . Because this information is not reflected in the anticipated posterior beliefs, the comparison of the mistakes implied by the elicited posteriors and anticipated posteriors would be unfair. In Appendix F we provide a way to eliminate the information orthogonal to a_2 in the posterior beliefs, and we show that the conclusion that subjects' elicited posterior beliefs are better calibrated than their anticipated posterior beliefs is robust.




Figure 6 – Quality of predictions and improvement in predictions

4.6 Treatment comparison

To see whether a change in the environment is an obstacle to learning, we first compute the actual informativeness of a_2 in both conditions. In the Same Tasks condition we find that $LR^q(a_2 = 1) = 5.08$ and $LR^q(a_2 = 0) = 0.07$, while in the Different Tasks condition we find that $LR^q(a_2 = 1) = 2.80$ and $LR^q(a_2 = 0) = 0.13$. The fact that the likelihood ratios are closer to one in the Different Tasks condition implies that a_2 is less predictive of a_4 than in the Same Tasks condition. This means that our treatment manipulation was successful at increasing the “distance” between the decision problems. We compute the normative posterior beliefs separately for each treatment group based on the likelihood ratios computed in this condition.

Table 4 contains the treatment comparison of key variables.¹⁹ There are no significant treatment differences in effort choice at date 4, the anticipated posterior, or the elicited posterior. We construct two measures that allow us to compare, from a normative viewpoint, the learning in both conditions without dividing each subsample again into the two subgroups $a_2 = 0$ and $a_2 = 1$. The first measure is the difference between the elicited posterior and the normative posterior, equal to $p_3^i(a_4 = 1) - p_3^N(a_4 = 1 | a_2^i)$. The second is a measure of under-reaction to information, equal to $p_3^i(a_4 = 1) - p_3^N(a_4 = 1 | a_2 = 0)$ for the subgroup $a_2^i = 0$ and equal to $p_3^N(a_4 = 1 | a_2 = 1) - p_3^i(a_4 = 1)$ for the subgroup $a_2^i = 1$. It measures by how much subjects under-update relative to the normative benchmark (in both directions). We find no significant difference in either of these measures between

¹⁹Table 8 in Appendix G shows that treatment groups are balanced according to gender, age, mathematical ability, and date 2 effort.

Variable	Same Tasks ($N = 79$)	Different Tasks ($N = 73$)	Difference	p-value
Effort (date 4)	45.6	45.2	0.4	>0.999
Prior	63.2	58.3	4.9	0.404
Elicited posterior	56.7	52.6	4.1	0.550
Normative posterior ($a_2 = 1$)	89.7	79.6	10.1	
Normative posterior ($a_2 = 0$)	10.5	15.8	-5.3	
Elic. post. - nor. post.	4.1	-0.8	4.9	0.256
Underreaction to information	5.12	-0.13	5.25	0.223
Anticipated posterior	59.1	57.1	2.0	0.742
Elic. post. - ant. post.	-2.4	-4.5	2.1	0.669
Underestimation of learning	5.78	5.85	-0.06	0.990

Table 4 – Treatment comparisons of key variables. For the comparison of effort choice the p-value is based on Fischer’s exact test, for all other comparisons a two-sided t-test was used.

the two treatment groups, which suggests that subjects learn equally well in both conditions.

To assess whether the underestimation of future learning is more severe in one condition, we construct two measures again. The first is the difference between the elicited posterior and the anticipated posterior, equal to $p_3^i(a_4 = 1) - p_3^{i,A}(a_4 = 1 | a_2^i)$. The second expression is a measure of the underestimation of future learning, equal to $p_3^i(a_4 = 1) - p_3^{i,A}(a_4 = 1 | a_2 = 0)$ for the subgroup $a_2^i = 0$ and equal to $p_3^{i,A}(a_4 = 1 | a_2 = 1) - p_3^i(a_4)$ for the subgroup $a_2^i = 1$. We find no significant difference in either of these measures between the two experimental conditions, which suggests that the underestimation of future updating affected participants in both conditions equally. We therefore conclude that the change in the environment that we implemented had no effect on subjects’ learning.

Result 5 (Treatment effect). *Subjects learn equally well and underestimate their learning to a similar degree in both conditions.*

4.7 Are belief elicitation used as a soft commitment?

It is possible that a sophisticated individual could use the belief elicitation as soft commitment devices. By stating a high belief, a subject makes it more expensive for her future self to only work a little, thereby incentivizing her to exert more effort. If belief elicitation are used as commitment, participants no longer state their true beliefs and what looks like naiveté may in fact be a sophisticated commitment strategy. [Augenblick and Rabin \(2018b\)](#) and [Fedyk \(2016\)](#) test for the strategic use of belief elicitation by varying the incentives of the elicitation. They

	(1)	(2)	(3)	(4)
Dep. Variable:	Effort (date 2)	Effort (date 2)	Effort (date 4)	Effort (date 4)
Belief paid	1.702 (1.675)	1.165 (1.431)	-0.259 (1.807)	-0.227 (1.186)
Commitment choice		0.602*** (0.0782)		0.858*** (0.0681)
Constant	18.68*** (1.142)	4.534*** (1.690)	17.18*** (1.415)	-0.362 (1.464)
Observations	168	168	168	168
R^2	0.006	0.282	0.000	0.542

Table 5 – OLS regressions of the determinants of effort choice with robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

argue that if elicitation were used as a soft commitment device, then stated beliefs should increase in the incentives. Both papers find no such effect.

We implement a different test of this potential confound. Before every belief elicitation, participants in our experiment were told that their beliefs would be payoff-relevant with probability 50%. After they stated their beliefs, they were then told whether they were randomly selected to be in the group for whom this belief elicitation was payoff-relevant. We can then compare two groups who are identical in terms of their beliefs, but whose beliefs differ in whether they constitute a monetary commitment to high effort. Then, if participants did in fact use their beliefs as a soft commitment, effort should be higher in the group whose beliefs were payoff-relevant.

Table 5 shows OLS regressions of effort choice on the payoff relevance of stated beliefs. The first column shows that effort at date 2 was not significantly higher if beliefs were paid. Column 2 confirms this, while controlling for participants' commitment preferences. Columns 3 and 4 focus on effort at date 4. The coefficients of the belief paid variable are now slightly negative and insignificant. If anything, we would have expected that the sophisticated deployment of the belief elicitation as a soft commitment would have increased after subjects had a chance to learn. As an additional step, we asked subjects in the post-experimental survey how their reported beliefs compared to their best prediction.²⁰ A vast majority (91.4%) indicated that they tried to report their best estimate, while only 6 subjects (3.6%) stated that they reported a larger likelihood of high effort than the one they had in mind. Among these 6 subjects only one justified this decision by the desire to

²⁰We asked the following question (translated from German): “In the past weeks we asked you several times to report the likelihood with which you believed you would complete more than a given number of screens in the future. What did you think when you answered these questions? 1. I tried to report my best estimate. 2. I did not think much of it and reported whatever answer came to my mind. 3. I reported a lower likelihood than the one I had in mind. 4. I reported a larger likelihood than the one I had in mind.”

affect future effort incentives. Our results therefore suggest that belief elicitation were not used as a soft commitment device.

5 Discussion

Naiveté about self-control is one of the best documented phenomena in behavioral economics. Our paper studies its evolution and asks whether individuals are able to learn from their past behavior to make better predictions about their future self-control.

Explanations for persistent naiveté. We find no evidence of an inferential bias that hinders learning. Instead, we show that individuals learn well and are able to transport their acquired self-knowledge from one environment to another environment. Therefore, our data does not resolve the puzzle of persistent naiveté. Neither is the puzzle resolved in theoretical work by [Ali \(2011\)](#) and [Hestermann and Le Yaouanq \(2018\)](#), who find that naiveté about self-control should be self-limiting if agents are Bayesian. A naive individual fails to commit to future effort and thereby exposes herself to the opportunity to learn about her self-control problem ([Ali, 2011](#)). Moreover, a naive individual may partially attribute her failure to exert effort to her current environment and therefore be compelled to change her environment which in turn enables her to see that her lack of self-control was the culprit all along ([Hestermann and Le Yaouanq, 2018](#)).²¹

Taken together with these papers, our findings greatly diminish the space of possible explanations for persistent naiveté. Future experiments could enrich the decision-making environment and thereby increase the scope for participants' misattribution of their failures to external factors.²² It is also possible that individuals' learning is hampered by noisy or biased recall. The role of memory can easily be accommodated in our framework by increasing the time lag between dates.

In order to elicit high-quality predictions of future effort, we restricted our attention to one learning period. The learning we observe in this one period implies that beliefs should quickly approach the true empirical frequencies of high effort.

²¹These theories are based on the common assumption of incidental learning. [Christensen and Murooka \(2018\)](#) show that naive present-biased agents might procrastinate forever at learning if this requires an active and costly decision.

²²[Gagnon-Bartsch et al. \(2018\)](#) explain mislearning by arguing that individuals might fail to attend to important data (e.g., their own behavior in the past) if this data is not valuable according to their incorrect model of the world. However, our purposefully stripped-down experimental design might make attending to all relevant data sufficiently simple to mute this mechanism of mislearning.

An experiment that contains several learning periods, perhaps by giving up on the elicitation of the complete distribution of prior beliefs over joint events, could test this conjecture.²³ Such an experiment could also ascertain the extent to which people learn about their ability to learn, i.e. analyze the stability of the non-belief in the propensity to learn.

Implications of the non-belief in the propensity to learn. While learning turns out to be normatively appropriate *ex post*, we show that participants fail to anticipate this *ex ante*. This novel result raises the important question of whether individuals experiment inefficiently little in their lives. For instance, workers might fail to try out different work arrangements (e.g., how to organize their daily schedule) and miss out on the opportunity to learn what makes them most productive.

The non-belief in the propensity to learn might also play a role in erroneous decisions that the literature has commonly attributed to naiveté only. Consider evidence in [DellaVigna and Malmendier \(2006\)](#) for the excess demand for annual gym memberships, relative to individuals' actual attendance. Naiveté does not by itself predict the take-up of an incorrect membership. Individuals who overestimate their future attendance but believe in their propensity to learn might indeed find it optimal to experiment with a pay-per-visit scheme before committing to a long-term contract. In this and other settings, the welfare loss due to insufficient experimentation can be measured by contrasting individuals' willingness to pay for a trial period before they make a long-term decision with the actual welfare improvement that would result from this experimentation.

Another important consequence of the non-belief in the propensity to learn might be the failure to activate self-regulation mechanisms that help individuals overcome self-control problems. [Ainslie \(1975\)](#) argues that some of people's most important means of self-control are internal. In choosing between a smaller sooner reward (being lazy, eating a fatty food, getting angry etc.) and a larger later reward (reaping the pecuniary benefits of our effort, being in good shape, and having good relationships etc.), our impulsivity might drive us to choose the smaller sooner reward. But if we are able to see our choice as a first in a long sequence of similar choices between smaller sooner rewards and larger later rewards and if we have the conviction that our making the impulsive choice contains a cue that we will end up

²³For a horizon with n work dates, the complete prior distribution is defined over a 2^{n-1} -dimensional simplex, and thus the complexity of the elicitation procedure rises quickly with the number of periods. Under the assumption that the marginal belief distributions are time-invariant, which was approximately true in our experiment, one could ask subjects to form beliefs about the number of future periods at which they will exert high effort. This method requires the elicitation of a prior probability distribution defined over a n -dimensional simplex.

choosing the smaller reward again in the future, then this bundling of rewards may tip the scale in favor of avoiding the impulsive choice. [Bénabou and Tirole \(2004\)](#) provide a theory for why our lack of self-control today should affect our prediction of similar choices in the future. They argue that people have imperfect recall when it comes to their deep preferences and have to infer them from their past behavior. For this reason, failing to exert self-control has a negative effect on our self-image and our attempts at future self-regulation.

Our results cannot speak to the causal effect of avoiding the present biased choice today on effort in the next period, nor do we measure whether participants' subjective model reflects such causality. However, a direct implication of the bundling of rewards à la [Ainslie \(1975\)](#) is that people should view their behavior in this period and the next as highly correlated. Put differently, a belief in the informativeness of behavior is a necessary condition for achieving self-control by bundling rewards. The near-complete absence of this belief in our subjects therefore indicates that this particular self-regulation mechanism will be difficult for people to deploy. Future work could test the causal effect of a non-belief in the propensity to learn on first period present bias by exogenously shifting beliefs about the autocorrelation of effort levels.

Applications for our methodology. Our technique for retrieving the perceived and actual information structure generated by naturalistic signals is well-suited to be applied in the field. For example, consider a longitudinal survey that elicits respondents' savings goals for the next two years, their subjective likelihood of meeting them before and after the first year, and whether goals were actually met. This dataset would allow a researcher to ask whether expectations were updated in a Bayesian fashion, which in turn would inform appropriate policy measures aimed at, for example, alleviating undersaving. A policy maker can trust a Bayesian population to shed biased beliefs about its saving behavior by itself and to eventually commit to save more. But if a population is found to underweigh the signal contained in its past savings decisions, then it may make sense to target it with information campaigns.

References

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* 82(4), 463.

- Ali, S. N. (2011). Learning self-control. *Quarterly Journal of Economics* 126(2), 857–893.
- Ariely, D. and K. Wertenbroch (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science* 13(3), 219–224.
- Augenblick, N., M. Niederle, and C. Sprenger (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics* 130(3), 1067–1115.
- Augenblick, N. and M. Rabin (2018a). Belief movement, uncertainty reduction, and rational updating. Working paper.
- Augenblick, N. and M. Rabin (2018b). An experiment on time preference and misprediction in unpleasant tasks. *Review of Economic Studies*. Forthcoming.
- Barron, K. (2016). Belief updating: Does the ‘good-news, bad-news’ asymmetry extend to purely financial domains? Working Paper.
- Bénabou, R. and J. Tirole (2004). Willpower and personal rules. *Journal of Political Economy* 112(4), 848–886.
- Benjamin, D. J. (2018). Errors in probabilistic reasoning and judgment biases. Forthcoming in the Handbook of Behavioral Economics.
- Bisin, A. and K. Hyndman (2018). Present-bias, procrastination and deadlines in a field experiment. Working paper.
- Buser, T., L. Gerhards, and J. van der Weele (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty* 56(2), 165–192.
- Chabris, C. F., D. Laibson, and J. P. Schuldt (2008). Intertemporal choice. In S. N. Durlauf and L. Blume (Eds.), *The New Palgrave Dictionary of Economics*. London: Palgrave Macmillan.
- Charness, G., R. Oprea, and S. Yuksel (2018). How do people choose between biased information sources? Evidence from a laboratory experiment. Working Paper.
- Christensen, E. G. B. and T. Murooka (2018). Procrastination and learning about self-control. Working paper.
- Cohen, J. D., K. M. Ericson, D. Laibson, and J. M. White (2016). Measuring time preferences. Working paper.

- Coutts, A. (2018). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*. Forthcoming.
- Cubitt, R. P. and D. Read (2007). Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics* 10(4), 369–389.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature* 47(2), 315–72.
- DellaVigna, S. and U. Malmendier (2006). Paying not to go to the gym. *American Economic Review* 96(3), 694–719.
- DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly journal of economics* 118(3), 909–968.
- Eil, D. and J. M. Rao (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics* 3(2), 114–38.
- Enke, B. and F. Zimmermann (2018). Correlation neglect in belief formation. *Review of Economic Studies*. Forthcoming.
- Ericson, K. M. (2017). On the interaction of memory and procrastination: Implications for reminders, deadlines, and empirical estimation. *Journal of the European Economic Association* 15(3), 692–719.
- Ericson, K. M. M. (2011). Forgetting we forget: Overconfidence and memory. *Journal of the European Economic Association* 9(1), 43–60.
- Ertac, S. (2011). Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback. *Journal of Economic Behavior & Organization* 80(3), 532–545.
- Eyster, E. and G. Weizsäcker (2010). Correlation neglect in financial decision-making. Working paper.
- Falk, A. and F. Zimmermann (2016). Consistency as a signal of skills. *Management Science* 63(7), 2197–2210.
- Fedyk, A. (2016). Asymmetric naivete: Beliefs about self-control. Working paper.

- Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature* 40(2), 351–401.
- Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2018). Channeled attention and stable errors. Working paper.
- Giné, X., D. Karlan, and J. Zinman (2010). Put your money where your butt is: A commitment contract for smoking cessation. *American Economic Journal: Applied Economics* 2(4), 213–35.
- Gotthard-Real, A. (2017). Desirability and information processing: An experimental study. *Economics Letters* 152, 96–99.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association* 1(1), 114–125.
- Hestermann, N. and Y. Le Yaouanq (2018). It's not my fault! Self-confidence and experimentation. Working paper.
- Hossain, T. and R. Okui (2018). Belief formation under signal correlation. Working paper.
- John, A. (2018). When commitment fails - evidence from a field experiment. *Management Science*. Forthcoming.
- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica* 77(2), 603–606.
- Kauffman, M. (2018). Projection bias in effort choice. Working paper.
- Kaur, S., M. Kremer, and S. Mullainathan (2015). Self-control at work. *Journal of Political Economy* 123(6), 1227–1277.
- Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature* 52(4), 1075–1118.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2), 443–478.
- Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review* 105(5), 267–72.

- Levy, G. and R. Razin (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review* 105(4), 1634–45.
- Loewenstein, G., T. O’Donoghue, and M. Rabin (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics* 118(4), 1209–1248.
- Möbius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2014). Managing self-confidence. Working paper.
- O’Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review* 89(1), 103–124.
- Ortoleva, P. and E. Snowberg (2015). Overconfidence in political behavior. *American Economic Review* 105(2), 504–35.
- Phillips, L. D. and W. Edwards (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology* 72(3), 346.
- Schwardmann, P. and J. van der Weele (2018). Deception and self-deception. Working paper.
- Spiegler, R. (2016). Bayesian networks and boundedly rational expectations. *Quarterly Journal of Economics* 131(3), 1243–1290.
- Toussaert, S. (2018). Eliciting temptation and self-control through menu choices: A lab experiment. *Econometrica* 86(3), 859–889.
- Zimmermann, F. (2018). The dynamics of motivated beliefs. Working paper.

Appendix

The appendix is organized as follows. Appendices [A](#) and [B](#) provide additional information on subjects' present bias and the stability of their behavior across time. Appendix [C](#) features histograms of prior, posterior and anticipated posterior beliefs. In Appendix [D](#), we provide a more stringent test of Bayesian updating that uses a finer information structure than the binary signal space considered in the paper. Appendix [E](#) reports the updating behavior of subjects who stated a misspecified prior at date 1. Appendix [F](#) establishes the robustness of results in section [4.5](#) and Appendix [G](#) demonstrates that treatment groups are balanced according to observable characteristics.

A Present bias

Here we document present bias using the number of screens (between 0 and 40) as our measure of commitment choices and effort. At date 1, participants commit to completing an average of 23.94 screens at date 2. But they complete only 19.53 screens, significantly fewer screens than they intended to (two-sided t-test, $p\text{-value} < 0.01$). We observe a similar time inconsistency between date 3 and date 4, where participants commit to an average of 20.44 screens in advance, but end up completing only 17.04 screens (two-sided t-test, $p\text{-value} < 0.01$). Figure [7](#) shows histograms of the present bias inherent in the effort choices. Our measure of present bias is obtained by subtracting the actual effort choice from the commitment choice. Out of our 168 participants, 79 exhibit at least some present bias between dates 1 and 2 and 75 between dates 3 and 4.

B Stability of behavior

The stability of behavior can be seen on the scatter plot of effort choice at date 2 and effort choice at date 4 depicted in Figure [8a](#). The correlation between the number of screens completed at the first work date and the number of screens completed at the second work date is 0.73. Figure [8b](#) exhibits the positive relationship between present bias at both dates. The correlation between the two variables is 0.34.

In the main text we show that effort a_2 is highly predictive of effort a_4 . It turns out that present bias between dates 1 and 2 is also predictive of present bias between dates 3 and 4. Indeed, coding present bias as a binary measure equal to 1 if the




Figure 7 – Distribution of present bias (commitment choice - effort).

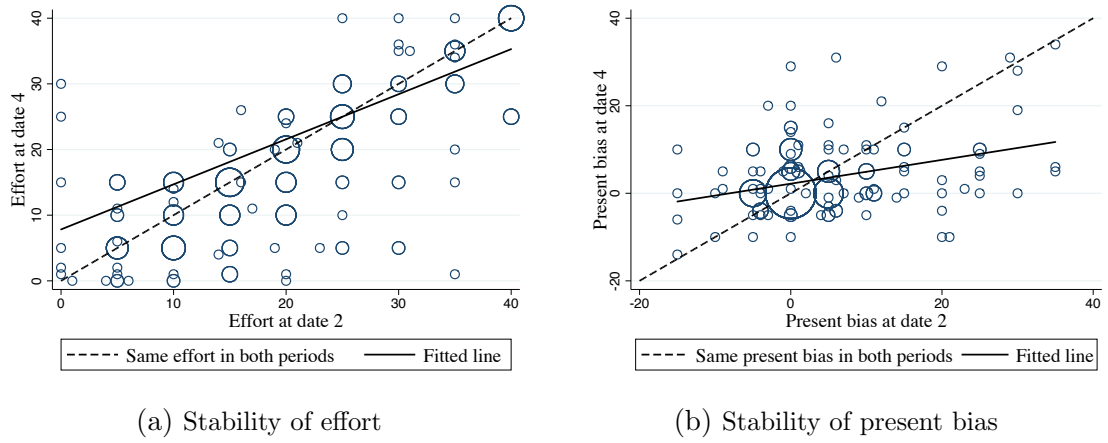


Figure 8 – Scatter plot: correlation between behavior at dates 2 and 4

commitment choice is strictly larger than the actual effort, we find that the fraction of subjects who are present-biased between dates 3 and 4 equals 44.6% in the full sample, while it goes up to 55.7% for subjects who exhibit present bias between dates 1 and 2, and down to 34.8% for subjects who are not present-biased between dates 1 and 2. A Fischer's exact test confirms that the probability of present bias between dates 3 and 4 differs between these two subgroups ($p = 0.008$).

C Distributions of beliefs




Figure 9 – Distributions of beliefs about effort at date 4. Belief distributions are described by the weight attached to $a_4 = 1$.

D A more stringent test of Bayesian learning

The construction of the normative posterior beliefs outlined in Section 2 can be generalized to a finer information structure than the one provided by the binary effort level a_2 . That is, for any observable individual-level event E realized at date 2, we can compute the actual likelihood ratio $q(E | a_4 = 1)/q(E | a_4 = 0)$ and construct the associated normative posterior belief conditional on E . For instance, we could in principle condition posterior beliefs on the exact number of screens completed by the participant at date 2. This would then allow us to retrieve the actual informational content of the precise effort level, which might contain more information about a_4 than the binary effort a_2 .

In practice, constructing normative posterior beliefs for the finest possible partition of the signal space would lead to underpowered statistical tests. We therefore partition date 2-effort into four categories only: Very low effort, moderately low effort, moderately high effort, and very high effort. From the empirical frequencies of a_4 in each of these four categories we construct the likelihood ratio of the date 2 effort and combine it with the average prior belief of the population to obtain the normative posterior beliefs. We can thus detect whether the near-Bayesian updating after $a_2 = 0$ uncovered in section 4.2 masks opposing anomalies in the reaction to a very low effort and moderately low effort.

For each effort level, Table 6 shows the corresponding likelihood ratio, the conditional frequency of date 4-effort, and compares the elicited posterior beliefs with the normative posterior. Three messages emerge. First, there is indeed more information in this finer partition than in the binary one used in section 4.2. In particular, a very high effort is a much stronger signal of subsequent high effort than a moderately high effort. Second, subjects learn equally well from a moderately high effort and from a very high effort. Third, the insignificant underreaction to a bad signal ($a_2 = 0$) uncovered in section 4.2 is driven by insufficient learning after very low effort. Completing fewer than 10 screens is a very negative signal of one's propensity to exert a high effort at date 4. But following this event, subjects report posterior beliefs which are more optimistic than the normative benchmark. This difference seems quantitatively important, but is not significant, which might be owed to the small sub-sample that this analysis is based on.

E Updating behavior after a zero-probability event

Here we report the updating behavior of the 16 subjects who are excluded from our analysis of learning as their prior beliefs p_1^i assign a probability zero to their

Date 2-effort (Number of screens)	Very low (0-9)	Moderately low (10-19)	Moderately high (20-29)	Very high (30-40)
Actual informativeness	0.11	0.14	2.31	7.37
Elicited posterior	20.3	18.6	75.2	94.2
Normative posterior	14.7	18.0	78.6	92.1
Frequency of $a_4 = 1$	7.7	9.6	64	85
Elic. post. - nor. post.	5.6	0.6	-3.4	2.1
p-value	0.389	0.892	0.443	0.324
N	26	52	50	40

Table 6 – Comparison of elicited and normative posterior with a 4-element signal space. P-values derive from a two-sided t-test under the null hypothesis that the difference between elicited and normative posterior is equal to zero.

actual effort a_2^i . 11 subjects exerted a low effort ($a_2^i = 0$) in spite of initially reporting prior beliefs that satisfy $p_1^i(a_2 = 1) = 100\%$. Their average prior beliefs $p_1(a_4 = 1)$ equal 96.4%, while their average posterior beliefs go down to 27.4%. 5 subjects exerted a high effort ($a_2^i = 1$) in spite of initially reporting $p_1(a_2 = 1) = 0\%$. Their average prior beliefs $p_1(a_4 = 1)$ equal 0%, while their average posterior beliefs equal 80%. In total, 11 of these 16 subjects reported posterior beliefs that put probability one on an effort a_4 equal to their realized effort a_2^i , while 4 subjects stated a posterior belief equal to their prior.

F Quality of predictions

Here we discuss further the issue mentioned in Footnote 18. Consider a rational agent with well-calibrated prior beliefs $p_1(a_2, a_4) = 0.25$ for all (a_2, a_4) , and who receives perfect information about a_4 just before date 3. The elicited posterior beliefs of this agent are more precise than her anticipated posterior beliefs, even though this agent has rational expectations. The issue is that the information received before date 3 is uncorrelated with a_2 and can therefore not be reflected in the prior beliefs.

To deal with this issue and provide a better comparison of the mistakes implied by the anticipated and actual learning, we eliminate any information orthogonal to a_2 in the elicited posterior beliefs. That is, we calibrate two hypothetical likelihood ratios, $LR^h(a_2 = 1)$ and $LR^h(a_2 = 0)$, such that, if all subjects in the subgroup a_2 updated their prior beliefs with the (common) likelihood ratio $LR^h(a_2)$, their resulting posterior beliefs would be equal on average to their elicited posterior beliefs. We then replace the elicited posterior belief of each subject by the posterior belief constructed from the prior of this subject and the calibrated likelihood ratio, which we call the *common LR posterior*, and we compute the associated mistake

at the individual level. This procedure eliminates the information orthogonal to a_2 by assuming that the rate of updating is the same for all subjects who exerted the same effort level a_2 .

Table 7 shows the average mistake implied by the different beliefs. The findings confirm that subjects' learning from a_2 itself is better than their anticipated learning from a_2 , as the mistake implied by the common LR posterior is smaller than the mistake implied by the anticipated posterior. These results confirm that participants learn well from their experience *ex post* but that their prior beliefs underestimate the magnitude of the information contained in a_2 .

Mistake 1		Mistake 2		Difference	p-value
Prior	31.7	Ant. post.	27.9	3.8	0.003
		Common LR. post.	23.6	8.1	<0.001
		Elic. post.	18.4	13.3	<0.001
Ant. post.	27.9	Common LR. post.	23.6	4.2	0.003
		Elic. post.	18.4	9.5	<0.001
Common LR. post.	23.6	Elic. post.	18.4	5.2	0.02

Table 7 – Mistakes implied by the different beliefs. For all comparisons a two-sided t-test was used.

G Baseline balance across treatment groups

Variable	Same Tasks ($N = 79$)	Different Tasks ($N = 73$)	Difference	p-value
Binary effort (date 2)	53.1	58.9	-5.7	0.516
Continuous effort (date 2)	19.7	20.9	-1.2	0.490
Female	64.6	61.6	2.9	0.738
Age	24.7	24.0	0.7	0.492
Mathematics score	4.2	3.9	0.3	0.656

Table 8 – Baseline balance across treatment groups. For the comparison of binary effort choice and gender (a binary variable equal to 1 for a female participant) the p-value is based on Fischer's exact test, for all other comparisons a two-sided t-test was used. The mathematics score was measured by the average mathematics grade in high school.