Wissenschaftszentrum Berlin
für Sozialforschung

Kai Barron
Christina Gravert

Confidence and career choices: An experiment

Discussion Paper

SP II 2018–301r2

January 2018 (2nd revision June 2020)

Research Area
Markets and Choice

Research Unit
Economics of Change

Wissenschaftszentrum Berlin für Sozialforschung gGmbH
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

Affiliation of the authors:

Kai Barron, WZB (kai.barron@wzb.eu)

Christina Gravert, University of Copenhagen (cag@econ.ku.dk)

Abstract

Confidence and career choices: An experiment[*]

Confidence in one's own abilities is often seen as an important determinant of being successful. Empirical evidence about how such beliefs about one's own abilities causally influence choices is, however, sparse. In this paper, we use a stylized laboratory experiment to investigate the causal effect of an increase in confidence on two important choices made by workers in the labor market: (i) choosing between jobs with a payment scheme that depends heavily on ability [high earnings risk] and those that pay a fixed wage [low earnings risk], and (ii) the subsequent choice of how much effort to exert within the job. We find that an exogenous increase in confidence leads to an increase in subjects' propensity to choose payment schemes that depend heavily on ability. This is detrimental for low ability workers due to high baseline levels of confidence.

Keywords: Overconfidence, experiment, beliefs, real-effort, career choices

JEL classification: C91, D03, M50, J24

---

[*]We would like to thank Teodora Boneva, Antonio Cabrales, Rustamdjan Hakimov, Zahra Murad, Heiner Schuhmacher, Robert Stüber, Georg Weizsäcker and participants at THEEM ESA European Meeting and Nordic Conference in Behavioral and Experimental Economics for helpful comments and suggestions. This research was funded by the Jan Wallanders and Tom Hedelius Foundation and the Tore Browaldhs foundation, grant number P2016-0051:1. Kai Barron gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft (through CRC TRR 190). Declarations of interest: none. All mistakes are our own.

1 Introduction

Confidence in one's own abilities is commonly thought to be an important determinant of success. A large body of work has studied the channels through which holding overconfident beliefs may be beneficial. For example, Bénabou and Tirole (2002, 2003) provide a discussion of how a higher level of self-confidence can motivate individuals to work harder, overcome obstacles, and take beneficial risks. Brunnermeier and Parker (2005) argue that individuals may hold upward-biased beliefs in order to enjoy the consumption value of a rose-tinted view of the future (and present and past), while Köszegi (2006) studies the behavior of individuals who derive ego-utility from overconfident beliefs. In complement to these theoretical discussions, a growing empirical literature has studied various mechanisms through which confidence can contribute to an individual's success. This literature presents evidence suggesting that more confident individuals are better at *persuading others* that they are of high ability (Burks et al., 2013; Schwardmann and Van der Weele, 2019; Solda et al., 2019), *work harder* (Puri and Robinson, 2007; Pikulina et al., 2017; Chen and Schildberg-Hörisch, 2019), and that overconfidence is an *adaptive evolutionary trait* (Bernardo and Welch, 2001; Heifetz et al., 2007; Johnson and Fowler, 2011).

This view is in stark contrast with the standard Bayesian rational agent perspective – for a rational Bayesian, more accurate beliefs are always better. Less accurate beliefs lead to more mistakes, which results in a loss of utility. A substantial body of theoretical and empirical evidence has documented some of the costly mistakes that overconfident individuals may make, such as exposing themselves to excessive risk (relative to their risk preferences) in financial markets (Odean, 1998; Barber and Odean, 2001), poor managerial decisions (Malmendier and Tate, 2005), or over-entry into competition (Camerer and Lovallo, 1999; Niederle and Vesterlund, 2007; Dohmen and Falk, 2011).

Despite the wealth of evidence documenting the widespread existence of overconfidence (see, e.g., Moore and Healy, 2008) and the prominence of overconfidence as one of the most commonly studied behavioral biases (see, e.g., Bénabou and Tirole, 2016), causal evidence relating an exogenous shift in overconfidence to a change in behaviour is relatively sparse.[1] Yet, without causal evidence of the effect of confidence on decisions in a particular domain, one can neither claim that increasing someone's confidence is beneficial, nor that it is detrimental. Any observed correlation between individual heterogeneity in overconfidence and behavior or outcomes could simply be driven by other unobserved characteristics that generate the behavior or outcomes. More broadly, a large set of economic models postulate that

[1]Some notable exceptions to this include papers that create exogenous variation in one's belief in oneself by varying the feedback individuals receive such as in Schwardmann and Van der Weele (2019) and Chen and Schildberg-Hörisch (2019) or by varying reasons for self-selection into the experiment e.g. Camerer and Lovallo (1999).

economic behavior in the presence of uncertainty operates as follows: new information arrives $\rightarrow_A$ beliefs are updated according to Bayes rule $\rightarrow_B$ choices are then prescribed by these posterior beliefs. There has been a considerable effort invested recently in investigating the first relationship ($\rightarrow_A$), namely in testing the descriptive validity of the way belief formation has been modeled traditionally, and trying to improve on it.[2] However, less attention has been paid to the second relationship ($\rightarrow_B$), namely whether a shift in beliefs actually translates into a shift in choices in the way that traditional models predict. Here, we attempt to contribute to this effort of studying experimentally how a shift in beliefs translates into changes in action choices.

In this paper, we use a stylized career choice game to study several channels through which an exogenous shift in beliefs may operate. We develop a simple theoretical framework and show empirically how an upward shift in confidence can causally influence decision-making regarding the subject's preferred payment scheme, effort provision, and resulting earnings.

The following illustrative example provides the intuition for how an increase in confidence can affect decision making in a labor market setting. Imagine a computer programmer who is about to graduate from college (let's call her "Taylor"). Taylor is well educated. She has the choice between two types of jobs: (1) a job at a mid-sized company that will pay her a fixed wage, and (2) working at a start-up where her earnings will depend heavily on her performance. She will earn far more at the startup if she is better than the average competing programmer who is also graduating from college and far less if she is worse.

Assume that Taylor, like many other people, believes that she is better than average (see, amongst others Kruger (1999); Burson et al. (2006); Healy and Moore (2007); Moore and Healy (2008); Benoît et al. (2015)). This belief may be influential for two of her decisions. First, she needs to choose the type of job that she thinks will be a good fit for her. Second, conditional on being in the job, she needs to decide how much effort to exert.

For the first choice of which job to take, most models would posit that the higher her confidence in her own ability, the more likely she is to choose to work at the start-up.[3] Choosing to work at the start-up is the payoff maximizing choice if her ability level is *actually* above average, but harmful, if she *wrongly believes* she is better than average.

With regards to her second choice of how much effort to exert within the job, however, an inflated level of confidence could motivate her to work harder in the start-up. The reason for this is that when individuals payment depends on their ability to produce, the returns

[2]See, for example, work by Grether (1980), Eil and Rao (2011), Möbius et al. (2014), Benjamin (2019), Coutts (2019), Heger and Papageorge (2018) and Barron (2020) amongst many others.

[3]However, risk aversion might temper this inclination to choose start-up work by pushing Taylor towards the less risky job, thereby offsetting the influence of her high confidence to some degree (see, e.g., Bonin et al. (2007)).

to effort are proportional to their ability. An inflated view of their own ability implies an upward shift in their perceived returns to effort. However, if this confidence turns out to be misguided, the individual's motivation to exert a high effort level may also be misguided and she may have been better off conserving some of her energy for other activities. It is of interest to evaluate empirically whether a shift in an individual's confidence does actually translate into a shift in the individuals actions in the way that these models typically describe it should - i.e. whether the common theoretical relationship between beliefs and choices provides an accurate descriptive picture of actual behavior.

Ideally, one would study these questions using survey data of real job and effort choices. However, this approach poses several challenges. One issue is that it is non-trivial to gain access to accurate measurements of the beliefs, effort choices and payment scheme preferences of job seekers, holding demand side factors constant. Further, even if one does have access to these high-quality data for these variables, it is not easy to find a natural experiment that provides an exogenous shock to these beliefs. We circumvent these issues using a laboratory experiment, which allows us full control over the environment. Doing this, we generate exogenous variation in beliefs to measure the causal effect of a shift in beliefs on (i) the selection into fixed [low earnings risk] or ability-contingent [high earnings risk] payment schemes, and (ii) effort exerted within a given incentive scheme. We derive our hypotheses for the experiment from a simple theoretical framework.

In the experiment, participants are divided into groups of ten. Each participant takes a test measuring their cognitive ability. This serves as our measurement of the participant's ability. They are then asked to estimate their belief about the probability that their IQ test performance was in the Top Half of their group of ten. We are interested in studying how a shift in the belief about their place in the ability distribution translates into decisions for job and effort choices made by the participants. We designed the experiment such that: (i) the ability distribution is held constant across treatments, and (ii) the influence of the individual's ability on her payoff is fixed prior to her payment scheme and effort choices. We therefore fix the participant's ability and belief about their ability (confidence) at the beginning of the experiment, and examine how it affects the decisions that follow.

After completing the IQ test, participants complete ten rounds of a mundane real-effort task that is chosen to reflect pure effort, and have little dependence on ability. In each round, except the first, participants must choose to be compensated for their effort according to one of two available payment schemes. Subjects can either choose to work for an ability-contingent piece rate or a fixed piece rate that does not depend on their ability. The ability-contingent piece rate pays a high wage if the subject is in the Top Half of her group in the IQ test at the beginning of the experiment and nothing if she is in the Bottom Half. The fixed risk-free piece rate increases in each round, but always lies below the high piece rate of the ability-contingent piece rate. Thus, if a subject is highly confident of being in the Top Half of her group, choosing the ability-contingent piece rate maximizes her earnings.

The exogenous variation in beliefs about relative ability is generated in the experiment by exposing the entire group of ten participants to either a harder or an easier version of the IQ test (participants only interact with other participants who faced the same test). Subjects randomly assigned to the easy test condition are expected to assess their relative rank in the IQ test to be higher than subjects assigned to the difficult test condition. This is commonly referred to as the 'hard-easy' effect (Kruger, 1999; Moore and Kim, 2003; Moore and Small, 2007; Healy and Moore, 2007; Dargnies et al., 2019). The underlying idea is that individuals fail to fully appreciate that when they find a test easy [difficult], the test is likely to also be easy [difficult] for all participants, not just for themselves. They therefore adjust their estimate of their own score more than they adjust their estimate of others' scores, which leads to a predictable (and biased) shift in their estimate of their relative rank.

A closely related concept to the hard-easy effect that can also generate overconfidence is "reference group neglect", where subjects under-respond to a change in the group's composition (Camerer and Lovallo, 1999; Schüssler et al., 2018). Both the hard-easy effect and reference group neglect assert that subjects neglect the fact that there is a shift in the background distribution against which they are being compared. The hard-easy effect refers to a neglect of the fact that changing the difficulty of the test will lead to a shift in the distribution of others' scores. In contrast, reference group neglect refers to neglect of a shift in the distribution of other's scores due to a particular subset of individuals selecting into participation.[4] Both of these mechanisms can result in the distorted belief that one is ranked higher than one actually is, which is commonly referred to as "overplacement", and which may then influence decision making in important contexts.[5]

In our study, we find evidence that a small shift in the difficulty of the test leads to a large shift in the average belief that subjects holds regarding their relative placement. The effect is particularly strong for subjects who are actually in the Bottom Half of the group. These subjects report higher beliefs on average in the easy treatment than in the hard treatment. The beliefs of those in the Top Half are less affected on average.

In terms of this shift in beliefs translating into actions, we find that this increase in confidence due to exposure to the easier test leads subjects to choose the ability-contingent piece rate more often. If randomly confronted with the hard test, subjects are more likely to

[4]Both these concepts fall into the broader important class of biases, where subjects neglect some feature of the data generating process, leading to biased inference and systematic mistakes in decision making. Notable examples include *correlation neglect* and *selection neglect* (see, e.g., Jehiel, 2018; Barron et al., 2019; Enke and Zimmermann, 2019; Enke, 2020). Reference group neglect could be viewed as a form of selection neglect.

[5]As opposed to "overprecision", overconfidence in the accuracy of one's beliefs, and "overestimation", the overconfidence in one's absolute ability (Moore and Healy, 2008).

choose the fixed piece rate. This shift in job choice occurs even though the incentives faced remain constant. This suggests that the way that knowledge is tested within an education system could have implications for the later choices made by individuals, holding the fundamental levels of their knowledge and abilities constant. With regard to effort, we find that in our experiment the motivation level of participants tends is high, and largely insensitive to their beliefs, implying we do not observe a shift in effort choices.

The shift in beliefs has important consequences for earnings of the Bottom Half of the group. This group earns only about a quarter of what the top group earns on average, but their average earnings are reduced even further, by about 40 percent, when their confidence has been exogenously increased. The reason for this is that overconfidence in relative ability is costly for below-average-ability individuals, as it increases their probability of choosing an ability-contingent incentive scheme, which is a mistake for these individuals.

An additional important contribution of this paper is that it provides evidence towards addressing one of the challenges in the literature studying beliefs: understanding the causal mapping of beliefs to economic decisions. Many papers take as given that beliefs have a simple linear causal mapping to decisions, and yet in the few papers that have investigated this relationship, the results are not so straight forward.[6] Costa-Gomes and Weizsäcker (2008), for example, show that subjects in their games fail to best-respond to their stated beliefs almost half of the time, providing evidence for the fact that beliefs do not directly map in a simple way into actions. However, in a later follow-up paper, the authors introduce exogenous variation in beliefs within a trust game to demonstrate that elicited beliefs have a causal impact on choices in that domain (Costa-Gomes et al., 2014). In the context of a repeated public goods game, Smith (2013) also provides evidence of a causal effect of beliefs on actions. In that context, Smith (2013) argues that the magnitude of the causal effect he estimates using an IV approach is about half the size of the effect that would be estimated if using OLS and not accounting for the endogeneity of the beliefs. Together, these papers show that the relationship between elicited beliefs and observed actions is not always as expected, and further empirical investigation is warranted across different domains. This paper contributes to this exercise.

Our paper also contributes to the literature that studies the relationship between beliefs about relative ability and stylized labor market choices, such as career choices, effort provision, willingness to compete and risk-taking behavior (e.g. Camerer and Lovallo, 1999; Niederle and Vesterlund, 2007; Dohmen and Falk, 2011; Bruhin et al., 2018; Cheung

[6] In a recent paper, Fischer and Sliwka (2018) distinguish between confidence in the level of one's ability and confidence in one's ability to learn (i.e. their ability production function). Using feedback to generate variation in beliefs, they show that an increase in one's belief about their ability to learn increases motivation to learn, while an increase in their belief about their ability level reduces motivation.

and Johnstone, 2017; Pikulina et al., 2016, 2017; Chen and Schildberg-Hörisch, 2019).[7] Dohmen and Falk (2011), for example, also study the choice between a variable and fixed-payment wage scheme. The authors show that individual characteristics, such as relative self-assessment and risk aversion are important predictors of how individuals sort into the different incentive schemes. Our paper differs from this body of work in a two important ways. First, we focus attention solely on the causal effect of a shift in beliefs, while holding the true ability distribution constant. Second, our design separates the measurement of ability and effort, allowing us to fix the individual's ability measurement and influence on her payoff prior to facing the different incentive schemes (i.e. we hold the ability component of the production function fixed). Our design therefore rules out the possibility of substituting lack of ability with effort, and allows us to isolate the influence of a belief shift on effort, without it being obscured by heterogeneous ability.

Our experimental design is related in spirit to Camerer and Lovallo (1999), who study entry into competition due to underestimating one's competitors. In their experiment, individuals are either informed or not informed during recruitment into the experiment that they will be competing against other subjects in a skill-based task. This information leads to more self-selection into the experiment based on performance level in the informed treatment and subsequent over-entry into a skill-based competition. All the high-skilled, highly competitive individuals select into the experiment, and then select into competing within the experiment. They fail to recognize that others are behaving in the same way and so they are not competing against a random draw from the population. Importantly, the results they observe are not necessarily caused by overconfidence, since overconfidence might co-vary with other individual characteristics that contribute both to the self-selection into the experiment and into competing.[8] This observation does not diminish the contribution of the study, since in many real world settings, overconfidence is likely to co-vary with these unobserved factors in the same way. We view our paper as providing complementary evidence to the insightful work cited above by focusing on the narrower question of isolating the role played by purely by confidence. Our paper also demonstrates a method that may be used to examine the causal role played by a shift in confidence on actions in other domains.

The remainder of the paper is structured as follows. Section 2 outlines the theoretical framework and hypotheses, section 3 describes the experimental design. Section 4 discusses the results, and section 5 concludes.

[7]Furthermore, Huck et al. (2018) consider how individuals might endogenously manage their own beliefs in order to motivate themselves to perform better in a real effort task.

[8]This is supported by the fact that a replication by Dankova and Servatka (2019) which extends Camerer and Lovallo (1999) by studying both men and women (since the original study focuses on men) finds that the results are highly sensitive to the participants gender.

2 Theoretical framework

In this section, we motivate the experimental design and hypotheses through means of a simple model where a worker chooses her payment scheme and then effort. The objective of the theoretical framework is to simply provide some discipline and precision to the ensuing discussion. To do this, we augment the model used by DellaVigna and Pope (2017) in their large-scale study of real effort and motivation.

Consider an individual i, who can earn money by performing a task that requires costly effort, e. She is either a high ability or low ability individual, $a \in \{a_L, a_H\}$. Prior to performing the task, the individual chooses between two incentive schemes: (i) one where high-ability individuals will earn a high wage, $w(a_H) = w_H$, and low-ability individuals will earn a low wage, $w(a_L) = w_L$, or (ii) one that pays a fixed wage to everyone, $\bar{w}$, where $w_H > \bar{w} > w_L$. After choosing her incentive scheme, she chooses the level of effort, e, she would like to exert. Following DellaVigna and Pope (2017) we allow the subject to derive some intrinsic utility from effort, denoted by s.[9]

A risk-neutral individual would choose her incentive scheme, $w \in \{\bar{w}, w(a)\}$, and effort level, e, by solving the following:

$$\max_{w\in\{\bar{w},w(a)\}} \max_{e\geq 0} \mathbf{E}_a[(s+w)\cdot e - c(e)] \tag{1}$$

where $c(e)$ is the cost of exerting effort, and is assumed to satisfy $c'(e) > 0$ and $c''(e) > 0$. The expectation operator, $\mathbf{E}_a$, denotes the expectation with respect to the individual's ability, and s represents the individual's intrinsic motivation for completing the task. Since uncertainty about one's own ability is only directly relevant for the wage schedule, we can rewrite equation 1:

$$\max_{w\in\{\bar{w},w(a)\}} \max_{e\geq 0} (s+\mathbf{E}_a[w])\cdot e - c(e) \tag{2}$$

Equation 2 shows that the individual's subjective belief regarding the likelihood that she is high-ability, $\hat{\pi} = \hat{\mathbf{P}}(a = a_H)$, is important both for her decision of which incentive scheme to take, and how much effort to exert if she chooses the ability-contingent incentives. Essentially, the choice of an incentive scheme involves a choice between being paid a certain piece rate of $\mathbf{E}_a[\bar{w}] = \bar{w}$, or an expected piece rate of $\mathbf{E}_a[w(a)] = \hat{\pi}\cdot w_H$ for each unit of effort. We normalise $w_H = 1$.

[9]As in DellaVigna and Pope (2017), we view this intrinsic motivation term as capturing, in reduced form, any non-monetary reward the workers derive from working on the task. In terms of the laboratory experiment described below, this is taken to include any sense of duty to, or gratitude towards the experimenter for the fixed show-up fee. DellaVigna and Pope (2017) argue that this non-monetary reward term is important for explaining the commonly observed non-zero effort in fixed-wage laboratory experiments.

The effort choice

Conditional on the choice of an incentive scheme, w, the individual chooses effort, e^*, according to the condition $c'(e^*) = s + \mathbf{E}_a[w]$. Under the certain piece rate (PR) incentive scheme she chooses $e^*_{PR} = c'^{-1}(s + \bar{w})$. Under the ability-contingent (AC) incentive scheme, she chooses $e^*_{AC} = c'^{-1}(s + \hat{\pi} \cdot w_H)$. If her confidence is sufficiently low (i.e. she believes the probability that she is high ability is below the following threshold: $\hat{\pi} < \frac{\bar{w}}{w_H}$), then the individual exerts more effort under the certain piece rate incentives. However, if she is sufficiently confident in her own ability, such that $\hat{\pi} > \frac{\bar{w}}{w_H}$ she expects a higher piece rate under ability-contingent incentives, and would work harder under these incentives. We therefore define a threshold belief, namely $\pi^e := \frac{\bar{w}}{w_H}$, such that:

- For low beliefs (i.e. if $\hat{\pi} \in [0, \pi^e]$), the individual exerts more effort under certain piece rate incentives than under ability-contingent incentives.
- For high beliefs (i.e. if $\hat{\pi} \in [\pi^e, 1]$), the individual exerts more effort under ability-contingent incentives than under certain piece rate incentives.

Importantly, for high levels of intrinsic motivation, s, differences in effort due to variation in monetary incentive schemes will be harder to detect. This is a common feature of laboratory real effort tasks (DellaVigna and Pope, 2017; Erkal et al., 2016; de Araujo et al., 2015). The time limit of the task constrains effort to be below a certain effort level, $e \leq \bar{e}$, which can be binding if $\bar{e} \leq e^*_{PR}$ and $\bar{e} \leq e^*_{AC}$ (e.g. this occurs when the intrinsic incentives are strong enough to maintain effort close to the boundary under both sets of monetary incentives). If this is the case, then the observed effort level chosen under both sets of incentives will be equal. As we will see in the results below, in our experiment the intrinsic motivation to exert effort appears to have been high and we do not observe substantial variation in effort across incentive schemes. While this reduces the informativeness of the lessons that we can draw about effort choices under different incentive schemes in our experiment, it has the advantage of allowing us to study the relationship between beliefs and the incentive choice decisions more cleanly, since effort choices are largely held constant.

The incentive scheme choice

When choosing between incentives schemes, the individual chooses the ability-contingent incentives whenever she expects to earn more per unit of effort under them than she would under the certain piece rate per unit of effort.

$$(s + \hat{\pi} \cdot w_H) \cdot e^*_{AC} - c(e^*_{AC}) \geq (s + \bar{w}) \cdot e^*_{PR} - c(e^*_{PR}) \quad (3)$$

This inequality holds whenever $\hat{\pi} \cdot w_H \geq \bar{w}$.[10] It holds even if the effort level chosen under both incentives schemes is the same (i.e. if $e^* = \bar{e}$). Under risk neutrality, the threshold for the choice of incentives, and the threshold for effort choices are equal (i.e. $\pi^i = \pi^e = \frac{\bar{w}}{w_H}$).

We can therefore summarise the influence of beliefs on choices as follows:

- a *low confidence* individual (i.e. one with a belief $\hat{\pi} \in [0, \pi^i]$) will (i) choose the certain piece rate incentives, and (ii) exert (weakly) lower effort under ability-contingent incentives than under certain piece rate incentives.
- a *high confidence* individual (i.e. one with a belief $\hat{\pi} \in (\pi^i, 1]$) will (i) choose the ability-contingent incentives, and (ii) exert (weakly) higher effort under ability-contingent incentives than under certain piece rate incentives.

In Appendix A.1.2, we relax the risk neutrality assumption and show that the threshold belief at which individuals will switch incentive choice differs from the one at which effort is affected by incentives.

2.1 Hypotheses

The theoretical framework above provides us with a set of hypotheses that we will test in the experiment. The central objective is to ask how a shift in confidence about one's own ability affects incentive scheme choices and effort choices. To do this, in the experiment, we use the well established Hard-Easy effect (see, e.g., Moore and Healy, 2008) to induce exogenous variation in subject's beliefs about their own ability, $\hat{\pi}$, keeping everything else constant (e.g. actual ability, a).

In the experiment, it is reasonable to expect heterogeneous beliefs. Therefore we consider a continuum of agents who hold beliefs, distributed on the unit interval, $\hat{\pi} \sim F(\hat{\pi})$, such that $f(\hat{\pi})$ is everywhere positive on $\hat{\pi} \in [0, 1]$. Since evaluating our main hypotheses relies on our experimental design generating exogenous variation in beliefs across treatments, our first hypothesis tests whether we observe a shift in beliefs due to the hard-easy effect in our experiment:

Hypothesis 1. *(Shift in Beliefs): Beliefs about one's own relative ability in the Easy treatment will be higher on average than beliefs in the Hard treatment.*

[10] To see this, notice that if $\hat{\pi} \cdot w_H > \bar{w}$, the individual could simply choose the ability-contingent incentives and set effort equal to the optimal effort level under certain piece rate incentives, $e = e^*_{PR}$, and receive a higher expected payoff than under the certain piece rate incentives, i.e.

$$(s + \hat{\pi} \cdot w_H) \cdot e^*_{PR} - c(e^*_{PR}) \geq (s + \bar{w}) \cdot e^*_{PR} - c(e^*_{PR})$$

Since $e = e^*_{AC}$ maximises the LHS of this inequality, equation 3 must also hold.

Our second hypothesis tests whether incentive scheme choices are affected by the hypothesized shift in beliefs between treatments. The logic behind this hypothesis is that individuals who hold a higher belief in their own ability are more likely to choose the ability-contingent incentives – i.e. an upward shift in $\hat{\pi}$ for all individuals implies that $\hat{\pi} \geq \frac{\bar{w}}{w_H}$ will hold for a greater fraction of individuals.

Hypothesis 2. *(Incentive Choices): An exogenous increase in confidence will lead to a higher fraction of individuals choosing the ability-contingent incentives.*

Third, we ask how a shift in beliefs affects effort choices. We consider two scenarios: one in which the intrinsic motive to exert effort is high and the elasticity with respect to extrinsic monetary incentives is low, and one in which the intrinsic motive is low and therefore the elasticity with respect to extrinsic monetary incentives is high.

In the scenario where the intrinsic motive to exert effort is sufficiently high (as it is in much of the experimental literature involving real effort in the lab) an increase in confidence will not change effort. In the scenario where subjects are predominantly motivated by extrinsic monetary rewards in their effort choices, the more confident individuals will exert more effort under ability-contingent incentives.

Hypothesis 3. *(Effort Choices): We will observe one of the following two patterns of behavior for effort choices. Either: (i) [low intrinsic motivation] For high confidence individuals, effort choices will be higher under the ability-contingent incentives than under the certain piece rate incentives. In this scenario, an upward shift in confidence will increase overall average effort. (ii) [high intrinsic motivation] Effort choices will not be influenced by the incentive scheme. In this scenario, an exogenous shift in confidence will not affect effort choices.*

Ultimately, we also want to examine the effect that an upwards shift in confidence has on earnings. Within each treatment in our experiment, individuals are classified into two groups – high and low ability. An upward shift in beliefs is likely to lead to a very different effect on outcomes individuals in these two groups. Essentially, individuals who are actually of high ability benefit from an upward shift in confidence as this leads to them switching towards the ability-contingent incentive scheme more often. The reverse is true for low-ability individuals. The boost in confidence can be harmful for them as they may switch to the ability-contingent incentive scheme even though it results in a loss of earnings. Appendix section A.1.1 contains a more detailed discussion of this intuition, and the Hypothesis 4 below summarizes the main testable implications:

Hypothesis 4. *(Earnings): While an increase in confidence will have an ambiguous effect on average earnings, the framework suggests that: (i) it will lead to weakly lower average earnings for low-ability individuals, (ii) it will lead to weakly higher average earnings for high-ability individuals, and (iii) it will increase earnings inequality overall.*

3 Experimental Design

Our experimental design aims, firstly, to assess how an agent's confidence in her relative ability causally affects her choice of incentive scheme for a subsequent real effort task. This mirrors the labour market decision of whether to pursue employment that is highly dependent on one's ability or not. Secondly, we evaluate the relationship between the agent's confidence and her effort provision under the chosen incentives.

An experiment with these objectives should have the following features:

i. An exogenous shift in subjects' beliefs about their relative ability,

ii. a separation of the role of ability and effort in the production function,

iii. a fine-grained measurement of a participant's willingness to switch from a fixed incentive scheme to an ability-contingent scheme

3.1 Overview of the experimental timeline

Our experiment is designed with these features in mind. Figure 1 outlines the timeline of the experiment. The first part of the experiment obtains a baseline measure of the participants' willingness to exert effort for a fixed wage. This stage is implemented prior to treatment, and therefore provides a control measure of effort that is unaffected by the hard-easy task assignment. Next, the ability task is the stage in which the treatment is introduced, with participants either exposed to a hard or easy version of the ability task. In each session, subjects are randomized into two groups of ten. One group completes the hard task, and one completes the easy task. The reason for this is to have within-session randomization as it is important to control for session fixed effects, but participants never interact with the participants in the other group. After the ability task, we elicit participants' beliefs about their relative placement in the group. Thereafter, they proceed to the main effort task. In this task, participants face ten rounds of the completing the effort task under either a fixed wage or an ability-contingent wage. In each round, except the first, participants first choose one of the two wage structures, and then exert effort. The ability-contingent wage stays constant in every round, while the fixed wage is ratcheted up in order to provide a fine-grained measure of the participant's indifference point. The experiment concludes with a measurement of subjects risk preferences and the questionnaire. We provide further details on each of the components of the experiment in the following sections.

Figure 1: Sequence of experimental parts

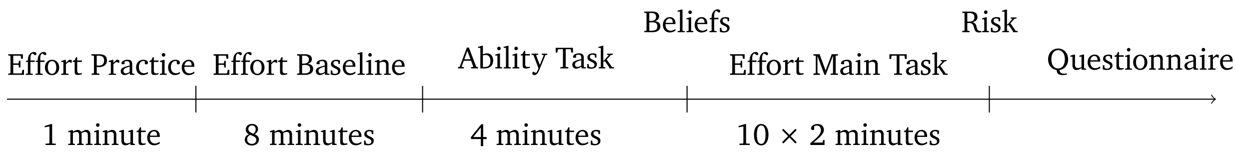

3.2 The ability and the effort measurement tasks

The main components of the experiment are the "Ability Task", used to measure a, and the "Effort Main Task", used to measure e. One challenge for an experiment of this nature is that it is non-trivial to measure ability[11] and effort[12] separately. We explicitly try to address this issue by using two separate tasks—one that we view as depending more on the individual's ability, and less on the effort she exerts; and one that depends more on effort, and less on her ability. We contend that this choice of tasks provides us with a reasonably clean measure of these two variables of interest.

The ability task

The "Ability Task" consists of a test that is often used to measure IQ, namely Raven Progressive Matrices.[13] Subjects have four minutes to solve as many matrices as they can. Subjects can go back and forth between the 12 matrices and can change their answers until the time is up. Every correct answer yields one point, and there are no negative points for wrong answers. We chose not to directly incentivize the task for two reasons. First, IQ tests tend to induce an intrinsic motive to perform well. Second, we wanted to limit the role of hedging in influencing the belief elicitation. Importantly, the motive to perform well faced by subjects should not differ between the two treatments, because at this point of doing the IQ task, subjects did not know that they would be incentivized for accurate beliefs at a later time—this is to prevent them from intentionally performing poorly.

The effort task

While the "Ability Task" was chosen such that the limiting factor in participants' performance is their ability, the "Effort Task" was chosen to be a task where ability plays a minor role

[11] We view *ability*, a, as being a fixed characteristic of the individual that she cannot change during the time frame of the experiment.

[12] We view *effort*, e, as a being a malleable object that the participant has full control over.

[13] A matrix consists of nine related patterns of which one is missing. Below the matrix, there are eight possible patterns to complete the set of nine. Subjects have to find the one piece that best completes the set. There is only one correct answer. Due to the short time frame for solving up to 12 matrices, performance depends mostly on cognitive ability. Subjects are either skilled at finding the missing pattern or not. Every matrix is shown on its own screen.

and all participants have a lot of control over their performance (i.e. performance depends predominantly on how motivated the individual is in exerting effort). For this purpose, we adopted the slider task by Gill and Prowse (2012). Using the mouse, participants move sliders on the screen from position zero to position 50. Sliders are shown in sets of 20. When all 20 sliders are set to 50 the subject can click the submit button and the sliders are reset to zero for a new round. In Section 3.5 below, we discuss the incentives subjects face.

Since the treatment is introduced during the "Ability Task" phase of the experiment, we also measure each individual's baseline effort level prior to treatment under fixed piece rate incentives. This serves two purposes. Firstly, it allows us to check for balance of effort in the slider task between treatment groups, prior to the treatment manipulation. Secondly, it allows us to control for baseline effort levels when assessing the impact of the treatment, thereby reducing unobserved individual level heterogeneity.

3.3 Treatment variation

The objective of the treatment variation is to exogenously shift confidence using a minimal intervention. Therefore, the two treatment conditions are completely identical except for a slight difference in the difficulty of the Ability Task. Within each session, subjects are randomly assigned to one of two groups. One entire group is exposed to a harder version, and the other group to an easier version of the Raven Progressive Matrices. Eight of the twelve puzzles are identical in both treatments. The remaining four are either slightly easier or slightly harder than the rest. Table 1 shows the precise sequence of Raven matrices faced in each of the treatment groups. In each of the four differences, switching a C-matrix for an E-matrix represents an increase in difficulty. Moving backwards and forwards between puzzles is allowed in the task. Subjects only interact with other participants who completed exactly the same test as them (i.e. in their treatment group of ten subjects), and they know this.[14]

Table 1: Treatment variation in ability task

	1	2	3	4	5	6	7	8	9	10	11	12
High (Easy)	C1	C3	D7	C6	D8	C2	D6	C5	C7	C12	C4	D4
Low (Hard)	E4	E12	D7	C6	D8	E6	D6	C5	C7	E1	C4	D4
same / different	$\neq$	$\neq$	$=$	$=$	$=$	$\neq$	$=$	$=$	$=$	$\neq$	$=$	$=$

This approach draws on the finding in the psychology literature that when the difficulty

[14]Our experiment uses within-session assignment to treatment in order to avoid the numerous potential issues associated with between-session randomization.

of a task increases, this causes a downward shift in an individual's confidence regarding their relative position in the ability distribution. Conversely, facing an easier task makes individuals more confident regarding their position in the distribution (Burson et al., 2006; Healy and Moore, 2007; Larrick et al., 2007; Moore and Healy, 2008; Bordley et al., 2014; Benoît et al., 2015). Importantly, these results assume a constant group composition, so there is no reason for the individual's actual rank to change when the difficulty of the test is shifted. An explanation for this result is that when the difficulty of a test is reduced, individuals find the test easier and adjust their assessment of their own performance upwards. However, they do not adjust their belief of the distribution of others' scores up as much. This results in a higher relative assessment of their own performance. Kruger (1999) shows that this miscalibration can lead to the majority of subjects evaluating themselves as worse-than-average in difficult tasks and better-than-average in easy tasks.

In our experiment, we therefore name the treatment in which subjects face the hard test, the "Low confidence" treatment, and the treatment in which subjects face the easier test, the "High confidence" treatment.

3.4 The belief elicitation task

After the "Ability Task" we elicit subjects' beliefs about their relative performance in comparison to the 9 other participants in their group who faced the same task as them.[15] More specifically, we asked subjects the following question: *"What do you think is the probability that you scored among the top 5 participants in the IQ picture task?"*. In order to give some guidance in thinking about probabilities, we provided the participants with a scale of possible answers ranging from *"0 - I am certain that I scored in the bottom half"* to *"100 - I am certain that I scored in the top half"*. Participants were free to state any number from 0 to 100. Their guess is incentivized using the quadratic scoring rule (Selten, 1998).[16] The quadratic scoring rule is explained in detail to them in the instructions provided, both on-screen and on paper. The scoring rule is designed to provide the highest expected payoff when subjects state their true beliefs. Maximum earnings are €2 for the belief elicitation task. The belief elicitation came as a surprise at this point in the experiment to prevent subjects artificially manipulating their earlier "Ability Task" scores in anticipation of the belief elicitation.

[15]Subjects knew that they were randomly assigned to one of two groups of ten within the session, and that all the other members of their group faced the same "Ability Task", however they didn't know that the other group faced a different "Ability Task". It was made clear to them that only their own group's performance was relevant for them, and there was no interaction at all between groups.

[16]We view the usage of both the QSR incentives, and the intuitive descriptions associated with different probability levels, as providing an appealing compromise approach to the belief elicitation. It allows us to maintain the appealing theoretical properties of a proper scoring rule, but also provides subjects with an intuitive and easily understandable guide for answering the question.

In addition, we asked participants to report their best guess of how many points they scored in the task and what they believe the 5th highest score in their group is (unincentivized, in order to avoid hedging).

3.5 The wage scheme choice

One central objectives of the experimental design is to evaluate how a shift in confidence affects subjects' choices between fixed wage schemes and ability-contingent wage schemes. The simplest way to do this would be to allow subjects to choose between a single fixed wage scheme and a single ability-contingent wage scheme. However, with just this binary choice, a shift in subjects' beliefs may change their valuation of the two options, but not enough to lead to a change in choices. In order to obtain a more fine grained measurement of subjects' precise relative valuation of the two payment schemes, we constructed a task where subjects make the choice ten times, but the fixed wage is gradually ratcheted up in each successive choice. Therefore, subjects face ten rounds of two-minute real effort tasks. In each of these rounds (except the first), subjects can choose whether they would prefer a fixed piece rate, or whether they want to work under the ability-contingent piece rate. Figure 2 below summarizes the two wage schemes available in each round. A Subject's switching point reveals the round in which they are indifferent between the two wage schemes.

Figure 2: Payment Scheme in Main Effort Task

Period Number (2 min each)	Option A Payment (per 20 sliders)	Option B Payment (per 20 sliders)
1	Can't Choose Option A	€1 if TOP HALF; €0 if BOTTOM HALF
2	€0.1	€1 if TOP HALF; €0 if BOTTOM HALF
3	€0.3	€1 if TOP HALF; €0 if BOTTOM HALF
4	€0.4	€1 if TOP HALF; €0 if BOTTOM HALF
5	€0.5	€1 if TOP HALF; €0 if BOTTOM HALF
6	€0.55	€1 if TOP HALF; €0 if BOTTOM HALF
7	€0.6	€1 if TOP HALF; €0 if BOTTOM HALF
8	€0.65	€1 if TOP HALF; €0 if BOTTOM HALF
9	€0.7	€1 if TOP HALF; €0 if BOTTOM HALF
10	€0.8	€1 if TOP HALF; €0 if BOTTOM HALF

Another important feature of the experimental design is that the ability component of the production function is held fixed, and the distribution is the same in the two treatments (i.e. 50% in the Top Half, and 50% in the Bottom Half). This is done by using the "Ability Task" to obtain a fixed measurement of ability before the participant reports her belief, and before she chooses between payment schemes. This implies that differences in beliefs, incentive-scheme choice and effort can be causally attributed to the shift in beliefs due to the treatment. It rules out potential issues that can arise if one were to allow participants to first make an incentive scheme choice, and then produce output that depends on both ability and effort. For example, in this case, one would not be able to distinguish a high ability individual who exerts low effort from a low ability individual who exerts high effort.

There are a few additional features of the payment scheme choice that are worth noting. First, to rule out learning effects (with respect to own ability), subjects receive no feedback about either their relative ability score or their performance in any components of the "Effort Task" tasks until the very end of the experiment. Second, in the initial period, all participants have to work for the ability-contingent piece rate. This feature allows us to assess how the shift in confidence affects effort provision when all subjects are forced to work under the ability-contingent piece rate, thereby avoiding endogenous selection effects—in all other rounds the incentive scheme that a subject faces is endogenous. Third, the fixed piece rate increases in each period from €0.15 per 20 sliders in the second period to €0.80 per 20 sliders in the last period. Once the expected earnings from the ability-contingent piece rate are lower than the fixed piece rate in that period, individuals should switch to the fixed piece rate and choose it for the remainder of the experiment, assuming risk neutral preferences.[17]

3.6 The risk elicitation

Finally, we elicit risk preferences by adapting the preferences module on risk taking by Falk et al. (2016) to our setting. The staircase procedure is essentially equivalent to a traditional multiple price list, presenting multiple choices between a sure payoff and a gamble, but simply requires fewer decisions on the part of the subject in comparison to a traditional price list by avoiding redundant choices. The staircase we use has four choices between a sure payment and a risky gamble. The outcomes of the risky gamble are always €0 or €1, each associated with a 50 percent chance of occurring. The sure payment value was varied across decisions to allow us to elicit the subject's point of indifference. One of the decisions

[17]In the end, five of the ten rounds are randomly chosen for payment. Subjects had to answer four control questions before starting the task to ensure comprehension of the payment scheme. To elicit the baseline motivation of moving sliders, subjects complete 9 minutes of the "Effort Task" at the start of the experiment, the first minute being an unincentivized practice round. In the baseline round, we pay €0.30 per 20 sliders, and all completed sets are paid out.The objectives of the baseline round were the following. Firstly, it allows subjects to familiarize themselves with the slider task, thereby ameliorating learning effects during the main effort task. Secondly, it allowed us to obtain a baseline measure of subjects' effort choices prior to the treatment variation under fixed incentives. This allows us to check for baseline balance in effort, and also to control for subjects' baseline effort at the individual level.

was randomly chosen for payment.[18]

At the end of the experiment, we administered a comprehensive questionnaire.

3.7 The procedure

The experiment was programmed in zTree (Fischbacher, 2007) and conducted at the WZB-TU experimental laboratory in 2017. Participants were solicited through an online database using ORSEE (Greiner, 2015) from a subject pool of mostly undergraduate students from all faculties. In total 100 subjects participated in 5 sessions, 20 in each. 47 of them were female, 49 male and four chose not to self-report their gender. Subjects received a show-up fee of €5 plus their earnings from the tasks. Mean earnings for the 60 minutes sessions amounted to €13.30. The relevant instructions were handed out to participants at the beginning of each stage and read out loud. Complete instructions as they appeared to participants are provided in the Online Appendix.

4 Results

4.1 Does the hard-easy treatment shift beliefs?

The main objective of our treatment manipulation is to exogenously shift participants' beliefs about their relative performance in the IQ test. In line with Hypothesis 1, Figure 3 shows that we find a significant difference in the participants' level of confidence in their own ability between the two treatment groups, where confidence refers to the individual's stated probability of being in the Top Half of their group (t-test, $p < 0.01$).[19] We refer to the easy task treatment as the High confidence treatment, and the difficult task treatment as the Low confidence treatment.

Result 1. *In line with the previous hard-easy effect literature, reducing the difficulty level of*

[18]We have two more staircases for which we use the subjects' own reported beliefs as the probability for the gamble. We do this as it allows us to compare their choices in the risk task with their incentive scheme choices. While the risk task involves a choice between a pure lottery and a fixed payment, the main task involves the choices of which payment scheme to work under. Both tasks share the same values of outcomes and subjective probabilities. Therefore, it allows us to assess whether there are any differences in how probabilities affect choice tasks and effort tasks. As a caveat, the risk elicitation always comes after the incentive scheme choices, which could motivate subjects to make consistent choices in the risk task. The third staircase is a mirror of the second and only used for keeping incentives fair to not favor individuals that stated high beliefs, as high beliefs increase the chance of winning the gamble.

[19]The top panel of Table 5, reported in the Appendices, shows that our treatment groups are balanced on observable characteristics (as would be expected given the within-session randomization). The second panel confirms that, by design, the ability score differs, but the relative ability distribution is identical between treatments.

the ability task increases the average confidence that participants have in their own relative performance.

Figure 3: Average stated beliefs by treatment

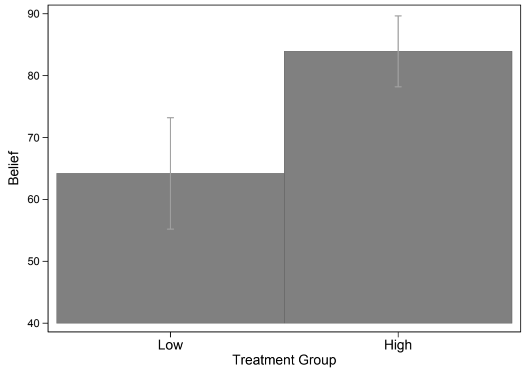

Note: (i) Vertical lines denote 95% CI around the mean.

4.2 The influence of beliefs on wage scheme choices

Next, we present evidence on whether the increase in confidence translates into actions by increasing the proportion of individuals choosing a high earnings risk wage scheme (i.e. the ability-contingent wage). This is a test of Hypothesis 2. Figure 4 shows that the ability-contingent wage is chosen significantly more often in the HIGH treatment condition than in the LOW treatment condition (diff. = 18pp; t-test, $p < 0.01$).[20] The first two columns of Table 2 show that this result is unaffected by session fixed effects.

Figure 4: Propensity to Choose Ability-Contingent Incentives

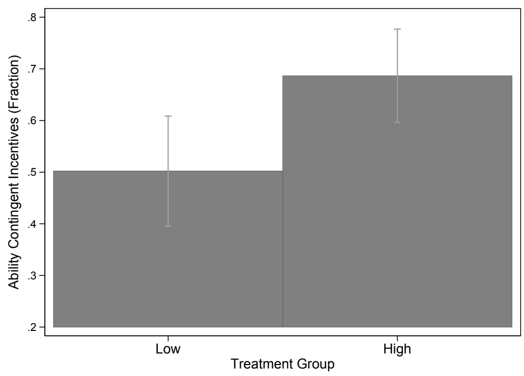

Note: (i) Vertical lines denote 95% CI around the mean.

[20]Wondering what Taylor would do if she were a man? In Appendix A.3, we analyze the gender heterogeneity in the stylized career choices we observe.

In addition to documenting the treatment effect on choices, it is informative to provide more direct evidence on whether this treatment effect operated via beliefs. To do this, columns (3) of Table 2 shows that the subjects' reported beliefs about their likelihood of being in the Top Half are highly predictive of their incentive scheme choices (a 1pp increase in a participants' belief is associated with choosing the ability-contingent incentives 0.86pp more often). However, this relationship may be endogenous. A nice feature of the experimental design is that we can use the treatment variation as an instrument for beliefs. Columns (4) and (5) report the results from this exercise, showing that the exogenous shift in beliefs does indeed translate directly into a change in wage choices. This cleanly demonstrates a causal relationship between beliefs and action choices in this context, showing that this result is not driven by other unobserved differences between individuals who hold high beliefs and low beliefs.

Result 2. *An increase in confidence results in a higher propensity to choose the ability-contingent wage scheme.*

Table 2: Propensity to Choose the AC Incentives

	OLS (1)	OLS (2)	OLS (3)	IV (4)	IV (5)
Treatment (HIGH=1)	0.18**	0.18**			
	(0.07)	(0.07)			
Subj Belief			0.86***	0.94***	0.93***
			(0.10)	(0.26)	(0.26)
Risk (CE p=50)					0.21
					(0.19)
Constant	0.50***	0.55***	-0.01	-0.07	-0.17
	(0.05)	(0.08)	(0.09)	(0.20)	(0.22)
Session Fixed Effects		✓	✓	✓	✓
Observations	100	100	100	100	100
Adjusted R^2	0.058	0.068	0.462		
First-Stage F				13.91	13.88

Notes: (i) In the IV Regressions, Subjective Beliefs are instrumented using the treatment dummy. (ii) Standard errors in parentheses. (iii) Dependent variable: fraction of AC choices in rounds 2 to 10.
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The discussion above has considered wage choices in all rounds pooled together. Figure 5 shows average wage choices in each round separately. As expected, the fraction of ability-contingent wage choices decreases across rounds in both treatments as the fixed wage rate increases. However, in every round the fraction is higher in the High confidence treatment, with the gap largest in the middle rounds – Table 7 in the Appendix shows that there is a significant difference in rounds 4 to 8 (t-test, $p < 0.05$). This is unsurprising since a large fraction of subjects would need to hold extreme beliefs to generate a large treatment difference in the first or last round. Overall, the evidence indicates that the shift in confidence resulted in a substantial change in behavior,

Figure 5: Propensity to Choose Ability-Contingent Incentives (by Round)

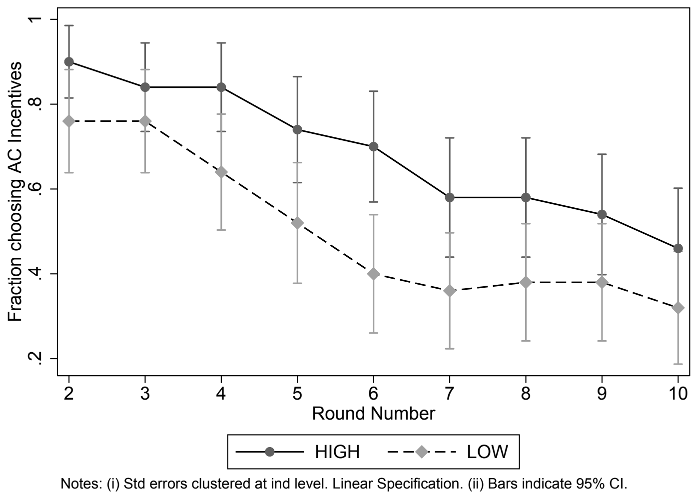

4.3 Influence of beliefs and incentive choice on effort

Once an individual has chosen her incentive scheme, the second choice she has to make is the choice of how much effort to exert. For this effort choice, our simple theoretical framework yields two sets of mutually exclusive predictions – one for scenarios where intrinsic motivation is high, and one for scenarios where intrinsic motivation is low.

Taken together, the data collected in our experiment is more consistent with the first scenario. We measure effort using the variable "effort per minute", which reflects the number of sliders completed during each minute within a particular round. We find no significant difference in average effort exerted between treatment groups (see Figure 6 and Table 7 in the Appendix). Further, we present two additional pieces of evidence in favor of the explanation that intrinsic motivation is high in the experiment.

Figure 6: Per minute effort in baseline and main task by treatment

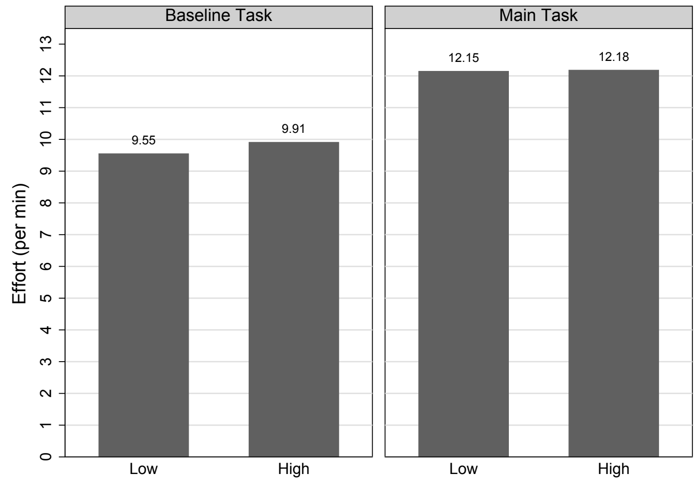

First, we show that effort is not strongly associated with the participant's expected wage rate. Figure 7 plots the average per minute effort exerted in the baseline round, as well as in every subsequent round. While we do see some initial learning, after the baseline round, there is very little change in effort exerted even though the value of the fixed piece rate increases from €0.1 to €0.8, and the fraction of individuals choosing this fixed piece rate incentives increases substantially in both treatment groups between round 2 and round 10.

Figure 7: Effort choices across rounds, by treatment

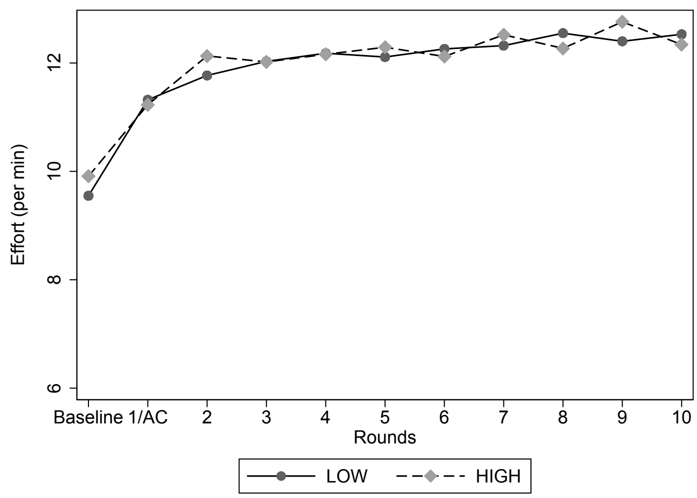

Second, we show that effort does not appear to be associated with the participant's beliefs even when she is exogenously assigned to the ability-contingent wage scheme. Figure 8 focuses on the first round in which all participants faced the ability-contingent incentive scheme. This figure shows that in both treatments, effort is highly unresponsive to beliefs.

While our treatment successfully shifted the beliefs of participants in the two treatments, it did not affect the relationship between beliefs and effort, which is rather flat.[21]

Figure 8: Per minute effort in first round under ability-contingent incentives

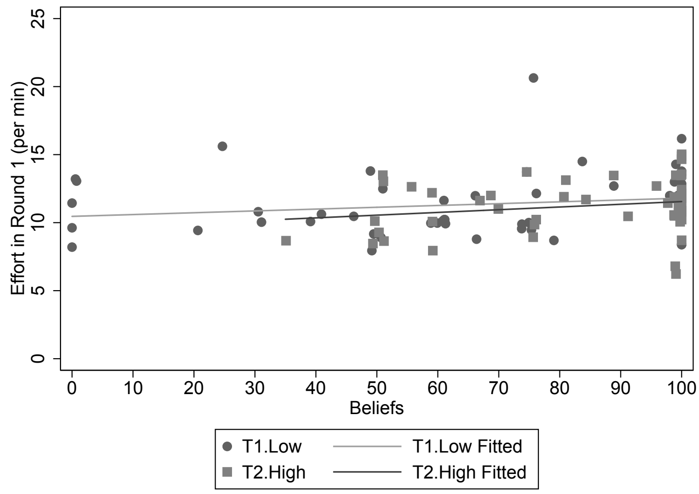

Both these results are interesting in light of the fact that we have already seen that there was a strong causal response of the wage scheme choice to a shift in beliefs. This means we can rule out the hypothesis that beliefs are generally not meaningful for action choices. Overall, this evidence is in line with the explanation that the intrinsic motive to exert effort in the task crowds out the extrinsic motive, resulting in a very low elasticity of effort to changes in expected monetary rewards.[22]

Table 3 reiterates these results by examining the correlates of effort in Round 1 (i.e., under the ability-contingent wage). Columns (1) and (2) confirm that there is no treatment difference in effort choices; columns (3) and (4) provide further evidence that there is no significant relationship between an individuals' belief and her effort choices. In columns (5) and (6), we include baseline effort, which was measured prior to the treatment variation being introduced. We again observe no significant relationship between effort and beliefs.

Result 3. *Effort choices are largely unresponsive to shifts in beliefs, and to the participant's choice of incentive scheme. The low elasticity of effort to change in the expected monetary rewards is indicative of a scenario where intrinsic motivation is high.*

[21]This finding is not surprising considering the recent literature on the unresponsiveness of effort to incentives in real-effort lab experiments (de Araujo et al., 2015; Corgnet et al., 2015; DellaVigna and Pope, 2017; Erkal et al., 2016). Most papers find that there is significant intrinsic motivation to work on the task, regardless of any incentives.

[22]It seems hard to imagine a more boring and tedious task than moving sliders for 20 minutes, but something even less intrinsically rewarding would be necessary to test our hypothesis without any intrinsic motivation.

Table 3: Effort Choice (per minute) Under AC Incentives (Round 1)

	OLS (1)	OLS (2)	OLS (3)	IV (4)	IV (5)	IV (6)
Treatment (HIGH=1)	-0.09	-0.09				
	(0.45)	(0.44)				
Subj Belief			1.11	-0.46	-2.29	-2.30
			(0.79)	(2.19)	(1.79)	(1.79)
Baseline Effort					1.00***	1.00***
					(0.16)	(0.16)
Risk (CE p=50)						0.50
						(1.37)
Constant	11.32***	11.97***	11.09***	12.27***	3.91***	3.68**
	(0.31)	(0.54)	(0.77)	(1.76)	(1.10)	(1.27)
Session Fixed Effects		✓	✓	✓	✓	✓
Observations	100	100	100	100	100	100
Adjusted R^2	-0.010	0.013	0.033			
First Stage F				13.92	13.09	13.03

Notes: (i) Dependent variable: Round 1 effort per minute. (ii) Higher values of risk variable (i.e. certainty equivalent for 50-50 gamble) imply risk loving. (iii) Standard errors in parentheses
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.4 Earnings

We now turn to the effect of increased confidence in earnings. Hypothesis 4 states that increased average confidence will lead to i) weakly lower earnings for low-ability individuals, ii) weakly higher earnings for high ability individuals and iii) result in a higher earnings inequality overall. To evaluate this hypothesis, we split the sample into the Top Half and Bottom Half ability groups and look at the impact of the treatment on each group. Figure 9 shows that Bottom Half individuals' earnings reduced by 40% from € 3.47 to € 2.11, and Top Half earnings almost unchanged by the treatment, at just above € 11. Table 4 shows the same pattern of results, indicating a significant drop in the earnings of the Bottom Half group ($p < 0.05$). The average effect when pooling ability types has a negative sign but is not significant.

Figure 9: Average earnings by ability and treatment

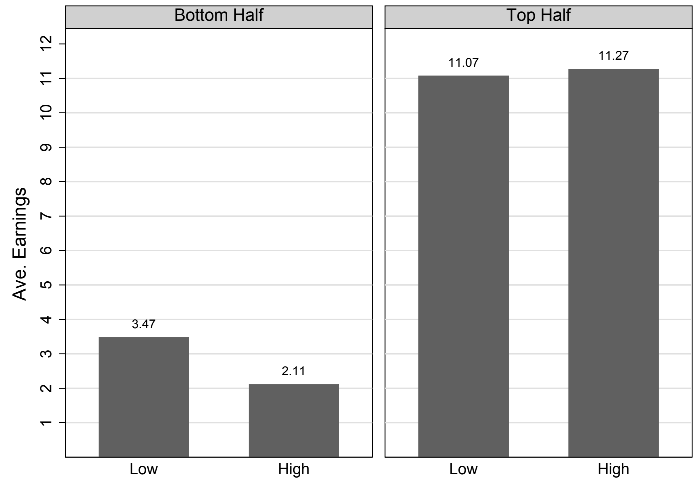

Table 4: Change in Earnings due to Exogenous Belief Shift

	All		Bottom		Top	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (HIGH=1)	-0.59	-0.57	-1.36**	-1.35**	0.19	0.20
	(0.99)	(1.01)	(0.63)	(0.64)	(0.82)	(0.81)
Constant	7.27***	10.39***	3.47***	5.81***	11.07***	11.69***
	(0.70)	(2.32)	(0.45)	(1.55)	(0.58)	(1.87)
Baseline Effort		✓		✓		✓
Risk CE (p=0.5)		✓		✓		✓
Session Fixed Effects		✓		✓		✓
Observations	100	100	50	50	50	50
Adjusted R^2	-0.007	-0.029	0.070	0.062	-0.020	-0.001

Notes: (i) Dependent variable: Main Task Earnings, (ii) Std Errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

This evidence is consistent with the predictions discussed in the theoretical framework section, showing that an increase in confidence leads to a drop in earnings for the low ability individuals who are already earning far less, and thereby moving in the direction of higher overall earnings inequality.[23]

[23] While the GINI coefficient increases from 0.275 in Low to 0.293 in HIGH, and Figure 15 provides suggestive evidence of higher inequality in the HIGH treatment by plotting the earnings histograms of both treatments, a Mann-Whitney ranksum test indicates that there is no significant difference between the earnings distributions in the two treatments. Furthermore, a difference-in-difference estimate of the change in the earnings between the TOP HALF and BOTTOM HALF within each group has a negative point estimate of -1.56 but is not significant at the 10% level. We are underpowered to detect the effect of treatment on inequality.

Result 4. *An increase in confidence leads to low ability individuals earning even less than their already low earnings, while high ability individuals are unaffected. This is suggestive of an increase in inequality with higher confidence, but our data does not permit us to estimate a significant change in inequality.*

One outstanding question is why we observe a relatively large decrease in earnings for the low ability individuals, but hardly any change in earnings for the high ability individuals. We consider discuss this question in more detail in section A.2 in the Appendices. In short, the impact of treatment on the earnings of the Bottom Half, but not the Top Half is driven by the fact that individuals in the Top Half are already highly confident in their ability, and frequently choose the ability-contingent wage, while the Bottom Half appear to hold more malleable beliefs and are willing to be convinced that they are in the Top Half when taking an easier test.

5 Concluding Discussion

In this paper, we have shown how easily confidence about one's relative ability can be shifted and how that confidence can causally influence the choices we make (e.g. the choice of which wage scheme one prefers). In the experiment, this upward shift in confidence results in a reduction in the amount earned by low-ability individuals, who are already earning very little. These individuals are therefore harmed by the confidence boost.

While one should exercise caution when extrapolating findings from a laboratory experiment to the real world, our findings, from our cleanly controlled setting, can be used to provide insights into the mechanisms that might be affecting career choices.

In many professions, talent or ability are essential for success. Thus, holding overconfident beliefs about one's ability may prove extremely costly if it leads to individuals mistakenly entering into these talent-intensive careers. For example, for artists, musicians, football players, and arguably for academic researchers, the reward schedule is highly ability-contingent. While exerting a high degree of effort is necessary in these professions, one can only compensate for a lack of ability through increased effort to a limited extend. Thus, entering into these professions can be a mistake for low-ability individuals. Similarly, entrepreneurs, where ability can be considered a combination of the business idea and the talent of the entrepreneur, have a high risk of toiling for years without any success.

Given how costly it is to choose the wrong profession, and how this choice seems to hinge on one's beliefs about oneself, it is worrying that the results of our experiment, as well as those in the previous hard-easy effect literature, demonstrate that one's beliefs about one's own abilities are highly malleable – particularly those of individuals of lower ability.

In our experiment, participants inference about their placement in the ability distribution is influenced by the perceived difficulty of the task. In a real world setting this could mean that, exposing children or students systematically to challenges that are "too easy" (e.g. "spoon-feeding" them) might result in them holding an artificially high sense of confidence in their own abilities. These inaccurate beliefs may be reinforced and intensified through other societal channels, e.g. by the well-meaning rose-tinted feedback of i) family and friends and ii) the education system. It is therefore possibly of concern that both of these channels seem increasingly tailored towards providing positively skewed feedback. First, in a controlled laboratory experiment, Gneezy et al. (2017) show that even strangers are usually unwilling to give negative feedback to another person face-to-face, even when it is costly to withhold this feedback. Second, looking at observational data, the past two decades have seen an enormous grade inflation at the university level, both in the US and many European countries (Rosovsky and Hartley, 2002). According to the Higher Education Statistics Agency (HESA, 2017), in 1994, only 7 percent of all students received a first class degree in the UK. In 2016, it is now more common to receive a first class degree than a lower second (24 percent vs. 21 percent).

The overarching lesson of this paper is that while there may certainly be benefits to building up confidence, when it comes to professions or tasks which rely heavily on talent that cannot be easily compensated for by increased effort, interventions aimed at increasing confidence might hurt exactly those people they are intended to help. It might be better to construct interventions that help individuals develop their abilities, but provide them with accurate feedback.

References

Barber, B. M. and T. Odean (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics 116*(1), 261–292.

Barron, K. (2020). Belief updating: Does the ‘good-news, bad-news’ asymmetry extend to purely financial domains? *Experimental Economics*.

Barron, K., S. Huck, and P. Jehiel (2019). Everyday econometricians: Selection neglect and overoptimism when learning from others. *Working Paper*.

Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics 117*(3), 871–915.

Bénabou, R. and J. Tirole (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies 70*(3), 489–520.

Bénabou, R. and J. Tirole (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives 30*(3), 141–64.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 69–186. Elsevier.

Benoît, J.-P., J. Dubra, and D. 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.

Bernardo, A. E. and I. Welch (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy 10*(3), 301–330.

Bonin, H., T. Dohmen, A. Falk, D. Huffman, and U. Sunde (2007). Cross-sectional earnings risk and occupational sorting: The role of risk attitudes. *Labour Economics 14*(6), 926–937.

Bordley, R., M. LiCalzi, and L. Tibiletti (2014). A target-based foundation for the hard-easy effect bias. *Working Paper*.

Bruhin, A., L. Santos-Pinto, and D. Staubli (2018). How do beliefs about skill affect risky decisions? *Journal of Economic Behavior & Organization 150*, 350–371.

Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Burks, S. V., J. P. Carpenter, L. Goette, and A. Rustichini (2013). Overconfidence and social signalling. *Review of Economic Studies 80*(3), 949–983.

Burson, K. A., R. P. Larrick, and J. Klayman (2006). Skilled or unskilled, but still unaware of it how perceptions of difficulty drive miscalibration in relative comparisons. *Journal of personality and social psychology 90*(1), 60.

Camerer, C. and D. Lovallo (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review 89*(1), 306–318.

Chen, S. and H. Schildberg-Hörisch (2019). Looking at the bright side: The motivation value of overconfidence. *European Economic Review*.

Cheung, S. L. and L. Johnstone (2017, February). True Overconfidence, Revealed through Actions: An Experiment. IZA Discussion Papers 10545, Institute for the Study of Labor (IZA).

Corgnet, B., R. Hernán-González, and E. Schniter (2015, Jun). Why real leisure really matters: incentive effects on real effort in the laboratory. *Experimental Economics 18*(2), 284–301.

Costa-Gomes, M. A., S. Huck, and G. Weizsäcker (2014). Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect. *Games and Economic Behavior 88*, 298–309.

Costa-Gomes, M. A. and G. Weizsäcker (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies 75*(3), 729–762.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics 22*(2), 369–395.

Dankova, K. and M. Servatka (2019). Gender Robustness of Overconfidence and Excess Entry. *Journal of Economic Psychology 72*, 179–199.

Dargnies, M.-P., R. Hakimov, and D. Kübler (2019). Self-confidence and unraveling in matching markets. *Management Science*.

de Araujo, F. A., E. Carbone, L. Conell-Price, M. W. Dunietz, A. Jaroszewicz, R. Landsman, D. Lamé, L. Vesterlund, S. Wang, and A. J. Wilson (2015). The effect of incentives of real effort: Evidence from the slider task.

DellaVigna, S. and D. Pope (2017). What motivates effort? evidence and expert forecasts. *Review of Economic Studies*.

Dohmen, T. and A. Falk (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. *American Economic Review 101*(2), 556–90.

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. (2020). What you see is all there is. *Quarterly Journal of Economics (forthcoming)*.

Enke, B. and F. Zimmermann (2019). Correlation neglect in belief formation. *Review of Economic Studies*.

Erkal, N., L. Gangadharan, and B. Han Koh (2016). What motivates participants in real-effort tasks: Is it the journey or the destination? *Working Paper*.

Falk, A., A. Becker, T. Dohmen, D. Huffman, and U. Sunde (2016, February). The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences. Working Papers 2016-003, Human Capital and Economic Opportunity Working Group.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Fischer, M. and D. Sliwka (2018). Confidence in knowledge or confidence in the ability to learn: An experiment on the causal effects of beliefs on motivation. *Games and Economic Behavior 111*, 122–142.

Gill, D. and V. Prowse (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review 102*(1), 469–503.

Gneezy, U., C. Gravert, S. Saccardo, and F. Tausch (2017). A must lie situation - avoiding giving negative feedback. *Games and Economic Behavior 102*, 445 – 454.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association 1*(1), 114–125.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly journal of economics 95*(3), 537–557.

Healy, P. J. and D. A. Moore (2007). Bayesian overconfidence.

Heger, S. A. and N. W. Papageorge (2018). We should totally open a restaurant: How optimism and overconfidence affect beliefs. *Journal of Economic Psychology 67*, 177–190.

Heifetz, A., C. Shannon, and Y. Spiegel (2007). The dynamic evolution of preferences. *Economic Theory 32*(2), 251.

HESA (2017). Qualifications obtained.

Huck, S., N. Szech, and L. M. Wenner (2018). More effort with less pay: On information avoidance, optimistic beliefs, and performance. *working paper*.

Jehiel, P. (2018). Investment strategy and selection bias: An equilibrium perspective on overoptimism. *American Economic Review 108*(6), 1582–1597.

Johnson, D. D. and J. H. Fowler (2011). The evolution of overconfidence. *Nature 477*(7364), 317–320.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association 4*(4), 673–707.

Kruger, J. (1999). Lake wobegon be gone! the below-average effec and the egocentric nature of comparative ability judgments. *Journal of personality and social psychology 77*(2), 221.

Kruger, J. and D. Dunning (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of personality and social psychology 77*(6), 1121.

Larrick, R. P., K. A. Burson, and J. B. Soll (2007). Social comparison and confidence: When thinking you are better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes 102*(1), 76 – 94. Special Issue on Social Comparison Processes.

Malmendier, U. and G. Tate (2005). Ceo overconfidence and corporate investment. *The journal of finance 60*(6), 2661–2700.

Möbius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2014). Managing self-confidence. *NBER Working paper*.

Moore, D. A. and P. J. Healy (2008). The trouble with overconfidence. *Psychological review 115*(2), 502.

Moore, D. A. and T. G. Kim (2003). Myopic social prediction and the solo comparison effect. *Journal of personality and social psychology 85*(6), 1121.

Moore, D. A. and D. A. Small (2007). Error and bias in comparative judgment: on being both better and worse than we think we are. *Journal of personality and social psychology 92*(6), 972.

Niederle, M. (2017). A gender agenda: A progress report on competitiveness. *American Economic Review 107*(5), 115–19.

Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? do men compete too much? *The Quarterly Journal of Economics*, 1067–1101.

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of finance 53*(5), 1775–1798.

Pikulina, E., L. Renneboog, and P. Tobler (2016). Do confident individuals generally work harder?

Pikulina, E., L. Renneboog, and P. N. Tobler (2017). Overconfidence and investment: An experimental approach. *Journal of Corporate Finance 43*, 175–192.

Puri, M. and D. T. Robinson (2007). Optimism and economic choice. *Journal of Financial Economics 86*(1), 71 – 99.

Rosovsky, H. and M. Hartley (2002). Evaluation and the academy: Are we doing the right thing.

Schüssler, K. et al. (2018). The influence of overconfidence and competition neglect on entry into competition. Technical report, CRC TRR 190 Rationality and Competition.

Schwardmann, P. and J. J. Van der Weele (2019). Deception and self-deception. *Nature Human Behavior 3*, 1055–1061.

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics 1*(1), 43–62.

Smith, A. (2013, Sep). Estimating the causal effect of beliefs on contributions in repeated public good games. *Experimental Economics 16*(3), 414–425.

Solda, A., C. Ke, L. Page, and W. Von Hippel (2019). Strategically delusional. *Experimental Economics*, 1–28.

van Veldhuizen, R. (2017). Gender differences in tournament choices: Risk preferences, overconfidence or competitiveness? Technical report, Discussion Paper.

A Appendices

A.1 Theoretical Framework: Additional Comments

A.1.1 Further details for Hypothesis 4

We define $F_L(\hat{\pi})$ to be the distribution of subjective beliefs of low ability individuals (a_L), and $F_H(\hat{\pi})$ to be the distribution of subjective beliefs of high ability individuals (a_H). This implies that $F_m(\pi^i)$, where $m \in L, H$, denotes the fraction of individuals with ability a_m that choose the certain piece rate incentives. Similarly, if we consider an upward shift in confidence, $\triangle\hat{\pi}$, then after the shift in beliefs, $F_m(\pi^i - \triangle\hat{\pi})$ is the fraction of individuals with ability a_m that choose the certain piece rate incentives.

Expressions 4 and 5 reflect the intuitive idea that an upward shift in confidence will be harmful to low ability individuals who switch to inappropriate ability-contingent incentives, and beneficial to high ability types who, absent the treatment, would not have chosen the ability-contingent scheme.[24]

Gain in Earnings for High Ability Individuals:

$$[F_H(\pi^i) - F_H(\pi^i - \triangle\hat{\pi})] \cdot (w_H \cdot e^*_{AC} - \bar{w} \cdot e^*_{PR}) \geq 0 \tag{4}$$

Loss of Earnings for Low Ability Individuals:

$$[F_L(\pi^i) - F_L(\pi^i - \triangle\hat{\pi})] \cdot (0 - \bar{w} \cdot e^*_{PR}) \leq 0 \tag{5}$$

The term $F_m^{SWITCH} = [F_m(\pi^i) - F_m(\pi^i - \triangle\hat{\pi})] \geq 0$ denotes the fraction of individuals of ability type a_m who switch from certain piece rate to ability-contingent incentives when there is an upward shift in confidence by $\triangle\hat{\pi}$. The magnitude of the change in earnings will depend on several factors: (i) the number of individuals who switch their incentive scheme choice, F_m^{SWITCH}, in each group, (ii) the change in effort between the two incentive schemes, and (iii) the incremental size of the gaps between the wages. Thus without first

[24]The expressions in the main text assume an interior solution for effort choices. However, if there is a binding constraint on effort choices then expressions 4 and 5 simplify to:

Gain in Earnings for High Ability Individuals

$$[F_H(\pi^i) - F_H(\pi^i - \triangle\hat{\pi})] \cdot (w_H - \bar{w}) \cdot \bar{e}$$

Loss of Earnings for Low Ability Individuals

$$[F_L(\pi^i) - F_L(\pi^i - \triangle\hat{\pi})] \cdot (-\bar{w}) \cdot \bar{e}$$

determining the magnitude of these factors, we cannot unambiguously predict the average change in earnings.

However, if an upwards shift in confidence leads to an increase in earnings for high ability individuals and a decrease for low ability individuals then earnings inequality will increase.

These ideas are summarized in hypothesis 4 in the main text.

A.1.2 Adding Risk Aversion to the Theoretical Framework

In the main text, we used a simple theoretical framework to organize thinking about what we might expect from the experimental data. In that discussion, we assumed that individuals are risk neutral. However most of the intuitions carry through when we relax the risk neutrality requirement. Below, we show that a similar logic applies if we allow for risk aversion over monetary payoffs. More specifically, consider an individual who chooses w and e to maximise:

$$\max_{w\in\{\bar{w},w(a)\}} \max_{e\geq 0} \mathbf{U}(w,e) = s\cdot e + \mathbf{E}_a[u(w\cdot e)] - c(e)$$

where $u'(\cdot) > 0$, $u''(\cdot) < 0$ and $u'(0)+s > c'(0)$. First, we consider the individuals choice between the two incentive schemes.

Choice of Incentive Scheme

As in the main text, let the optimal effort levels under the ability-contingent and certain piece rate incentives be denotes by e^{**}_{AC} and e^{**}_{PR} (assuming an interior solution). Given these optimal effort choices under the two incentive schemes, we have the following condition for the individual choosing the ability-contingent incentives. The individual prefers ability-contingent incentives if:

$$s\cdot e^{**}_{AC} + \hat{\pi}\cdot u(w_H\cdot e^{**}_{AC}) - c(e^{**}_{AC}) \geq s\cdot e^{**}_{PR} + u(\bar{w}\cdot e^{**}_{PR}) - c(e^{**}_{PR})$$

This inequality implies that there is a threshold subjective belief, $\pi^{i,r}$, such that the individual prefers to choose the certain piece rate incentives for lower subjective beliefs, i.e. $\hat{\pi} \leq \pi^{i,r}$, and prefers to choose ability-contingent incentives for larger subjective beliefs, i.e. $\hat{\pi} \geq \pi^{i,r}$. In general, we can specify $\pi^{i,r}$ as follows:

$$\pi^{i,r} := \frac{u(\bar{w}\cdot e^{**}_{PR})}{u(w_H\cdot e^{**}_{AC})} + \frac{s\cdot(e^{**}_{PR} - e^{**}_{AC}) - (c(e^{**}_{PR}) - c(e^{**}_{AC}))}{u(w_H\cdot e^{**}_{AC})}$$

Notice that if there is a binding constraint on effort in the lab, and optimal effort under both incentive schemes is given by $\bar{e}^{**}$, this expressions simplifies to:

$$\pi^{i,r} = \frac{u(\bar{w} \cdot \bar{e}^{**})}{u(w_H \cdot \bar{e}^{**})}$$

The discussion above can be summarised with the following decision rules for incentive choices:

i. The individual chooses the certain piece rate incentives if she has low self-confidence (i.e. beliefs on the interval: $\hat{\pi} \in [0, \pi^{i,r}]$).

ii. The individual chooses the risky ability-contingent incentives if she has high self-confidence (i.e. beliefs on the interval: $\hat{\pi} \in [\pi^{i,r}, 1]$).

Choice of Effort

In the main text, under risk neutrality, the threshold belief at which the individual is indifferent between choosing the ability-contingent and certain piece rate incentives was the same as the threshold at which the individual would choose the same level of effort under both incentives schemes. With risk aversion, these two thresholds for (i) incentive scheme schoice, and (ii) effort choice, are no longer the same.

In this section, we first consider the case where there is an internal solution for the optimal effort level. Later, we return to the case where effort is constrained in the lab (e.g. due to time constraints). Effort choices under the two incentive schemes are obtained by solving the following two equations (assuming interior solutions):

$$\hat{\pi} w_H \cdot u'(w_H \cdot e) = c'(e) - s \text{ gives } e^{**}_{AC} \tag{6}$$

$$\bar{w} \cdot u'(\bar{w} \cdot e) = c'(e) - s \text{ gives } e^{**}_{PR} \tag{7}$$

It is clear that in general e^{**}_{AC} may be either larger or smaller than e^{**}_{PR}, depending on the individual's subjective belief about her own ability, $\hat{\pi}$. In particular, when $\hat{\pi} = 0$, the individual believes that she is low ability for sure, and she chooses more effort under the piece rate incentives ($e^{**}_{AC} < e^{**}_{PR}$). The reverse is true when she believes that she is the high ability type for sure ($\hat{\pi} = 1 \Rightarrow e^{**}_{AC} > e^{**}_{PR}$).

However, since individuals are given a choice of which incentive scheme to choose, it is of interest to know whether all individuals who choose the ability-contingent incentives choose a higher effort level than the individuals who choose the fixed piece rate.

In order to answer this question, firstly, notice that the RHS of equations 6 and 7 is identical. Since the LHS of each equation is decreasing in e, considering the ratio of the

LHS of the two expressions, $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)}$, can help us to assess which optimal effort choice will be higher.

Secondly, since $w_H > \bar{w}$ and $u''(\cdot) < 0$, the ratio $\frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)}$ is increasing as effort increases. This implies that $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)}$ is also increasing in e. One can think of this term as the ratio of the marginal shift in utility due to money of an additional unit of effort, under the two incentive schemes. For ease of exposition, we will refer to the "monetary component" of the utility function as the part that excludes the intrinsic reward for effort, and the cost of effort (i.e. $\mathbf{E}_a[u(w \cdot e)]$).

Thirdly, the optimal effort level under thecertain piece rate incentives doesn't depend on the individual's subjective belief about her ability - all individuals choose the same optimal effort level under certain piece rate, e^{**}_{PR}, irrespective of their belief, $\hat{\pi}$. This means that irrespective of $\hat{\pi}$, we can evaluate the ratio of marginal utilities of money, $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)}$, at the point of optimal effort under the certain piece rate incentives, e^{**}_{PR} (which doesn't depend on $\hat{\pi}$) to determine the threshold value of $\hat{\pi}$ for which the optimal effort choice switches from being higher under certain piece rate incentives to being higher under ability-contingent incentives.

Define $\pi^{e,r}$ to be the value of $\hat{\pi}$ at which the ratio of marginal utilities of money, $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)}$, evaluated at e^{**}_{PR}, is equal to 1. In particular:

$$\pi^{e,r} := \frac{\bar{w}}{w_H} \cdot \frac{u'(\bar{w} \cdot e^{**}_{PR})}{u'(w_H \cdot e^{**}_{PR})} \tag{8}$$

With these definitions for $\pi^{i,r}$ and $\pi^{e,r}$ in hand, we can assess the relationship between the individual's choice of incentive scheme and her effort choice.

Proposition 1. *If the individual chooses the ability-contingent incentives, then she exerts more effort than she would have if she had faced the fixed piece rate incentives. In particular,*

$$\hat{\pi} \geq \pi^{i,r} \Rightarrow \hat{\pi} \geq \pi^{e,r}$$

If she chooses the piece rate incentives, then she chooses more effort under certain piece rate incentives if $\hat{\pi} \leq \pi^{e,r}$, and more effort under ability-contingent incentives if $\hat{\pi} > \pi^{e,r}$.

Essentially, we show that there are three possible intervals for $\hat{\pi}$ that specify incentive choice and effort choice behaviour:

- If $\hat{\pi} \in [0, \pi^{e,r}]$, the individual chooses the certain piece rate incentives and exerts more effort under certain piece rate incentives than ability-contingent incentives.

- If $\hat{\pi} \in (\pi^{e,r}, \pi^{i,r})$, the individual would prefer the certain piece rate incentives, but if she faced the ability-contingent incentives, then she would exert more effort than under the certain piece rate incentives.

- If $\hat{\pi} \in [\pi^{i,r}, 1]$, the individual chooses the ability-contingent incentives and exerts more effort under ability-contingent incentives than certain piece rate incentives.

Case 1: First we consider $\hat{\pi} \in [0, \pi^{e,r})$, or equivalently, $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e^{**}_{PR})}{u'(\bar{w} \cdot e^{**}_{PR})} < 1$. This means that the monetary marginal utility of effort is lower under ability-contingent incentives than under certain piece rate incentives at $e = e^{**}_{PR}$. Therefore, since $\bar{w} \cdot u'(\bar{w} \cdot e^{**}_{PR}) = c'(e^{**}_{PR}) - s$, we have:

$$\hat{\pi} w_H \cdot u'(w_H \cdot e^{**}_{PR}) < \bar{w} \cdot u'(\bar{w} \cdot e^{**}_{PR}) = c'(e^{**}_{PR}) - s \tag{9}$$

But since $u''(\cdot) < 0$ and $c''(\cdot) > 0$, there cannot be an $e > e^{**}_{PR}$ such that $\hat{\pi} w_H \cdot u'(w_H \cdot e) = c'(e) - s$, since the left hand side of equation 6 is decreasing in e and the right hand side is increasing in e. This implies that when $\hat{\pi} \in [0, \pi^{e,r})$, we must have $e^{**}_{AC} < e^{**}_{PR}$. Essentially, this says that if the monetary marginal utility of effort is larger under PR, evaluated at $e = e^{**}_{PR}$, then the individual will choose a higher effort level under certain piece rate incentives than under ability-contingent incentives. Notice, also, that since this ratio of monetary MUs is increasing in e, $\forall e \leq e^{**}_{PR}$, we must have: $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e)}{u'(\bar{w} \cdot e)} < 1$. Therefore, the monetary marginal utility $\forall e \leq e^{**}_{PR}$ is larger under the certain piece rate incentives than ability-contingent incentives, and the non-monetary component of the utility function is identical. Therefore, integrating the marginal utilities over e implies that $\forall e \leq e^{**}_{PR}$:

$$s \cdot e + \hat{\pi} \cdot u(w_H \cdot e) - c(e) \leq s \cdot e + u(\bar{w} \cdot e) - c(e)$$

and since $e = e^{**}_{PR}$ maximises the RHS, and $e^{**}_{AC} < e^{**}_{PR}$ it must be the case that:

$$s \cdot e^{**}_{AC} + \hat{\pi} \cdot u(w_H \cdot e^{**}_{AC}) - c(e^{**}_{AC}) < s \cdot e^{**}_{AC} + u(\bar{w} \cdot e^{**}_{AC}) - c(e^{**}_{AC}) \leq s \cdot e^{**}_{PR} + u(\bar{w} \cdot e^{**}_{PR}) - c(e^{**}_{PR})$$

In summary, if $\hat{\pi} \in [0, \pi^{e,r})$, then: (i) $e^{**}_{AC} < e^{**}_{PR}$, and (ii) the individual chooses the certain piece rate incentives (i.e. $\hat{\pi} \in [0, \pi^{i,r})$).

Case 2: Second, we consider $\hat{\pi} \in [\pi^{e,r}, 1]$, or equivalently, $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e^{**}_{PR})}{u'(\bar{w} \cdot e^{**}_{PR})} \geq 1$. Now, the monetary marginal utility of effort is higher under ability-contingent incentives than under certain piece rate incentives at $e = e^{**}_{PR}$. This implies that $\hat{\pi} w_H \cdot u'(w_H \cdot e^{**}_{PR}) > \bar{w} \cdot u'(\bar{w} \cdot e^{**}_{PR}) = c'(e^{**}_{PR}) - s$, which means that the overall marginal utility of effort under the ability-contingent incentives is positive at $e = e^{**}_{PR}$ (i.e. $\left.\frac{\partial \mathbf{U}_{AC}(w,e)}{\partial e}\right|_{e=e^{**}_{PR}} > 0$. Therefore, the optimal effort level under the ability-contingent incentives is higher than under the piece rate incentives $e_{AC}^{**} > e_{PR}^{**}$. However, it is important to note that $\frac{\hat{\pi} w_H}{\bar{w}} \cdot \frac{u'(w_H \cdot e_{PR}^{**})}{u'(\bar{w} \cdot e_{PR}^{**})} \geq 1$ does not necessarily imply that the individual chooses the ability-contingent incentives. In Case 2, we can have either choice of incentives.

In summary, if $\hat{\pi} \in [\pi^{e,r}, 1]$, then: (i) $e_{AC}^{**} \geq e_{PR}^{**}$, and (ii) the individual chooses the ability-contingent incentives if $\hat{\pi} \in [\pi^{i,r}, 1]$.

Together, Case 1 and Case 2 show that for the interval of subjective beliefs $\hat{\pi} \in [0, \pi^{e,r}]$, the individual would choose certain piece rate incentives, and would choose higher effort under certain piece rate incentives than she would under ability-contingent incentives.

For the interval $\hat{\pi} \in [\pi^{i,r}, 1]$, the individual would choose the ability-contingent incentives, and would choose higher effort under ability-contingent incentives than she would under certain piece rate incentives. And for the interval $\hat{\pi} \in (\pi^{e,r}, \pi^{i,r})$, the individual prefers the certain piece rate incentives, but would choose higher effort if she were to face the ability-contingent incentives.

As a caveat, the discussion above refers to interior solutions for effort choices. Of course, if effort choices in the lab face a binding constraint (e.g. the time limit), then the discussion above does not apply, and effort choices are the same under the two incentives schemes.

Influence on Earnings

The discussion in the main text for the risk-neutral agent essentially maps directly to the case of the risk-averse agent, replacing: (i) π^i with $\pi^{i,r}$, (ii) e_{PR}^* with e_{PR}^{**}, and (iii) e_{AC}^* with e_{AC}^{**}. All the main intuitions remain the same. Therefore, we have:

<u>Gain in Earnings for High Ability Individuals Due to Upward Shift in Confidence</u>

$$[F_H(\pi^{i,r}) - F_H(\pi^{i,r} - \triangle\hat{\pi})] \cdot (w_H \cdot e_{AC}^{**} - \bar{w} \cdot e_{PR}^{**}) \geq 0 \tag{10}$$

<u>Loss of Earnings for Low Ability Individuals Due to Downward Shift in Confidence</u>

$$[F_L(\pi^{i,r}) - F_L(\pi^{i,r} - \triangle\hat{\pi})] \cdot (0 - \bar{w} \cdot e_{PR}^{**}) \leq 0 \tag{11}$$

A.2 Why is there a larger impact on the earnings of Bottom Half individuals?

The discussion in the main text showed that while the increase in confidence reduced the earnings of the Bottom Half individuals, it had no impact on the Top Half individuals. Since effort levels, within ability type, are relatively unresponsive to treatment, the lack of an impact on earnings of the Top Half individuals must be due to fewer of them switching their incentive scheme. In this section, we look at how the beliefs of the different ability types are affected by the treatment.

Firstly, Figure 10 plots the CDFs of beliefs in each of the treatment groups. The figure shows that the entire belief distribution is shifted to the right between the Low and the High treatment groups.[25] However, in order to understand why we only observe a shift in earnings for the Bottom Half ability individuals, we need to consider the belief distributions of each ability type separately (as indicated in equations 4 and 5 by the $F_H(\cdot)$ and $F_L(\cdot)$ functions).

Figure 10: CDF of Beliefs, by Treatment

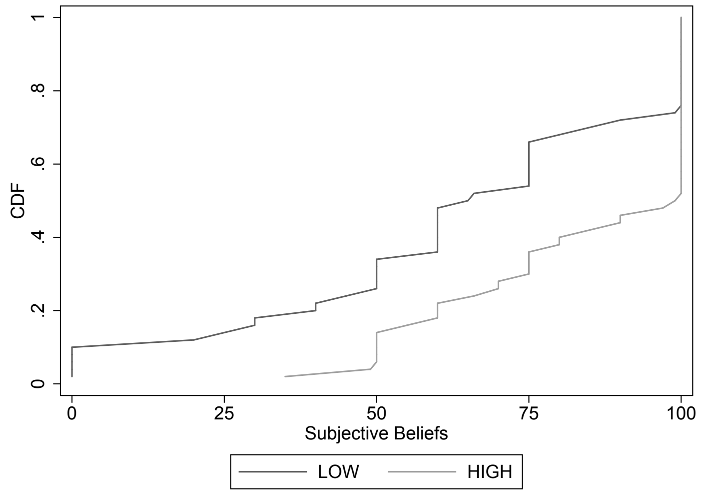

Figure 11 displays these belief CDFs for each ability type separately, comparing treatments. It is immediately apparent from these figures that the majority of the shift in beliefs between treatments is due to the shift in beliefs among individuals in the Bottom Half of the distribution. One reason for this is that on average there is relative overconfidence even in the Low treatment, with the Top Half individuals holding very high beliefs, leaving little room for their beliefs to increase. Essentially, the Top Half individuals are always confident that they are in the Top Half, and the treatment does little to shift this. In contrast, the

[25]A Mann-Whitney ranksum test indicates that the beliefs in the two treatments are drawn from different distributions ($p < 0.01$)

Bottom Half individuals appear to hold more malleable beliefs about their ability. When faced with an easier test, they adjust their level of confidence upwards which leads to costly mistakes in incentive choices.[26]

Figure 11: CDF of Beliefs of Top Half and Bottom Half, by Treatment

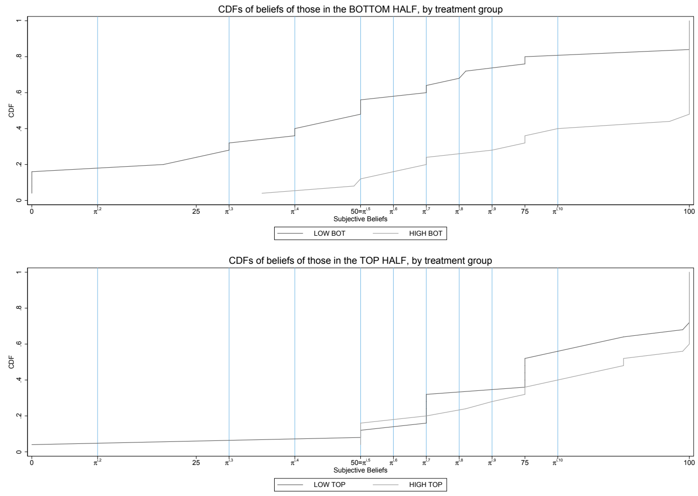

The vertical blue lines in Figure 11 refer to the beliefs thresholds, $\pi^{i,r}$, that indicate the optimal incentive choice for the risk-neutral individual in each round, r. For example, $\pi^{i,5} = 0.5$ is the threshold for round 5. The risk neutral individual should choose ability-contingent incentives in round 5 if her belief is higher than this threshold. Therefore, we can directly read off the fraction of individuals who should choose ability-contingent incentives in each round, under risk neutrality, given their beliefs. This serves to illustrate equations visually 4 and 5, and to demonstrate how the shift in beliefs among the Bottom Half translates into differences in incentive choices, which are less pronounced among Top Half individuals.

In summary, the impact of treatment on the earnings of the Bottom Half, but not the Top Half is driven by the fact that the Top Half are already highly confident in their ability, and choosing ability-contingent incentives, while the Bottom Half hold more malleable

[26]At first glance, this finding might remind the reader of the Dunning-Kruger effect (Kruger and Dunning, 1999). The Dunning-Kruger effect claims that low-ability individuals do not have the means to understand that they are low ability and thus grossly overestimate their relative ability, while high ability individuals can correctly assess their position or are even a bit underconfident in their relative abilities. If the Dunning-Kruger effect would be dominant in our experiment, we should have seen no effect of the treatment or potentially even the opposite. If more difficult tasks make it harder for low ability individuals to estimate their position in a relative ability ranking because they lack the knowledge to evaluate how well they did, then we should have seen *higher* average beliefs of the Bottom Half individuals in the HARD test than in the EASY test treatment. And yet we see the opposite. It seems more likely that the perceived level of difficulty indicates how well they think they performed. A task that *feels* easy creates the belief that one is good at it and especially better than others.

beliefs and are willing to be convinced that they are in the Top Half when taking an easier test.

A.3 Gender heterogeneity in confidence and decision making

There is the large body of evidence documenting that there tends to be a gender gap in confidence in one's own ability (see, e.g., Niederle and Vesterlund (2007), van Veldhuizen (2017) and Niederle (2017)). It is therefore of interest to ask whether we observe this gender gap in our context. We assume that conditional on ability type, men have a higher confidence level than women (i.e. $\forall \hat{\pi} : F_m^M(\hat{\pi}) \leq F_m^W(\hat{\pi})$, where $F_m^G(\cdot)$ denotes the subjective belief CDF for individuals of gender, $G \in \{M, W\}$, and ability type a_m, with $m \in L, H$). This leads to the following hypothesis:

Hypothesis 5. *(Gender Differences): Conditional on ability level, the average man will: (i) hold higher beliefs about his ability, (ii) is more likely to choose the ability-contingent incentives, (iii) will choose higher effort (if non-binding), in comparison to the average woman.*

Below, we test these hypotheses and describe how gender is correlated with our outcomes variables. Our subjects are balanced on gender between treatments and sessions. We show being male is associated with a similar magnitude upward shift in confidence as our treatment effect, and the pattern of outcomes observed in our treatment-control comparison is similar to the pattern of outcomes observed in our male-female comparison (except effort choices).

A.3.1 Gender differences in ability

To eliminate ability differences as an explanation for potential gender differences in beliefs and payment scheme choices, we selected an ability task that was not gendered and with no evidence of gender effects in previous experiments. Table 6 shows that men and women are almost identical in their average scores (8.98 and 8.92), and their probability of being in the Top Half of their group (0.51 and 0.49). Neither difference is statistically significant. The similarity in the performance of men and women implies that ability should not explain any observed differences in beliefs between men and women.

A.3.2 Gender differences in beliefs and incentive choices

In line with the literature documenting the gender confidences gap, Figure 12 shows that in both our treatment conditions women state lower average beliefs about being in the Top Half. Table 6 shows that the gender-confidence gap is on average 11 percentage points and is significant at the 5% level.[27]

[27] If we test for the gender-confidence gap within each treatment separately, the ttest has a p-vale of 0.054 for the High treatment and 0.29 for the Low treatment, but our sample size is insufficient for robust analysis at this level of disaggregation.

Figure 12: Average beliefs about being in the Top Half by gender

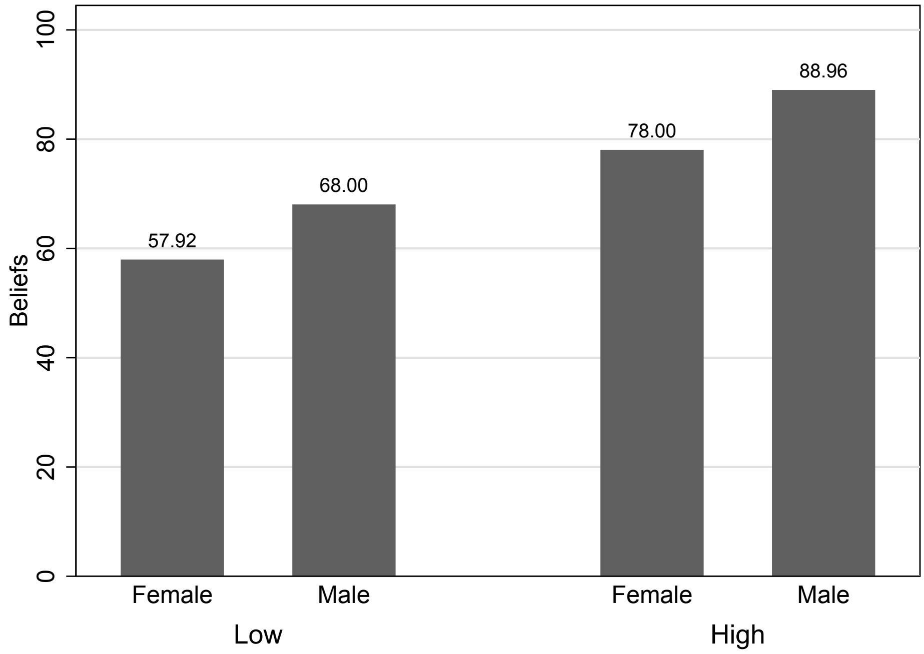

Regarding incentives choices, following their beliefs, men choose the ability-contingent incentives more often than women, despite not being more likely to be in the Top Half and benefit from these incentives. Figure 13 shows the choices of payment scheme separately by gender and treatment. The pattern is similar to the gender-treatment pattern observed for beliefs in 12. Table 6 shows that the average woman chooses ability-contingent incentives 50% of the time, while the average man chooses ability-contingent incentives 68% of the time ($p < 0.05$). Table 8 shows that men choose ability-contingent incentives more often in every individual round, with a significant difference in six of the nine rounds. The table is suggestive of a larger gender gap in choices for later rounds.

More specifically, there appear to be some men who are very unwilling to switch away from the ability-contingent incentives to the certain piece rate incentives. A striking illustration of this is that even in round 10 when the certain piece rate incentive piece rate is €0.8, approximately 50% of men prefer to gamble on being in the Top Half and getting a piece rate of €1, and earning €0 if they are wrong.

Figure 13: Propensity to Choose Ability-Contingent Incentives

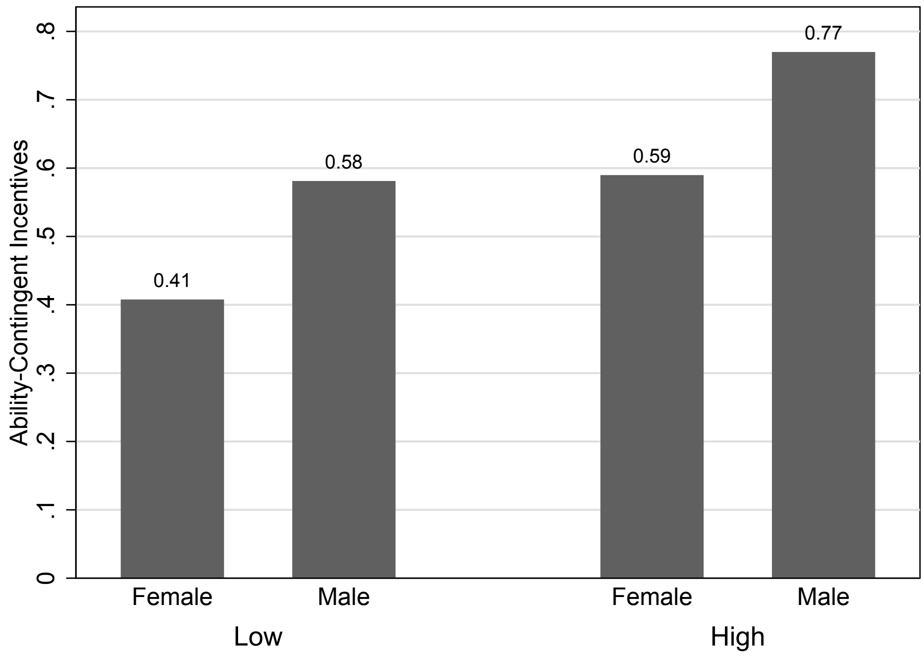

A.3.3 Gender differences in earnings

In section 4.4 we showed that an increase in average confidence had no significant effect on average earnings, but hurts the low-ability individuals. In this section, we consider whether the gender confidence gap translates into gender differences in earnings. Table 6 shows that on average, there is no significant difference between the earnings of men (€ 6.95) and women (€ 7.10).

However, Table 6 also shows that men are exerting significantly higher effort than women ($p < 0.01$). Since earnings are determined by both effort and incentive choices, we need to consider the contributions of effort and choices to earnings separately to understand the mechanisms. To do this, we remove the role played by effort choices by constructing a variable that stands for the *earnings per unit of effort*. Essentially, at the individual level, we calculate how much an individual earns for each set of 20 sliders she completes. This variable allows us to measure how optimal the incentive choice decisions are given an individual's level of effort.

Figure 14: Earnings per unit effort, by gender and treatment

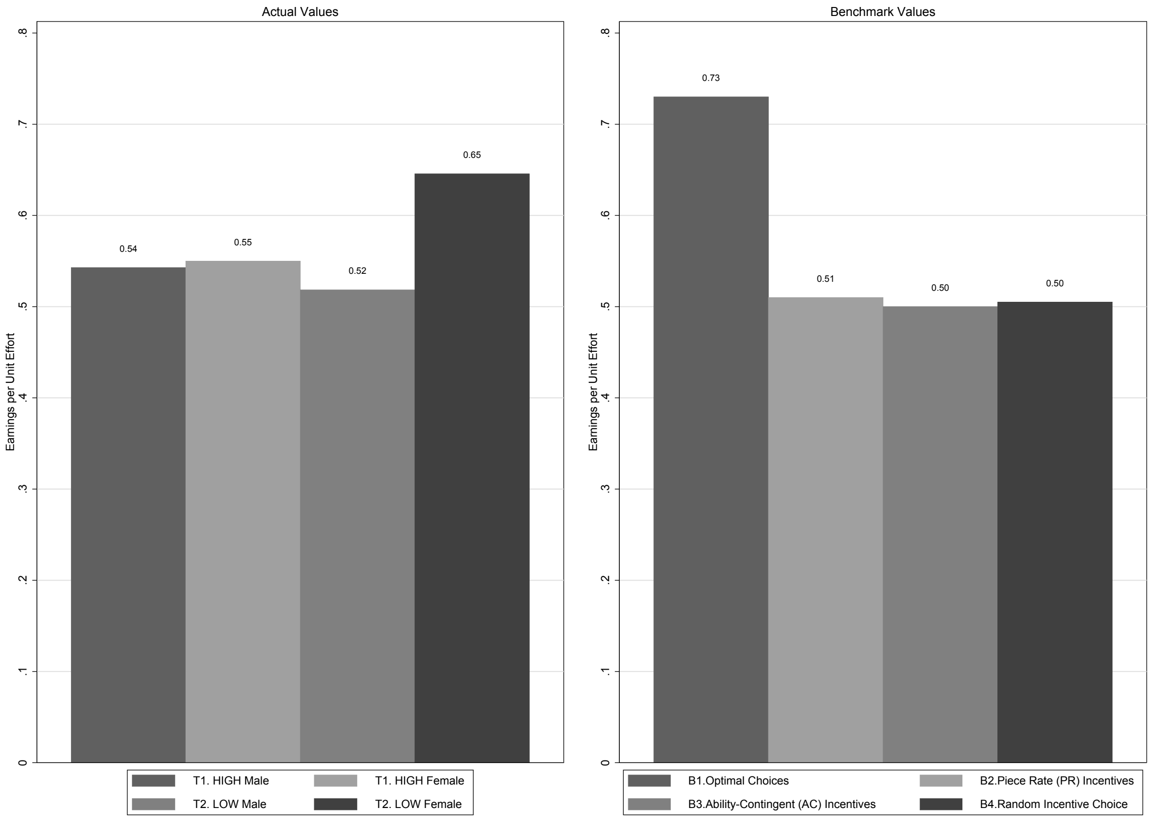

The left panel of Figure 14 reports the average value of this "earnings per unit of effort" variable for each gender-treatment group. The interesting result here is that women in the Low confidence treatment group are choosing the correct incentives far more often than any other group. The intrinsically lower beliefs of women together with the exogenously triggered lower beliefs by the difficult task lead to more efficient choices by women in this group. To provide a benchmark for how well these women do, in the right-hand panel we plot four benchmark earnings per unit effort possibilities: (1) average earnings for a group who all choose completely optimally, (2) average earnings for a group who all always choose the certain piece rate incentives, (3) average earnings for a group who all always choose the ability-contingent incentives, (4) average earnings for a group who always choose randomly between incentive schemes.

Except women in the Low treatment, participants are on average performing rather poorly in terms of incentive choices, scoring between €0.52 and €0.55 per unit of effort. This amount is only a little more than they would earn if they chose completely randomly (€0.50). However, women in the Low treatment earn €0.65 per unit effort, which is closer to the first best value of €0.73 than random choice.

A.3.4 Summary of gender heterogeneity in this context

In our results, there is no evidence that risk aversion had a mediating effect for either gender. This result adds to the evidence that gender differences in payment scheme choices are significantly affected by beliefs about relative ability rather than just competitive preferences or risk aversion (van Veldhuizen, 2017). Since it is not a competition, anxiety, fear or thrill should not play a role during the real effort task. There are no externalities imposed on the other participants when an individual exerts a high effort, so other regarding preferences are not relevant here. Women chose the ability-contingent piece rate when they believe that they are in the upper half of the group. In our design, the most successful subjects were the on average less confident women because they were better at selecting into the top and bottom group according to their actual ability.

A.4 Additional Tables

Table 5: Comparison of means by treatment

	LOW	HIGH
Male (=1)	0.48	0.54
	(0.51)	(0.50)
Age	25.20	24.62
	(6.55)	(3.88)
Effort (Baseline Task, per min)	9.55	9.91
	(1.80)	(1.70)
Risk (CE p=50)	0.48	0.49
	(0.16)	(0.14)
Treatment Variables		
Treatment (High =1)	0	1
	(0.00)	(0.00)
Ability Score	6.92	10.90***
	(2.79)	(1.20)
Ability Top Half (=1)	0.50	0.50
	(0.51)	(0.51)
Outcome Variables		
Subj Belief (%)	64.20	83.92***
	(31.67)	(20.16)
AC Incentive Choice (Frac)	0.50	0.69***
	(0.37)	(0.32)
Effort (Main Task, per min)	12.15	12.18
	(2.29)	(2.04)
Earnings (Main Task)	7.27	6.69
	(4.78)	(5.14)
Risk (CE p=Subj Belief)	0.56	0.71***
	(0.23)	(0.22)
N	50	50

Notes: (i) standard deviations in parentheses, t-tests: *=10%, **=5%, ***=1%

Table 6: Comparison of means by gender

	Female	Male
Male (=1)	0	1
	(0.00)	(0.00)
Age	23.96	25.16
	(3.61)	(4.37)
Effort (Baseline Task, per min)	9.46	10.11*
	(1.62)	(1.79)
Risk (CE p=50)	0.48	0.48
	(0.16)	(0.13)
Treatment Variables		
Treatment (High =1)	0.49	0.55
	(0.51)	(0.50)
Ability Score	8.98	8.92
	(2.71)	(3.23)
Ability Top Half (=1)	0.51	0.49
	(0.51)	(0.51)
Outcome Variables		
Subj Belief (%)	67.74	79.55**
	(29.55)	(26.38)
AC Incentive Choice (Frac)	0.50	0.68**
	(0.36)	(0.34)
Effort (Main Task, per min)	11.65	12.79***
	(1.89)	(2.17)
Earnings (Main Task)	7.10	6.95
	(4.15)	(5.62)
Risk (CE p=Subj Belief)	0.56	0.70***
	(0.24)	(0.22)
N	47	49

Notes: (i) standard deviations in parentheses, t-tests: *=10%, **=5%, ***=1% (ii) The 3 missing women, and 1 missing man answered “Prefer not to answer” to the gender question in the survey.

Table 7: Comparison of mean incentive choices by treatment

	LOW	HIGH
Treatment (High =1)	0	1
	(0.00)	(0.00)
Subj Belief (%)	64.20	83.92***
	(31.67)	(20.16)
Ability Top Half (=1)	0.50	0.50
	(0.51)	(0.51)
<u>Incentive Choices</u>		
AC Incentive Choice (All Rounds)	0.50	0.69***
	(0.37)	(0.32)
AC Incentives Round 1 (=1)	1	1
	(0.00)	(0.00)
AC Incentives Round 2 (=1)	0.76	0.90*
	(0.43)	(0.30)
AC Incentives Round 3 (=1)	0.76	0.84
	(0.43)	(0.37)
AC Incentives Round 4 (=1)	0.64	0.84**
	(0.48)	(0.37)
AC Incentives Round 5 (=1)	0.52	0.74**
	(0.50)	(0.44)
AC Incentives Round 6 (=1)	0.40	0.70***
	(0.49)	(0.46)
AC Incentives Round 7 (=1)	0.36	0.58**
	(0.48)	(0.50)
AC Incentives Round 8 (=1)	0.38	0.58**
	(0.49)	(0.50)
AC Incentives Round 9 (=1)	0.38	0.54
	(0.49)	(0.50)
AC Incentives Round 10 (=1)	0.32	0.46
	(0.47)	(0.50)
N	50	50

Notes: (i) Ability Top Half: Reflects fraction of individuals in top half. Equals 0.5 by construction. (ii) AC Incentive Choice (All Rounds): individual level variable, averaged across all an individual's choices. (iii) standard deviations in parentheses, t-tests: *=10%, **=5%, ***=1%

Table 8: Comparison of mean incentive choices by gender

	Female	Male
Treatment (High =1)	0.49	0.55
	(0.51)	(0.50)
Subj Belief (%)	67.74	79.55**
	(29.55)	(26.38)
Ability Top Half (=1)	0.51	0.49
	(0.51)	(0.51)
Incentive Choices		
AC Incentive Choice (All Rounds)	0.50	0.68**
	(0.36)	(0.34)
AC Incentives Round 1 (=1)	1	1
	(0.00)	(0.00)
AC Incentives Round 2 (=1)	0.74	0.92**
	(0.44)	(0.28)
AC Incentives Round 3 (=1)	0.70	0.88**
	(0.46)	(0.33)
AC Incentives Round 4 (=1)	0.70	0.80
	(0.46)	(0.41)
AC Incentives Round 5 (=1)	0.60	0.67
	(0.50)	(0.47)
AC Incentives Round 6 (=1)	0.40	0.67***
	(0.50)	(0.47)
AC Incentives Round 7 (=1)	0.34	0.59**
	(0.48)	(0.50)
AC Incentives Round 8 (=1)	0.34	0.59**
	(0.48)	(0.50)
AC Incentives Round 9 (=1)	0.36	0.55*
	(0.49)	(0.50)
AC Incentives Round 10 (=1)	0.28	0.49**
	(0.45)	(0.51)
N	47	49

Notes: (i) Ability Top Half: Reflects fraction of individuals in top half. (ii) AC Incentive Choice (All Rounds): individual level variable, averaged across all an individual's choices. (iii) standard deviations in parentheses, t-tests: *=10%, **=5%, ***=1%

A.5 Additional Figures

Figure 15: Distribution of Earnings between Treatments

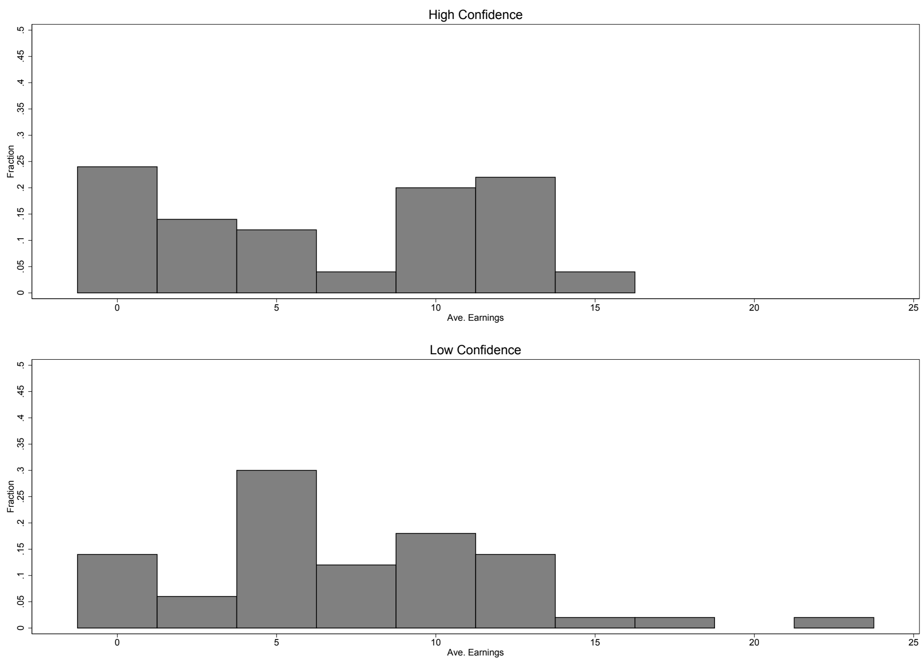

Figure 16: Propensity to Score in the Top Half

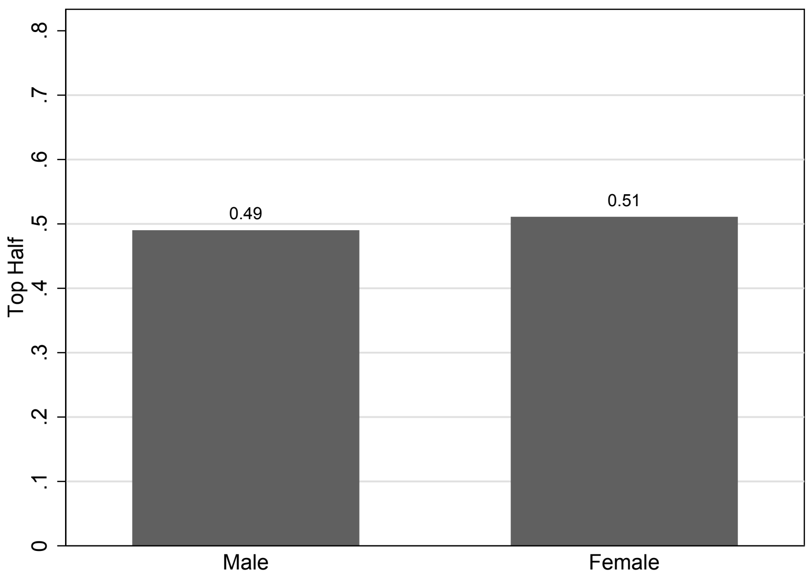

Discussion Papers of the Research Area Markets and Choice 2018

Research Unit: **Economics of Change**

Kai Barron, Christina Gravert SP II 2018-301
Confidence and career choices: An experiment

All discussion papers are downloadable:
http://www.wzb.eu/en/publications/discussion-papers/markets-and-choice