I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 14835

Reference Points and the Tradeoff between Risk and Incentives

Thomas Dohmen
Arjan Non
Tom Stolp

NOVEMBER 2021

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 14835

Reference Points and the Tradeoff between Risk and Incentives

Thomas Dohmen
IZA, University of Bonn and Maastricht University

Arjan Non
Erasmus University Rotterdam

Tom Stolp
Maastricht University

NOVEMBER 2021

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Reference Points and the Tradeoff between Risk and Incentives*

We conduct laboratory experiments to investigate basic predictions of principal-agent theory about the choice of piece rate contracts in the presence of output risk, and provide novel insights that reference dependent preferences affect the tradeoff between risk and incentives. Subjects in our experiments choose their compensation for performing a real-effort task from a menu of linear piece rate and fixed payment combinations. As classical principal-agent models predict, more risk averse individuals choose lower piece rates. However, in contrast to those predictions, we find that low-productivity risk averse workers choose higher piece rates when the riskiness of the environment increases. We hypothesize that reference points affect piece rate choice in risky environments, such that individuals whose expected earnings would exceed (fall below) the reference point in a risk-free environment behave risk averse (seeking) in risky environments. In a second experiment, we exogenously manipulate reference points and confirm this hypothesis.

JEL Classification: D81, D91, M52

Keywords: incentive, piece-rate, risk, reference point, laboratory experiment

Corresponding author:
Thomas Dohmen
University of Bonn
Lennestrasse 43
53113 Bonn
Germany
E-mail: t.dohmen@uni-bonn.de

* The authors gratefully acknowledge financial support from NWO Vidi grant 452-10-006. Dohmen and Non gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project A01). Dohmen also acknowledges funding by DFG under Germany's Excellence Strategy – EXC 2126/1– 390838866. We thank seminar audiences at the University of Bonn and Maastricht; conference participants of the KVS New Paper sessions 2019 in The Hague and IAAEU Workshop on Labour Economics 2021 in Trier; Matthias Kräkel, Dirk Schindler and three anonymous referees for valuable comments and suggestions.

1. Introduction

A tenet of principal-agent theory is the trade-off between risk and incentives. In principal-agent relationships the principal can write a contract that specifies a piece rate for observable output to solve the moral hazard problem that arises when agents' effort is non-contractible but costly to the agent. Typically, observed output does not depend on effort alone but also on risky factors that are beyond the control of the agent. As risk averse agents demand a compensation for exposure to risk, principals optimally offer them a lower piece rate and a higher base payment to reduce risk exposure, given unlimited liability (Budde and Kräkel, 2010). Since the risk premium that risk averse agents demand rises as output risk increases, the optimal piece rate is lower the higher output risk is. This induces a negative relationship between risk and incentives in these models.

We conduct two laboratory experiments to probe the predictions of standard contract models. Our findings point to the importance of the reflection effect (Kahneman and Tversky, 1979) in contract choice and indicate that expanding standard incentive theory to include reference points can better explain the tradeoff between risk and incentives. In our first laboratory experiment, we test the two main predictions of standard theory concerning agents' choice of piece rate contracts, namely (1) whether agents prefer piece rate contracts with weaker incentives, i.e. with a higher base payment but a lower piece rate, the more risk averse they are, and (2) whether they prefer weaker incentives as output risk increases.

A set of earlier studies based on experimental data documents that more risk averse workers prefer weaker incentives (e.g. Cadsby et al., 2007; Dohmen and Falk, 2011; Fehrenbacher and Pedell, 2012). However, these studies do not assess how changes in exogenous risk affect contract choice. Recent experimental studies by Corgnet and Hernán-González (2019) and Chowdhury and Karakostas (2020) manipulate output risk by introducing additive random shocks to production in order to examine whether principals respond to changes in output risk by adjusting the fixed and variable components of the contracts they offer to agents. Consistent with classical principal-agent models, both studies find that principals offer lower variable payment when output risk increases, particularly when principals believe that the agent is risk averse. Closely related to these latter two studies, we also exogenously manipulate output risk in our experiment. A crucial difference is that we investigate agents' preferences rather than principals' contract offers.

In our experiment, individuals can choose from a menu of linear piece-rate contracts to decide how they want to be rewarded for performance on a real-effort task. Contracts consist of a fixed component and a payment per unit of output. Mimicking the principal's contract design, we introduce a trade-off between the base payment and the piece rate: contracts that specify a higher base payment offer a lower payment per unit of output. Output is determined by individuals' production and a risky component that is operationalized as a randomly chosen percentage that is added to or subtracted from their production. This percentage is drawn from a known discrete uniform distribution that is symmetric around zero with incremental steps of 10%. We implement three treatments that differ in the range of the aforementioned distribution. In the high risk treatment, production shocks range from -100% to +100% with increments of 10%. This implies that in the best case production is doubled, while it is nullified in the worst case. In the low risk treatment, production shocks range between -40% and +40%; whereas in the no risk treatment no additional risk is added. We elicit subjects' risk preferences from their choices in incentivized lotteries and their self-assessed willingness to take risks.

In line with the above cited experimental studies on contract choice we find that more risk averse individuals prefer weaker incentives. A new insight of our study, however, is that individuals do not choose lower piece rates on average in a riskier environment. Remarkably, the average treatment effect masks important heterogeneity: In line with standard models, highly productive individuals choose lower piece rates when output risk is high, but contrary to the predictions of classic principal-agent models, low-productivity individuals choose higher piece rates in a more risky environment. This pattern of behavior is consistent with the idea that individuals value reaching an earnings target. Such a target might be socially determined (e.g., earn at least what the average person earns) or based on earnings expectations (e.g., earn at least what participants typically earn in a one hour lab experiment). In our setting, low-productivity workers might always fall short of this target in the absence of production risk. They might attain or approach their target, however, when they give leverage to the output shock by choosing a high piece rate. In case of a sufficiently large output shock, this high piece rate might enable them to reach the target despite low productivity. High productivity workers, who expect to earn more than their earnings target, might instead want to insure themselves against bad output shocks by reducing exposure to shocks (i.e. by choosing a lower piece rate). In other words, while low-productivity workers feel that they have little to lose and much to gain by choosing a high piece rate in a risky environment, high-productivity workers on the contrary perceive that they have little to gain and much to lose by doing so.[1]

In essence, such an explanation holds that reference points drive behavior. In order to investigate the role of reference points, we designed a second experiment that builds on the first experiment. Individuals perform the same real-effort task as in the first experiment and choose how they want to be rewarded from the same menu of linear piece-rate contracts. The second experiment features a 2x2 between-subjects design: we manipulate risk (no risk and high risk treatments) and reference points (low and high reference point treatments). We manipulate reference points by introducing salient counterfactual earnings: Either individuals are rewarded according to the piece rate they select, as in the first experiment, or they receive a fixed payment, of which the amount differs by treatment. We posit that potentially receiving a high fixed payment causes subjects to have a higher reference point than when this amount is low. The events of selecting a payoff relevant piece rate or receiving the fixed payment are equally likely and randomly determined by the computer. Individuals learn about the outcome directly before they choose their piece rate. This procedure confronts individuals with an amount they *could have earned*, which is likely to be regarded as a salient reference point. We hypothesize that the treatment effect of a high risk environment depends on the reference point. If the reference point is low, then individuals perceive their earnings as gains and select lower piece rates in response to risk. If the reference point is high, then individuals consider their earnings as a loss relative to the reference point. Therefore, they respond to risk in the environment by choosing higher piece rates.

In line with our hypotheses, individuals in the low reference point treatment choose lower piece rates when risk is introduced, while individuals in the high reference point treatment tend to choose higher, albeit not

[1] Gill et al. (2019) show how one's rank in the distribution of performance affects behavior. In comparison to middle-ranking individuals, those who rank first or last in the performance distribution will exert more effort during the next performance evaluation. First-ranking individuals may want to assure their first place by subsequently exerting greater effort while last-ranking individuals exert greater effort to assure they don't rank last again. Such an explanation corresponds with our hypothesis: high-productivity individuals choose lower piece rates to "remain on top" while low-productivity individuals select higher piece rates with the intent to climb up the ladder. Note, however, that we do not provide relative performance feedback.

significantly higher, piece rates in response to risk.[2] Taken together our findings indicate that reference points affect agents' preferences over piece rate contracts in the presence of output risk. If individuals perceive their expected earnings as gains, then the predictions of the standard principal-agent model apply. In particular, individuals react to greater risk by selecting lower piece rates. If individuals expect to earn less than the reference point, then they become more risk tolerant and select higher piece rates in the presence of risk. This finding calls for an extension of the principal-agent model to take reference dependence into account.[3]

Our finding concerning the role of the reflection effect for contract choice also offers another potential explanation for why the empirical evidence on the relationship between risk and incentives has remained inconclusive, even almost two decades after Prendergast (2002) concluded that "[] empirical research has not shown a convincing relationship between pay for performance and observed measures of uncertainty." (Prendergast, 2002, p. 1071-1072). One prominent explanation offered in the literature is that risk averse workers sort into relatively safe environments, obscuring the relationship between risk and incentives. In that sense, this study complements field studies on the trade-off between risk and incentives that find that increased risk is associated with weaker incentives after correcting for endogenous matching. For example, Ackerberg and Botticini (2002) find evidence that rent contracts of farmers in medieval Tuscany depend on the type of crop grown, i.e. the risk of crop failure. In particular, after taking into account that more risk-averse farmers prefer to grow less risky crops, weaker incentives are observed when the risk of crop failure is high. Likewise, Hilt (2008) also finds evidence for a negative tradeoff between risk and incentives in the US whaling industry, taking into account that more risk averse sailors sort into less risky whaling voyages.[4] The relationship between risk and incentives might

[2] Choosing a compensation scheme with steep incentives clearly implies risk-taking behavior. In that sense, our findings are consistent with other studies on the impact of reference points on risk-taking behavior, such as Linde and Sonnemans (2012) and Schwerter (2013). In contrast to these studies, our experiment is designed to study implications for contract theory.

[3] Other studies have investigated the effect of reference points on labor market outcomes, such as effort provision and performance (Abeler et al., 2011, Bartling et al., 2015, De Quidt et al., 2017, Ockenfels et al., 2015), labor supply (Fehr and Goette, 2007), and perceived attractiveness of labor contracts (De Quidt, 2018). These studies do not point at the importance of reference points for sorting into incentive schemes in risky environments.

[4] Not controlling for the riskiness of the environment, we would observe that risk averse workers are likely to receive stronger incentives, because they tend to sort into relatively safe environments. Prasad and Salmon (2013) conduct a lab experiment with stated-effort to shed light on endogenous matching. In their setup, principals can contract two tasks to the agent by making different offers (in terms of a fixed payment and a bonus payment). The tasks differ in earnings risk to the principal. They find that in comparison to the safe task, bonus payments are higher and fixed payments are lower when the task is risky. Moreover, more risk-averse workers are less likely to choose the risky task.

also be obscured in naturally occurring data because other factors that correlate with the riskiness of the environment drive the choice for output-based incentives, as discussed by Prendergast (2002).[5] Our findings are obtained in a more controlled lab environment, which has the advantage that we can rule out endogenous matching or third factors. In addition we can clearly define and manipulate output risk as well as the reference points, in order to identify how references points affect the nexus of risk aversion, risk and incentives.

The paper is organized as follows. In the next section we describe the design and hypotheses of the first experiment. In section 3 we present the results. In section 4 we discuss possible interpretations of the findings, and how they motivate us to conduct the second experiment. Section 5 describes the setup of the second experiment, section 6 presents the results. Section 7 concludes.

2. Design of first experiment

2.1. Hypotheses

A common assumption in principal-agent theory is that agents are risk averse and that the relation between agents' effort and their output is stochastic, i.e., the environment is risky in the sense that agents' performance is prone to random shocks. Incentivizing agents by paying them according to their performance therefore implies exposing agents to income risk. Risk averse agents perceive this as a cost, for which they demand a risk premium as compensation. The stronger the link between pay and performance and the larger the variance of random shocks, the larger the required risk premium. Principal-agent theory therefore predicts the following:

Hypothesis 1.1: More risk averse individuals choose lower piece rates when the environment is risky.

Hypothesis 1.2: Higher risk in the environment induces risk averse individuals to choose lower piece rates.

To provide a more formal derivation of these hypotheses, assume that individuals have the following utility function:

[5] For example, Prendergast (2002) explains that delegation of responsibility to employees is more likely when it is not clear how to best solve a task, i.e., when the environment is more risky. At the same time output-based pay is more likely when workers have more discretion in order to motivate them to perform the most productive tasks. In practice, we might therefore observe high-powered output-based incentives in high-risk environments.

$$U = u(w - c(e))$$

where *w* denotes monetary compensation and *c(e)* captures effort costs. Individuals are risk averse, i.e. $u'(w) > 0$, $u''(w) < 0$, and costs of effort are convex, i.e. $c'(e) > 0$, $c''(e) > 0$. Monetary compensation *w* consists of a base salary *s* and an amount *b* per unit of output. Output consists of two components: individuals' productivity and a random error term ε, which is normally distributed around zero with variance σ^2 and enters the production function multiplicatively. For simplicity, assume that productivity is simply equal to effort *e*, hence output is $e(1 + \varepsilon)$. To summarize, individuals' compensation is given by:

$$w = s + be(1 + \varepsilon)$$

Individuals maximize their expected utility by choosing effort and combination of base salary *s* and piece rate *b* from a given choice set. To ensure an interior solution, the choice set is such that a higher piece rate comes at the cost of a lower base salary, and the higher the piece rate the stronger the reduction in base salary: $\frac{ds}{db} < 0$ and $\frac{ds}{d^2b} < 0$. We show in Appendix A that the first-order condition that describes the optimal piece rate choice can be written as:

$$\frac{dE(U)}{db} = e + \frac{ds}{db} - rbe^2\sigma^2 = 0,$$

where $r = -\frac{E[u''(\cdot)]}{E[u'(\cdot)]} > 0$ is a measure of global risk aversion. In the absence of risk, the optimal piece rate balances the additional revenue from higher variable payment and the corresponding reduction in base salary. Clearly, when risk is introduced ($\sigma^2 > 0$) and agents are risk averse ($r > 0$), it is optimal to choose a lower piece rate in order to reduce the exposure to risk. In appendix A, assuming r is constant, we show formally that the optimal piece rate is decreasing in risk aversion (r) and the variance of random shocks (σ^2).[6]

This first order condition above also suggests an interaction effect between risk (σ^2) and risk attitude (r). It is important to understand what this implies. First, it is important to note that hypothesis 1.2 only holds for risk averse individuals: we would not expect that risk neutral or risk seeking individuals choose lower piece rates when risk in the environment increases. We will therefore investigate whether the effect of risk in the environment differs

[6] This prediction also follows from the widely used exponential utility model with additive shocks, as described in textbooks such as Cahuc and Zylberberg (2004) and Bolton and Dewatripont (2005, Ch. 4.2.). The model presented here implies that the same predictions hold if we allow for multiplicative shocks, as in our experimental setup.

by risk attitude. Second, it is easily verified that risk attitudes are not relevant when there is no risk in the environment ($\sigma^2 = 0$), which yields the prediction that Hypothesis 1.1. does not hold in the No-risk treatment. However, testing this experimentally is not straightforward as individuals may feel unsure about how much they will be able to produce, and consequently still perceive some risk in the No-Risk treatment. More generally, the interaction between risk (σ^2) and risk attitude (r) does not imply that we should expect a stronger relation between risk attitude and piece rate choice in more risky environments. The reason is that, theoretically, risk averse individuals choose a low piece rate when the environment is highly risky, so changes in risk aversion also have a small effect in absolute value. This offsets the intuitive effect that differences in risk attitudes are more relevant when the environment is risky. In the remainder of this section, we will describe the setup of the laboratory experiment we designed to test those hypotheses.

2.2. Description of the experiment

The outcome of interest of the experiment is subjects' decision on how they want to be rewarded for their output in a real-effort task. Figure 1 provides an overview of the stages of the experiment. The task consists of typing sequences of numbers 1 to 9 in ascending order, where each number entered has to be confirmed by a mouse-click. Each completed sequence of numbers 1-9 represents one unit of production. After reading the instructions and completing a sample sequence, subjects perform the task for 5 minutes. This stage is not financially incentivized, as doing so could influence subjects' choice of the piece rate contract in the subsequent stage, e.g. due to status quo bias. However, we inform subjects that "it is in their best interest to complete as many sequences as possible, because it will facilitate their decision-making in the experiment". As such, subjects are motivated to do their best as they will be better informed during the experiment. We provide feedback on their productivity at the end of this stage. This is essential information later in the experiment when subjects choose their preferred payment scheme. In the remainder of the paper, we refer to this period as the productivity stage. Subjects' productivity in this stage serves as an ability indicator. Also, by tracking the development of productivity during this stage we can verify that learning or fatigue effects are absent. As such, we minimize the inherent riskiness of the task so that we can effectively manipulate risk in the environment. Details of the task can be found in Online Appendix C.

After the productivity stage, subjects select their preferred compensation scheme. Subjects are informed that they will be paid for performing the *exact same* task again for 5 minutes. We refer to this stage as the work stage. They can choose between various contracts that each consist of a fixed amount and a piece rate per unit of output. Contracts with a higher piece rate offer a lower fixed amount. Figure 2 displays the menu of contracts that was offered. We mimic risk in the environment by adding or subtracting a fraction of the output. The percentage added or subtracted is randomly drawn from a discrete uniform distribution with zero mean.[7] That is, when s denotes the base salary and b the piece rate, individual's payoff w is described by

$$w = s + bq \ ,$$

where output $q = x(1+\epsilon)$ is defined as individuals' production x multiplied by the random shock 1+ ε. We inform subjects about the distribution of the random variable before they choose their compensation scheme.

We assign subjects randomly to one of three treatments, which differ with respect to the variance of the distribution of output shocks. Subjects choose their preferred payment scheme in one treatment condition only, i.e. we use a between-subjects design.[8] In the low risk treatment, at most 40% is added to or subtracted from individuals' production in the real-effort task. Specifically, the distribution of ε ranges from minus 40% to plus 40% in steps of 10%, i.e. -40%, -30%, …., 30%, 40%, where a negative number indicates that the number will be subtracted. All numbers are equally likely to be chosen. In the high risk treatment, the range is extended to 200%, i.e. -100%, -90%, …, 90%, 100%. Finally, we implement a no risk treatment where no shock is implemented (i.e. ε has zero mean and zero variance). The only possible source of uncertainty in this treatment is subjects' uncertainty about their own performance.

[7] Note that shocks are multiplicative, which is different from standard-text book models in which shocks enter the production function additively, i.e. $q = x + \epsilon$. The empirical predictions are the same, as long as the costs of effort *c(e)* can be measured in monetary terms, i.e. agent's utility is of the form U=u(w-c(e)). In our experimental application, multiplicative shocks ensure that earnings of subjects never fall below zero.

[8] There are two reasons why we use a between-subjects rather than a within-subjects design. First, a within-subjects design reduces the stakes proportional to the number of decisions. Second, a within-subjects design arguably leads to a strong experimenter demand effect. When the variance of the random shock changes with the treatment, subjects are prompted to think about the role of the variance in their decision, and how they are supposed to react to this change in the environment.

Before subjects are informed about their earnings in the work stage, we elicit their attitudes towards risk and losses by means of five incentivized binary lottery-choice tasks. The first two tasks are standard measures of risk preferences based on Dohmen et al. (2011) and Holt and Laury (2002), respectively. In the first task, subjects choose between a lottery and a safe payment in 15 choice situations. The lottery is the same in all situations: subjects earn 400 or 0 points with a 0.5 probability. The safe payment increases each choice situation with 25 points, starting at 25 points up to 375 points. Naturally, subjects prefer the lottery for low values of the safe payment, but switch to the safe payment when its value increases. The switch point is an indicator for risk tolerance.

The second task, inspired by Holt and Laury (2002), consists of choosing ten times between two lotteries. The first lottery pays either 200 or 160, the second lottery pays either 385 or 10. In both lotteries, the probability of receiving the highest outcome increases from 0.1 to 1 over the ten choice situations. The first lottery is therefore very attractive in the first choice situation, as it pays 160 rather than 10 with 0.9 probability, but strictly dominated in the final choice situation (200 instead of 385 for sure).

In the third task, we elicit loss aversion as in Fehr and Goette (2007). Subjects are asked six times whether they want to participate in a lottery or receive zero payment. In each lottery, subjects can earn 350 points with a 0.5 probability or loose a number of points. The possible loss increases from 50 points in the first decision problem to 300 points in the sixth decision. The number of lotteries accepted is an indicator of tolerance towards losses.

The fourth and fifth task jointly measure preferences for certainty as proposed by Callen et al. (2014). This measure is based on the idea that individuals weigh probabilities in a non-linear way, where absolute certainty is strongly overvalued compared to linear weighting (Kahneman and Tversky, 1979). Both decision tasks have a similar structure as the tasks above. In task 4, subjects choose ten times between two lotteries. The first lottery yields a gain of 300 points or 0 points. The probability of gaining 300 points increases from 0.1 to 1. The second lottery yields 300 points or 100 points, always with fifty/fifty probabilities. Task 5 is identical, except that the second lottery is replaced by a safe payment of 100 points instead of a fifty/fifty probability of gaining 100 or 300 points. Since the first set of lotteries are equal in task 4 and 5, and the alternative is a lottery in task 4 and a safe payment in task 5, we can compare the answers to obtain a measure of individual preference for certainty.

We incentivized decision making by randomly choosing one task for payment at the end of the experiment. Subsequently, the computer randomly selects one choice situation, and determines the outcome of the lottery if the subject has chosen to play a lottery in that particular choice situation. Finally, in the survey at the end of the experiment, we ask subjects how willing they are to take risks in general. Dohmen et al. (2011) validate this measure. Throughout the paper, our measure of risk attitude is the first principal component of five measures: the four lottery-choice tasks capturing risk aversion, plus the subjective willingness to take risks. We do not include the measure for loss aversion (task 3) in the principal component analysis.

Finally, we elicit measures of subjects' IQ and personality. The IQ test consists of ten Raven matrices of increasing difficulty level (Raven, 1962).[9] Subjects have ten minutes to solve the matrices. We control for IQ in our analysis in order to rule out that risk preferences reflect cognitive ability (see Dohmen et al., 2018). After the IQ test, subjects answer several questions about their personality and attitudes. In addition to the subjective risk attitude measure, we include the following three sub-traits of the big five personality index: vulnerability, anxiety, achievement-striving (NEO-PI-R, Costa and McCrae, 1992). Additionally, we selected items from different personality questionnaires that are related to effort costs. [10] In our analysis, we use the first principal component that results from a factor analysis of those items, explaining 32% of the variation.

2.3. Experimental procedures

The experiment was conducted at the Behavioral and Experimental Economics Laboratory (*BEElab*) at Maastricht University beginning of May 2015. We recruited undergraduate students from various disciplines. We excluded economics students in their third or fourth year, as they may be acquainted with principal-agent theory. In total, 194 students participated in the experiment. We randomized students into low risk and high risk treatments within sessions to minimize the impact of session-specific effects. The no risk treatments were conducted in two separate sessions (one on Friday morning, and one on Monday afternoon). The reason for doing so is that the instructions in the no risk treatment are significantly shorter, as there is no need to explain that individuals'

[9] Further information about the matrices is available upon request.

[10] Items are taken from the item lists that measure the following personality constructs: 1) 'liveliness' by Lee and Ashton (2004); 2) 'activity level' by Costa and McCrae (1992) and 3) 'conscientiousness' by John and Srivastava (1999). The items are available upon request.

production is affected by random shocks. Randomizing within the session would therefore lead to long waiting times for subjects in the no risk treatment, which might affect their decision making. For example, one may be concerned that bored individuals are more inclined to take risks. We informed subjects that the experiment was double-blind, i.e. that none of their actions or answers can be linked to their name or student-id. Any form of communication and use of electronic devices (calculators, phones etc.) was strictly forbidden throughout the experiment. Subjects were allowed to make paper-pencil calculations.

A typical session lasted about one hour and the average payoff was 18.85 Euros. There was no official show-up fee, but all subjects were awarded 200 points for completion of the IQ-test. The conversion rate from points to euro's was €0.021 per point. Subjects were paid out in private at the end of the experiment.

3. Results of first experiment

In this section, we first show descriptive statistics. In section 3.2, we present the analysis and results of testing our main hypothesis. Section 3.3 describes heterogeneous treatment effects. Finally, in section 3.4 the results are interpreted which

3.1. Descriptive statistics

Table 1 gives an overview of summary statistics of the main variables by treatment. Subjects in the different treatments are roughly comparable in terms of demographics such as age, gender, year and field of study.[11] We find no significant differences between observable characteristics across treatments, which suggests that our randomization was successful. Moreover, we find that individuals in different treatments perform an equal amount of sequences in the productivity phase. Also, productivity is higher in the incentivized work stage than in the productivity stage, suggesting that financial incentives further improve performance. Individuals' piece rate choice is therefore also a choice of a real-effort level, and not simply a choice between financial lotteries. Online

[11] Field of study statistics are not reported. A multinomial test for each field of study shows that the distribution of students over treatments does not significantly differ from the distribution predicted by random assignment, except for economics. For the economics students we find that there are significantly fewer students in the no risk treatment (N=2) than would be predicted by random assignment (p=0.02).

Appendix C details the characteristics of the sequence task and shows that the distribution is relatively narrow – such that risk concerns play a greater role in decision-making – and that individuals are likely able to accurately predict their performance in the work stage.

To minimize measurement error, we measure risk attitudes by using the first principal component of the five risk-preference measures described in Section 2.2.[12] Higher values indicate higher willingness to take risks. Excluding subjects who make inconsistent choices in at least one of the lottery-choice tasks leads to a loss of 24 observations: 6 in the no risk treatment, and 9 in the low and high risk treatments. As depicted in Table 1, we observe no differences in risk attitude across treatments.

3.2. Analysis and estimation results

We first examine whether individuals who are more willing to take risks select higher piece rates (and a corresponding lower base salary), by regressing individuals' chosen piece rate (0-10) on their risk attitudes and indicator variables for treatment condition.[13] Column (1) of Table 2 reports the estimation results where both risk attitudes and productivity are standardized for reasons of comparison. Consistent with hypothesis 1.1, individuals who are more willing to take risks choose higher piece rates. The effect is statistically significant at the 5 percent level.[14] In column (2) of Table 2, we add control variables to correct for subjects' productivity, demographics (gender, age, year of study, field of study, nationality), effort costs, Raven IQ score, loss aversion, and required certainty premium.[15] The coefficient estimates of risk attitude are quantitatively similar to those reported in the first column. A comparison of the coefficient estimates for the standardized measures of risk attitude and

[12] The first principal component has an eigenvalue of 2.82 and explains approximately 55% of the total variance of the five variables. The loadings equal approximately 0.5 for each of the lottery tasks and 0.3 for the subjective question. We only use the first component, which is the latent variable that captures most variation, and ignore higher order components. This is justified by Kaiser's criterion (Kaiser, 1960), which stipulates that only components with an eigenvalue exceeding 1 should be retained.

[13] The offered piece rates increase in discrete steps by one unit from 0 to 10. Since the data are cardinal, we prefer OLS. Alternatively, one could estimate an ordered probit model, which gives qualitatively similar results.

[14] We do not account for within-session dependence in estimating standard errors of regression parameters. As argued by Cameron et al. (2008, p. 414), "with a small number of clusters the cluster-robust standard errors are downward biased". In line with this argument, when we use the typical adjustment for clustering as proposed by White (1984) we often observe *smaller* standard errors as compared to the OLS estimates without adjustment. We therefore conservatively prefer not to cluster our standard errors, although doing so does not affect any of the main results.

[15] Note that one additional observation is omitted from the sample, as one individual in the estimation sample did not have a unique switching point in the lotteries measuring loss aversion.

productivity reveals that the effect of risk attitude on piece rate choice is about half the effect of productivity. In Online Appendix A, we replicate the analysis for each of the five measures of risk preference separately to assess the robustness of this result. For all measures, we find a positive association between risk attitude and piece rate choice (see Table A1).

The second hypothesis is that higher risk in the environment induces risk averse individuals to choose lower piece rates. A first impression of the results is provided by Figure 3, which shows kernel density estimates of the distribution of compensation choices by treatment. The clearest difference between treatments is that the distribution of choices in the low risk treatment is less dispersed than in the other treatments. Importantly, and contrary to prior expectations, the distributions of the low and high risk treatment are not shifted to the left as compared to the no risk treatment. If anything, the distribution of choices in the high risk treatment has a thicker right tail than the distribution of choices in the no risk treatment, contrary to what we would expect.

We examine this association in closer detail by the regression of compensation choice (i.e. subjects' chosen piece rate) on treatment dummies and individuals' risk attitudes reported in column (1) of Table 2. Consistent with the first impression based on Figure 3, we find no statistically significant treatment effects. For both the low and high risk treatments, we find that the estimated coefficients are positive, but small and statistically insignificant (p=0.73 and p=0.53, respectively).[16] Unreported regressions show that pooling both treatments and defining a treatment dummy that equals 1 if the treatment is either the low or the high risk treatment gives a similar insignificant result (p=0.54). As shown by estimation results in column (2) of Table 2, the results are not affected by including control variables.

When estimating the regressions reported in Table 2, we omit 25 observations from the analysis due to inconsistent behavior in the lottery choice tasks. To estimate the effects of both risk aversion and production risk for the whole sample, we impute values for inconsistent subjects and redo the analyses shown in Table 2. As the experiment contains four lottery tasks, and all subjects behave consistently in at least one lottery task, we use consistent lottery choices and the subjective risk measure to predict behavior in lottery choice tasks in which

[16] Throughout the paper, reported p-values are based on two-sided tests.

subjects behaved inconsistently.[17] Replicating the analyses of Table 2 using imputed values, we find that the results are quantitatively very similar. Results are reported in Table A2 in Online Appendix A.

3.3. Heterogeneous treatment effects

Since theory predicts that risk averse workers choose lower piece rates in response to risk while risk neutral or risk-loving workers do not, we estimate the treatment effect for risk averse workers separately.[18] The choices made in the four lottery-choice tasks allows us to distinguish between risk averse subjects on the one hand, and risk neutral or risk-loving subjects on the other hand.[19] Subjects behave risk averse if they are willing to forego expected earnings to decrease the variance of potential earnings. Considering all four lottery tasks, only 4% of the subjects are consistently risk neutral or risk seeking. Moreover, 36% of subjects are consistently risk averse.[20] We confine the analysis to individuals who are consistently risk averse in all four lottery-choice tasks. Column (3) of Table 2 reports OLS estimates of the effect of risk on piece rate choice for the aforementioned sample. We do not find negative treatment effects for the risk averse subset of our sample. The estimated treatment effects remain positive and insignificant.

We generalize this analysis by estimating interaction effects between risk attitude and treatment. Intuitively, we would expect that the effect of risk attitude is close to zero when the environment is very stable, as

---

[17] In particular, we regress the lottery outcome (i.e., switching point) on both the subjective measure of risk preference and the other lottery outcomes for subjects who are consistent in all lotteries. Then, we use the estimated coefficients to predict the lottery outcome for those who have chosen inconsistently. For example, suppose a subject shows inconsistent behavior in the lottery-choice task by Holt and Laury (2002) and consistent behavior in the other tasks. The estimated regression coefficients and subjects' values of the other tasks (including the subjective risk preference measure) are then multiplied to predict the response in the inconsistent lottery-choice task.

[18] We acknowledge that the experiment is not designed to test for heterogeneous treatment effects and therefore lacks statistical power to estimate such effects. We nevertheless check for heterogeneous effects to examine whether patterns in the data exist that point to other models of decision-making under risk.

[19] We cannot distinguish risk neutral from risk seeking subjects because the lottery choice tasks do not allow subjects to indicate a strict preference for one of two lotteries in each decision row. As such, switching at the row where the expected values of two lotteries are equal indicates that the subject is either risk neutral or (slightly) risk-seeking.

[20] There are several reasons why subjects' risk preferences may appear inconsistent over the lottery-choice tasks they perform. Firstly, the middle row, which represents a different gamble for each lottery-choice task, may be a focal point that affects decision-making. Likewise, the row of the switching point in the previous task might become salient, such that risk-taking behavior differs by lottery-choice task. Moreover, in one set of choice tasks subjects choose between a lottery and a certain outcome, whereas in the other set of tasks subjects choose between two lotteries. Potentially, these differences in framing affect risk-taking behavior.

in the no risk treatment, but becomes larger when risk in the environment increases.[21] The estimation results are reported in column (4) of Table 2. We find no significant interaction effects between the treatments and risk attitude. Note that although the effect of risk attitude is no longer statistically significant, the point estimate is unaffected: the effect is just estimated less precisely. We therefore find no indication that risk attitudes are more important when the environment is more risky, or, stated otherwise, that the treatment effects are concentrated among risk averse subjects.

We also explore heterogeneous treatment effects by productivity. Theoretically, there are two opposing effects of productivity on the impact of risk on compensation choice. On the one hand, risk aversion may be lower at higher wealth levels, so that individuals who expect to earn more are less sensitive to the risk treatments. On the other hand, since we implemented risk multiplicatively, a larger absolute number is added or subtracted when individuals are more productive. Consequently, more productive individuals may perceive the treatment as more risky.

In the fifth column of Table 2, we report estimation results of a regression that includes interaction effects between productivity and treatment. Productivity is standardized for ease of interpretation. Interestingly, we find significant negative interaction effects between productivity and the treatments. Note that these effects are not driven by differences in risk attitudes across productivity levels, because we do not find an interaction between risk attitudes and treatment, as discussed above. Adding those interactions therefore does not change the results.[22] The heterogeneous productivity effects are illustrated by Figure 4, which shows the estimated difference in the selected piece rate between the control and risk treatments conditional on individuals' productivity.[23] The risk treatment pools the low and high risk treatments. Subjects of average productivity do not show a statistically

[21] It should be noted that it is theoretically not clear whether the positive association between willingness to take risks and the piece rate is stronger in more risky environments. The reason is that, theoretically, risk averse individuals already choose a low piece rate in risky environments, so the effect of stronger risk aversion is also small in absolute value. This offsets the intuitive effect that differences in risk attitudes are more relevant when the environment is risky.

[22] In addition to OLS regressions, we have also estimated quantile regressions at the 25^{th}, 50^{th}, and 75^{th} percentile of the piece rate distribution. The results are broadly consistent with the OLS results: as productivity increases, introducing risk has increasingly negative effects on piece rate choice.

[23] The predicted treatment effects by productivity are based on coefficients of a regression model as in the fifth column of Table 2 where the risk treatments are pooled. Confidence intervals are estimated by the delta method.

significant response to treatment. However, contrary to theoretical predictions, low-productivity individuals choose significantly higher piece rates when risk is introduced. Specifically, in the risk treatment, the effect is statistically significant at the 10% level for the 33% least productive individuals. The most productive individuals choose lower piece rates, but the estimates are not significantly different from zero. As the interaction coefficients in column (5) of Table 2 are approximately equal for both the low and the high risk treatment, the patterns are qualitatively similar for both treatments.

4. Interpretation of results and refined hypotheses

The finding of the first experiment that highly productive individuals tend to select lower piece rates when risk is introduced, while less productive individuals tend to select higher piece rates can be explained by reference-dependent preferences. In models with reference-dependent utility, outcomes are perceived as gains or losses depending on how they compare to a reference level, also known as the reference point. Such models can explain our findings if individuals with low productivity are more likely to expect earnings that fall below their reference point than individuals with high productivity.[24]

Based on prospect theory (Kahneman and Tversky, 1979), we extend the simple model of section 2 by assuming that income is evaluated relative to a reference point R, that the utility function is symmetric in the gain and the loss domain, and that the utility function exhibits diminishing sensitivity to gains and losses. Expected utility can be written as:

[24] An alternative explanation is that low-productivity individuals are more likely to suffer from self-control issues that worsen when the environment becomes riskier, since there is less control over output. They therefore commit to exert effort by choosing a higher piece rate. Since we asked subjects to motivate their choice, we can shed light on this by investigating treatment differences in subjects' agreement with the statement that "a higher piece rate forces them to work harder". Concentrating on the bottom half of the productivity distribution, we find no statistically significant difference in the percentage of subjects that agree between the no risk treatment and the pooled risk treatments (82% vs 90%, p=0.27). We subsequently asked subjects to rank all statements they agree with in order of their importance for their piece rate choice. We find that in this sample of relatively low-productivity individuals, the motivation effect is considered less important in the pooled risk treatments than in the no-risk treatment. This suggests that low-productive subjects' higher piece rate choice is not driven by commitment problems.

$$E(U) = \int_{\varepsilon^*}^{\bar{\varepsilon}} u(w - c(e) - R)\varphi(\varepsilon)d\varepsilon - \int_{\underline{\varepsilon}}^{\varepsilon^*} u\left(R - \left(\mathrm{w} - c(e)\right)\right)\varphi(\varepsilon)d\varepsilon, \qquad (1)$$

where $\varphi(\varepsilon)$ is the probability distribution of the error terms, $\bar{\varepsilon}$ and $\underline{\varepsilon}$ are the highest and lowest possible realizations, and ε^* is defined as the shock that equates income net of effort costs with the reference point, i.e. $s + be(1 + \varepsilon^*) - c(e) - R = 0$. We assume that the distribution of shocks is symmetric around zero, implying that $\bar{\varepsilon} = -\underline{\varepsilon}$. The first part of equation (1) represents the gain domain ($w - c(e) > R$), while the second part represents the loss domain ($w - c(e) < R$). Note that we assume diminishing sensitivity to gains and losses, but we do not assume loss aversion (which would imply multiplying the second term by $\lambda > 1$).[25]

It is instructive to first analyze behavior in the absence of risk. In the absence of risk, the optimal choice of effort is described by:

$$\frac{dE(U)}{de} = u'(\cdot)[b - c'(e^*)] = 0, \qquad (2)$$

Individuals maximize their income: they equate the revenues of additional effort with the marginal costs of effort. The reference point affects $u'(\cdot)$, but does not affect the optimal effort choice. Likewise, the optimal piece rate is described by:

$$\frac{dE(U)}{db} = u'(\cdot)\left[\frac{ds}{db} + e^*\right] = 0, \qquad (3)$$

which is independent of the reference point.[26] Reference points therefore do not affect the choice of piece rates in the absence of risk. The effect of introducing risk in the environment depends on how expected earnings net of effort costs compare to the reference point.

[25] Loss aversion reinforces the tendency to behave risk averse in the gain domain, because it increases the incentive to avoid losses. At the same time, loss aversion reduces the incentive for risk seeking behavior when expected outcomes fall short of the reference point, as losses receive a higher weight than gains. The total effect is therefore ambiguous, depending on the magnitude of the loss aversion parameter and the probability that shocks push outcomes from the loss to the gain domain and vice versa.

[26] This result follows from the simplifying assumption that effort costs are evaluated the same way as money. Models that assume that agents are loss averse and that money and effort costs are separable, i.e. $U = u(w - R) - c(e)$, predict that agents exert more effort in the loss domain than in the gain domain in the absence of risk, keeping $|w - R|$ constant. The reason is that the marginal returns to money are higher in the loss domain due to loss aversion (see e.g. De Quidt et al., 2017, for an empirical test of this mechanism and De Quidt, 2018, and Corgnet and Hernán González, 2019, for theoretical models in a setting with binary shocks). This higher effort is reflected in higher piece rate choice. However, when there is risk in the environment and the utility function exhibits diminishing returns to gains and losses, such models still predict risk seeking behavior in the loss domain and risk averse behavior in the gain domain, provided changes in effort are not too large (since risk averse workers increase effort when the environment becomes more risky, see Sloof and Van Praag, 2010).

Hypothesis 2.1: When individuals have a low reference point, such that their expected earnings net of effort costs exceed the reference point for any given piece rate, they respond to the introduction of risk by choosing lower piece rates.

We provide a formal derivation in Appendix A. The intuition is that individuals who expect to earn more than the reference point perceive their earnings as gains. Being risk averse in the gain domain, they reduce their risk exposure by exerting less effort and selecting a lower piece rate than the one that maximizes their earnings. By similar logic, the model predicts the following:

Hypothesis 2.2: When individuals have a high reference point, such that their expected earnings net of effort costs fall below the reference point for any given piece rate, they respond to the introduction of risk by choosing higher piece rates.

Since individuals have diminishing sensitivity to losses, negative shocks have less impact on utility than positive shocks of equal magnitude, providing an incentive to take risks by exerting more effort and choosing higher piece rates. Finally, since behavior in the absence of risk is independent of the reference point, we predict the following:

Hypothesis 2.3: If there is no risk in the environment, reference points do not affect the choice of piece rates.

Hypothesis 2.4: If production risk is present, it follows from Hypotheses 2.1-2.3 that individuals choose higher piece rates when they have a high reference point than when they have a low reference point.

The key assumption in this model is that individuals exhibit diminishing sensitivity to gains and losses. It is also possible to explain our findings with alternative models of reference-dependent preferences that do not rely on this assumption. For example, it would also be possible to assume that individuals experience a jump in utility when their payoff exceeds a certain reference point (as in Diecidue and Van de Ven, 2008). This jump in utility represents the joy of meeting a particular income threshold. Such a discontinuity in the utility function motivates relatively low-productivity subjects to choose a higher piece rate when risk is introduced, because positive shocks allow them to reach the target level with higher probability. Relatively productive subjects have the opposite incentive: they do not consider the introduction of risk as an opportunity, but as a threat.

As a result, our aim is not to identify what particular model of reference-dependent utility best explains behavior in our experiment, but rather to assess whether reference-dependent preferences can explain our experimental findings and, more generally, whether they are important as a determinant of compensation choice.

Similarly, we do not aim to disentangle where references points come from and how they are formed.[27] Rather, we exogenously manipulate both risk and reference points in a second experiment in order to test the hypotheses derived above.

5. Design of the second experiment

The basic set-up and parametrization of the second experiment is comparable to the first experiment. Subjects become acquainted with the sequence task, choose a compensation scheme for the upcoming work stage, answer questions about their motivation, and then work for five minutes. Compared to the first experiment, we change the setup in two important ways. The main difference is that we manipulate reference points: subjects are randomized into a high or low reference point treatment. The second difference is that we conduct two instead of three risk treatments. Specifically, we drop the low risk treatment, so that there is either no risk or high risk (i.e. production shocks that range between -100% to +100%). We therefore have four treatments in total:

1. Low reference point, no risk;
2. High reference point, no risk;
3. Low reference point, high risk;
4. High reference point, high risk.

A further minor difference is that we collect fewer measures of risk preferences, since risk aversion is no longer the main variable of interest. Specifically, we collect the subjective risk question and the binary lottery-

[27] There are different potential determinants of reference points. For example, individuals may have certain expectations of what they will earn in the experiment. Reference points may therefore reflect income expectations. Alternatively, individuals may base their reference point on what others earn (Schwerter, 2013). Although subjects in our experiment do not receive explicit information about the earnings of other subjects, they may have an accurate idea of their relative productivity, and can compare their earnings to those of others.

choice task developed by Dohmen and Falk (2011). Figure 5 provides a detailed overview of the experimental setup. In what follows, we discuss the reference point manipulation and the manipulation checks.

5.1. Reference point manipulation

We manipulate reference points by introducing salient counterfactual earnings. At the start of the experiment, all subjects are informed that their reward for the work stage is either a constant amount of points, or the points they earn according to their performance and the piece rate scheme they select. The computer randomly determines for each subject which of the two payment systems applies, both with the same probability. Subjects are informed about the payment system that the computer has chosen right before they enter the piece rate selection stage. Thus, either subjects receive an amount regardless of their production in the work stage and their choice of piece rate is not payoff relevant (i.e. hypothetical), or they choose a piece rate according to which they are paid. Hence, subjects whose choice of piece rate is payoff relevant are aware of what they *could have earned*. [28] The latter amount is likely to serve as a reference point against which other potential earnings are compared.[29] In the analysis, we focus on these subjects and exclude subjects who made a hypothetical choice.

To manipulate the reference point, the constant payment can be either low or high. In the low reference point treatment, the constant payment equals 100 points. All subjects on a piece rate contract can expect to earn more than 100 points, so that they will perceive their earnings as a gain. Their earnings fall below 100 points only when they select a high piece rate (and corresponding low fixed payment) and their production suffers from a

[28] Our approach is similar to Schwerter (2013). Other approaches are the manipulation of expected earnings, as in Abeler et al. (2011), or contract framing, as in De Quidt (2018). The manipulation in Abeler et al. (2011) differs from ours in that they resolve uncertainty about which payment scheme applies at the end of the experiment. Although this approach has the advantage that all subjects make a payoff relevant choice, it transforms the piece rate choice into a choice over compound lotteries. As such, we cannot attribute treatment effects to changes in subjects' reference point. Manipulating the contract framing is straightforward when subjects earn a bonus if their performance exceeds a threshold, as in De Quidt (2018), but doing so is not straightforward in our setting with linear piece rates. We assume that subjects' endogenous formation of reference points (see Köszegi and Rabin, 2006) is (partly) adaptive, which is also confirmed by our manipulation checks (see section 5.1).

[29] We make this amount salient in two ways. First, we inform subject already in the beginning of the experiment that they will either receive a constant amount or will be rewarded according to the payment scheme they select themselves. Second, the probability on each outcome is 0.5, which ensures that subjects can reasonably expect to receive the constant amount, and perceive it as the amount they could have earned. Note that we need to balance saliency with a large enough sample size of incentivized choices.

severe negative shock.[30] Note that they can always select a piece rate of zero and receive the fixed payment of 370 points. In the high reference point treatment, the constant payment is 600 points. This amount is calibrated to be well above expected earnings for any level of productivity observed in experiment 1. Subjects will therefore perceive expected earnings as a loss relative to the constant payment they could have earned. Productive individuals who receive a sufficiently positive shock to production and select a high piece rate can earn more than the constant payment of 600 points.[31] The high reference point is therefore not so high that it might be considered irrelevant: it is above expected earnings, but not totally out of reach.

We conducted three manipulation checks based on mood changes, earnings aim, and earnings satisfaction. First, if our manipulation is successful, we expect that subjects who select a payoff relevant piece rate and do *not* receive 600 points experience smaller mood increases than those who do not receive 100 points. Second, we expect that subjects in the low reference point treatment aim for lower earnings than subjects in the high reference point treatment. Third, when we ask subjects to indicate their satisfaction with different earnings from the work stage, we expect that the increase in satisfaction from 0-200 points is greater for subjects in the low reference point treatment, while the increase from 500-700 points is larger in the high reference point treatment. Taken together, the three checks indicate that we successfully manipulated reference points. We refer to online Appendix D for a detailed description of the manipulation checks and the results.

[30] In the first experiment, the lowest earnings in the work stage was 318 points. Although theoretically possible, it is highly unlikely that individuals earn less than 100.

[31] In the first experiment, the highest earnings in the work stage was 527 points. Although earning more than 600 points is possible based on observed productivity in period 1, it is a challenging target that can only be reached by choosing a high piece rate and the realization of a positive random shock.

5.2. Experimental procedures

The second experiment was conducted at the end of November 2018. Just like the first experiment, the second experiment was conducted at the BEElab at Maastricht University. The experimental procedures were the same as those of the first experiment. A total of 263 students participated.[32]

To minimize the impact of session-specific effects, we randomized the high and low reference point treatments within sessions. Since the high risk treatment requires an additional explanation of how random output shocks affect compensation, we randomized the high and no risk treatments across sessions. Randomizing the risk treatments within sessions would otherwise have resulted in excessively long waiting times for subjects in the no risk treatment. The duration of a session was on average 50 minutes with an average payout of €15.82 per subject.

6. Results of second experiment

6.1. Descriptive Statistics

Table 3 shows descriptive statistics for all experimental subjects and for the estimation sample – i.e. the sample of subjects whose piece rate choice is payoff relevant –, the latter split out by treatment. As expected, subjects in the second experiment are very similar to subjects in the initial experiment in terms of age, gender, and study progress. [33] The estimation sample consists of subjects who are randomly selected to be rewarded according to the piece rate they have chosen and their output in the work stage. Differences in observable characteristics between the estimation sample and other experimental subjects are therefore small and, with the exception of

[32] To predict the sample size required for sufficient statistical power, we performed a power analysis based on the following assumptions. Firstly, assuming that relatively productive individuals have a low reference point and low-productivity individuals a high reference point, the data from experiment 1 show that chosen piece rates change by approximately 1 unit when risk is introduced: a decrease when the reference point is low and an increase when it is high. The standard deviation is approximately 2 leading to anticipated effect sizes of 0.5 and -0.5 standard deviations. We require the power (β) of the test to be 0.8 and set the significance level (α) to 0.05. Moreover, we assume an intracluster correlation coefficient of 0 – i.e. no variation in piece rates is explained by variation between sessions. This leads to a minimum of 126 subjects to perform the two tests (Cohen, 1988). To manipulate the reference point, we need to double this amount such that at least 252 subjects participate. As some subjects typically do not show up to the experiment, we invited somewhat more than 252 subjects, leading to the ultimate sample of 263 subjects.

[33] In the second experiment significantly more economics and business students participate and significantly less law students and students from other disciplines.

productivity and field of study, not statistically significant. Ignoring differences in risk attitude and choice of piece rate, which might be caused by experimental treatment, we find three statistically significant differences at the 10% level in observable characteristics between treatments.[34] As we conduct in total 24 unpaired two-sided t-tests (4 variables and 6 pairwise comparisons per variable), we would expect 2 to 3 differences to be significantly different by chance at a 10 percent significance level (i.e. 24 x 0.10 = 2.4).[35] This expectation is a lower bound, as it ignores dependence between variables (e.g. study year and age). We find 3 significant differences suggesting that our randomization is successful. We control for the abovementioned characteristics in our analysis.

6.2. Analysis and estimation results

As a starting point, we first estimate an OLS-regression of piece rate choice on risk treatment, controlling for productivity (standardized), risk attitude (standardized), gender, nationality (Dutch, German, Belgian, other), field of study (Business, Economics, other), and study year. In all analyses, we only include individuals who made a payoff-relevant choice. Results are reported in Table 4 and are in line with our first experiment: individuals in the high risk treatment do not choose significantly lower piece rates. Since we experimentally manipulated the reference point, we can now investigate whether the absence of an average effect masks heterogeneity by reference point.

Next, we estimate OLS regressions to analyze the effect of the treatments on piece rate choice. Table 5 reports estimates of the average effect of the four treatments: low reference point, high risk (β_{lh}); low reference point, no risk (β_{ln}); high reference point, no risk (β_{hn}) and; high reference point, high risk (β_{hh}). Column (1) and (2) report treatment effects with and without control variables. In column (2), we correct for individuals' productivity (standardized), risk attitude (standardized), gender, nationality (Dutch, German, Belgian, other), field

[34] Specifically, we find that the fraction of women is significantly different at the 10% level in treatment (1) in comparison to treatment (2), subjects' year of study is significantly different between treatments (2) and (3), and subjects are more productive in treatment (4) than in treatment (3).

[35] Study choice is not considered as it deals with multinomial outcomes. We find that study choice does not significantly predict treatment which thereby provides further evidence that the subjects from different treatments are comparable.

of study (Business, Economics, other), and study year.[36] Naturally, the estimation sample is restricted to individuals who make incentivized choices (N=134).[37]

Hypotheses 2.1 – 2.4 state that the ordering of estimated treatment coefficients is as follows:

$\beta_{lh} < \beta_{ln} = \beta_{hn} < \beta_{hh}$.

A first glance at Table 5, in which the low reference point, high risk (β_{lh}) treatment is the baseline in our regression, corroborates this hypothesis, as the coefficient estimates follow the ordering above. Low reference points and high risk cause individuals to select relatively low piece rates whereas high reference points and high risk induce individuals to select relatively high piece rates. Next, we confirm or reject each hypothesis in turn by testing the equality of coefficients.

Hypothesis 2.1 predicts that risk in the environment induces individuals with a low reference point to choose lower piece rates. The point estimate of β_{ln} is greater than β_{lh} in both columns. After adding explanatory variables in column (2) the difference between β_{ln} and β_{lh} is -1.15 and statistically significant at the 5 percent level, i.e. individuals in the high risk treatment choose roughly one-point lower piece rates. This finding supports our hypothesis that individuals who perceive their potential earnings as a gain make risk averse choices by selecting lower piece rates.

Hypothesis 2.2 states that greater risk causes individuals with a high reference point to select higher piece rates. Consistent with the hypothesis, we find that β_{hh} is greater than β_{hn} by about 0.6 piece rate points. However, the effect is not statistically significant in the regression models without (p=0.20) or with controls (p=0.41).

Hypothesis 2.3 indicates that in a predictable environment, differences in piece rate choice between high and low reference point treatments are small. We find that β_{hn} is somewhat greater than β_{ln}. In column (1) and (2) the results are not significantly different (p = 0.29 and p = 0.32, respectively). This finding corresponds with De Quidt et al. (2017), who do not find an effect of contract framing on effort in a predictable environment.

[36] Additionally, controlling for either time of the day or day of the week does not affect the results.

[37] Including individuals who receive the fixed amount and therefore make hypothetical choices also does not affect the qualitative results.

Hypothesis 2.4 predicts that in a risky environment individuals choose higher piece rates when they have a high reference point. Our estimates are in line with this prediction: Individuals in the high reference point treatment choose 2.2-point higher piece rates than individuals in the low reference point treatment. This is the logical implication of hypotheses 2.1 and 2.2: The reaction to risk depends on the reference point. Our findings therefore suggest that reference points crucially influence piece rate choice when the environment is risky. While we do not have statistically significant evidence that a high reference point actually induces risk seeking choices, the response to risk is sufficiently different in the high and low risk treatment to conclude that reference points matter for piece rate choice. In Table A3 in Online Appendix A we replicate the results of Table 5 by redefining dummy variables and using a Diff-in-Diff formulation, i.e. interact reference point condition and risk treatment. This specification immediately shows that the interaction between risk and reference point is statistically significant.

7. Concluding remarks

Our experiments show that reference points can influence the choice of incentives in ways that are not captured by standard principal-agent model with risk averse agents. Our findings therefore call for incorporating reference-dependent preferences into the principal-agent model in order to better understand the tradeoff between risk and incentives. This tradeoff is relevant in many settings and contractual environments, such as principal-agent relations in insurance, franchise, or labor markets.

Our findings may extend beyond our specific setting, where individuals self-select into linear piece-rate schemes, to sorting into incentive schemes more generally. While previous literature has shown that individuals sort into incentive schemes based on their productivity and risk attitude (Cadsby et al., 2007; Dohmen and Falk, 2011; Bernard et al. 2019), this literature largely overlooks the relevance of reference points, which could be important for a number of reasons. First, the available choice alternatives might create or shape reference points. Second, the probability of earning a given reference amount may differ across choice alternatives, hence affecting sorting into incentive schemes and behavior they induce. Financial incentives may therefore not be effective in attracting high-productivity workers. For example, contracts that reward relative performance (i.e. a promotion tournament) may also attract relatively low-productivity workers in high risk environments, since they may feel that they have little to lose and everything to gain by accepting such a risky contract. Third, most empirical support for self-selection on the basis of productivity relies on well-defined routine tasks such as installing windshields in cars (Lazear, 2000) or simple laboratory tasks (Dohmen and Falk, 2011). Our findings suggest that productivity sorting will be diluted in uncertain environments, not only because individuals differ in risk attitudes, but also because they may differ in the reference point against which they evaluate their income.

Our findings also imply that workers are willing to accept substantial income risks. This insight is particularly important when the formation of reference points is based on social comparisons, so that low-paid workers are relatively likely to earn less than their reference point. This implies that workers at the bottom of the income distribution are willing to take substantial risks to improve their situation. Employers may therefore shift risks to workers without having to compensate them for bearing those risks. Since those workers may be in a relatively bad position to bear those risks, this may be undesirable from a social point of view.

Our findings also have implications for contract design in settings where agents can influence the variance of possible outcomes. In such a setting, high output targets potentially make reference points salient, prompting incumbent agents to take excessive risks that are not in the interest of the principal. Examples include bonus contracts for traders in derivative markets. Contracts in these markets often rewarded upside risk but entailed little downside risk. Our findings suggest that redesigning contracts to entail more downside risk might not deter agents from excessive risk taking when by doing so agents may bring their earnings close to their reference level.

Those diverse implications underline the importance of further research on what determines reference points in the context of compensation choice. For example, individuals may base their reference point on social comparisons, status quo, or a personal income target independent of social comparison. Another important direction for further research is to assess the robustness and generality of our findings. Experimental variations include the way risk is introduced in the environment and the incentive schemes individuals can choose from.

References

Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman (2011). Reference Points and Effort Provision. *American Economic Review*, *101*(2): 470-92.

Ackerberg, Daniel A. and Maristella Botticini (2002). Endogenous Matching and the Empirical Determinants of Contract Form, *Journal of Political Economy*, 110(3): 564-591.

Bartling, Björn, Leif Brandes, and Daniel Schunk (2015). Expectations as Reference Points: Field Evidence from Professional Soccer. *Management Science* 61(11): 2646-2661.

Bernard, Mark, Thomas Dohmen, Arjan Non, and Ingrid Rohde (2019). Menus of Contracts Determine Sorting Patterns. *Journal of Economic Psychology*, 72: 293-311.

Bolton, Patrick, and Mathias Dewatripont. *Contract Theory.* MIT press, 2005.

Budde, Jörg and Matthias Kräkel (2011). Limited Liability and the Risk-Incentive Relationship. *Journal of Economics*, 102(2): 97-110.

Cadsby, Bram, Fei Song, and Francis Tapon (2007). Sorting and Incentive Effects of Pay-for-Performance: An Experimental Investigation. *The Academy of Management Journal*, 50(6): 387–405.

Cahuc, Pierre, and André Zylberberg (2004). Labor Economics. MIT Press.

Callen, Michael, Mohammad Isaqzadeh, James Long, and Charles Sprenger (2014). Violence and Risk Preference: Experimental Evidence from Afghanistan. *American Economic Review*, 104(1): 123-48.

Cameron, Colin, Jonah Gelbach, and Douglas Miller (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *The Review of Economics and Statistics,* 90(3): 414-427.

Chowdhury, Subhasish, and Alexandros Karakostas (2020). An Experimental Investigation of the 'Tenuous Trade-off' between Risk and Incentives in Organizations. *Theory and Decision*, 88(1): 153-190.

Cohen, Jacob (1988). *Statistical Power Analysis for the Behavioral Sciences, Second Edition.* Mahwah, NJ: Lawrence Erlbaum Associates.

Corgnet, Brice, and Roberto Hernán-González (2019). Revisiting the Trade-off between Risk and Incentives: The Shocking Effect of Random Shocks? *Management Science*, 65(3): 955-1453.

Costa, Paul and Robert McCrae (1992). Four Ways Five Factors are Basic. *Personality and Individual Differences,* 13(6): 653-665.

De Quidt, Jonathan (2018). Your Loss is My Gain: A Recruitment Experiment with Framed Incentives. *Journal of the European Economic Association,* 16(2): 522-559.

De Quidt, Jonathan, Francesco Fallucchi, Felix Kölle, Daniele Nosenzo, and Simone Quercia (2017). Bonus versus Penalty: How Robust are the Effects of Contract Framing? *Journal of the Economic Science Association*, *3*(2): 174-182.

Diecidue, Enrico, and Jeroen Van De Ven (2008). Aspiration Level, Probability of Success and Failure, and Expected Utility. *International Economic Review,* 49(2): 683-700.

Dohmen, Thomas, and Armin Falk (2011). Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender. *American Economic Review,* 101(2): 556-90.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2018). On the Relationship Between Cognitive Ability and Risk Preference. *Journal of Economic Perspectives*, 32(2): 115-134.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner (2011). Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences, *Journal of the European Economic Association*, 9(3): 522-550.

Fehr, Ernst, and Lorenz Goette (2007). Do Workers Work More When Wages are High? Evidence from a Randomized Field Experiment, *American Economic Review*, 97(1): 298- 317.

Fehrenbacher, Dennis, and Burkhardt Pedell (2012). Disentangling Incentive Effects from Sorting Effects: An Experimental Real-Effort Investigation, Working Paper.

Gill, D., Kissová, Zdenka., Lee, Jaesun, and Viktoria Prowse. (2019). First-Place Loving and Last-Place Loathing: How Rank in the Distribution of Performance Affects Effort Provision. *Management Science*, *65*(2), 494-507.

Hilt, Eric (2008). The Negative Trade-Off between Risk and Incentives: Evidence from the American Whaling Industry. *Explorations in Economic History,* 45(4): 424-444.

Holt, Charles, and Susan Laury (2002). Risk Aversion and Incentive Effects. *American Economic Review,* 92(5): 1644–1655.

John, Oliver, and Sanjay Srivastava (1999). The Big Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives. *Handbook of Personality: Theory and Research,* 2(1999): 102-138.

Kahneman, Daniel and Amos Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263-292.

Kaiser, Henry (1960). The Application of Electronic Computers to Factor Analysis. *Educational and Psychological Measurement* 20, 141–151.

Köszegi, Botond, and Matthew Rabin (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics,* 121(4): 1133-1165.

Lazear, Edward (2000). Performance Pay and Productivity. *American Economic Review* 90(5): 1346-1361.

Lee, Kibeom, and Michael Ashton (2004). Psychometric Properties of the HEXACO Personality Inventory. *Multivariate Behavioral Research,* 39(2): 329-358.

Linde, Jona, and Joep Sonnemans (2012). Social Comparison and Risky Choices. *Journal of Risk and Uncertainty,* 44(1): 45-72.

Ockenfels, Axel, Dirk Sliwka, and Peter Werner. "Bonus Payments and Reference Point Violations." *Management Science* 61.7 (2015): 1496-1513.

Prasad, Kislaya, and Timothy Salmon (2013). Self-Selection and Market Power in Risk Sharing Contracts. *Journal of Economic Behavior & Organization*, 90: 71-86.

Prendergast, Canice. (2002). The Tenuous Trade-off between Risk and Incentives. *Journal of Political Economy,* 110(5): 1071-1102.

Raven, John C. (1962). Advanced Progressive Matrices. London, H. K. Lewis & Co. Ltd.

Schwerter, Frederik (2013). Social Reference Points and Risk Taking. *Bonn Econ Discussion Papers* No. 11/2013. http://hdl.handle.net/10419/92981

Sloof, Randolph, and Mirjam Van Praag (2010). The Effect of Noise in a Performance Measure on Work Motivation: A Real Effort Laboratory Experiment. *Labour Economics* 17(5): 751-765.

Varian, Hal R. (1992). Microeconomic Analysis (3rd Edition). New York: Norton.

White, Halbert (1984). Asymptotic Theory for Econometricians. Orlando, Academic Press.

Figures and Tables

EXPERIMENT ONE

Figure 1 - Overview of the first experiment

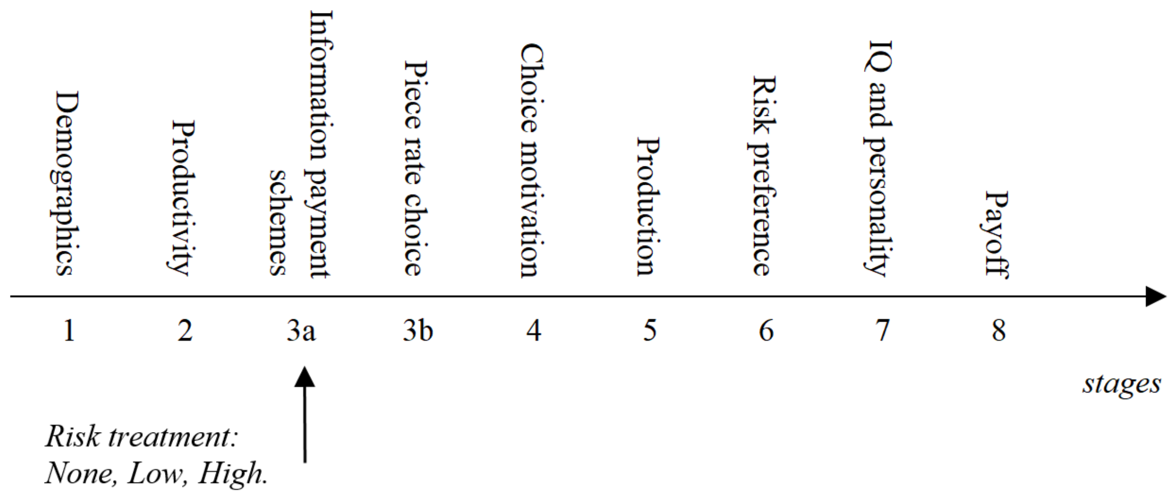

Figure 2 - Menu of piece rate contracts: screenshot

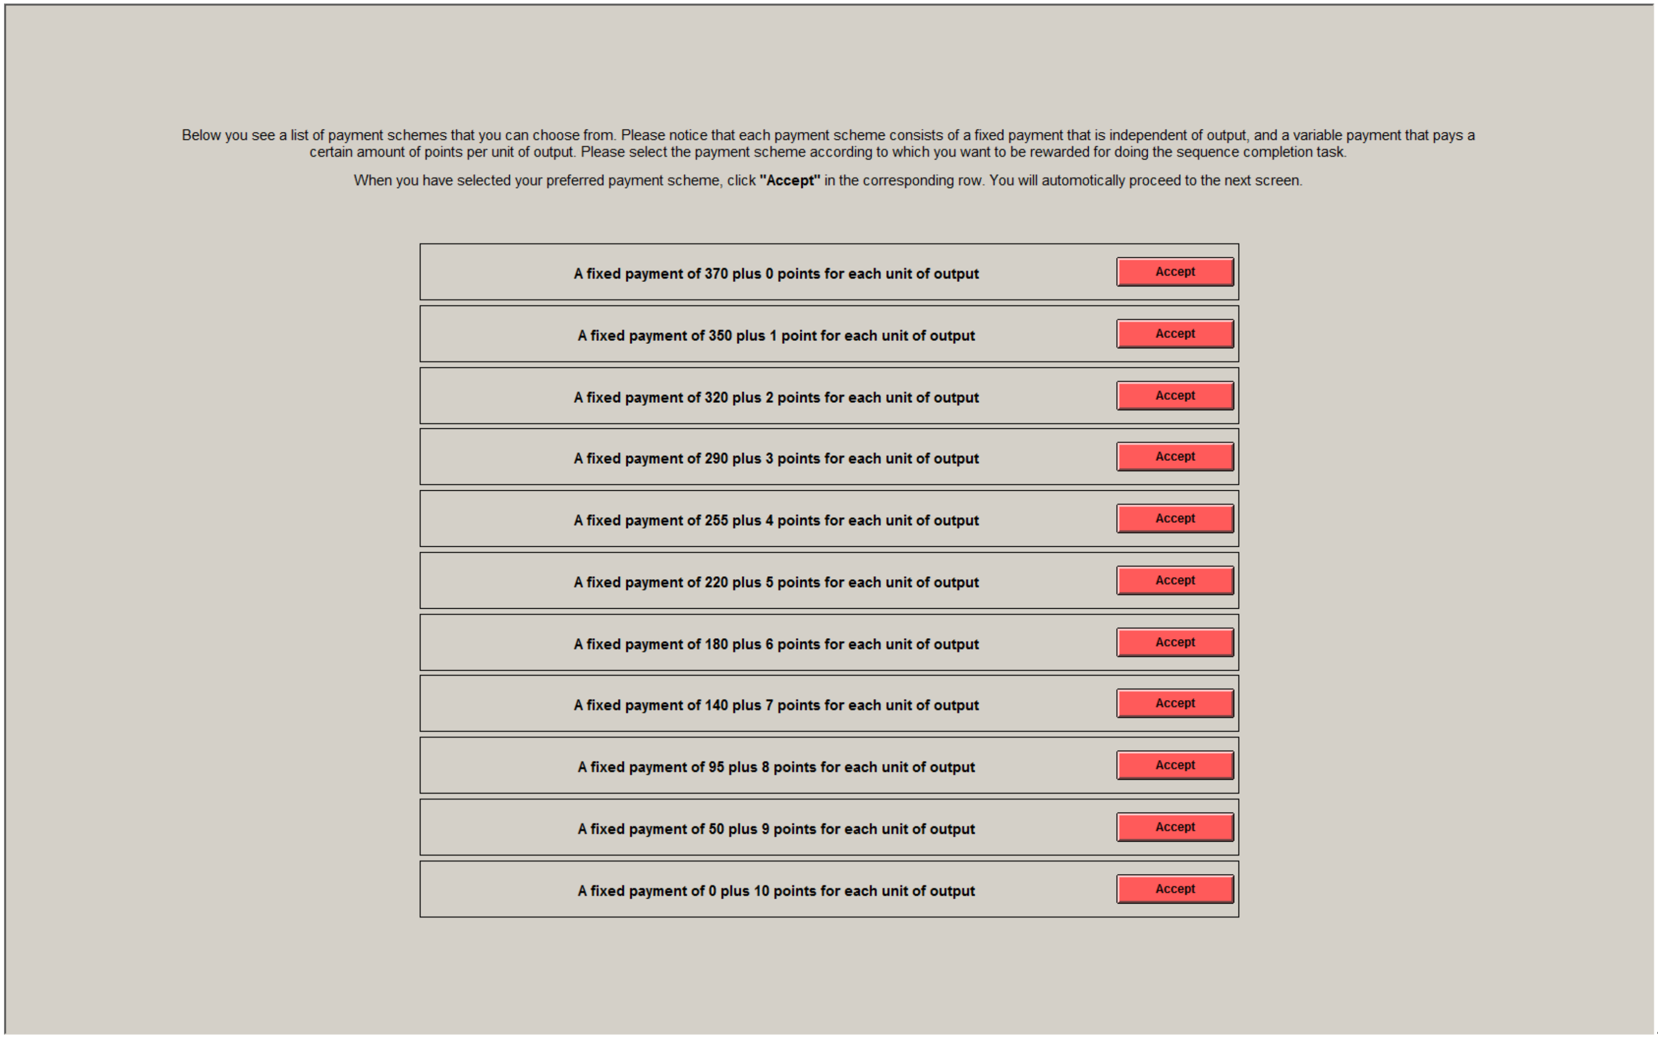

Figure 3 - Distribution of compensation choices by treatment: kernel density estimates

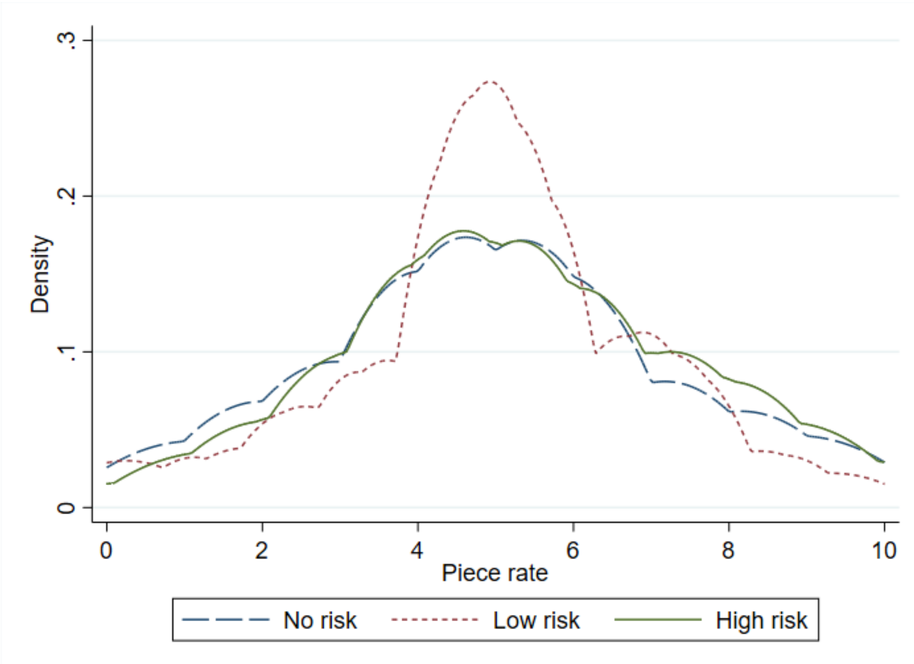

NOTE: The figure depicts kernel density estimates of piece rate choices by treatment.

Figure 4 - Marginal effect of risk treatment on piece rate choice by productivity

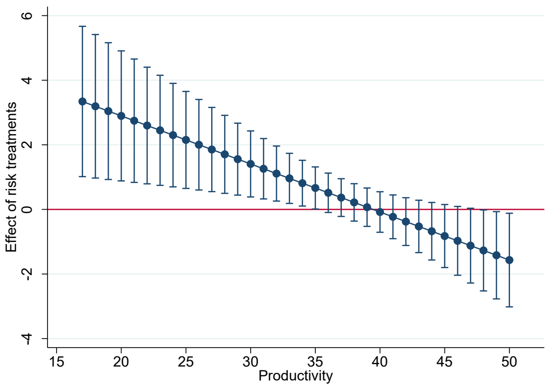

NOTE: The figure depicts the estimated effect of the risk treatment (i.e. the low and high risk treatment pooled) on piece rate choice relative to the no risk treatment. Bars show 90% confidence intervals.

Table 1 - Descriptive statistics first experiment

	TREATMENT			
	No risk	Low	High	Total
Age	20.67	21.33	21.34	21.15
	(0.27)	(0.29)	(0.30)	(0.17)
Women (shares)	0.60	0.59	0.71	0.64
	(0.07)	(0.06)	(0.05)	(0.03)
Study year	1.91	1.94	1.89	1.91
	(0.11)	(0.10)	(0.08)	(0.06)
Risk attitude (principal	0.09	-0.10	0.03	0.00
component)	(0.19)	(0.24)	(0.22)	(0.13)
Productivity	36.95	36.36	36.90	36.72
	(0.70)	(0.67)	(0.63)	(0.38)
Productivity work stage	41.16	39.55	38.97	39.80
	(0.70)	(0.87)	(0.99)	(0.51)
Piece rate choice	4.95	4.97	5.24	5.06
	(0.32)	(0.25)	(0.27)	(0.16)
Raven IQ	5.16	4.81	4.94	4.96
	(0.28)	(0.25)	(0.27)	(0.15)
Perceived effort costs	0.07	-0.13	0.08	0.00
(principal component)	(0.29)	(0.20)	(0.19)	(0.13)
Study (shares)				
International Business	0.47	0.30	0.33	0.36
Economics	0.04	0.12	0.17	0.11
Law	0.11	0.20	0.11	0.14
Psychology	0.04	0.04	0.03	0.04
Culture studies	0.02	0.04	0.01	0.03
Other	0.33	0.29	0.34	0.32
Observations	55	69	70	194

NOTE: Table 1 reports means by treatment condition. The standard error of the mean is shown between parentheses. Productivity refers to the number of completed sequences during the productivity stage. Risk attitude is the first principal component of the four lottery choice tasks. Higher values indicate a greater willingness to take risk. Subjects who choose inconsistently in at least one of the lottery tasks are excluded from the risk attitude measure such that statistics are shown for 49, 60, and 61 subjects in the no, low, and high risk treatment, respectively.

Table 2 – Effect of treatment and risk attitude on compensation choice

	(1) Piece rate	(2) Piece rate	(3) Piece rate	(4) Piece rate	(5) Piece rate
Treatment:					
No risk	Baseline	Baseline	Baseline	Baseline	Baseline
1 if low risk	0.15	0.27	0.83	0.25	0.36
	(0.44)	(0.40)	(0.85)	(0.40)	(0.40)
1 if high risk	0.28	0.23	0.77	0.22	0.33
	(0.43)	(0.41)	(0.80)	(0.41)	(0.41)
Risk attitude (PCA, std.)	0.44**	0.47**	0.68	0.56	0.52***
	(0.17)	(0.19)	(0.49)	(0.38)	(0.19)
Productivity (std.)		1.11***	1.58***	1.07***	1.67***
		(0.16)	(0.37)	(0.17)	(0.30)
Low risk X Risk attitude				-0.21	
				(0.27)	
High risk X Risk attitude				0.12	
				(0.27)	
Low risk X Productivity					-0.16**
					(0.07)
High risk X Productivity					-0.13*
					(0.08)
Constant	4.96***	2.94	0.71	2.43	2.99
	(0.32)	(2.28)	(4.05)	(2.31)	(2.26)
Controls	No	Yes	Yes	Yes	Yes
Observations	170	169	66	169	169
R-squared	0.04	0.32	0.46	0.33	0.35

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

NOTE: OLS estimates. The dependent variable in all columns is the chosen piece rate contract (0-10). Standard errors are in parentheses. Stars indicate significance: *** p<0.01, ** p<0.05, * p<0.1. Risk attitude is the principal component of the four lottery choice tasks and is standardized. Higher values indicate a greater willingness to take risk. Productivity is standardized. Controls include effort costs, Raven IQ, loss aversion, certainty premium, study field, study year, nationality, gender, and age. The estimation sample in column (3) is restricted to individuals who are consistently risk averse in all four lottery-choice tasks.

EXPERIMENT 2

Figure 5 - Overview of the second experiment

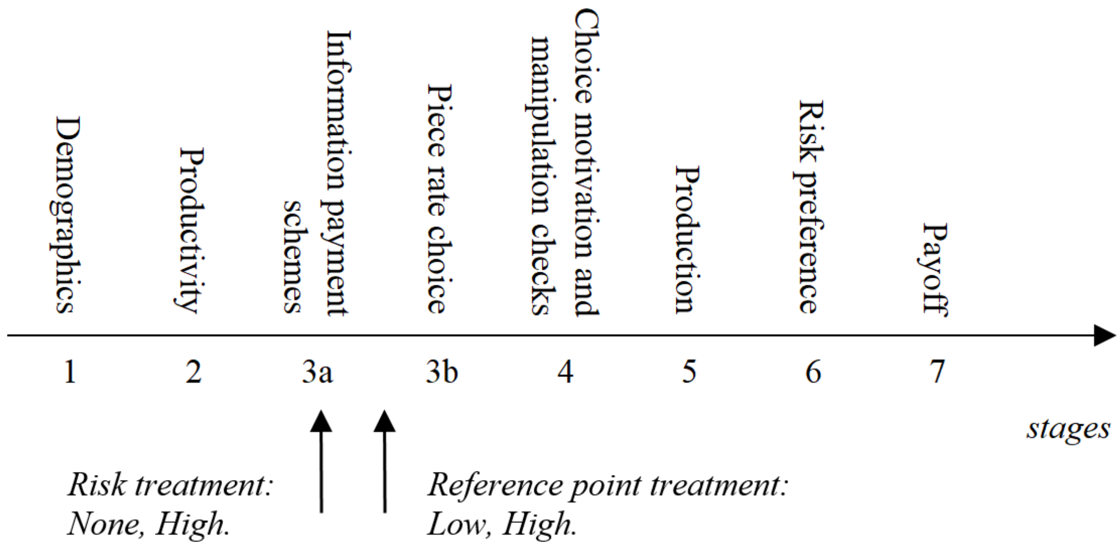

Table 3 - Descriptive statistics second experiment

	TREATMENT					
	Low reference point, no risk	High reference point, no risk	Low reference point, high risk	High reference point, high risk	Estimation sample	Total
Age	20.65	20.89	20.76	20.27	20.63	20.78
	(0.31)	(0.43)	(0.34)	(0.34)	(0.04)	(0.13)
Women (shares)	0.70	0.44	0.54	0.61	0.58	0.61
	(0.08)	(0.10)	(0.08)	(0.09)	(0.04)	(0.03)
Study year	1.92	2.33	1.78	1.94	1.97	2.05
	(0.21)	(0.23)	(0.14)	(0.21)	(0.10)	(0.07)
Risk attitude (principal	-0.22	-0.16	0.33	0.27	0.06	0.00
component)	(0.14)	(0.23)	(0.12)	(0.19)	(0.09)	(0.07)
Productivity	33.79	32.15	31.25	34.70	32.98	34.32
	(1.20)	(1.30)	(1.20)	(1.20)	(0.62)	(0.40)
Output work stage	38.00	40.33	37.46	39.97	38.81	38.39
	(1.08)	(1.21)	(1.44)	(0.80)	(0.59)	(0.48)
Piece rate choice	4.49	5.15	3.73	5.97	4.78	4.89
	(0.46)	(0.41)	(0.34)	(0.48)	(0.22)	(0.15)
Study (shares)						
International Business	0.54	0.59	0.49	0.61	0.55	0.54
Economics	0.19	0.22	0.22	0.15	0.19	0.24
Law	0.08	0.07	0.08	0.12	0.09	0.06
Psychology	0.00	0.00	0.00	0.00	0.00	0.01
Culture studies	0.00	0.00	0.00	0.00	0.00	0.01
Other	0.20	0.11	0.22	0.12	0.16	0.13
Observations	37	27	37	33	134	263

NOTE: Table 3 reports means by treatment condition. Standard error of the mean is shown between parentheses. The estimation sample consists of subjects who make an incentivized choice. Risk attitude is the principal component of the lottery choice task (Dohmen et al., 2011) and stated willingness to take risks. Higher values indicate a greater willingness to take risk. Productivity refers to the number of completed sequences during the productivity elicitation stage.

Table 4 - OLS-regression of piece rate choice on treatment

	(1) Piece rate
Treatment:	
No risk	Baseline
High risk	-0.32
	(0.41)
Risk attitude (principal component)	0.53**
	(0.24)
Productivity (std.)	0.80***
	(0.18)
Constant	7.19***
	(0.75)
Controls	Yes
Observations	134
R-squared	0.29

NOTE: OLS estimates. The dependent variable in all columns is the chosen piece rate contract (0-10). Standard errors are in parentheses. Stars indicate significance: *** p<0.01, ** p<0.05, * p<0.1. Productivity is standardized. Risk attitude is the standardized principal of the subjective risk attitude measure and the lottery task (Dohmen et al., 2011). Higher values indicate a higher willingness to take risk. Controls include gender, nationality, field of study, and study year.

Table 5 - OLS-regressions of piece rate choice on treatment

	(1) Piece rate	(2) Piece rate
Treatments:		
Low reference point, high risk (β_{lh})	Baseline	Baseline
Low reference point, no risk (β_{ln})	0.76	1.15**
	(0.57)	(0.52)
High reference point, no risk (β_{hn})	1.42**	1.72***
	(0.63)	(0.57)
High reference point, high risk (β_{hh})	2.24***	2.19***
	(0.59)	(0.52)
Productivity (std.)		0.69***
		(0.17)
Risk attitude (principal component)		0.56**
		(0.23)
Constant	3.73***	4.75***
	(0.41)	(0.62)
Controls	No	Yes
Observations	134	134
R-squared	0.11	0.38

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

NOTE: OLS estimates. The dependent variable in all columns is the chosen piece rate contract (0-10). Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1. Productivity is standardized, risk tolerance is the standardized first principal component of the subjective risk attitude measure and the lottery task. Higher values indicate a higher willingness to take risk. Controls include gender, nationality, field of study, and study year.

Appendix A

Proof of hypotheses 1.1 and 1.2.

As described in the main text, the utility function is described by:

$$U = u\big(w - c(e)\big)$$

where

$$w = s + be(1 + \varepsilon)$$

Optimal piece rate choice follows from the first-order condition:

$$\frac{dE(U)}{db} = \int_{\underline{\varepsilon}}^{\overline{\varepsilon}} u'\big(s + be(1 + \varepsilon) - c(e)\big)\left(e(1 + \varepsilon) + \frac{ds}{db}\right)\varphi(\varepsilon)d\varepsilon = 0,$$

Using that the expected product of two random variables X and Y can be written as $E[XY] = E[X]E[Y] + cov[X; Y]$, we can rewrite this expression as:

$$\frac{dE(U)}{db} = E[u'(\cdot)]E\left[e(1 + \varepsilon) + \frac{ds}{db}\right] + cov\left[u'(\cdot); e(1 + \varepsilon) + \frac{ds}{db}\right] = 0$$

Varian (1992, p. 380) argues that when random variables X and Y are bivariate normally distributed, $cov[U(X), Y] = E[u''(X)]cov[X, Y]$. Exploiting this property we get:

$$\frac{dE(U)}{db} = E[u'(\cdot)]E\left[e(1 + \varepsilon) + \frac{ds}{db}\right] + E[u''(\cdot)]cov[be\epsilon; e\epsilon] = 0,$$

which can be rewritten as

$$\frac{dE(U)}{db} = e + \frac{ds}{db} - rbe^2\sigma^2 = 0,$$

where $r = -\frac{E[u''(\cdot)]}{E[u'(\cdot)]}$ is a measure of global risk aversion. In what follows, assume that r is constant. $r = -\frac{E[u''(\cdot)]}{E[u'(\cdot)]}$Likewise, we derive the first-order condition for optimal effort:

$$\frac{dE(U)}{de} = b - c'(e) - rb^2 e\sigma^2 = 0$$

To show that $\frac{db}{d\sigma^2} < 0$ and $\frac{db}{dr} < 0$, we totally differentiate both first-order conditions and solve for db. We obtain after rewriting:

$$\frac{db}{d\sigma^2} = rbe \frac{-c''(e)e - b(1 - rbe\sigma^2)}{(-c''(e) - rb^2\sigma^2)\left(\frac{ds}{d^2b} - re\sigma^2\right) - (1 - 2rbe\sigma^2)^2} < 0$$

$$\frac{db}{dr} = \sigma^2 be \frac{-c''(e)e - b(1 - rbe\sigma^2)}{(-c''(e) - rb^2\sigma^2)\left(\frac{ds}{d^2b} - re\sigma^2\right) - (1 - 2rbe\sigma^2)^2} < 0$$

From the first-order conditions, we obtain that $rbe\sigma^2 < 1$, implying that the numerator is negative. The denominator is the discriminant of the system of equations, and should be positive for an optimum. Hence, $\frac{db}{d\sigma^2} < 0$ and $\frac{db}{dr} < 0$. Note that a change in σ^2 has an ambiguous effect on the magnitude of $\frac{db}{dr}$, so that it is not clear that risk attitudes have a larger effect on b in a risky environment.

Proof of hypotheses 2.1-2.2.

To analyze how behavior depends on the reference point when risk is introduced, we assume that individuals either have a high or low reference point. The reference point is defined as high when the reference point exceeds the expected earnings net of effort costs for any given piece rate.

Hypothesis 2.1: When individuals have a low reference point, such that their expected earnings net of effort costs exceed the reference point for any given piece rate, they respond to the introduction of risk by choosing lower piece rates.

Proof: First, we analyze how effort of workers changes when risk is introduced. The first-order condition for optimal effort is given by:

$$\frac{dE(U)}{de} = \int_{\varepsilon^*}^{\overline{\varepsilon}} u'(w - c(e) - R)\big(b(1+\varepsilon) - c'(e)\big)\varphi(\varepsilon)d\varepsilon +$$

$$\int_{\underline{\varepsilon}}^{\varepsilon^*} u'(R - \big(w - c(e)\big))\big(b(1+\varepsilon) - c'(e)\big)\varphi(\varepsilon)d\varepsilon = 0$$

As can easily be seen from the first order condition, optimal effort in the absence of risk is described by $b - c'(e) = 0$. Substituting $b - c'(e) = 0$ into the first order condition to compare with the risk-free situation, we can write the first-order condition as follows:

$$\int_{\varepsilon^*}^{0} u'(s+be(1+\varepsilon)-c(e)-R)b\varepsilon\varphi(\varepsilon)d\varepsilon + \int_{0}^{-\varepsilon^*} u'(s+be(1+\varepsilon)-c(e)-R)b\varepsilon\varphi(\varepsilon)d\varepsilon$$

$$+\int_{\underline{\varepsilon}}^{\varepsilon^*} u'\left(R-\left(s+be(1+\varepsilon)-c(e)\right)\right)b\varepsilon\varphi(\varepsilon)d\varepsilon$$

$$+\int_{-\varepsilon^*}^{\overline{\varepsilon}} u'(s+be(1+\varepsilon)-c(e)-R)b\varepsilon\varphi(\varepsilon)d\varepsilon$$

The first two parts capture relatively small negative and positive shocks, respectively, so that income is always in the gain domain. The sum of those two parts is negative. Each positive shock is matched with a negative shock of equal magnitude, but negative shocks receive a larger weight since $u'(\cdot)$ in case of a positive shock is smaller than when a shock of the same magnitude is negative. The other two parts capture larger negative and positive shocks, respectively. In this case, negative shocks lead to losses relative to the reference point. To compare the marginal utilities of income, we exploit that $E(w)-c(e)>R$. This condition implies that losses are smaller than gains when positive and negative shocks are of equal magnitude:

$$R-\left(s+be(1-\varepsilon)-c(e)\right) < s+be(1+\varepsilon)-c(e)-R$$

By concavity of the utility function, $u'(\cdot)$ in case of a positive shock is smaller than when the shock is negative, and the final two terms are also negative. This implies that $b-c'(e)>0$ to satisfy the first-order condition, hence effort needs to be reduced as compared to optimal effort when there is no risk.

By similar reasoning, we can show that, for given effort, individuals choose lower piece rates than in the absence of risk. The first-order condition for optimal piece rate is given by:

$$\frac{dE(U)}{de} = \int_{\varepsilon^*}^{\overline{\varepsilon}} u'(w-c(e)-R)\left(\frac{ds}{db}+e(1+\varepsilon)\right)\varphi(\varepsilon)d\varepsilon +$$

$$\int_{\underline{\varepsilon}}^{\varepsilon^*} u'(R-\left(w-c(e)\right))\left(\frac{ds}{db}+e(1+\varepsilon)\right)\varphi(\varepsilon)d\varepsilon = 0$$

Comparing with the optimal piece rate in the absence of risk by substituting $e=-\frac{ds}{db}$, and rewriting to distinguish between small and large shocks, we obtain:

$$\int_{\varepsilon^*}^{0} u'(s+be(1+\varepsilon)-c(e)-R)e\varepsilon\varphi(\varepsilon)d\varepsilon + \int_{0}^{-\varepsilon^*} u'(s+be(1+\varepsilon)-c(e)-R)e\varepsilon\varphi(\varepsilon)d\varepsilon$$

$$+\int_{\underline{\varepsilon}}^{\varepsilon^*} u'\left(R-\left(s+be(1+\varepsilon)-c(e)\right)\right)e\varepsilon\varphi(\varepsilon)d\varepsilon$$

$$+\int_{-\varepsilon^*}^{\overline{\varepsilon}} u'(s+be(1+\varepsilon)-c(e)-R)e\varepsilon\varphi(\varepsilon)d\varepsilon$$

This expression is negative since $u'(\cdot)$ in case of a positive shock is smaller than when a shock of the same magnitude is negative. Hence, to satisfy the first-order condition, $\frac{ds}{db} + e > 0$, implying that the optimal b is lower than in the absence of risk, assuming effort does not change. Since the introduction of risk also lowers effort for given incentives, the tendency to choose lower piece rates is reinforced.■

Hypothesis 2.2: When individuals have a high reference point, such that their expected earnings net of effort costs fall below the reference point for any given piece rate, they respond to the introduction of risk by choosing higher piece rates.

Proof: the proof is the same as for Hypothesis 2.1. Since the utility function is concave in losses, $u'(\cdot)$ in case of a positive shock is larger than when the shock is negative, implying that the optimal effort and piece rate are higher than in the absence of risk.■