More effort with less pay: On information avoidance, belief design and performance

by Steffen Huck, Nora Szech, Lukas M. Wenner

WORKING PAPER SERIES IN ECONOMICS

More Effort with Less Pay:

On Information Avoidance, Belief Design and Performance[*]

Steffen Huck   Nora Szech   Lukas M. Wenner

September 25, 2015

Abstract

In a tedious real effort task, subjects know that their piece rate is either low or ten times higher. When subjects are informed about their piece rate realization, they adapt their performance. One third of subjects nevertheless forego this instrumental information when given the choice — and perform stunningly well. Agents who are uninformed regarding their piece rate tend to outperform all others, even those who know that their piece rate is high. This also holds for enforced instead of self-selected information avoidance. All our findings can be captured by a model of optimally distorted expectations following Brunnermeier and Parker (2005).

Keywords: Optimal Expectations, Belief Design, Performance, Real Effort Task, Coarse Incentive Structures, Workplace Incentives

JEL-codes: D83, D84, J31, M52

---

1 Introduction

Orthodox economic theory posits that agents have a non-negative willingness to pay for instrumental information, that is, information that may affect their subsequent choices. They have no reason to refuse information that comes for free. In the workplace, for example, agents would want to know their precise wage. Knowing the pay schedule allows agents to adjust their performance optimally, balancing costs and expected rewards.

Recently, this view has been challenged. When anticipations matter, agents may have an incentive to avoid information (see, e.g., Caplin and Leahy, 2001, 2004; Bénabou and Tirole, 2002; Kőszegi, 2003; Brunnermeier and Parker, 2005; Schweizer and Szech, 2014). Optimal expectations will balance the psychological benefits of designed beliefs about the future with the material costs of making suboptimal choices. This recent literature studies how agents actively control their expectations. Optimal beliefs often turn out to be coarse. If an agent decides just for him- or herself, anticipatory utility can provide a reason to avoid information structures that are too precise.

Oster et al. (2013) demonstrate the power of information avoidance and belief design in the context of medical testing for the hereditary Huntington's Disease. They show that a large fraction of people who are at risk shy away from medical testing despite costs of testing being small and behavioral adjustments to test results being large. Their data also show that people who avoid getting tested seem to do pretty well in life. This suggests that information avoidance may be beneficial, at least for a substantial fraction of people. Oster et al. (2013) conclude that a model of belief design in the spirit of Brunnermeier and Parker (2005) captures observed behavior much more accurately than orthodox neoclassical approaches. In the latter, the only reason for information avoidance can be found in prohibitive costs of obtaining information. Furthermore, self-selection into information avoidance may play an important role for the findings as people decide themselves if they want to get tested for the disease or not.[1]

We want to explore whether preferences for information avoidance can also be found in a less extreme but more familiar economic setting: the workplace. We conduct a real-effort experiment (with a strenuous task). Across all treatments, subjects know that the piece rate

---

[1] In a less drastic context, when offering tests for herpes infections, Ganguly and Tasoff (2015) also find a preference for information avoidance in a substantial fraction of subjects.

is either high (1 EUR) or low (0.1 EUR), with equal probabilities. We first establish that information about this piece rate is instrumental. Our data show substantially greater effort and output for the high piece rate compared to the low piece rate when subjects are informed about their piece rate at the outset. Nevertheless, we observe a sizable fraction of subjects (robustly around one third of all subjects) who prefer not to receive precise and costless information about their piece rate when given the choice. We will refer to these subjects as information avoiders. When asked in a post-experimental questionnaire why they decided to stay uninformed, two main reasons emerge. Some information avoiders state that they did not want to be demotivated by a low piece rate. Others mention their aversion to too much pressure caused by a high piece rate. To the best of our knowledge, this presents the first laboratory evidence for the prevalence of information avoidance in an economic context in which information is instrumental.

In line with Oster et al. (2013), our data further reveal that subjects who avoid information perform highly significantly better than those who decide to learn their piece rate before starting to work. Our context allows us to explore whether self-selection into information avoidance is crucial for this finding. In a no-information treatment, we force subjects into not knowing their piece rate realizations. Performance results are again stunning and not different from those of self-selected information avoiders. Overall, performances under information avoidance are not only higher than performance results under complete information about the piece rates - they also tend to be higher than performance results of subjects who know that they are going to receive the high piece rate for sure. From a designer's perspective, it thus pays to equip subjects with coarse information structures.

As a theoretical explanation, we propose a model in the spirit of Brunnermeier and Parker (2005) incoporating heterogeneous agents.[2] This variation captures potential differences in the way agents react to performance schemes. Whereas for some agents higher rewards may unambiguously increase motivation and effort, for others such high rewards may have adverse effects. Specifically, when the piece rate is high, some agents may choke under pressure, a phenomenon that has received considerable attention in the psychology literature since Baumeister's seminal (1984) paper and that has more recently also been documented in a

---

[2]Answers from our ex-post questionnaire suggest a heterogeneity in fear of choking under pressure. For a homogenous population with some fear of choking, the Brunnermeier and Parker (2005) model captures all our findings straight away.

number of economic studies (see, for example, Dohmen, 2008; Ariely et al., 2009; Apesteguia and Palacios-Huerta, 2010).

In a population with heterogeneous agents, some uncertainty in incentives may prove superior to any fixed reward system. Coarse information structures may allow different types of agents to adjust their expectations in different ways, according to their personal preferences, and to bias beliefs according to their individual needs. Those who value the motivation-enhancing effect of high wages may bias their beliefs towards optimistic piece-rates, which increases their output. Those who dread choking may bias their beliefs in less drastic ways, thereby enhancing their performance. Optimal incentive design might, thus, make deliberate use of uncertainty. Indeed, tournament incentives which are widely used in the workplace might exploit this very mechanism. Likewise, the coarseness of payment contracts involving shares of the employing company may affect performance results even in workers on lower levels of a company's hierarchy.

Our paper is organized as follows. Section 2 describes the design, the procedures of our experiment, and the hypotheses. Section 3 presents our results and sets them into the context of the related literature. Section 4 proposes a variation of the Brunnermeier and Parker (2005) model allowing for heterogeneous agents, and Section 5 concludes.

2 Experimental Design

We conduct three main treatments, FULL INFO, NO INFO, and INFO CHOICE. These treatments are identical except that information about piece rates varies—exogenously or endogenously.

In all main treatments, subjects know that they receive an either a high (1 EUR) or a low (0.1 EUR) piece rate, with equal probabilities, for working on a tedious task.[3] They have 60 minutes to enter lines of strings, containing numbers, upper case and lower case letters, *backwards* into the computer interface. For each correctly entered string, they are going to receive their piece rate. Each string consists of 60 characters. For example, one of the strings used in the experiment looks as follows.

NXgCX7JHxYZj2cfBSd8JtkYp3LPcyDX8y8NNQhrzJfg22S2ACjC85EQ43B7L

---
[3]At the time of the experiment, 1 EUR $\approx$ 1.37 USD.

Each task consists of one of these randomly generated strings and all tasks are identical for all subjects.

We vary across treatments how much information subjects have about their piece rate while working on the task. In treatment FULL INFO, every subject gets to know his or her piece rate immediately, before working. In treatment INFO CHOICE, subjects can choose whether to know their piece rate immediately, or whether they only want to find out about it after they have finished working on the task. Thus subjects can avoid information if they want to. In treatment NO INFO, subjects are forced into information avoidance, but know the probability distribution over wages. At the end of the experiment, we ask subjects to provide us with some basic demographic information about themselves. In treatment INFO CHOICE we also ask them to state their reasons for choosing (not) to obtain information about their piece rate before working on the task.

We hypothesized that performance would depend on the piece rate, such that subjects would perform better under the high piece rate than under the low one. Treatment FULL INFO allows us to find out whether information about the piece rate is indeed instrumental. Furthermore, in line with Oster et al. (2013) and Ganguly and Tasoff (2015), we expected that nevertheless, a substantial fraction of subjects may decide to avoid information about the piece-rate in INFO CHOICE. Further, as suggested in Oster et al. (2013), performance may not suffer from information avoidance. If so, treatment NO INFO allows us to uncover potential causal effects of self-selection into information avoidance.

We also implement another treatment, MEDIUM WAGE, where every subject earns a piece rate of 0.55 EUR. This treatment allows us to disentangle whether subjects stick to their Bayesian prior when avoiding information, or if there is a tendency that a substantial fraction of subjects distort their beliefs behaviorally as suggested in Brunnermeier and Parker (2005). This treatment informs our modeling approach in Section 4.

Implementing the different information structures across treatments. The detailed procedures in our main treatments, FULL INFO, NO INFO, and INFO CHOICE, are as follows. When entering the lab, each subject is randomly allocated a red or a black chip by one of the experimenters. Half of the chips are black, the other half red. Each subject is then told to take a seat at a computer terminal where the screen shows a square with the color

corresponding to the color of his or her chip.[4] Subjects know that depending on the color (red or black), they can either earn 0.1 EUR or 1 EUR for each correctly entered string, with equal probability.

In order to determine which color corresponds to a high piece rate and which to a low one, we use the following procedure. We prepared two pieces of cardboard which look identical from the outside when folded, but inside either show a red or a black square. After showing the cardboard pieces (from outside and inside) to all participants, they are folded, secured with paper clips, placed into a small bag and shuffled. Another experimenter then draws one of the two folded cardboard pieces. The color of the drawn piece determines which color is associated with the high wage for this session.

In treatment FULL INFO, the cardboard is unfolded and the color is revealed to all subjects immediately. Thus subjects know whether they are going to receive the high piece rate, or the low one.

In treatment NO INFO, the folded cardboard is placed onto a white board at the front of the room where it remains for the whole duration of the experiment and is revealed to all participants once the allowed time for the task (60 minutes) is up. Hence subjects do not know whether they earn the high or the low piece rate when working on the task.

In treatment INFO CHOICE, subjects are asked on their computer screen whether they would like to receive the information about the color now, or wait until the end of the experiment. After clicking the button corresponding to their choice, another screen appears which states the subject's decision. After all subjects have made their choice, the experimenter walks through the lab and privately reveals the color inside the cardboard to those subjects who decided to see it. As in the NO INFO treatment, the folded cardboard is then placed onto a white board and revealed afterwards. Thus, in this treatment, subjects choose whether they want to know their piece rate beforehand or not. Then the real effort task starts.

After each string that subjects enter, they learn whether they entered it correctly. They can then click on a button to continue. Subjects are informed about the time that remains. Subjects are not allowed to use any electronic devices, but are each given a copy of a well-known German weekly, called DER SPIEGEL. This magazine has a weekly circulation of more than one million. It contains all sorts of articles, from investigative journalism over

---

[4]The procedure makes sure that subjects can see that there are chips of two colors and that the procedure is entirely random.

reports on politics to articles about scientific discoveries and information on cultural events and sports. Subjects are explicitly told that they can make use of the magazine "...whenever, during the experiment, [they] would like to take a break or pass time". Thus, no subject has to feel obliged to work on the task if he or she prefers to spend their time otherwise.

In the MEDIUM WAGE treatment, there is no need for any randomization in the beginning and subjects immediately start working on the task after reading the instructions. In order to keep the context as comparable as possible, we also inform subjects that other participants could earn either 0.1 EUR or 1 EUR for the same task in previous experiments, with equal probability.

In total, our sample consists of 238 subjects. All treatments were run at the WZB-TU Laboratory in Berlin between November 2013 and April 2014. There were no restrictions imposed on the invited participants regarding gender, subject of study, or previous experience with experiments. We used z-tree (Fischbacher, 2007) as the experimental software and ORSEE (Greiner, 2004) to recruit subjects. Participants received a show-up fee of 5 EUR and average earnings over all treatments amounted to 14.29 EUR. Each session lasted 80 to 90 minutes.[5]

3 Results

In treatment FULL INFO, subjects know their piece rate realization before starting to work on the task. Our data show that subjects perform significantly better under the high versus the low piece rate. Subjects working for the low piece rate of 0.1 EUR solve 20.67 tasks on average, whereas subjects working for the high piece rate of 1 EUR solve 26.21 tasks correctly. This difference is significant (p=0.043)[6] and confirms that information about the piece rate is instrumental as hypothesized.[7]

We expected that nevertheless, in treatment INFO CHOICE, a substantial fraction of subjects would avoid information about the piece rate and postpone learning the piece rate to the end of the experiment. This is exactly what we find. 30 out of 95 subjects (31.6%)

---

[5]We also ran two sessions without the weekly magazine. We report the findings from these sessions as well. Overall, results are quite similar to those in our main treatments, though effects are slightly less pronounced.

[6]Unless indicated otherwise, all p-values are calculated using a Wilcoxon Rank Sum test.

[7]The other possible measure of effort would be the number of attempted tasks. Arguably, our task is prone to errors and for some subjects this measure might more correctly reflect the actual effort put in. Others might choose a more risky strategy and tolerate more errors. The two measures are highly correlated ($\rho = 0.8763$, p-value=0.0000) and the results are very similar.

decide to avoid information. They prefer to learn their piece rate only after they have worked. These subjects thus decided to forgo information that turned out to be of instrumental value as shown in treatment FULL INFO.[8]

In a post-experimental questionnaire, we ask information avoiders about their motives for not getting informed about their piece rates. There are basically two types of answers that subjects provide. Several subjects tend to argue that they wanted to avoid being demotivated in case of having a low piece rate. Other subjects stressed that they were afraid of the pressure in case of learning that they receive a high piece rate.[9]

Looking at performance, information avoiders do extremely well. They solve 30 tasks correctly on average, while information receivers solved 21.31 tasks correctly (p=0.0002). Strikingly, information avoiders even tended to outperform the subgroup of subjects in INFO CHOICE who received the information that their piece rate was high (30 versus 25.53, p=0.0573).

Potentially, as the answers to the ex post questionnaire suggest, information avoiders work under beliefs that motivated their performance in better ways than those subjects who opted for information. Yet as we argued above, treatment INFO CHOICE does not allow us to rule out potential effects of self-selection on unobservables that correlate with performance. For example, high-ability subjects might have self-selected into information avoidance, which would explain the observed performance differential. In order to identify the role of self-selection, we ran treatment NO INFO in which subjects have no option to learn their piece rates before working. As in the other main treatments, subjects knew that piece rates were either 1 EUR or 0.1 EUR, with equal probability. Yet subjects knew that their piece rate was only disclosed to them at the end of the experiment. Subjects' information structure in NO INFO is, hence, similar to the one that information avoiders chose in INFO CHOICE, with the only difference that self-selection was not possible in the NO INFO treatment. Our data reveal that performance results between these two groups of subjects are almost identical and statistically indistinguishable (30 versus 28.02, p=0.3710).[10] The data show that not knowing the piece rate realization enhances performance, even if subjects do not freely opt

---

[8]As a robustness check, we analyze the two sessions that where identical to INFO CHOICE except that subjects did not have access to the magazine. There, 15 out of 44 subjects (34.1%) decided not to acquire information, an effect of almost identical magnitude.

[9]We analyse the subjects' statements in more detail below.

[10]There is no significant difference either for any of the other two groups that receive information about their piece rates, at any conventional level of statistical significance.

piece rate	FULL INFO		INFO CHOICE		NO INFO		MEDIUM WAGE	
	mean (s.d.)	median	mean (s.d.)	median	mean (s.d.)	median	mean (s.d)	median
0.1	20.67 (10.49) N=24	22	17.69 (11.37) N=35	17				
0.55							24.66 (9.58) N=47	24
1	26.21 (8.75) N=24	26	25.53 (9.86) N=30	23				
unknown			30 (9.35) N=30	29	28.02 (8.41) N=48	28		
aggregate	23.44 (9.17) N=48	25	24.05 (11.44) N=95	23	28.02 (8.41) N=48	28	24.66 (9.58) N=47	24

Table 1: Mean and median performance across treatments

for this coarse information structure. As self-selection does not lead to statistically different performance results, we can pool the data for the different information structures under which subjects worked. From a designer's perspective, opting for the NO INFO policy leads to significantly better performance results than FULL INFO (23.44 versus 28.02, p=0.0274).

Information avoidance enhances performance, even if it is enforced. When asked, subjects argue with controlling their beliefs while working. Both, avoiding demotivation from a potentially low piece rate as well as avoiding too much pressure from a potentially high piece rate are stated as motives for information avoidance. This suggests that there may be a heterogeneity in our subjects, and that, at least for some of them, their performance varies non-monotonically with the wage they expect to receive. If all subjects were to use their Bayesian prior when they do not know their exact wage, their expected wage is 0.55 EUR. If this wage is sufficiently high to avoid demotivation and sufficiently low to prevent choking under pressure, the performance at a fixed wage of 0.55 EUR should not differ much from the performance under information avoidance. However, if, for at least some subjects, their performance-maximizing wage is different from 0.55 EUR, information avoidance may allow them to optimally design their expectations as described in Brunnermeier and Parker (2005), away from the Bayesian prior. With room for this kind of behavioral belief distortion, information avoidance may even outperform any kind of fixed reward system in a heterogenous

piece rate	POOLED		p-value of pairwise test		
	mean (s.d.)	median	0.1	0.55	1
0.1	18.90 (11.03)	19			
0.55	24.66 (9.58)	24	0.0091		
1	25.83 (8.56)	25	0.0003	0.2073	
unknown	28.78 (8.78)	28	0.0000	0.0015	0.0645

Table 2: Pairwise comparisons of performance results across treatments

population. In order to understand whether there is evidence for behavioral belief design, we ran treatment MEDIUM WAGE. In this treatment, subjects know that they work for a piece-rate of 0.55 EUR. They also know that other subjects worked for either 1 EUR or 0.10 EUR with equal probability, in order to keep the context comparable. We find that subjects solve on average 24.66 tasks correctly in this treatment.

Table 1 presents an overview of all our results per piece rate and treatment. Table 2 displays the performance results for the different information structures subjects had while working on the task, pooling the data of self-chosen and enforced information structures. We find that the result from above regarding the effort level for the known wages 0.1 EUR and 1 EUR carries over: Subjects perform significantly better at the higher wage (18.90 vs. 25.83 correct tasks on average, p=0.0003).[11] On the other hand, for subjects who did not know whether their piece rate is 0.1 EUR or 1 EUR, we obtain an average performance of 28.78 correctly solved tasks. In other words, a participant who does not know his wage solves about 3 more tasks than a subject who knows that he or she receives the high wage of 1 EUR per solved task (p=0.0645).[12] Figure 1 plots the empirical distribution functions for the four cases. Visually, the distribution for the case in which the piece rate is unknown almost first-order stochastically dominates the distributions for all other treatments, in which piece rates were deterministic. Performance results under information avoidance are also statistically different from performance results under the secure 0.55 EUR piece rate (p=0.0015). Therefore, while it may be that a fraction of subjects sticked to the undistorted Bayesian

---

[11]It appears that most subjects did not take longer breaks during the one hour working time. Even under the low piece rate of 0.1 EUR, less than 30 percent of subjects spend more than ten minutes with any task. Looking at the other piece rates (including unknown), no subject paused for a longer interval.

[12]Using attempted tasks as the effort measure we find a similar effect: when the wage is unknown subjects attempt on average 40.51 tasks as compared to 35.46 at the certain wage of 1 EUR (p=0.0052).
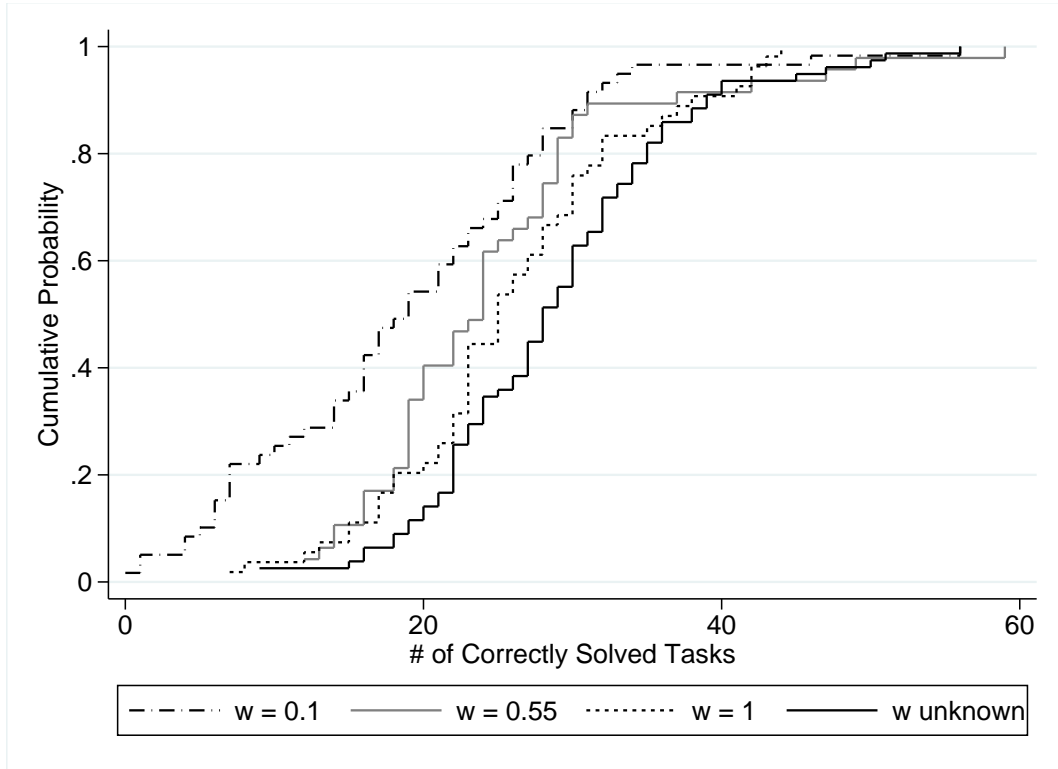
Figure 1: CDFs of effort choices, by wage.

expectation of 0.55 EUR under information avoidance, a substantial fraction of subjects are likely to have distorted their beliefs in behavioral ways.

All in all, from a designer's perspective, giving people the choice to avoid information does not harm aggregate performance compared to a full information policy. Subjects who self-select into information avoidance perform extremely well. This result is in line with findings from Oster et al. (2013). Moreover, enforcing information avoidance significantly increases performance results. This effect must be partly due to behavioral belief distortion as described in Brunnermeier and Parker (2005). If subjects mostly tended to stick to the Bayesian prior under information avoidance, performance results would have to be roughly equal to the MEDIUM WAGE treatment, which is not the case.

Our data further suggest that a policy of information avoidance tends to outperform paying the high piece rate for sure, even though the latter is much more costly from a designer's perspective. An explanation for this, as put forward by several subjects themselves, could be a fear of choking under pressure in a substantial fraction of subjects. This potential heterogeneity informs our modeling approach in Section 4. A model in the spirit of Brunnermeier and Parker (2005) allowing for heterogeneous preferences can capture our observations. Of

course, it may also be that subjects homogeneously distort their beliefs to some moderately optimistic belief, such as a piece rate of 0.80 EUR. If all subjects were to some extent afraid of choking under pressure, performance results could turn out higher than performance results for the high piece rate of 1.00 EUR. As several subjects tended to argue either with choking under pressure or with avoiding potential demotivation from a low piece rate, we will elaborate the case of heterogenous agents in a population in Section 4.

Discussion. We would first like to point out that our results on performance and information avoidance are inconsistent with classical models of expected utility in which anticipatory feelings do not matter. Consider a utility function that is separable in monetary payoffs and effort costs, and make the natural assumption that the cost of effort is increasing. Then, irrespective of the agents' risk attitudes, the optimal level of effort is predicted to be higher in the case where the agent works for a fixed piece rate compared to the case where he or she faces a lottery over this piece rate and any lower piece rate.

At the end of the experiment, after subjects finished working on the task, we asked them to state the reasons for obtaining (avoiding) information about their wage. Among those who decided to obtain information, the vast majority of subjects (49 out of 65) state that they wanted to know their piece rate in order to adjust their effort accordingly. These subjects thus acknowledge the instrumental role of information for them.[13] A minority of people (13 out of 65) explicitly state preferences for information, e.g. "curiosity", as their reason for obtaining information. Looking at subjects who avoided information, a more diverse picture emerges. On the one hand subjects state that by not knowing their wage they wanted to ensure that they would be sufficiently motivated in light of the risk of receiving a low wage (7 out of 30), and on the other hand subjects reveal that they are afraid of being under too much pressure when knowing for certain that their wage is high (8 out of 30).[14] Interestingly, another group of subjects (9 out of 30) explicitly state that both potential demotivation as well as pressure of a high wage influenced their decision, or that they felt they could work best without being influenced too much by a certain wage. All these answers point towards an important role for belief design in our setting. Subjects consciously reflect on they way

---

[13]A typical statement was, for example, "I wanted to know whether putting in effort would be worthwhile, if the wage would have been low, I would not have bothered".

[14]For the first group, a typical statement reads "with the hope of the higher payment, I wanted to keep up my motivation", whereas those potentially choking under the pressure of a high wage gave answers like "I chose not to learn the color in order to take the pressure off myself. I would have been even more error-prone".

their performance will be affected by the wage distribution and deliberately choose to avoid information in order to optimize their performance.

Related findings in the literature. Our finding that roughly a third of participants choose to avoid instrumental information does not have a precedent in the literature in so far as settings as the workplace are concerned. Oster et al. (2013) document evidence for information avoidance in the domain of health outcomes. The authors find that only 7 percent of individuals that can be classified as being at-risk of contracting the hereditary Huntington Disease decide to undergo testing that provides them with certainty about their health status. This effect of information avoidance is significantly larger than ours, but unlike in our setting obtaining information is not costless. Furthermore, effects on anticipatory feelings are arguably much more drastic than in our setting, such that information avoidance may become even more attractive.

In a laboratory study Ganguly and Tasoff (2015) give subjects the option to avoid being tested for herpes at a cost of 10 USD. Depending on the type of virus they are tested for, 5.2% to 15.6% of individuals are willing to give up 10 USD in order not to be tested. While in this setting information about the health status is instrumental, Ganguly and Tasoff (2015) also look at demand for non-instrumental information. Giving subjects the option to learn (or avoid) information about potential monetary payments at some cost, they find that some subjects are willing to pay money to avoid learning about the outcome of a lottery whereas others are willing to pay for early resolution of uncertainty (each of the two groups roughly comprises between 30% and 40% of the sample).

Eil and Rao (2011) find evidence for information avoidance when they elicit subjects' willingness-to-pay of learning their relative rank in terms of IQ and attractiveness. The authors show that subjects who believe themselves to be below average in a category are willing to pay for not learning their true rank in the ability distribution. However, in their setting when subjects were given the option to learn their rank, this information did not have an impact on their earnings in the experiment, and is thus non-instrumental.

Likewise, Eliaz and Schotter (2010) and Falk and Zimmermann (2014) analyze the demand for non-instrumental information. Both studies find evidence that subjects have a preference for obtaining non-instrumental information. In Eliaz and Schotter (2010), a majority of subjects choose to pay a fee to obtain information that will only alter their confidence

about their decision, but not their decision itself. In Falk and Zimmermann (2014), subjects can choose the timing when to find out about whether they are to receive small electrical shocks and a large majority chooses to find out immediately, and can thus be classified as "curious", i.e. having a strict preference for information. Compared to our setting, these results indicate that we might be underestimating the magnitude of the effect of anticipatory utility and the desire for belief design. For subjects in our sample with a high degree of curiosity (we can identify some of them via their answers in the post-experimental questionnaire, see above) or desire for confidence, these motives might outweigh the benefits from being able to optimally design their beliefs when the piece rate is unknown.

In relation to our findings in the domain of effort choices, Ariely et al. (2009) find in experiments conducted in the U.S. and in India for a variety of different tasks that high incentives might backfire and reduce performance. For the majority of the tasks administered in their study, performance varies non-monotonically with the compensation offered; moderate incentives typically deliver the highest performance level. Dohmen (2008) and Apesteguia and Palacios-Huerta (2010) document a similar effect for professional athletes. Building on Baumeister (1984), all these studies identify the phenomenon that individuals may choke under too much pressure induced by high rewards. Therefore, the fear of choking under pressure that a fraction of subjects expressed, in our study, may be rather plausible.

Shen et al. (2015) document that in certain real-effort situations a small reward that is uncertain and either higher or lower (e.g. 1 USD or 2 USD with equal probability, or a smaller versus a larger amount of candy) may generate better performance results than the fixed higher reward (e.g. 2 USD). In all the settings studied by Shen et al. (2015), overall rewards are small. The authors suggest that the uncertainty about these rewards may increase subjects' excitement, and subsequently, their motivation with which they engage in the task. While this may be true when stakes are low, we consider it unlikely that excitement alone drives our results, in which overall stakes are rather high (about 30 EUR if the piece rate is 1 EUR). Various papers show that most agents are risk-averse in economic settings, and try to avoid overly risky lotteries (e.g. Holt and Laury, 2002). As discussed above, in our ex-post questionnaire, many subjects argue that they decided to avoid information in order to prevent demotivation from a low piece rate, while others stated that they wanted to avoid feelings of pressure. Excitement from the lottery does not seem to play a major role

in the explanations subjects gave. In the next section, we adapt the model of Brunnermeier and Parker (2005) on belief design in order to account for potentially heterogeneous goals in behavioral belief design across agents.

4   Theory

All our results can be captured by the model of behaviorally distorted beliefs of Brunnermeier and Parker (2005) (BP henceforth). In the following, we are going to elaborate on the case of a heterogenous population of agents in which some agents are afraid of choking under pressure if the piece rate, or wage w henceforth, is high. We are motivated to do so as in our ex-post questionnaire, there seemed to appear a heterogeneity in preferences across subjects. For a homogenous population in which all agents are a bit afraid of choking under pressure, all our findings can be explained by the BP model directly.

In their model, agents optimally choose their beliefs as to balance benefits from anticipatory emotions and costs in decision making due to biased beliefs. In our experimental setting, subjects either know the wage they are working for (in treatments FULL INFO, MEDIUM WAGE, and if they opted for information to be revealed in INFO CHOICE) or they face uncertainty about whether it is the high or the low wage (NO INFO and INFO CHOICE, if they decided to stay uninformed). Whereas the former case leaves no room for manipulation of beliefs, in the latter subjects might hold subjective beliefs that do not treat the high and the low wage as being equally likely (which would be the Bayesian prior).

In our setting, an agent derives utility from the payment she receives for solving tasks but has to bear the cost of effort. We model effort directly as the number of correctly solved tasks, e, and assume risk neutrality throughout. The (expected) payment is we, where w is the (expected) wage. As described in the previous section, we aim to develop a model that can capture the possibility of choking under pressure, i.e. the phenomenon that an agent's performance might be adversely affected if the (expected) wage for the task is high. We therefore allow the cost of effort not only to depend on e, but also (potentially negatively) on w for these agents. In the specific case we look at below, such a cost function delivers an optimal effort level that is hump-shaped in the wage, i.e. effort is maximized at an intermediate wage. In cases where there is uncertainty about the wage, we interpret w as the expected wage, potentially distorted by optimal belief design by the agent. Assuming

additive separability of monetary payments and effort costs the agents' consumption utility is then given by $u(e, w) = we - c(e, w)$.

At time 0, subjects in treatment INFO CHOICE decide whether to learn their wage or not. Subjects in the other treatments either know their wage by default or are forced into information avoidance. At time 1, subjects decide how much effort to exert when working on the task and they experience *anticipatory utility* based on their expected consumption utility which materializes at time 2. Following BP, we assume that agents who do not know their wage can optimally design their beliefs $\pi \in [0, 1]$, where π denotes the belief that the wage is high, w_H, rather than low, w_L (where $w_H > w_L > 0$).[15] The chosen belief affects anticipatory utility but has no direct effect on consumption utility. Indirectly, however, it does affect consumption utility through the agent's effort choice which will be based on the subjective wage $w(\pi) = \pi w_H + (1 - \pi) w_L$.

Following BP, agents' overall well-being is given by a weighted sum of anticipatory utility and consumption utility where the relative weight of anticipatory utility is denoted by $\delta \geq 0$. Given a belief π and the corresponding expected wage $w(\pi)$, our agent chooses effort e to maximize $w(\pi)e - c(e, w(\pi))$, i.e. the anticipated consumption utility. Expecting the effort choice $e^*(w(\pi))$, the agent maximizes overall well-being by choosing the belief π that balances anticipatory feelings and final consumption utility. In our setting, the belief choice directly corresponds to a choice of the subjective wage. Hence, we can therefore suppress the dependence on π and allow the agent to choose $w \in [w_L, w_H]$ directly. Defining $\bar{w} = \frac{1}{2}(w_H + w_L)$, our agent then maximizes

$$U(w|\bar{w}) = \delta \left[we^*(w) - c(e^*(w), w) \right] + \bar{w}e^*(w) - c(e^*(w), \bar{w}). \tag{1}$$

The optimal choice of beliefs has to consider the following trade off: an agent may distort beliefs away from $\bar{w}$ in order to increase her anticipatory utility. By being more optimistic about her odds to be paid the high wage, she manipulates herself into exerting more effort since effort is determined by the subjective expected wage. However, choosing a belief different from $\pi = 0.5$ may come at a cost because the agent will exert more effort than what is optimal given $\bar{w}$. In general, choking and non-choking agents differ in their choice

---

[15]Note that allowing for the possibility that beliefs can be freely (and therefore in a non-Bayesian manner) chosen is what makes this model more suitable for our purposes than, for example, Bénabou and Tirole (2002). In addition, the latter model focuses on how agents forget or suppress bad news through selective memory, rather than optimally distorting beliefs in the presence of uncertainty.

of how much to distort their beliefs. Under our assumptions on the cost function, by being more optimistic, non-choking agents unambiguously increase their anticipatory utility and the more they care about it the more optimistic they will be. Agents who choke at high wages, however, will prefer to distort their beliefs less in order to work at a (subjective) expected wage that is lower than w_H.

It is worth noting that the wage that enters the cost function differs between the anticipatory utility term and the consumption utility term. When the agent experiences "true" consumption utility, her cost of effort is given by the undistorted expected wage $\bar{w}$. Nevertheless, when anticipating these costs we assume that they are based on the subjective wage w.

In order to demonstrate how this variation of BP can account for our experimental findings, we formalize the above intuition using the family of cost of effort functions given by

$$c(e, w) = \frac{1}{2}\alpha e^2 + \gamma e f(w)$$

with $\alpha > 0$, $\gamma \geq 0$, and $f(w) > 0, f'(w) > 0, f''(w) > 0$.

It is straightforward to see that for a given w, optimal effort is given by

$$e^*(w) = \frac{1}{\alpha}\left(w - \gamma f(w)\right).$$

Optimal effort e^* is strictly concave in the (expected) wage and if γ is sufficiently large, e^* is decreasing for sufficiently large w. We add two more assumptions which essentially impose that γ is not too high. We assume that $w_L > \gamma f(w_L)$ and $w_H > \gamma f(w_H)$ so that optimal efforts e^* are positive over $[w_L, w_H]$. Moreover, we assume $\gamma f'(w_L) < 1$ so that optimal efforts are increasing in wage for small wage levels, though not necessarily for high wages.

As desired, our model delivers an optimal effort level that can be hump-shaped in the wage. Specifically, we define $\hat{w}$ as the unique wage that maximizes effort. If an interior solution $\hat{w} < w_H$ exists, it solves $\gamma f'(\hat{w}) = 1$. This means that for any agent with $w_H > \hat{w}$ effort is maximal at a wage level lower than w_H. We think of these agents as choking under the pressure of a wage that is too high. To simplify the exposition below, we also assume that $\hat{w} > \bar{w}$. If there is no $\hat{w}$ solving $\gamma f'(\hat{w}) = 1$, we know that e^* is increasing and $\hat{w} = w_H$. Define $\hat{\gamma}$ as $\hat{\gamma} = 1/f'(w_H)$. Then, the agents who do not choke under pressure, that is, those

with $\hat{w} = w_H$, can be characterized as the agents with $\gamma \leq \hat{\gamma}$. These agents always exert more effort as the wage increases.

Before presenting the solution for the full optimization problem, let us consider an agent who only cares about anticipatory utility. This agent will choose w to maximize $we^*(w) - c(e^*(w), w)$. By the envelope theorem, the first order condition for this problem is given by

$$e^*(w)\frac{de^*(w)}{dw} = 0. \tag{2}$$

Since e^* is positive, the unique maximizer of anticipatory utility is the wage $\hat{w}$ which maximizes effort. This holds both for the interior solution of the choking agents and for the corner solution $w = \hat{w}$ of the non-choking agents.[16]

Next, consider the other extreme, an agent who does not care about anticipatory utility at all, $\delta = 0$. He or she chooses the wage w to maximize $\bar{w}e^*(w) - c(e^*(w), \bar{w})$, which yields the following first order condition:

$$\frac{de^*(w)}{dw}\left(e^*(\bar{w}) - e^*(w)\right) = 0.$$

It is straightforward to check via the second derivative that $w = \bar{w}$ maximizes the objective function. However, for choking agents there might be a second maximizer due to the hump-shaped nature of effort in wage, namely the effort level that induces the same effort level as $\bar{w}$, if it exists. For our purposes it is immaterial which one of the two the agent chooses. A "standard" agent who is not affected by anticipatory utility chooses the same effort level as an agent who faces a sure wage of $\bar{w}$, like in our MEDIUM WAGE treatment. Thus, for $\delta = 0$ our model nests the neoclassical model since under risk neutrality effort choices should only depend on the average wage.

Putting these two effects together for intermediate values of δ, we obtain the FOC of the full objective function in (1) as:

$$\frac{de^*(w)}{dw}\left(e^*(\bar{w}) - (1-\delta)e^*(w)\right) = 0.$$

Analyzing this expression allows us to derive the solution w^* to the maximization problem.

---

[16]It is straightforward to check that this equation has more than one solution, but that only $w = \hat{w}$ is indeed a maximum, provided that $\hat{w} < w_H$. Hence, for choking agents, the wage that maximizes anticipatory utility is interior and equal to the effort-maximizing wage, whereas for non-choking agents the optimal wage is w_H since for $\gamma < \hat{\gamma}$ sufficiently low the LHS in (2) is positive everywhere on $[w_L, w_H]$.

Proposition 1. *Define $\delta^* = 1 - \frac{e^*(\bar{w})}{e^*(\hat{w})} \in (0,1]$. (a) For all $\delta < \delta^*$, the optimal wages w^* are implicitly defined through the equation*

$$e^*(\bar{w}) = (1 - \delta)e^*(w^*)$$

All solutions induce the same effort and satisfy $w^* > \bar{w}$. (b) For $\delta \geq \delta^*$, $w^* = \hat{w}$.

Proof. See Appendix. □

The proposition shows that in the case where $\gamma < \hat{\gamma}$ (i.e. where the agent does not choke and therefore $\hat{w} = w_H$) a sufficiently large δ implies that the agent will choose $\pi = 1$, that is, he or she chooses effort and receives anticipatory utility under the (distorted) belief that the piece rate will be w_H with probability 1. These non-choking agents then exert the same effort in the case where they do not know their wage and in the case where they have found out that their wage is w_H per task solved. Non-choking agents with a lower value of δ, will adopt interior $\pi \in [0.5, 1)$ because they care relatively more about the upward distortion of consumption utility induced by over-exerting effort.

Agents that choke under pressure, as represented by a positive value of γ, do not distort beliefs in such an extreme way. Since their optimal level of effort is strictly below $e^*(w_H)$, they will, provided they care enough about anticipatory utility (δ is large enough), distort beliefs only up to the point where they exert the maximum level of effort, $e^*(\hat{w})$. Hence, the model delivers the result that these choking agents exert a strictly higher effort in the case where they do not know their piece rate, compared to where they know for sure that they will be paid according to w_H.

Corollary 1. *Consider a group of agents consisting of two types of agents with different parameters $\gamma_2 > \hat{\gamma} > \gamma_1$. If $\delta \geq \delta^*$ then average effort (= average number of correctly solved tasks) of the group will be higher when the agents do not know whether their wage is w_L or w_H than in the case when they all know that that their wage will be w_H.*

Proof. Under the assumptions on γ_1 and γ_2 (which ensures that the effect of choking is sufficiently large as to guarantee that agents of type 2 ("choking agents") put in maximal effort at a wage lower than w_H) and δ (that anticipatory utility is large enough), $w_1^* = w_H$ and $w_2^* = \hat{w}$ holds for the case of unknown wages. Furthermore, agents of type 1 exert effort of $e_1^*(w_H)$ whereas agents of type 2 exert effort of $e_2^*(\hat{w})$. Under a known wage of w_H, the respective effort levels are given by $e_1^*(w_H)$ and $e_2^*(w_H) < e_2^*(\hat{w})$. Choking agents thus exert

18

higher effort when the wage is unknown whereas standard agents exert the same effort level as for a known high wage, proving the statement. □

Our variant of BP can also explain why in treatment INFO CHOICE some agents choose not to obtain information about their wage. An agent who decides whether to learn her wage faces a tradeoff between optimally choosing her effort after having eliminated uncertainty about the wage, but also forgoes the opportunity to benefit from being able optimally to choose her belief and benefit from distorting anticipatory utility upwards. Formally, an agent decides not to learn the wage if

$$\delta \left[w^* e^*(w^*) - c(e^*(w^*), w^*) \right] + \bar{w} e^*(w^*) - c(e^*(w^*), \bar{w}) \geq$$
$$\frac{1}{2}(1 + \delta) \left[w_H e^*(w_H) - c(e^*(w_H), w_H) + w_L e^*(w_L) - c(e^*(w_L), w_L) \right] \qquad (3)$$

Proposition 2. *There exists a $\hat{\delta} \geq 0$ such that all agents with $\delta > \hat{\delta}$ prefer not to know whether their wage is w_H or w_L.*

Proof. Dividing both sides by $(1 + \delta)$ ensures that the RHS of the condition in (3) stays constant once we increase δ. The LHS is then simply a weighted average between anticipatory utility and consumption utility. Using a standard envelope theorem argument, we then see that increasing δ strictly increases the LHS because $w^* e^*(w^*) - c(e^*(w^*), w^*) \geq \bar{w} e^*(w^*) - c(e^*(w^*), \bar{w})$. Hence, for large enough δ, inequality (3) is satisfied. □

To conclude this section, consider again the experimental results described in the previous section. An orthodox model of effort choice without anticipatory utility or choking ($\delta = \gamma = 0$), would predict that for all treatments where the wage is known, average effort is increasing in the wage, $e^*(w_L) < e^*(\bar{w}) < e^*(w_H)$ and that under risk neutrality agents who do not know their wage choose $e^*(\bar{w})$. Also, we should not see anybody rejecting the information about the wage. Our results do not conform to this. Introducing anticipatory utility can remedy this: agents who do not know their wage, optimally choose more optimistic beliefs, inducing themselves to exert more effort. If they value anticipatory utility sufficiently much, they will actively choose to stay uninformed. Observe, however, that such a model predicts that the highest effort level chosen by the uninformed subjects is at most $e^*(w_H)$. It cannot explain our finding that in the aggregate subjects with an unknown wage perform better than $e^*(w_H)$. Hence, in order to fully explain our results, we must incorporate the concept of "choking" into the BP model. We posit that some agents' productivity may be highest at

a wage strictly below w_H. For a known wage these agents' performance may still be as in the standard model, but their effort will be highest at a subjective wage $\bar{w} < w < w_H$. Not knowing the wage might therefore induce them to be most productive. Moreover, if they care enough about anticipatory utility, they will also choose not to learn the wage before they start working.

5   Conclusion

We considered a real-effort task in which agents can choose to receive instrumental information about their piece rate before starting to work. Our data show that more than 30 percent of subjects deliberately decide to forgo learning their piece rate, revealing a preference for information avoidance. Furthermore, agents avoiding information achieve considerably better performance results than agents opting for information.

In order to uncover potential causal effects of self-selection on performance results, we run a treatment in which subjects are forced to stay uninformed about their piece rate. Performance again turns out to be higher than under the certain high wage. Information avoidance, even if enforced instead of chosen, significantly enhances performance. Moreover, performance under information avoidance tends to be better than performance under the certain high piece rate. This is at odds with basic predictions from orthodox economic theory.

Looking into reasons why subjects choose to avoid information suggests that there may be different types of subjects: Some avoid information in order to avoid potential demotivation if the piece rate turns out to be low; others say that learning about a high piece rate could make them feel stressed and lead to choking under pressure. Both effects have been documented in the both the psychology and economics literature.

In an otherwise standard Brunnermeier and Parker (2005) model, we incorporate heterogeneity of agents and the possibility of choking under pressure. This extended version of BP captures all key patterns of our data. A substantial part of subjects appears to have biased their beliefs about the piece rate considerably upwards in order to stay motivated in the task. Other subjects may have behaviorally distorted their beliefs in order to avoid feelings of pressure.

By and large, our study documents that giving agents room to design their beliefs may

not only be beneficial in contexts such as health (as has been documented by Oster et al., 2013 and Ganguly and Tasoff, 2015), but also in economic settings of paid effort exertion and performance in the workplace. While a direct randomization over piece rates may be unpopular with workers and unions, there are other more subtle and perfectly accepted ways of introducing uncertainty about effective pay in a firm. Any type of incentive scheme that introduces interdependencies between workers' payments generates scope for beneficial belief distortion and rampant overprovision of effort in contests may have one of its root causes in anticipatory utility.[17] Additionally, in many, typically large, firms workers are often paid in company shares as part of their salary. Since the individual worker probably has a negligible effect on the evolution of the share's value, spiecifically when he or she works on a lower hierarchy level within the company, this will introduce exogenous uncertainty into the worker's compensation. Our results suggest that this might be an effective way for a firm to increase workers' efforts. Employees may distort their beliefs about the firms' future performance according to their needs, which increases their effort.

The possibility of inducing more effort with less pay is tantalizing. We believe that it offers much scope for further research.

Appendix

Proof of Proposition 1

Using the specific functional form, the first order condition for maximizing (1) can be written via equation (4) as

$$\frac{de^*(w)}{dw}\left(e^*(\bar{w}) - (1-\delta)e^*(w)\right) = 0.$$

Solutions of this condition are either $\hat{w}$, since $\frac{de^*(w)}{dw}|_{w=\hat{w}} = 0$, provided it exists, or the w that solve the term in brackets. In order to solve the term in brackets, $e^*(w)$ must be larger than $e^*(\bar{w})$ and δ must not be too large. The boundary value δ^* is derived from $e^*(\bar{w}) = (1-\delta^*)e^*(\hat{w})$. A solution w^* which sets $e^*(\bar{w}) - (1-\delta)e^*(w)$ to zero exists whenever $\delta \leq \delta^*$. Since e^* is continuous and increasing over $[0, \hat{w}]$, there is always a solution $w^* \in [\bar{w}, \hat{w}]$. There may be further solutions $w^* \in [\hat{w}, w_H]$ which induce the same wage.

---

[17]For a detailed survey of the experimental literature on contests, see Dechenaux et al. (2014) who document evidence for overprovision of effort in a large number of different settings.

Consider the second derivative, given by

$$-(1-\delta)\left[\frac{de^*(w)}{dw}\right]^2 + \frac{d^2 e^*(w)}{dw^2}\left(e^*(\bar{w}) - (1-\delta)e^*(w)\right).$$

If the solution to the problem is given by a solution of $e^*(\bar{w}) = (1-\delta)e^*(w^*)$, this is a maximum because such a solution can only exist for $\delta < \delta^* < 1$, so that the second derivative is negative.

Given the definition of $\hat{w}$, $e^*(w)$ is maximized at $\hat{w}$ if $\hat{w} < w_H$. $\hat{w}$ sets the first term in the second derivative to zero. Furthermore, for any $\delta > \delta^*$, $\frac{d^2 e^*(w)}{dw^2}|_{w=\hat{w}} < 0$ and the term in brackets then is positive. We thus have a maximum here as well. For $\delta < \delta^*$, $\hat{w}$ will be a minimum because in this case $e^*(\bar{w}) - (1-\delta)e^*(\hat{w})$ will be negative. For $\delta = \delta^*$, the second derivative is zero, but for $\varepsilon > 0$, the first derivative at $\hat{w} - \varepsilon$ is positive and at $\hat{w} + \varepsilon$ it is negative. In the case where $\hat{w} = w_H$, and there is no w to satisfy $e^*(\bar{w}) = (1-\delta)e^*(w^*)$, the first derivative in (4) will be strictly positive and thus $w^* = w_H$.

References

APESTEGUIA, J. AND I. PALACIOS-HUERTA (2010): "Psychological Pressure in Competitive Environments: Evidence from a Randomized Natural Experiment," *American Economic Review*, 100, 2548–64.

ARIELY, D., U. GNEEZY, G. LOEWENSTEIN, AND N. MAZAR (2009): "Large Stakes and Big Mistakes," *The Review of Economic Studies*, 76, 451–469.

BAUMEISTER, R. F. (1984): "Choking under pressure: Self-consciousness and paradoxical effects of incentives on skillful performance," *Journal of Personality and Social Psychology*, 46, 610–620.

BÉNABOU, R. AND J. TIROLE (2002): "Self-Confidence and Personal Motivation," *The Quarterly Journal of Economics*, 117, 871–915.

BRUNNERMEIER, M. K. AND J. A. PARKER (2005): "Optimal Expectations," *American Economic Review*, 95, 1092–1118.

CAPLIN, A. AND J. LEAHY (2001): "Psychological Expected Utility Theory and Anticipatory Feelings," *The Quarterly Journal of Economics*, 116, 55–79.

——— (2004): "The supply of information by a concerned expert," *The Economic Journal*, 114, 487–505.

DECHENAUX, E., D. KOVENOCK, AND R. M. SHEREMETA (2014): "A survey of experimental research on contests, all-pay auctions and tournaments," *Experimental Economics, forthcoming*.

DOHMEN, T. J. (2008): "Do professionals choke under pressure?" *Journal of Economic Behavior & Organization*, 65, 636 – 653.

EIL, D. AND J. M. RAO (2011): "The good news-bad news effect: asymmetric processing of objective information about yourself," *American Economic Journal: Microeconomics*, 3, 114–138.

ELIAZ, K. AND A. SCHOTTER (2010): "Paying for confidence: An experimental study of the demand for non-instrumental information," *Games and Economic Behavior*, 70, 304–324.

FALK, A. AND F. ZIMMERMANN (2014): "Beliefs and Utility: Experimental Evidence on Preferences for Information," Working paper.

FISCHBACHER, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

GANGULY, A. AND J. TASOFF (2015): "Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future," Working paper.

GREINER, B. (2004): "An online recruitment system for economic experiments," in *Forschung und Wissenschaftliches Rechnen*, Gesellschaft für wissenschaftliche Datenverarbeitung, vol. 63, 79–93.

HOLT, C. A. AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 1644–1655.

KŐSZEGI, B. (2003): "Health anxiety and patient behavior," *Journal of Health Economics*, 22, 1073 – 1084.

OSTER, E., I. SHOULSON, AND E. R. DORSEY (2013): "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease," *American Economic Review*, 103, 804–30.

SCHWEIZER, N. AND N. SZECH (2014): "Optimal Revelation of Life-Changing Information," Working paper.

SHEN, L., A. FISHBACH, AND C. K. HSEE (2015): "The Motivating-Uncertainty Effect: Uncertainty Increases Resource Investment in the Process of Reward Pursuit," *Journal of Consumer Research*, 41, 1301–1315.

Working Paper Series in Economics