

More Effort with Less Pay:
On Information Avoidance, Optimistic Beliefs, and
Performance*

New and Preliminary Version

Steffen Huck Nora Szech Lukas M. Wenner

March 30, 2017

Abstract

Recent behavioral models argue in favor of avoidance of instrumental information. We explore the role of information avoidance in a real-effort setting. Our experiment offers three main results. First, we confirm that preferences for avoidance of instrumental information exist, studying information structures on performance pay. Second, information avoiders outperform information receivers. This result holds independently of effects of self-selection. Third, the findings support theories on information avoidance that favor an optimistic belief design, such as Brunnermeier and Parker (2005), instead of a Bayesian rational approach such as Bénabou and Tirole (2002). Subjects do not stick to their Bayesian expectations, but bias them optimistically under coarse information.

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

JEL-codes: D83, D84, J31, M52

*We are grateful to Kai Barron, Sibille Hackel, Alex Imas, Terri Kneeland, Henrik Orzen, Jonathan Parker, Gautam Rao, Frank Rosar, Andy Schotter, Nikolaus Schweizer, and Florian Zimmermann. for helpful comments. Steffen Huck: Wissenschaftszentrum Berlin (WZB) and Department of Economics, University College London. Email: steffen.huck@wzb.eu. Nora Szech: Department of Economics, Karlsruhe Institute of Technology (KIT), WZB, and CESifo. Email: nora.szech@kit.edu. Lukas M. Wenner: Department of Economics, University of Cologne. Email: lukas.wenner@uni-koeln.de

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. There is no reason to refuse information that comes for free. When working, for example, agents would want to know the exact payment scheme, and adapt their performance accordingly. Knowing the pay schedule allows agents to adjust their performance optimally, balancing costs and expected rewards.

Recently, this view on preferences for instrumental information has been challenged. Information avoidance seems to be a phenomenon that is surprisingly pervasive in a variety of contexts. Two recent surveys of the literature (Golman et al., 2017; Hertwig and Engel, 2016), highlight a number of situations in which agents actively “decide not to know.” Instead, agents prefer to remain ignorant about information which would affect their behavior considerably. Besides health (Oster et al., 2013; Ganguly and Tasoff, 2016), evidence comes from contexts relevant for one’s self image (Eil and Rao, 2011; Dana et al., 2007), but also from very different contexts such as portfolio investment decisions (Karlsson et al., 2009; Sicherman et al., 2016).

Recent models from behavioral theory on information avoidance can be divided into two groups. Either, they assume that agents who avoid information stick to the Bayesian expectation, such as Bénabou and Tirole (2002) or Caplin and Leahy (2001, 2004). Or they assume that agents bias their expectations optimistically under coarse information structures, as in Brunnermeier and Parker (2005) or Gollier and Muermann (2010). Under the latter approach, uncertainty thus provides some leeway for unrealistic optimism.

The goal of this paper is three-fold. First, we explore whether preferences for information avoidance can also be found in work situations. To this purpose, we consider a real-effort task and vary how much information subjects have on the performance pay. Second, we analyze how information avoidance affects performance in the real-effort task. Our design can distinguish between behavior in self-selected information structures versus enforced information structures. Hence, we can isolate the causal role of information structures from effects of self-selection. Third, we analyze whether Bayesian or optimistically biased beliefs provide the better suited model framework for describing the behavior of information avoiders in our context.

A number of applied theoretical papers have explored situations in which agents may want to avoid instrumental information, e.g., if anticipations influence utility. One prominent example is medical testing, which is often avoided by patients (compare Caplin and Leahy, 2004; Schweizer and Szech, 2016). Kőszegi (2003) considers a related problem in which doctors bias information delivered to patients, who accordingly infer that they are part of a cheap talk game. In all these approaches, optimal expectations balance the psychological benefits of deliberately coarse information with the material costs of making suboptimal choices. Yet the empirically open question remains regarding what expectations information avoiders stick to. Do they basically hold on to the Bayesian correct expectations? Or do they build biased, potentially unrealistically optimistic expectations instead, and perform accordingly? Our results indicate that the latter is the case.

As one specific example, 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.¹ The authors find that a large fraction of people at risk shy away from the genetic test, even though direct costs of testing are small and subsequent adjustments to test results are typically large. In addition, their data also display that patients who avoid the medical test seem to do pretty well in life. This suggests that information avoidance may be beneficial, at least for a substantial fraction of people. The authors suppose that information avoiders follow a Brunnermeier and Parker (2005) approach, i.e., that test avoiders tend to believe that chances of getting Huntington's are much smaller than the Bayesian correct probability. Yet Oster et al. (2013) can only analyze self-selected information structures. Their findings may be driven by people with a more positive attitude towards life, self-selecting into avoiding the test. In our less drastic set-up, we are able to control for effects of self-selection.

In the work context we study, we find that first, preferences for information avoidance exist. To the best of our knowledge, this presents the first highly controlled evidence for the prevalence of information avoidance in an economic context in which information is instrumental. Second, we find that no matter whether enforced or self-chosen, coarse information structures lead to better performance results than precise information. The effect of self-selection is non-significant in our data. Coarse information thus motivates subjects better

¹Huntington's is a genetically caused, severe disorder that breaks out around age 40, and leads to a drastic physical and mental decline.

than precise information. Third, agents do not stick to the Bayesian correct prior, but seem to bias their beliefs optimistically under coarse information.

All in all, our results indicate that uncertainty can have a significant performance-enhancing effect. Indeed, tournament incentives which are widely used at 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 who is unlikely to influence a company’s share-value through his or her behavior. If there is some room for optimism, it may well be the case that workers bias their beliefs into that direction. For a principal, it may thus turn out profitable to leave some coarseness in payment schemes. Likewise, for employees, it may be difficult to realize when searching for a new job would be advantageous. It may be that chances for promotion or high bonus payments are often perceived too optimistically at the current employer.

Our paper is organized as follows. Section 2 describes the design of our experiment and formulates our key predictions. Section 3 presents the results and how they relate to orthodox theory, to behavioral approaches on information avoidance assuming Bayesian beliefs, and to behavioral approaches assuming optimistic belief distortion. Our results support the last approach. In our context, preferences for information avoidance occur, and agents do not seem to stick to Bayesian expectations, but to bias their expectations optimistically. Section 4 situates our paper in the literature, discusses an extension of the Brunnermeier and Parker approach accounting for heterogeneities in optimistically biased beliefs, and discusses alternative explanations for our empirical results. Section 5 concludes.

2 Experimental Design and Hypotheses

We conduct four main treatments, FULL INFO, NO INFO, INFO CHOICE, and FIXED WAGE. The first three treatments are identical except that information about piece rates varies—exogenously or endogenously. The FIXED WAGE treatment allows us to separate between behavioral models on information avoidance assuming Bayesian rationality, such as Bénabou and Tirole (2002), and models assuming optimistic belief distortion a la Brunnermeier and Parker (2005). This treatment is therefore designed differently, as explained in detail below.

In FULL INFO, NO INFO, and INFO CHOICE, subjects know that they receive either a high (1 EUR) or a low (0.1 EUR) piece rate, with equal probabilities, for performance results

in a tedious task.² Subjects 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 receive a piece rate. Each string consists of 60 characters. For example, one of the strings used in the experiment looks as follows.

NXgCX7JHxYZj2cfBSd8JtkYp3LPcyDX8y8NNQhrzJfg22S2ACjC85EQ43B7L

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 they want 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 while working. In treatment NO INFO, subjects are forced into information avoidance, but know that both piece rate realizations are equally likely. They thus know exactly as much as self-selected information avoiders in INFO CHOICE. At the end of the experiment, we ask subjects for some basic demographics about themselves. In treatment INFO CHOICE, we also ask them in an open form about their reasons for choosing (not) to obtain information about their piece rate before working on the task.³

In our work context, treatment FULL INFO allows us to find out whether information about the piece rate is indeed instrumental. This is the case if agents perform better under the 1 EUR piece rate than under the 0.1 Euro piece rate. In line with recent behavioral theory, it may nevertheless be the case that preferences for information avoidance exist in this context.

Prediction 1. *Preferences for information avoidance exist, though information has instrumental value.*

Both, Bayesian models on information avoidance such as Bénabou and Tirole (2002) (henceforth BT), as well as approaches assuming optimistic belief distortion such as Brunnermeier and Parker (2005) (henceforth BP), can predict this outcome. All these models

²At the time of the experiment, 1 EUR \approx 1.37 USD.

³The exact wording of the question was “Please explain why you decided to look at the colored cardboard at the beginning of the experiment, or respectively, why not.” The colored cardboard was used to implement information revelation, see below.

rationalize information avoidance, albeit via different channels.⁴ Here, we provide a written exemplary description of how the BT and the BP model can respectively apply to our context, and what distinguishes them from each other. In the appendix, we provide the formal versions.

In BT, an agent who has to decide about whether to learn the piece rate or not, could face a problem of temptation. From an ex-ante perspective, he or she would like to exert an optimal level of effort, conditional on the piece rate to be received. Yet if he or she anticipates that a low piece rate could have an unwanted effect of demotivation, avoiding to learn the piece rate early could be advantageous. Avoiding information thus provides a kind of commitment device. If the Bayesian expected piece rate is sufficiently attractive, the agent may decide to stick to this expectation instead of receiving full information and running the risk that motivation severely decreases.

In BP, an agent derives part of the utility from his or her beliefs. Applied to our real effort setting, the model would assume that agents optimistically distort their beliefs about their piece rate under uncertainty, and perform accordingly. Agents thus trade off the benefits from biased beliefs about the payments against the costs of ex-post suboptimal effort choices. Hence, if agents place sufficiently large weight on their belief utility, they will avoid precise information, and prefer to leave some room for optimistic belief distortion. This may happen consciously, or sub-consciously.

The crucial difference between the two modeling approaches for our context is thus whether agents optimistically distort their expectations and perform accordingly, or whether their performance suggests that they stick to the Bayesian correct expectation instead. Both approaches support that coarse information can boost performance compared to precise information. This leads us to our second main hypothesis.

Prediction 2. *Coarse information causally fosters performance results compared to precise information, even though Bayesian expected piece rates are identical in both settings.*

Both classes of models, à la BT and à la BP, can explain such a preference for information avoidance, and an according boost in performance results. The difference in predictions comes in when analyzing performance results under coarse information in more detail. Ac-

⁴See also Bénabou (2015) for a comparison of the two approaches. For other applied studies documenting preferences for information avoidance, see the overviews by Golman et al. (2017) and Hertwig and Engel (2016).

According to Bayesian rational approaches such as BT, agents' performance hinges on the Bayesian expectation of the piece rate. In contrast, in models of optimistic belief distortion such as BP, agents hold behaviorally biased expectations and perform accordingly.

We implement the FIXED WAGE treatment in order to test for these different predictions. In FIXED WAGE, agents know that they earn a piece rate of 0.55 EUR for sure. This is the Bayesian correct expectation from the NO INFO treatment. Following BT, agents should have this Bayesian expectation in mind while working in NO INFO. Performance results should then be comparable to those from FIXED WAGE. If there was a difference in performance, it should be that FIXED WAGE outperforms NO INFO as agents typically tend to be risk-averse.

In contrast, under BP, agents should hold too optimistic beliefs regarding the future payments under coarse information. They could thus perform significantly better (and for sure not worse) in NO INFO than in FIXED WAGE. If so, from a principal's point of view, keeping the piece rate uncertain could causally foster performance compared to a certain payment of same expected value.

Prediction 3. *If agents behaviorally distort their beliefs under uncertainty (as, e.g., in BP), performance results in NO INFO can be better than in FIXED WAGE. If the opposite holds true, the data support Bayesian approaches of belief design (such as, e.g., BT) instead.*

In contrast to field studies, e.g., from the literature on medical testing, our experimental design can further shed light on whether there is a *causal* effect of coarse versus precise information on behavior. This is possible because part of our subjects are forced into specific information structures, as is the case in FULL INFO as well as in NO INFO. Our set-up therefore allows us to isolate the causal role of the information structure on work performance. To measure potential additional effects of self-selection, we compare performance from NO INFO and FULL INFO to performance in the corresponding, self-selected sub-groups from treatment INFO CHOICE.

In the rest of this section, we describe the implementation of the treatments in detail. We present the results from our study in Section 3.

Implementing the different information structures across treatments. The detailed procedures in the three treatments FULL INFO, NO INFO, and INFO CHOICE, were 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.⁵ 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-

⁵The procedure makes sure that subjects are aware that there are chips of two colors and that the procedure is entirely random.

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 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 FIXED 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, 2015) 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.⁶

3 Results

In treatment FULL INFO, subjects know their piece rate realization before starting to work on the task. Our data display that subjects perform significantly better under the high than under 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 ($p = 0.043$),⁷ and confirms that information about the piece rate is instrumental as hypothesized.⁸ The result furthermore shows that the combination of the task we employed and the difference between the low and the high

⁶We 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.

⁷Unless indicated otherwise, all p-values are calculated using a Wilcoxon Rank Sum test.

⁸The 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 ($\rho = 0.8763$, $p = 0.0000$) and the results are very similar.

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

Table 1: Mean and median performance across treatments

wage is suitable to uncover meaningful responses to incentives, despite the relatively large number of solved tasks in the low wage condition.

Having established that information is instrumental, we can turn to Prediction 1. In support of all theories discussed which predict information avoidance, 30 out of 95 subjects (31.6%) prefer not to know their piece rate while working. These subjects thus decide to forgo information of instrumental value as shown in treatment FULL INFO.⁹ Our data thus confirm Prediction 1.

Next, we turn to Prediction 2. The findings from INFO CHOICE support Prediction 2. Information receivers solve 30 tasks correctly on average, while information avoiders solve only 21.31 tasks correctly ($p = 0.0002$). Controlling for potential effects of self-selection, we find a similar result. NO INFO leads to significantly better performance results than FULL INFO (23.44 versus 28.02, $p = 0.0274$). We thus confirm Prediction 2.

Furthermore, we can also analyze how the information avoiders compare to the sub-group of information receivers who found out that their piece rate is 1 EUR. We find that the latter solve 25.53 tasks on average, which is (weakly) significantly less than the 30 tasks solved by information avoiders ($p = 0.0573$). This reinforces our finding from above that information

⁹As a robustness check, we analyze the two sessions that were 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.

avoidance can have large effects on performance.

Moreover, while coarse information significantly enhances performance, effects of self-selection are not significant in our data.¹⁰ Therefore, in the following, we also investigate data that is pooled over self-selected and enforced information, e.g. in order to have a larger data base for plotting distributions in performance.

So far, the data is in line both with Bayesian approaches such as BT, and with models of optimistic belief distortion such as BP. The difference between the two model approaches is reflected in Prediction 3. According to BT, under coarse information, subjects should perform worse or equally well as under the Bayesian piece rate payed out for sure. Following BP, to the contrary, subjects may perform better under uncertainty than under the secure piece rate matching the Bayesian correct expectation. This is what we find.

In FIXED WAGE, 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 as comparable as possible. We find that subjects solve on average 24.66 tasks correctly in this treatment. This is less than the average of 28.78 tasks if the wage is unknown, yet 0.55 in Bayesian expectation. The difference is significant ($p = 0.0015$). With regard to our two classes of models, this increase in performance under uncertainty is only compatible with approaches of optimistic belief distortion such as BP. Under Bayesian models such as BT, we would have to see the opposite effect if agents tend to be risk-averse. Even with risk-neutral agents, under BT, there is no reason why NO INFO should outperform FIXED WAGE. Yet we see that this is significantly the case. For a principal, implementing a situation of coarse information fosters performance compared to providing the expected piece rate with certainty.

Table 1 presents an overview of all performance results, per piece rate and per 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 as there are no significant differences between the two at any conventional level. 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 piece rate (18.90

¹⁰To make this statement precise, we test for all wages (0.1 EUR, 1 EUR and “unknown”) whether the distributions differ between the self-selected case (treatment INFO CHOICE) or the forced condition (treatment FULL INFO or NO INFO, respectively). In all cases $p > 0.1$.

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 piece rates.

vs. 25.83 correct tasks on average, $p = 0.0003$).¹¹

Figure 1 plots the empirical distribution functions for the four cases of information structure. We see that under the uncertain piece rate, performance results are best. All in all, we obtain a strong indication that uncertainty substantially fosters performance compared to paying the expected piece rate with certainty. Visually, the distribution for the case in which the piece rate is unknown even almost stochastically dominates the distribution for the 1 EUR piece rate. This parallels the finding in the INFO CHOICE treatment: The difference between the two distributions is weakly significant, also for the pooled data ($p = 0.0645$). We address this aspect further in the discussion section below.

Thus, from a principal’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 can be explained by optimistic belief distortion as, e.g., suggested by BP.

3.1 Calibrating the Brunnermeier and Parker Model

Our findings suggest that models of distorted beliefs such as BP are a good fit for our data. Therefore, the next logical step is to calibrate the extent to which subjects may have biased their expectations under uncertainty.

To this end, we implemented a variant of the treatment FIXED WAGE, but changed the

¹¹It 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.




Figure 1: CDFs of effort choices, by wage.

piece rate from 0.55 EUR to 0.8 EUR. Otherwise, the treatment remained exactly the same. We accordingly call this treatment the FIXED WAGE 0.8 treatment. The choice of 0.8 EUR is motivated as follows. 0.8 is a value less distorted from 0.55 than 1, but still considerably higher than 0.55.¹² For many subjects, 0.8 may serve as a better motivation than 0.55, and thus reflect a more optimistic interpretation of an uncertain situation (in which the Bayesian correct estimate would be 0.55).

FIXED WAGE 0.8 turns out to match the findings from the situations of uncertainty, exactly as approaches à la BP predict. Performance results under the uncertain piece rate compare to those under a fixed piece rate, yet one that is considerably higher than the Bayesian correct estimate of 0.55.¹³ The distribution for the 0.8 fixed wage treatment is visually and statistically similar to the distribution of performance if the wage is unknown.

¹²We ran this treatment in November 2016 with 47 new subjects from the same subject pool, at the same location as before.

¹³Performance results for 1 EUR and 0.8 EUR turn out to be statistically comparable ($p = 0.1953$).

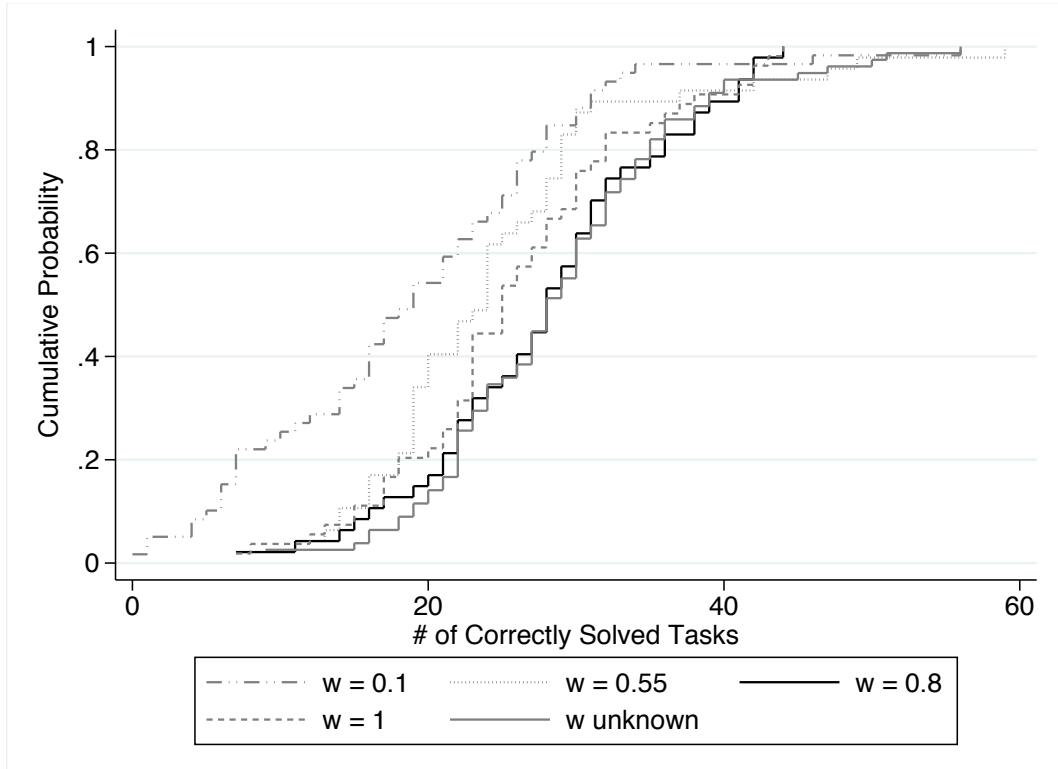


Figure 2: CDFs of effort choices, by wage, including additional treatment

4 Discussion

In this Section, we address four points. First, we explore the idea that there may be heterogeneities in what piece rates subjects may want to believe in if there is uncertainty. This discussion is motivated by the fact that in our ex post questionnaire, subjects provided different motivations for avoiding information. Second, we discuss in detail how our findings relate to the existing literature. Also from this literature, a heterogeneity in preferences seems plausible. Third, we extend the BP approach, incorporating a heterogeneity regarding optimally distorted beliefs. If subjects bias their expectations to different extents, uncertain incentive structures can even lead to better incentive structures than any fixed wage scheme. In our data, we see a marginally significant tendency for this. Forth, we share some thoughts regarding the design of incentives.

4.1 Heterogeneity in Optimally Biased Beliefs

The BP approach captures our findings very well. Yet in our ex-post questionnaire, we see that subjects provide different reasons for why they avoided information. While the majority of agents possibly like to tell themselves that they will get a very high piece rate, others may

want to hold on to more moderate beliefs.¹⁴ Two arguments for information avoidance came up frequently. Most often subjects argued that they wanted to avoid knowing that the piece rate realized low instead of high in order to feel more motivated. Yet others argued that they intended to avoid feeling stressed and a potential choking under pressure. This suggests that subjects have different motives for information avoidance, and may thus bias their expectations differently.

Therefore, we extend the BP approach to the case of heterogeneous agents. As illustrated in detail in the appendix, we find that with heterogeneous agents, coarse information structures can outperform any kind of fixed wage scheme. The intuition behind is that agents who work best if incentives are high have room to believe payments will be large. At the same time, agents afraid of choking under pressure have enough room for sticking to moderate expectations, thereby avoiding situations of choking.

In our data, we see that performance results under coarse information tend to outperform performance for 1 EUR. For this, revisit Table 2. Performance under uncertainty (but with Bayesian expected wage of 0.55) tends to be better than under a certain payment of 1 ($p = 0.0645$). This supports the idea that there may be a heterogeneity in subjects. Believing in a high piece rate of 1 may lead to good performance results for several, but not for all subjects.

4.2 Related Literature

We find that that about one third of our subjects prefer to avoid instrumental information. For various but arguably very different contexts, preferences for information avoidance have been documented by Golman et al. (2017) and Hertwig and Engel (2016). The authors report preferences for information avoidance when it comes to self-image, health, or investment decisions.¹⁵ As one stark example, Oster et al. (2013) report that only 7 percent of individuals at risk of the hereditary Huntington’s Disease opt for the medical test.

In a laboratory study Ganguly and Tasoff (2016) give subjects the option to avoid being tested for milder virus infections at a cost of 10 USD. Depending on the type of virus subjects are tested for, 5.2% to 15.6% of individuals are willing to give up 10 USD in order not to be tested, even though the diseases can be treated quite well. Ganguly and Tasoff (2016) also

¹⁴See the appendix for details.

¹⁵A related phenomenon exists also in the morally relevant domain, compare Dana et al. (2007) and Falk and Szech (2014) for empirical evidence.

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 comprise between 30 and 40 percent 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 to be below average in a category are willing to pay for not learning their true rank in the distribution. Rank is not connected to any direct monetary incentive in that study, but of course, knowing one's rank may contain instrumental information for future decisions outside the laboratory.

Eliaz and Schotter (2010) and Falk and Zimmermann (2016) analyze the demand for 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 (2016), subjects can choose when to find out about whether they are to receive mild electric shocks. Without anything interesting to do, a large majority chooses to find out immediately, and can thus be classified as "curious." This is maybe why in our setting, when given the choice, the majority of subjects opt for information, though information avoidance leads to better performance results. More research in this direction would be interesting.

Ariely et al. (2009) demonstrate 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 stress the phenomenon that individuals may choke under too much pressure induced by high rewards. This renders the fear of choking under pressure that a fraction of our subjects expressed 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, overall rewards are small. The authors suggest that the uncertainty about the 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). In our ex-post questionnaire, many subjects argue that they decided to avoid information in order to prevent demotivation from a low piece rate realization, while others state 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. Nevertheless, it would obviously be important to explore in future research the role of different mechanisms for why uncertain incentives can foster performance results. Possibly, the roles of the different mechanisms also depend on context.

4.3 Alternative Explanations

The reader may have thought of other potential explanations for our findings, and we did so as well. In the following, we discuss such attempts.

Satisficing. As argued by Simon (1955), agents may be satisficers instead of utility maximizers. That is, once payoffs reach some critical level, agents do not find it worthwhile to exert more effort. Coupled with optimally designed beliefs, satisficing could be an alternative explanation to choking: Agents may not exert as much effort when working for a wage of 1 EUR once they are “satisfied” with the number of tasks they have completed.

Two reasons speak more to the direction of choking, though. First, several subjects mention a fear of choking in the questionnaire, while no subject mentions ideas of satisficing. Second, there is no evidence for satisficing in the data either. For wages larger than 0.1 EUR, subjects never seem to stop working before the allotted time of 60 minutes is up. (In contrast, for a wage of 0.1 EUR, some subjects stopped working.) Stopping once one has earned enough, however, would be a behavior expected from satisficers, especially as participants in our experiment do not have certainty about whether they solve a question correctly, which makes it harder to ex-ante aim for a certain level of pay and adjust effort

accordingly.

Regret Aversion. Another possibility would be that agents feel regret when they work under wage uncertainty. This could lead them to perform extremely well in order to reduce regret should the low wage be revealed. The problem with this argument, and related concepts such as disappointment or loss aversion, is that working hard to avoid regret leads to a high spread of possible payments before the lottery is resolved. This is unlikely to be an attractive feature for agents. In addition, neither regret nor loss aversion can offer an explanation for the observed information avoidance. The model by Gollier and Muermann (2010) combines disappointment aversion and belief distortions, but, as the authors show, the former works as an opposing force to optimistic beliefs. Their model therefore predicts that agents have a preference for early resolution of uncertainty, contrary to what we find.

4.4 Some Thoughts on Optimal Payment Schemes and Information Structures

Our results have implications for principals who want to design payment schemes for their agents. The high performance when the wage is uncertain suggests that such wage schedules are profitable for principals. They either increase output while keeping the expected wage constant (uncertain wage vs. 0.55 EUR), or keep performance at a similar level while lowering the payment to workers (uncertain wage vs. 0.8 EUR). This, however, begs the question whether a principal could do even better than using any of the payment schemes we investigated. We leave an in-depth investigation of this topic for future research, but, nevertheless, offer some speculative thoughts.

Given that we find clear evidence for active belief design, it remains an intriguing possibility whether one could lower the expected wage of the lottery by making the low wage more likely than the high wage, while still obtaining similar levels of performance. Especially agents for whom anticipatory feelings are important, might respond by distorting beliefs to similar levels, thereby reducing effort only a little, but at the same time allowing the principal to save money (in expectation).

Another promising area to explore would be whether partial information release has positive effects on performance. For the context of medical testing, in which agents typically shy away from the test in fear of ending up in a situation of no hope, Schweizer and Szech

(2016) explore possibilities of partial revelation. They find that for the medical context such as testing for Huntington's, tests that perfectly reveal the good, but shroud the bad state may be a promising alternative for the vast majority of agents who do not dare to take the perfectly revelatory test, but have an interest in better information. This is even the case under Bayesian correct beliefs. Similar types of information structures could be explored in the work context.

5 Conclusion

We studied a real effort task in which agents decided whether they wanted to receive instrumental information about their piece rate before starting to work, or not. Our data display that about one third of subjects deliberately decide to forgo information, thus revealing a preference for information avoidance. Furthermore, agents avoiding information achieve considerably better performance results than agents opting for information.

In order to rule out potential causal effects of self-selection on performance results, we ran another treatment in which subjects were forced to stay uninformed about their piece rate. Performance again turned out to be significantly higher than under complete information. Information avoidance, even if enforced instead of self-chosen, significantly increases performance.

Two classes of behavioral models, those assuming Bayesian rationality such as BT, and those assuming optimistic belief distortion such as BP, can explain the results up to this point. In order to differentiate between the two, we ran another treatment in which subjects received the expected piece rate with certainty. The data show that subjects in this treatment perform worse than under uncertainty. This result supports the approaches à la BP, under which agents optimistically distort their beliefs from the Bayesian correct estimate. In a follow-up treatment, we find indeed that agents perform comparably well as under uncertainty if they receive 0.8 EUR as a secure piece rate. This fits well to the models à la BP, and sets the findings apart from the predictions generated under the Bayesian model approach such as BT.

Not all our findings are at odds with orthodox economic theory. Specifically, without uncertainty, standard incentives work pretty well in the sense that a higher piece tends to foster performance results. The predictions of standard theory only turn out non-suitable

when uncertainty becomes part of the incentive scheme. All in all, our data therefore supports models of optimistic belief distortion under uncertainty.

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, 2016), but also in economic settings of paid effort exertion and performance. 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. Rampant overprovision of effort in contests may have one of its roots in optimistic belief distortion.¹⁶ 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, specifically 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.

¹⁶For a detailed survey of the experimental literature on contests, see Dechenaux et al. (2015) who document evidence for overprovision of effort in a large number of different settings.

Appendix

A.1 Formal Derivation of the Predictions

We consider an agent who receives a piece rate $w \in \{w_L, w_H\}$ for working on a task. Completing e tasks comes at a convex effort cost $c(e)$. Hence, in this formulation, effort directly translates into performance, i.e. output. At period 0, the agent faces a lottery where each wage is equally likely to realize and he decides whether to find out his wage or not. In period 1, the agent works on the task and incurs his effort cost. He is paid $w \cdot e$ in period 2 and does not discount payments. We assume risk neutrality throughout, but emphasize how risk aversion would change the results where appropriate.

Bénabou and Tirole (2002)

The key ingredient of this model is a temptation problem of the agent. In period 0, for a given wage w , the optimal effort level $e^*(w)$ solves

$$w = c'(e^*(w))$$

If, however, the wage is low, $w = w_L$, the agent does not exert effort according to $e^*(w_L)$ in period 1. Instead, the agent gives in to the temptation of “slacking off” and chooses effort level $\hat{e}(w_L)$ according to:

$$\beta w_L = c'(\hat{e}(w_L))$$

Since $\beta < 1$ and $c'(\cdot)$ is increasing by assumption, the agent’s effort at the low wage is lower than ex-ante optimal, $e^*(w_L) > \hat{e}(w_L)$.

Now, consider the agent’s decision whether to obtain information about his wage. If he decides to acquire the information, his expected utility will be given by:

$$\frac{1}{2}(w_L \hat{e}(w_L) - c(\hat{e}(w_L))) + \frac{1}{2}(w_H e^*(w_H) - c(e^*(w_H))) \quad (\text{A.1})$$

If instead, he does not obtain information, his expected utility is given by

$$\bar{w} e^*(\bar{w}) - c(e^*(\bar{w}))$$

where $\bar{w} = \frac{1}{2}(w_L + w_H)$. For any w , the agent’s utility $U(w) = w e^*(w) - c(e^*(w))$ is convex in w , hence it follows that for $\beta = 1$, that is in the absence of a temptation problem, the

agent prefers to know his wage because $\frac{1}{2}U(w_L) + \frac{1}{2}U(w_H) > U(\bar{w})$. For $\beta < 1$, however, the agent's utility is given by (A.1) which decreases as β decreases. Hence, if β is sufficiently low and $U(w)$ not too convex,

$$\frac{1}{2}(w_L \hat{e}(w_L) - c(\hat{e}(w_L))) + \frac{1}{2}(w_H e^*(w_H) - c(e^*(w_H))) < \bar{w} e^*(\bar{w}) - c(e^*(\bar{w}))$$

which implies that the agent prefers to avoid information about his wage. Note, however, that effort when facing the lottery and when facing a sure wage of \bar{w} is in both cases given by $e^*(\bar{w})$. Hence, this model predicts that information avoiders perform no better than agents in our 0.55 EUR (FIXED WAGE) treatment. While risk neutrality would predict that efforts are exactly equal, it is straightforward to show that risk averse agents would exert higher effort under the fixed wage because they dislike the lottery which increases in their variance as the number of completed tasks increases.

Brunnermeier and Parker (2005)

Using the same notation as above, we can write the agents utility in this model as

$$\theta [w(\pi)e^*(w(\pi)) - c(e^*(\pi))] + (1 - \theta) [\bar{w}e^*(\bar{w}) - c(e^*(\bar{w}))] \quad (\text{A.2})$$

The interpretation of this expression is as follows: total utility is a weighted average of anticipatory utility and realized (“consumption”) utility, where $0 \leq \theta < 1$ denotes the relative weight of the former. 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} .

We analyze the agent's problem by starting with the effort choice. Whenever he works under wage uncertainty (whether forced or chosen), he will choose his beliefs π , the likelihood of receiving the high wage. From this, we find $w(\pi) = \pi w_H + (1 - \pi)w_L$ as the subjective wage. Hence, when the agent works on the task he chooses his effort level taken this distorted

wage as given. Importantly, this effort is chosen to maximize the anticipatory utility:

$$w(\pi) = c'(e^*(w(\pi)))$$

When choosing his optimal belief π^* , however, he maximizes the expression in (A.2), that is, he takes into account his anticipatory utility (which only depends on his belief $w(\pi)$) as well as his consumption utility which is calculated based on the objective expected wage \bar{w} .

The first order condition for the belief choice can then be written as:

$$\theta e^*(w(\pi^*)) + (1 - \theta) \frac{de^*(w(\pi^*))}{dw(\pi^*)} (\bar{w} - w(\pi^*)) = 0 \quad (\text{A.3})$$

Due to the convexity of $c(\cdot)$, e^* is increasing in the wage, hence, for $\theta = 0$, we find that $\bar{w} = w(\pi^*)$ and for $\theta > 0$ that $\bar{w} < w(\pi^*)$. Thus, whenever the agent cares about anticipatory emotions, he distorts his beliefs upwards. We also note that as $\theta \rightarrow 1$, the agent will maximally distort beliefs, i.e. select $\pi^* = 1$.

What is left to show is that such an agent would find it worthwhile to avoid information. If, on the contrary, he opts for information, he does not distort beliefs and simply chooses effort optimally according to $w = c'(e(w))$. We therefore need to have that:

$$\theta U(w(\pi^*)) + (1 - \theta) [\bar{w} e^*(w(\pi)) - c(e^*(w(\pi)))] > \frac{1}{2} U(w_L) + \frac{1}{2} U(w_H)$$

A standard “envelope” argument implies that the L.H.S. is increasing in θ . Hence, this condition will eventually be satisfied for a large enough θ . This implies that the model can predict information avoidance and at the same time makes a clear prediction regarding the performance level we should observe. Since the agent distorts beliefs his effort is given by $e^*(w(\pi^*))$ which will be larger than the effort he exerts at a fixed wage of \bar{w} .

A.2 A Formal Model of Choking and Belief Design

We follow the notation in the previous section, with the modification that instead of the assumption that effort $e(w)$ directly translates into the agent’s output, we write output $y(w) = e(w)f(w)$, that is, the product of effort and productivity. The latter can be thought of as the agents’ efficiency in working on the task (i.e. the propensity of making mistakes), potentially depending on the wage she faces. This means that the optimal effort level is

chosen as to maximize

$$w(\pi)ef(w(\pi)) - c(e)$$

for a given (potentially distorted) wage $w(\pi)$. We thus find the first order condition as:

$$w(\pi)f(w(\pi)) = c'(e^*(w(\pi)))$$

We capture the idea of choking via the productivity $f(w)$. Specifically, we assume that for “chokers” $f(w)$ is decreasing for wages higher than some threshold wage $w_c \in (\bar{w}, w_H)$, and constant otherwise. The decreasing productivity for high wages can be, for example, understood as a higher propensity of making mistakes, or more generally as stress- or pressure-induced adverse effects on performance. In addition to $f'(w) < 0$, we assume that $f(w)$ is log-concave, i.e. the ratio $f'(w)/f(w)$ is decreasing, for all $w > w_c > \bar{w}$.¹⁷ The assumption that $f'(w)/f(w)$ is decreasing implies that there exists a $w = \hat{w}$ which satisfies $f'(\hat{w})\hat{w} + f(\hat{w}) = 0$. If $w_c < \hat{w} < w_H$, this captures the essence of the agent’s choking behavior. \hat{w} is the wage which maximizes the agents’ effort.

As before, the agent now maximizes the weighted sum of anticipatory utility and consumption utility. In our setting, the belief choice directly corresponds to a choice of the subjective wage. To ease exposition, we therefore suppress the dependence of w on π and allow the agent to choose $w \in [w_L, w_H]$ directly. Our agent then maximizes

$$\theta [we^*(w)f(w) - c(e^*(w))] + (1 - \theta) [\bar{w}e^*(w)f(w) - c(e^*(w))]$$

We then obtain the first order condition as:

$$\frac{f'(w)}{f(w)}w + \theta - \left(\frac{f'(w)}{f(w)} + \frac{e'(w)}{e(w)} \right) (1 - \theta)(w - \bar{w}) = 0$$

From this expression we notice the following: If there is no choking, i.e. $f(w)$ is constant for all $w \leq w_H$, the condition collapses to the first order condition in (A.3) above. $\theta = 0$, by assumption $f'(\bar{w}) = 0$, hence $w = \bar{w}$. From this it follows as well that the agent would never choose $w < \bar{w}$. Also, if there is choking and $\theta \rightarrow 1$, the optimal belief selects $w = \hat{w}$, since

¹⁷A more natural condition would be to simply assume that $f(w)$ is decreasing and concave, directly capturing the intuition that the choking behavior is exacerbated as the wage increases. Our assumption of log-concavity implies concavity, of course, and is all we need for our results, hence we use the less restrictive condition. The assumption that there is no choking below the average wage \bar{w} allows us to focus attention on upward belief distortions which we see in the data.

\hat{w} solves $f'(\hat{w})\hat{w} + f(\hat{w}) = 0$. As mentioned above, this \hat{w} also is the wage which maximizes effort. Hence, the agent will never distort beliefs above \hat{w} , because for any $w > \hat{w}$ there is a $w > \hat{w}$ which induces the same effort at a higher productivity. We finally note that output $y(w)$, which we defined as $y(w) = e(w)f(w)$ is decreasing at \hat{w} . Hence, it must be the case that $y(\hat{w}) > y(w_H)$, as well as for some other wages in $[w_c, \hat{w}]$.

Finally, we can use this result to argue why, in a heterogeneous population where only some agents choke at high wages, coarse information structures can do better than any fixed wage $w \in [w_L, w_H]$. For simplicity consider the case where θ is close to 1. With wage uncertainty, non-choking agents optimally distort their beliefs such that $w(\pi) \approx w_H$, whereas choking agents choose beliefs such that $w(\pi) \approx \hat{w}$. Hence, while both types of agents distort their beliefs differently, the aggregate output of a group of agents consisting of both types will be higher than it could be under a fixed wage of either w_H or \hat{w} , or indeed any other fixed wage.

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. (2015): “The economics of motivated beliefs,” *Revue d’économie politique*, 125, 665–685.
- 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.
- DANA, J., R. A. WEBER, AND J. X. KUANG (2007): “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness,” *Economic Theory*, 33, 67–80.
- DECHENAUX, E., D. KOVENOCK, AND R. M. SHEREMETA (2015): “A survey of experimental research on contests, all-pay auctions and tournaments,” *Experimental Economics*, 18, 609–669.
- 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 N. SZECH (2014): “Diffusion of being pivotal and immoral outcomes,” Working paper.

- FALK, A. AND F. ZIMMERMANN (2016): “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 (2016): “Fantasy and dread: the demand for information and the consumption utility of the future,” *Management Science*.
- GOLLIER, C. AND A. MUERMANN (2010): “Optimal choice and beliefs with ex ante savoring and ex post disappointment,” *Management Science*, 56, 1272–1284.
- GOLMAN, R., D. HAGMANN, AND G. LOEWENSTEIN (2017): “Information Avoidance,” *Journal of Economic Literature*, 55, 96–135.
- GREINER, B. (2015): “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- HERTWIG, R. AND C. ENGEL (2016): “Homo ignorans: Deliberately choosing not to know,” *Perspectives on Psychological Science*, 11, 359–372.
- HOLT, C. A. AND S. K. LAURY (2002): “Risk Aversion and Incentive Effects,” *American Economic Review*, 92, 1644–1655.
- KARLSSON, N., G. LOEWENSTEIN, AND D. SEPPI (2009): “The ostrich effect: Selective attention to information,” *Journal of Risk and Uncertainty*, 38, 95–115.
- KÓSZEGI, 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 (2016): “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.
- SICHERMAN, N., G. LOEWENSTEIN, D. J. SEPPI, AND S. P. UTKUS (2016): “Financial Attention,” *The Review of Financial Studies*, 29, 863–897.
- SIMON, H. A. (1955): “A behavioral model of rational choice,” *The quarterly journal of economics*, 69, 99–118.