Costly information acquisition: The influence of stakeholder earnings*

Timo Heinrich†, Bindu Arya‡, Alexander Haering§ and Sven Horak**

February 2022

Abstract

Information is often acquired within organizations that generate earnings for employees and stakeholders. In this paper we analyze the causal effects of inequality on information acquisition performance and vary the pay of agents relative to the earnings of passive stakeholders. Our experimental results reveal that disadvantageous inequality does not have a negative effect on agents' performance.

Keywords: Information acquisition, pay inequality, stakeholders.

JEL Codes: C91, D83, M52.

---

* We thank the editor-in-chief Eldad Yechiam and two anonymous referees for very helpful comments. We are also grateful to Jeannette Brosig-Koch, Alina Ferecatu and Conny Wunsch for their valuable feedback. We also benefited from the comments of seminar participants at Durham University and those of conference participants at the 11th Nordic Conference on Behavioral and Experimental Economics (Oslo), the 17th GEABA meeting (Basel) and the 4th Global Entrepreneurship and Innovation Conference (Tianjin). Franziska Then and Stephanie Zimmer provided excellent research assistance. Financial support from the BMBF through the IN-EAST School of Advanced Studies (Grant No. 01UC1304) is gratefully acknowledged.

† Timo Heinrich, Institute for Digital Economics, Hamburg University of Technology, Am Schwarzenberg-Campus 2 (B), 21073 Hamburg, Germany; email: timo.heinrich@tuhh.de (corresponding author).

‡ Bindu Arya, College of Business Administration, University of Missouri–St. Louis; email: bindua@umsl.edu.

§ RWI – Leibniz-Institute for Economic Research Essen and Fresenius University of Applied Sciences; email: alexander.haering@rwi-essen.de.

** Peter J. Tobin College of Business, St. John's University; email: horaks@stjohns.edu.

1 Introduction

Weitzman (1979) motivates his classical paper on optimal search with an example of the research and development process of a large organization. The organization's researchers consider new production technologies and can sequentially develop them to resolve uncertainty about the rewards they deliver. The risks and rewards of these technologies differ, and only one technology can be used for production. Weitzman then derives the optimal sequence of developing production technologies for an individual decision-maker under a range of conditions.

As in Weitzman's example, research is usually conducted within the social settings of organizations and not by individuals in isolation. In this paper, we ask how the earnings generated for stakeholders affect the agents' choices in information acquisition. This question is highly relevant, for example, to venture capitalists who set up contracts with entrepreneurs or to business owners who need to incentivize employees in research and development. Does higher performance pay for agents improve outcomes *per se*? Or do owners need to consider agents' pay alongside their own earnings?

We study information acquisition under controlled laboratory conditions that allow us to systematically vary the pay structure and observe the quality of decision making. We build on a sequential search task as studied by Weitzman (1979) in which decision-makers can explore different options with known risks and rewards. This task has a straightforward optimal decision sequence, which allows us to compare the performance of all subjects to the same objective benchmark.

Behavior in the task by Weitzman (1979) has already been studied experimentally by Slonim (1994) and Gabaix et al. (2006). Both papers only consider individual choices made in isolation. Slonim (1994) focuses on experience and on the method to elicit choices. He argues that the cognitive effort associated with identifying the optimal search path may decrease with experience. In fact, he finds that behavior gets closer to the optimal path when the tasks are played repeatedly. Gabaix et al. (2006) propose the directed cognition algorithm based on myopic cost-benefit calculations as a boundedly rational model for information acquisition behavior. They find that the directed cognition algorithm predicts aggregate information acquisition patterns quite well. Whether the path suggested by directed cognition differs from the optimal path depends on the parameterization of the tasks. When both paths diverge, they find that directed cognition does a better job of matching laboratory evidence.

To our knowledge, no study has considered information acquisition performance in social settings. In laboratory settings, where total payoffs of agents are not transparent, a preferences for merit-based pay has been observed (Fochmann et al., 2019), yet when pay inequality is exogenously imposed it has a detrimental effect on the quality of work (Greiner et al., 2011; Bracha et al., 2015). Also, in field experiments, a negative effect on job satisfaction (Card et al., 2012) and job performance has been observed (Cohn et al., 2014; Breza et al., 2018). In a recent field experiment, Cullen & Perez-Truglia (2021) not only consider horizontal pay inequality but also the differences in pay between workers and their managers. In their experiment the authors vary the information workers receive. They observe that workers who learn that peers earn more than they thought, exert less effort as measured by the number of emails sent, the working hours and generated sales. Yet, when they find out that their managers earn more, these workers exert *more* effort.[1]

In our experimental analysis, we employ the same parameterization as Gabaix et al. (2006) but create a setting, in which a stakeholder benefits from the agent's effort.[2] Thus, we study a sequential search problem without discounting and with different binary outcomes with known probabilities. An agent and a passive stakeholder are paid based on the decisions made by the agent. The agent enters choices which can be observed (but not influenced) by the stakeholder. Note that this approach avoids any uncertainty about inequality in effort costs simply because the stakeholder does not exert any effort. We abstract away from principal-agent interaction or free-riding within groups. Instead, we always consider a setting in which incentives are perfectly aligned. However, we exogenously vary absolute and relative pay levels of the agent and the stakeholder focusing on disadvantageous inequality.

[1] There are also a number of non-experimental studies that correlate pay dispersion with performance across different tasks (see Pfeffer, 2007; Shaw, 2014; and Downes & Choi, 2014, for general overviews). With respect to research and innovation pay dispersion appears to have a detrimental effect (see Pfeffer & Langton, 1993, Yanadori & Cui, 2013, Wang et al., 2015 and Amore & Failla, 2020).

[2] Gabaix et al. (2006) also conduct a second experiment that studies a more-complex choice problem in which the optimal path cannot be determined analytically. We only consider the parameters of their first experiment.

2 Experimental design

We focus on the sequential search problem presented by Weitzman (1979) and adopt the notation by Gabaix et al. (2006) in the following. Within a particular task, an agent has to choose between three projects i. Each project has a known probability p_i of generating a known payoff V_i, if the respective project is successful and zero otherwise. For a payment of c_i, the agent can acquire information about project i and learns whether it was successful. Once the agent has decided to stop exploring, she can pick the highest-paying project and earns the respective payoff minus total search costs. The agent can only select projects where she already learned about the success or failure. Optimally, the risk-neutral agent would always explore projects based on the index Z_i. For projects with unknown payoffs, the index is calculated as:

$$Z_i = (p_i V_i - c_i)/p_i. \tag{1}$$

The index is equal to V_i for a successful project and 0 for a failed project. If the project with the highest value Z_i has an unknown outcome, the agent would acquire information about it. If the project has a known outcome, she would pick this project and stop exploring. Note that the index implies that if two projects have the same expected net gain, the one with the smaller probability of success will be selected – low-probability, high-payoff projects should be investigated first, because they will end the search.

Gabaix et al. (2006) also propose an alternative way of selecting the sequence of projects as a boundedly rational model for actual choice behavior. They propose the "directed cognition algorithm", which suggests that people explore projects in the sequence of decreasing expected gains

$$G_i = p_i(V_i - S) - c_i \tag{2}$$

where S is the value of the best currently known winning project. If the highest gain is provided by a project with known outcome, the algorithm stops, and this project is selected. Based on this heuristic, agents maximize expected gains myopically and choose as if each choice was the last to be made.

Subjects in all treatments of the experiment acquire information about their projects in ten tasks. The tasks are ordered randomly, subjects face each task only once and one of the tasks is randomly selected for payment at the end of the experiment. We use the same parameterization as Gabaix et al. (2006), so that in five of the tasks the sequence of optimal choices diverges from the directed cognition sequence (tasks A to E). In the other five, the two coincide (tasks F to J). See Online Appendix A for the parameters of all ten tasks, an example and a screenshot.

Subjects face the tasks A to J in pairs. The first subject acts as the decision-making agent entering the decisions. She is matched to a second subject who witnesses project characteristics, choices and outcomes and acts as the stakeholder. This player is neutrally introduced as “Participant 2”. Both subjects cannot communicate but they know that the choices of the agent affect the payoffs of both. See Online Appendix B for the instructions. The experiment was implemented via z-Tree (Fischbacher, 2007).

Table 1: Experimental treatments

	Pay differentials		
Treatment	Agent	Stakeholder	N
LOW	$\delta = 0.7$	$\delta = 0.7$	72
HIGH	$\delta = 1.3$	$\delta = 1.3$	72
UNEQUAL	$\delta = 0.7$	$\delta = 1.3$	72

We ran three experimental treatments that systematically vary the pay within pairs of subjects, as shown in Table 1.[3] The treatments are designed to disentangle self-interested and inequality averse motivations in search. In the **LOW** and the **HIGH** treatment, both the agent and the stakeholder earn the same (equal) prize V_K where K is the project selected in one randomly chosen round. However, we vary the stakes both can earn. Either both receive 30% less than the nominal payoff (δ = 0.7) or both receive a bonus of 30% (δ = 1.3). In the **UNEQUAL** treatment, we create

[3] We also ran a control treatment (δ = 1) not reported here, replicating the first experiment in Gabaix et al. (2006). In this treatment, agents made choices individually and not within pairs. For example, we observe very similar frequencies of subjects choosing optimally in their first moves: 34% in tasks A to E of both experiments and 82% in tasks F to J of our experiment, compared to 74% in Gabaix et al. (2006).

disadvantageous inequality from the agent's point of view: Her stakeholder receives a bonus of 30% on top of her payoff V_{jK} (δ = 1.3), while she earns 30% less (δ = 0.7).[4]

Upon entering the laboratory, subjects drew a ball from an urn which assigned them to a cubicle. This way subjects were also randomly assigned to a treatment and a role. In each session we ran all three treatments in order to eliminate potential session effects. The distribution of female and male agents was balanced across treatments by using two different urns. Only for one subject we had to deviate from this procedure because of several no-shows in one session. Thus, we have an equal number of male and female agents in treatments **LOW** and **UNEQUAL** and 47% female participants in **HIGH**. After taking a seat in their cubicle, the subjects read the experimental instructions. Questions about the instructions were answered in private. Before the actual series of search tasks started, subjects participated in a test of understanding. In this test, they made choices in one exemplary search task, which was not payoff relevant. After making their choices, they had to calculate possible payoffs from this task themselves (which all were able to do).

The actual series of payoff-relevant search tasks was followed by a preference elicitation task and a brief questionnaire.[5] Earnings were converted from experimental currency units (ECU) to Euro. One ECU was worth Euro 1.50 to subjects. Then subjects were paid in private based on the random draws and the decisions they made in the search task and the preference elicitation task. Finally, they left the laboratory individually. Note that all stakeholders were seated in closed sound-proof cubicles while agents were seated in open cubicles. Stakeholders only left their cubicles once all agents were paid. The experiment was conducted at the Essen Laboratory for Experimental Economics (elfe) in Germany. We implemented a between-subjects design, so each participant only participated in one treatment. In total, 216 student subjects took part in our experiment. They earned Euro 18.92 on average from a session that lasted approximately 90 minutes.

[4] We chose this parameterization to create a realistic representation of pay differentials and a salient unfairness at the same time. In our experiment, those earning less receive 35 percent of the total pay. In standard ultimatum games using typical stakes, for example, offers in the range between 30 and 40 percent of the pie have been found to be rejected 30 percent of the time by List & Cherry (2000) and 42 percent of the time by Slonim & Roth (1998).

[5] The estimates of risk aversion and loss aversion we obtained using this particular task may be confounded by probability weighting (Kahneman & Tversky, 1979) or a diminishing sensitivity to zero (e.g. Erev et al., 2008).

3 Hypothesis

We assume that, with increasing cognitive effort, people improve their choices and are more likely to follow the optimal path, i.e., to choose projects in decreasing order of Z_i (cf. Slonim, 1994). We assume that agents form rational expectations about the additional pay from exerting additional effort and from identifying a path with higher-expected payoff. In other words, we assume that there is an unobserved search cost (in addition to c) in identifying the optimal path.[6]

We follow Breza et al. (2018) in adapting the framework by DellaVigna et al. (2016) and assume that morale effects can influence the choice of effort and thereby performance (see Online Appendix C for details). A purely self-interested agent is not influenced by morale effects and will increase her effort if pay increases. She will not be influenced by the size of stakeholder pay. Thus, a wage increase results in a higher optimal effort choice. This suggests more optimal decisions in **HIGH** than in **LOW** and **UNEQUAL**, as effort is increasing in one's own pay.

An agent averse to disadvantageous inequality, as modeled by Fehr & Schmidt (1999), will be influenced by moral effects. She will also exert more effort when her pay increases. However, an increase in the stakeholder's pay will reduce the agent's effort. This relationship is described by our main hypothesis:

Hypothesis: *An increase of the stakeholder's pay leads to a lower propensity to acquire information optimally.*

[6] This approach follows Bolton & Faure-Grimaud (2009) in that we assume that people deal optimally with their cognitive limitations (but see also the discussion by Conlisk, 1996, and Lipman, 1991, 1999, on the infinite regress problem). It differs from their approach in that we assume people to correctly anticipate the costs and benefits of exerting additional effort. Our approach is also similar to the papers on search behavior by Caplin et al. (2011) and Reutskaja et al. (2011) in that we treat search as a real-effort task in as much as we cannot control for the subjects' effort costs. We also take a reduced-form approach: We only consider the resulting payoffs of the dynamic decision task and do not model each individual decision, i.e., each step of the decision path within a task. Such an analysis could consider further behavioral drivers such as coherence effects that have been observed to influence information search (see, e.g., Mischkowski et al., 2021).

This suggests fewer optimal decisions in **UNEQUAL** than in **LOW** due to the agent's disutility from being paid less than the stakeholder.

4 Results

In the following, we first treat performance as a binary variable indicating whether an agent followed the optimal decision sequence. We calculate the optimal decision sequence for each task separately by considering a subject's previous choices within a task. Also note that for the analyses presented here, we consider the optimal path to be the optimal path under risk neutrality and, thus, to be the same for all participants. This way our results can be readily compared to the analyses of Gabaix et al. (2006). Also, from a management perspective, the normative prediction under risk-neutrality is arguably the more informative benchmark.

In addition, we judge performance based on the opportunity costs incurred in expectation to account for payoff differences across tasks. For this purpose, we calculate the expected payoff of an agent's strategy based on the observed choices and under the assumption that all unobserved choices would have been made optimally (conditional on observed choices and outcomes). We then subtract it from the expected payoff of the optimal decision sequence. Note that we do not apply the payoff differentials δ for this calculation to keep the variation in pay exogenous (cf. Online Appendix C).

Behavior across tasks

In tasks A to E, the optimal sequence diverges from the path described by the directed cognition algorithm. It demands agents to explore the low-probability, high-payoff project first, even though its expected payoff is lower than that of the high-probability, low-payoff project. In tasks F to J, both sequences coincide, i.e., the optimal sequence in these tasks is characterized by choosing projects in the order of decreasing expected payoffs.

The findings by Gabaix et al. (2006) suggest that many people use the directed cognition algorithm as a heuristic. Aggregating the data from our three treatments, we observe a similar pattern for decisions in social settings, as shown in Figure 1: In tasks A to E, only 38% of the agents choose the optimal path, while 53% follow the directed cognition algorithm. In tasks F to J, 81% choose the optimal path, but both strategies cannot be distinguished. The share of optimal play is significantly larger in tasks F to J than in tasks A to E ($p < 0.001$, Wilcoxon signed-rank test).[7]

Thus, in line with the results by Gabaix et al. (2006), the share of optimal information acquisition behavior is significantly larger in tasks that can be solved optimally by the directed cognition algorithm. Note, however, that performance may not be directly comparable between both sequences. Payoffs from the optimal path tend to be larger in sequence F to J than in sequence A to E. With respect to expected opportunity costs, there is more money left on the table in tasks F to J than in tasks A to E ($p < 0.001$, Wilcoxon signed-rank test). In the former, the expected opportunity costs amount to ECU 0.873, in the latter to ECU 0.568. That means, the higher share of optimal choices in Tasks F to J does not compensate for the loss in expected payoffs generated by those behaving sub-optimally.

Figure 1: Share of optimal choices across tasks (all treatments pooled, 95% confidence intervals)

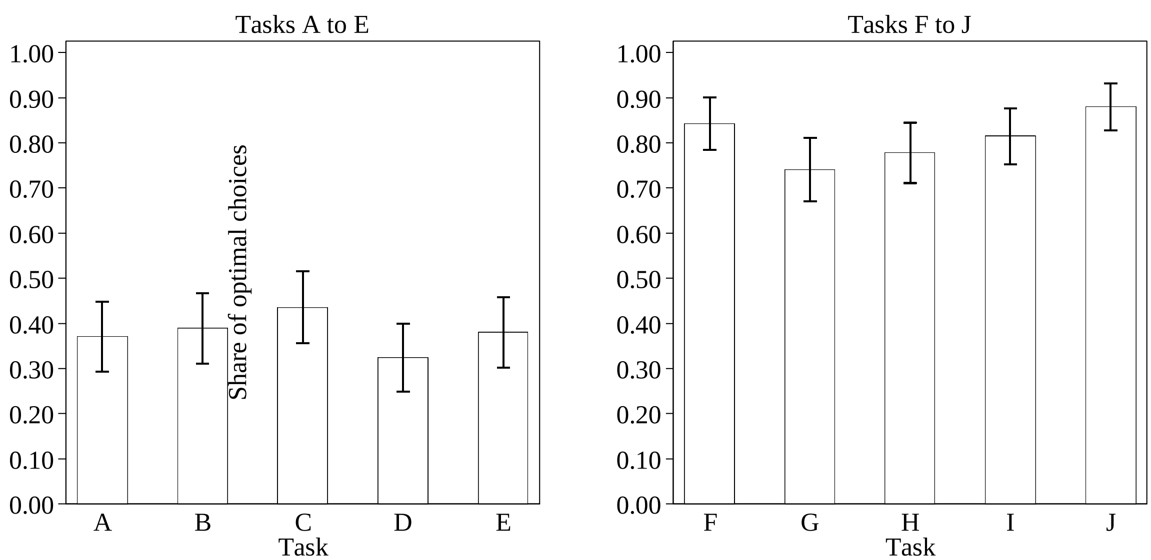

[7] All reported p-values are based on two-tailed tests.

Treatment comparisons

Figure 2 displays the share of optimal choices averaged across tasks separately by treatment. Average performance over tasks F to J ranges from 79.4% to 82.7% while the average performance over tasks A to E ranges from 28.3% to 46.7%. The tasks A to E are played optimally less often than tasks F to J within each of the treatments ($p < 0.001$, Wilcoxon signed-rank tests). As expected, performance also varies more strongly across treatments in tasks A to E than in tasks F to J, in which performance is quite close to the optimum already.

When comparing the share of optimal choices across the three treatments, we do not find any significant differences with respect to tasks F to J ($p \geq 0.368$, Mann-Whitney-*U* tests). When comparing performance in tasks A to E the difference between the **UNEQUAL** treatment (in which agents perform best) and the **LOW** treatment (in which they perform worst) is significant ($p = 0.033$) while the other two treatment comparisons are insignificant ($p \geq 0.209$).

As a robustness check, we run the ordinary least squares (OLS) regressions shown in Table 3. They take the share of tasks an agent played optimally as a dependent variable and are run separately for tasks A to E and tasks F to J. Depending on the model specification, they also take the agent's gender and age as explanatory variables next to the treatment dummies **LOW** and **HIGH** (with **UNEQUAL** as the baseline category). The regression results are in line with the pattern described above: While there are significant treatment differences in all models for tasks A to E, we find no significant differences in tasks F to J across regression models (1), (2) and (3) ($p \geq 0.579$, Wald tests). With respect to tasks A to E, the OLS regressions confirm that performance in **LOW** is worse than in the **UNEQUAL** treatment ($p \leq 0.037$). Based on the estimates the share of optimal choices is around 18% lower in **LOW**. There is no significant difference between **UNEQUAL** and **HIGH** or between **LOW** and **HIGH** ($p \geq 0.201$).

We also compare the share of choices adhering to the directed cognition algorithm across treatments in tasks A to E. On average 47.2% of choices in **UNEQUAL**, 51.1% of choices in **HIGH** and 60.6% of choices in **LOW** were made based on the directed cognition algorithm. However, there are no significant differences between treatments ($p \geq 0.148$, Mann-Whitney-*U* tests). Figures and regression results are provided in Online Appendix E.

When comparing the expected opportunity costs between treatments in tasks A to E, we also observe a better performance (i.e. the lowest opportunity costs) in treatment **UNEQUAL** (ECU 0.399) than in treatments **LOW** (ECU 0.610) and **HIGH** (ECU 0.696). But the differences are insignificant ($p \geq 0.059$, Mann-Whitney-*U* tests). Figures and regression results are provided in Online Appendix F.

Our main hypothesis suggests that an inequality averse agent will perform worse under disadvantageous inequality. But, if anything, we observe better performance in **UNEQUAL** than in **LOW**: Based on the risk neutral benchmark, performance measured as the number of optimal choices even increases with stakeholder pay. Accordingly, we formulate the following observation:

Observation: *Unilaterally increasing the pay of an agent's stakeholder does not have a negative effect on the agent's performance.*

Figure 2: Share of optimal choices across tasks by treatment (95% confidence intervals)

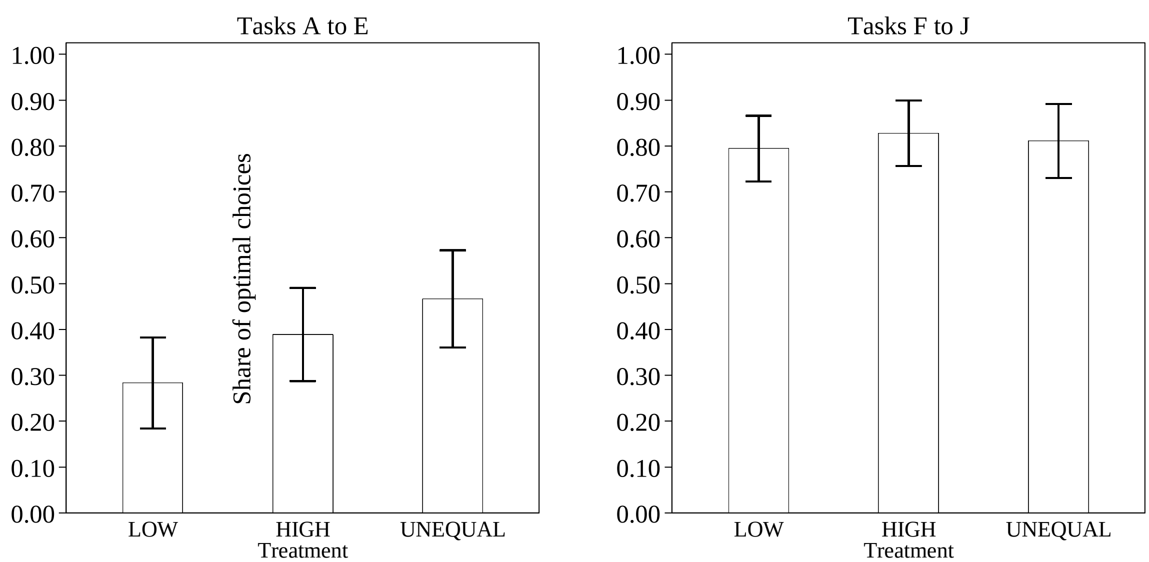

Table 3: OLS regressions, dependent variable: Share of optimal choices

	Tasks A to E			Tasks F to J		
	(1)	(2)	(3)	(1)	(2)	(3)
LOW	-0.183*	-0.183*	-0.184*	-0.017	-0.017	-0.018
	(0.086)	(0.087)	(0.087)	(0.064)	(0.063)	(0.064)
HIGH	-0.078	-0.075	-0.076	0.017	0.014	0.013
	(0.087)	(0.086)	(0.086)	(0.064)	(0.064)	(0.064)
Female		0.085	0.086		-0.088	-0.081
		(0.070)	(0.072)		(0.051)	(0.051)
Age			0.001			0.006
			(0.011)			(0.008)
N	108	108	108	108	108	108
R^2	0.042	0.056	0.056	0.003	0.031	0.037

Standard errors in parentheses, * p<0.050, ** p<0.010, *** p<0.001.

5 Conclusion

First, we confirm the result by Gabaix et al. (2006), who find that their directed cognition algorithm describes actual behavior quite well. We observe a similar decision pattern for agents acting in a setting with a passive stakeholder. In addition, when considering the share of optimal decisions as a measure of effort, we observe that disadvantageous inequality in pay improves the quality of decisions. This effect is present for tasks that cannot be solved optimally by applying the directed cognition algorithm. When considering expected opportunity costs as a measure of effort, we still observe the best performance under disadvantageous inequality, but the difference is not significant. Overall, we find that disadvantageous inequality in pay does not have a negative effect on performance.

The observation that an increase in pay does not necessarily lead to more effort is in line with the results of several empirical studies that consider individual-effort provision in real-effort settings in the laboratory (see, e.g., Gneezy & Rustichini, 2000; Pokorny, 2008; Ariely et al., 2009) or in the field (see, e.g., Camerer et al., 1997; Crawford & Meng, 2011; Fehr & Goette, 2007). Somewhat puzzling is the increase in performance with unequal pay we find with respect to the share of optimal decisions. This behavior may of course be specific to our sample or the procedures we have used. But interestingly, it is in line with recent findings by Cullen & Perez-Truglia (2021) who observe that in the field, workers increase their effort once they learn that their managers earn more than they thought. Their results suggest that negative morale effects are absent with respect to vertical comparisons: Workers satisfaction with their jobs, their salary and their firm's wage differentials are unchanged when they learn that their superiors earn more. Even though our setting is neutrally framed, the asymmetry between agents and stakeholders may also explain the absence of negative morale effects.

References

Amore, M. D., & Failla, V. (2020). Pay dispersion and executive behaviour: Evidence from innovation. *British Journal of Management*, 31(3), 487-504.

Ariely, D., Gneezy, U., Loewenstein, G. & Mazar, N. (2009): Large stakes and big mistakes, *Review of Economic Studies*, 76(2), 451-469.

Bolton, P. & Faure-Grimaud, A. (2009): Thinking ahead: The decision problem, *Review of Economic Studies*, 76(4), 1205-1238.

Bracha, A., Gneezy, U. & Loewenstein, G. (2015): Relative pay and labor supply, *Journal of Labor Economics*, 33(2), 297-315.

Breza, E., Kaur, S. & Shamdasani, Y. (2018): The morale effects of pay inequality, *The Quarterly Journal of Economics*, 133(2), 611-663.

Camerer, C., Babcock, L., Loewenstein, G. & Thaler, R. (1997): Labor supply of New York City cabdrivers: One day at a time, *Quarterly Journal of Economics*, 112(2), 407-441.

Caplin, A., Dean, M. & Martin, D. (2011): Search and satisficing, *American Economic Review*, 101(7), 2899-2922.

Card, D., Mas, A., Moretti, E. & Saez, E. (2012): Inequality at work: The effect of peer salaries on job satisfaction, *American Economic Review*, 102(6), 2981-3003.

Cohn, A., Fehr, E., Herrmann, B. & Schneider, F. (2014): Social comparison and effort provision: Evidence from a field experiment, *Journal of the European Economic Association*, 12(4), 877-898.

Conlisk, J. (1996): Why bounded rationality?, *Journal of Economic Literature*, 34(2), 669-700.

Crawford, V. P. & Meng, J. (2011): New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational expectations targets for hours and income, *American Economic Review*, 101(5), 1912-1932.

Cullen, Z., & Perez-Truglia, R. (2021). How much does your boss make? The effects of salary comparisons. *Working paper*.

DellaVigna, S., List, J. A., Malmendier, U. & Rao, G. (2016): Estimating social preferences and gift exchange at work, *Working paper*.

Downes, P. E. & Choi, D. (2014): Employee reactions to pay dispersion: A typology of existing research, *Human Resource Management Review*, 24(1), 53-66.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. Journal of Behavioral Decision Making, 21(5), 575-597.

Fehr, E. & Goette, L. (2007): Do workers work more if wages are high? Evidence from a randomized field experiment, *American Economic Review*, 97(1), 298-317.

Fehr, E. & Schmidt, K. M. (1999): A theory of fairness, competition, and cooperation, *Quarterly Journal of Economics*, 114(3), 817-868.

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

Fochmann, M., Sachs, F., & Weimann, J. (2019). Managing wages: Fairness norms of low-and high-performing team members. Working paper.

Gabaix, X., Laibson, D., Moloche, G. & Weinberg, S. (2006): Costly information acquisition: Experimental analysis of a boundedly rational model, *American Economic Review*, 96(4), 1043-1068.

Gneezy, U. & Rustichini, A. (2000): Pay enough or don't pay at all, *Quarterly Journal of Economics*, 115(3), 791-810.

Greiner, B., Ockenfels, A. & Werner, P. (2011): Wage transparency and performance: A real-effort experiment, *Economics Letters*, 111(3), 236-238.

Kahneman, D. & Tversky, A. (1979): Prospect theory: An analysis of decision under risk, *Econometrica*, 47(2), 263-291.

Lipman, B. L. (1999): Decision theory without logical omniscience: Toward an axiomatic framework for bounded rationality, *Review of Economic Studies*, 66(2), 339-361.

Lipman, B. L. (1991): How to decide how to decide how to...: Modeling limited rationality, *Econometrica*, 59(4), 1105-1125.

List, J. A. & Cherry, T. L. (2000): Learning to accept in ultimatum games: Evidence from an experimental design that generates low offers, *Experimental Economics*, 3(1), 11-29.

Mischkowski, D., Glöckner, A., & Lewisch, P. (2021). Information search, coherence effects, and their interplay in legal decision making. *Journal of Economic Psychology*, 87, 102445.

Pfeffer, J. (2007): Human resources from an organizational behavior perspective: Some paradoxes explained, *Journal of Economic Perspectives*, 21(4), 115-134.

Pfeffer, J. & Langton, N. (1993): The effect of wage dispersion on satisfaction, productivity, and working collaboratively: Evidence from college and university faculty, *Administrative Science Quarterly*, 382-407.

Pokorny, K. (2008): Pay - but do not pay too much: An experimental study on the impact of incentives, *Journal of Economic Behavior & Organization*, 66(2), 251-264.

Reutskaja, E., Nagel, R., Camerer, C. F. & Rangel, A. (2011): Search dynamics in consumer choice under time pressure: An eye-tracking study, *American Economic Review*, 101(2), 900-926.

Shaw, J. D. (2014): Pay dispersion, *Annual Review of Organizational Psychology and Organizational Behavior*, 1(1), 521-544.

Slonim, R. & Roth, A. E. (1998): Learning in high stakes ultimatum games: An experiment in the Slovak Republic, *Econometrica*, 66(3), 569-596.

Slonim, R. (1994): Learning in a search-for-the-best-alternative experiment. *Journal of Economic Behavior & Organization*, 25(2), 141-165.

Wang, T., Zhao, B. & Thornhill, S. (2015): Pay dispersion and organizational innovation: The mediation effects of employee participation and voluntary turnover, *Human Relations*, 68(7), 1155-1181.

Weitzman, M. L. (1979): Optimal search for the best alternative, *Econometrica*, 47(3), 641-654.

Yanadori, Y. & Cui, V. (2013): Creating incentives for innovation? The relationship between pay dispersion in R&D groups and firm innovation performance, *Strategic Management Journal*, 34(12), 1502-1511.

Online Appendix to

"Costly information acquisition:

The influence of stakeholder earnings"

Timo Heinrich[1], Bindu Arya[2], Alexander Haering[3] and Sven Horak[4]

[1] Timo Heinrich, Institute for Digital Economics, Hamburg University of Technology, Am Schwarzenberg Campus 2 (B), 21073 Hamburg, Germany; email: timo.heinrich@tuhh.de (corresponding author).

[2] Bindu Arya, College of Business Administration, University of Missouri–St. Louis, USA; email: bindua@umsl.edu.

[3] RWI – Leibniz-Institute for Economic Research Essen and Fresenius University of Applied Sciences, Germany; email: alexander.haering@rwi-essen.de.

[4] Peter J. Tobin College of Business, St. John's University, USA; email: horaks@stjohns.edu.

Appendix A – Parameters of the tasks, example and screenshot

Table A1: Task parameters based on Gabaix et al. (2006)

	V_1	p_1	V_2	p_2	V_3	p_3
Task A	21	0.09	10	0.76	1	1
Task B	19	0.11	10	0.79	1	1
Task C	23	0.09	13	0.72	1	1
Task D	18	0.12	10	0.81	1	1
Task E	20	0.12	12	0.85	1	1
Task F	22	0.48	11	0.74	1	1
Task G	24	0.34	9	0.70	1	1
Task H	18	0.52	11	0.74	1	1
Task I	25	0.39	9	0.70	1	1
Task J	10	0.09	8	0.85	1	1

Note: Costs c_i are always 1.

Figure A1 contains a screenshot of the agent's screen with the three projects. It shows that probabilities and payoffs are known to subjects. Payoff amounts are given in experimental currency units (ECU). For each task we observe one of the possible decision paths. To illustrate the possible decision paths, let us consider this as an example: Project 1 is the safe option, guaranteeing a payoff of ECU 1, because it is always successful. Project 2 has a payoff of ECU 21 in case of success and a probability of success of 9%. Project 3 yields ECU 10 in case of success. Its success rate is 76%. Projects 2 and 3 yield no payoff if they fail. Would a subject acquire information about the outcome of the two uncertain projects? Probably, because, in expectation, both yield a payoff that is larger than the payment of $c = 1$. But which one should be explored first?

Calculating the index Z_i for the three projects as in formula (1) yields the optimal sequence. The values are $Z_1 = 1$, $Z_2 = (0.09*21-1)/0.09 = 9.89$ and $Z_3 = (0.76*10-1)/0.76 = 8.68$. Therefore, project 2 should be explored first and taken if successful, because its outcome is certain and $Z_2 = 21$. If not, project 3 should be explored and selected if successful. Only if neither is successful, project 1 is selected. The decision sequence will be different for someone applying the directed cognition algorithm. This heuristic only considers the gain from the next move being made. In the beginning, the best-known successful project is project 3, therefore $S = 1$. Based on (2), the gains from the three projects are G_1 = 1(1-1)-0 = 0, G_2 = 0.09(21-1)-1 = 0.8 and G_3 = 0.76(10-1)-1 = 5.84, so now project 3 is considered first. If it is successful, the gains change to G_1 = 1(1-10)-0 = -9, G_2 = 0.09(21-10)-1 = -0.01 and G_3 = 1(10-10)-0 = 0, and the sequence ends. If project 3 fails, project 2 will be explored.

Figure A1: Screenshot of agent's screen (task A)

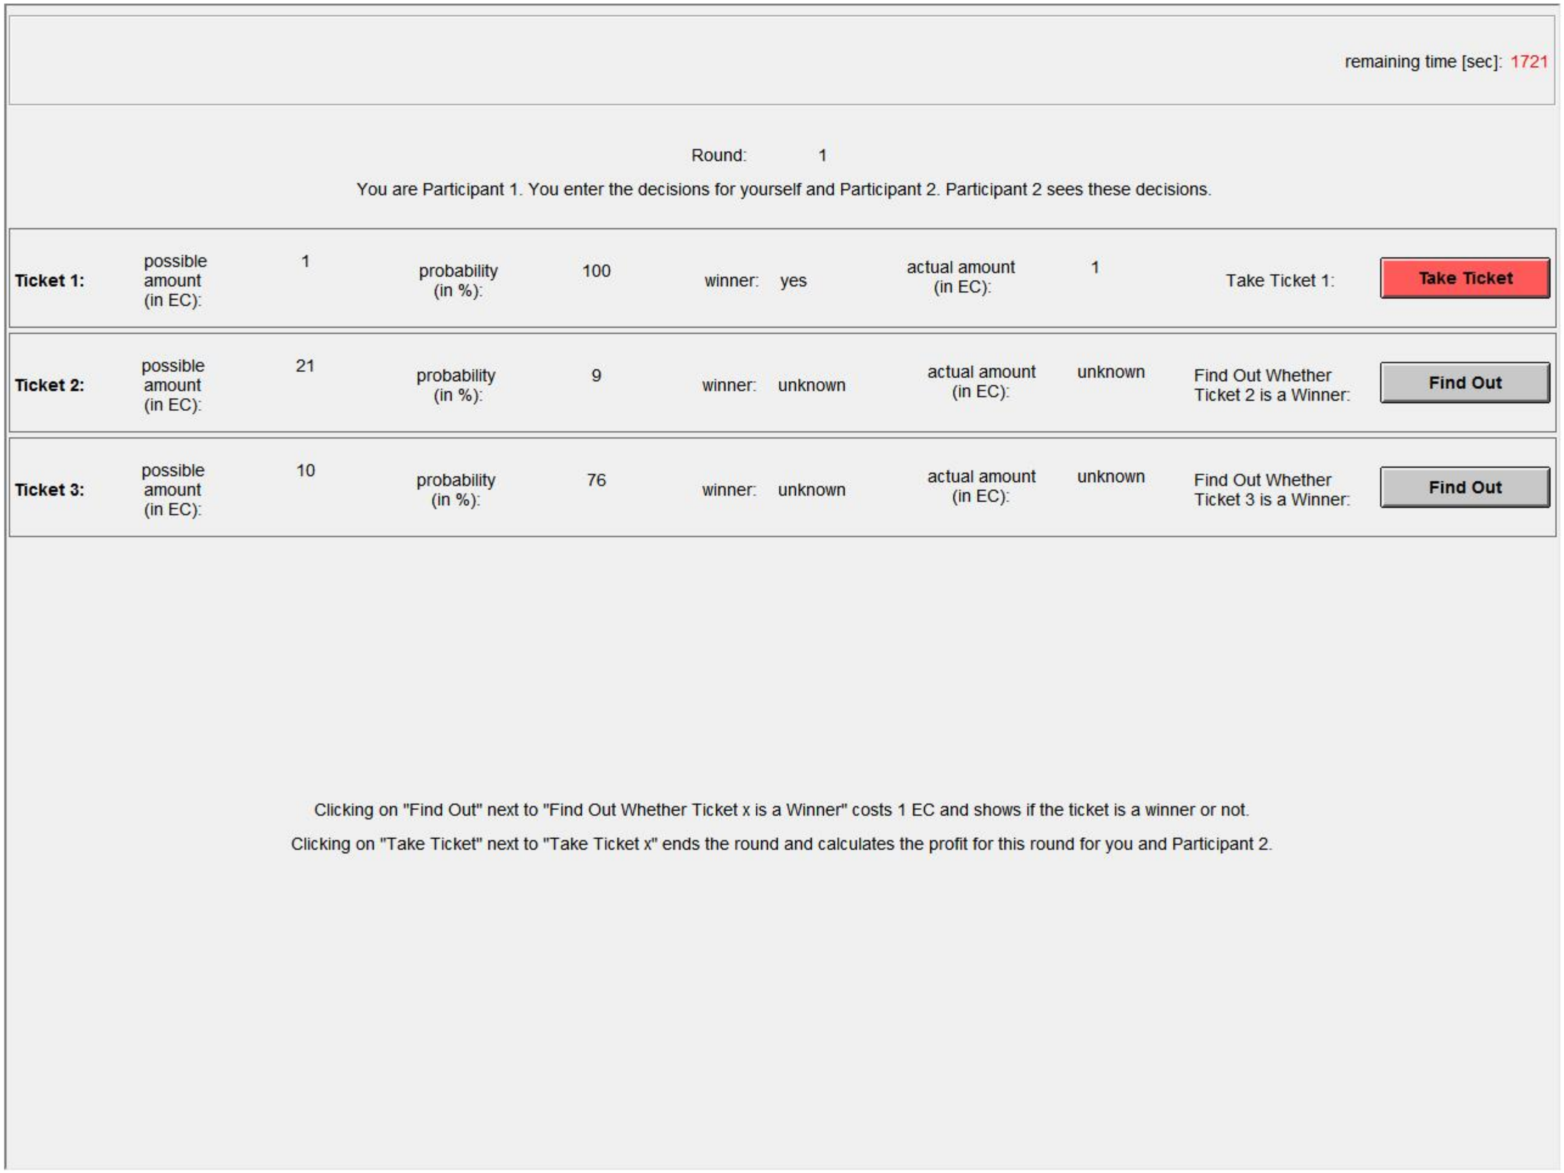

Appendix B – Instructions (translated from German)

Welcome to the experiment!

Foreword

You are taking part in a study of decision behavior in experimental economic research. During the study you and the other participants are asked to make decisions. By doing so you can earn money. How much money you earn, will depend on your decisions. At the end of the experiment you will be paid your total earnings in cash.

None of the participants will be informed about the identity of the other participants during the experiment.

Instructions

Please read the following instructions carefully. About 5 minutes after we hand out the instructions, we will come to you to answer questions. If you have any questions during the experiment, you can raise your hand at any time. We will then come to you.

During the experiment you will take part in 10 lottery rounds.

Assignment of partner and role

At the beginning of the experiment you will get assigned to your partner and randomly receive one of two roles: Participant 1 or Participant 2. During the whole experiment you will play together with the same partner and keep the same role. You will not learn anything about the identity of your partner.

Description of the lottery games

During the lottery rounds of the experiment you will be asked to choose among a group of "lottery tickets". Tickets have different probabilities of paying off. All amounts are given in Experimental Currency (EC).

Participant 1 enters the decisions for himself and Participant 2. Participant 2 only sees these decisions.

Consider the following set of tickets:

Ticket 1:	possible amount (in EC): 20	probability (in %): 50	winner: unknown	actual amount (in EC): unknown	Find Out Whether Ticket 1 is a Winner:	Find Out
Ticket 2:	possible amount (in EC): 10	probability (in %): 75	winner: unknown	actual amount (in EC): unknown	Find Out Whether Ticket 2 is a Winner:	Find Out
Ticket 3:	possible amount (in EC): 1	probability (in %):100	winner: yes	actual amount (in EC): 1	Take Ticket 3:	Take Ticket

Here, ticket 1 has a 50 % chance of paying 20 EC. To be more precise, it pays off 20 EC with a 50 % chance and pays off 0 EC with a 50 % chance. Similarly, ticket 2 pays off 10 EC with a 75 % chance and 0 EC with a 25 % chance. Ticket 3 always pays 1 EC.

For 1 EC you can find out whether an "unknown" ticket is a winner or not by clicking on "Find Out" next to "Find Out Whether Ticket is a Winner" in the far right hand column. You may investigate as many tickets as you wish, at a cost of 1 EC each.

At any point you may stop learning and end the round by choosing ONE ticket. In order to do so Participant 1 clicks on "Take Ticket" in the far right hand column. The "Take Ticket" option will be available for all tickets that are known to be a winner or not. Even for those that have been revealed as losing tickets. Once you select a ticket, we will calculate your winnings for that round as the value of the selected ticket minus

1 EC for each “unknown” ticket that you investigated. We will tell you your profit for the round and then reset the screen for the next round.

<u>Payment</u>

At the end of the experiment you will get paid your winnings (net of learning costs) in 1 of the 10 rounds selected at random. [**LOW:** <u>Participant 1 and Participant 2 receive a malus of 30 % in addition to the profit of this round. I.e. their profit is reduced by 30 %.</u> **HIGH:** <u>Participant 1 and Participant 2 receive a bonus of 30 % in addition to the profit of this round. I.e. their profit is increased by 30 %.</u> **UNEQUAL:** <u>Participant 1 receives a malus of 30 % in addition to the profit of this round, while Participant 2 receives a bonus of 30 %. I.e. their profit is decreased by 30 % or increased by 30 %.</u>]

1 EC is equivalent to 1.50 Euro.

Example: *Assuming Participant 1 and Participant 2 have investigated 2 tickets in one round and take a ticket with an actual amount of 10 EC. In this case the profit for this round is for each of the both participants:*

Profit for this round = 10 EC – 2 EC = 8 EC

[**LOW:**

*However, Participant 1 and Participant 2 receive a malus in addition to the profit. The malus is calculated as the absolute value of 30 % of the profit for this round, in this round 8 EC * 30 % = 2.4 EC. That is, the actual profits that are paid are:*

Profit for Participant 1 = 8 EC – 2.4 EC = 5.6 EC

Profit for Participant 2 = 8 EC – 2.4 EC = 5.6 EC

]

[**HIGH:**

*However, Participant 1 and Participant 2 receive a bonus in addition to the profit. The bonus is calculated as the absolute value of 30 % of the profit for this round, in this round 8 EC * 30 % = 2.4 EC. That is, the actual profits that are paid are:*

Profit for Participant 1 = 8 EC + 2.4 EC = 10.4 EC

Profit for Participant 2 = 8 EC + 2.4 EC = 10.4 EC

]

[**UNEQUAL:**

*However, Participant 1 receives a malus in addition to the profit and Participant 2 a bonus. Malus and bonus are calculated as the absolute value of 30 % of the profit for this round, in this round 8 EC * 30% = 2.4 EC. That is, the actual profits that are paid are:*

Profit for Participant 1 = 8 EC – 2.4 EC = 5.6 EC

Profit for Participant 2 = 8 EC + 2.4 EC = 10.4 EC

]

You will start with a practice round that illustrates the process. Then the experiment will begin. The experiment will consist of 10 rounds (in addition to the practice round).

On the next page you will find the screens for one exemplary round.

Screen for Participant 1

Screen for Participant 2

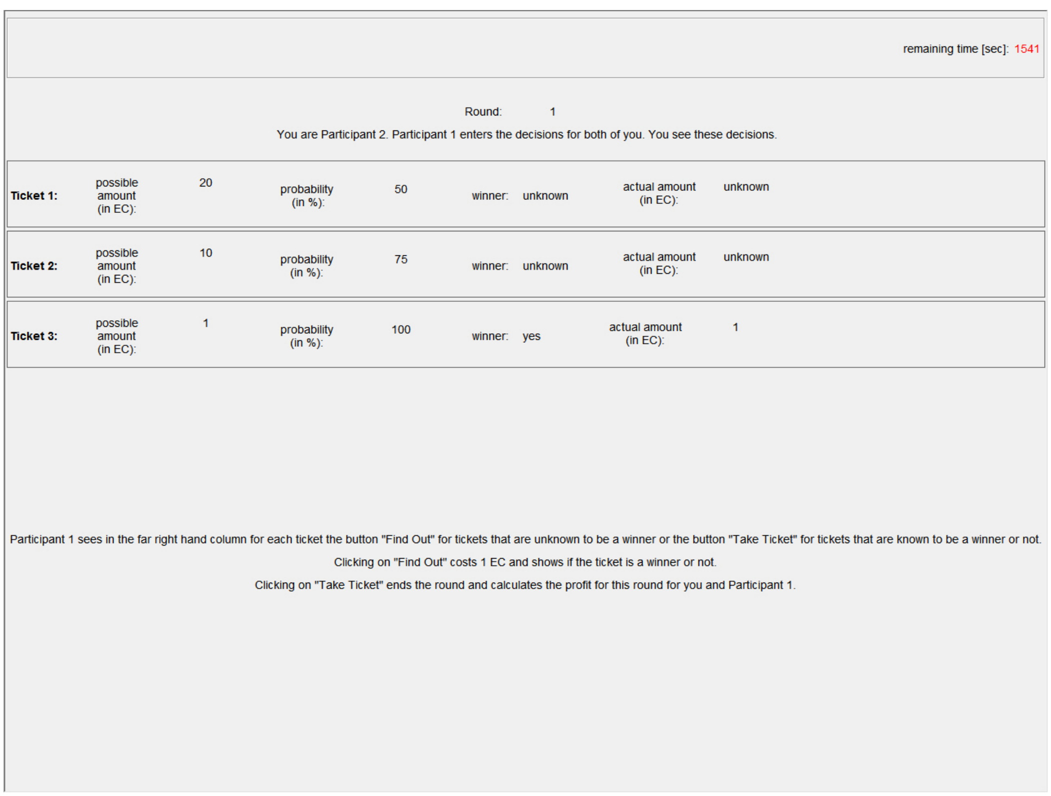

Lottery decisions

In the following we ask you to make 12 lottery decisions for which you will receive additional earnings.

Please read the following instructions carefully. About 5 minutes after we hand out the instructions, we will come to you to answer any questions. If you have any questions during the experiment, you can raise your hand at any time. We will then come to you.

Description of the lottery decisions

One of the 12 lottery decisions will be randomly selected and the corresponding profit will be added to your previous earnings. Please note that your decisions influence your profit only and not that of the other participant.

During the 12 decisions you can choose between a lottery and a secure (or "safe") payment. You will have an additional 2.00 Euro at your disposal. According to your decisions and the lottery results (which is randomly selected by a computerized random number generator), this amount can be reduced or increased.

Example: *You can choose between a secure payment of 0.00 Euro and a lottery. If you choose the lottery you receive a profit of 1.00 Euro with a probability of 50% and with a probability of 50% a loss of -0.50 Euro. Do you prefer the secure payment or the lottery?*

Payment

At the end of the experiment the profit of the lottery decisions will be paid out to you, together with the previous earnings.

The profit of the lottery decisions is composed of the endowment of 2.00 Euro and – depending on your decision – of the secure payment or the lottery result.

Appendix C – Derivation of main hypothesis

We follow Breza et al. (2018) in adapting the framework by DellaVigna et al. (2016) and assume that morale effects can influence the choice of effort. An agent i's utility function is given by:

$$U(e_i, w_i, y) = e_i w_i - g(e_i) - M(e_i, w_i, y) \tag{3}$$

where e_i is her cognitive effort and w_i is her pay, i.e., the amount by which her payoff increases from additional cognitive effort. The effort costs are convex and given by $g(e_i)$ with $g'(e_i) > 0$, $g''(e_i) > 0$ and $g(0) = 0$. The morale effect $M(e_i, w_i, y)$ depends on the agent's pay w_i, the effort e_i and some reference point y.

In our basic model, there is no morale effect. The optimal effort is simply derived from the first-order condition $w_i - g'(e_i) = 0$ as:

$$e_i^* = g'^{-1}(w_i) \tag{4}$$

where $g'^{-1}(\ )$ is the inverse of the effort cost function and monotonically increasing. Because $-g''(e^*) < 0$, this is a maximum.

We now consider how the prediction above changes once we introduce morale effects. Like Breza et al. (2018), we assume that the reference point y is formed based on the pay per unit of effort received by stakeholders. In our case, the reference pay is equal to the pay the stakeholder earns from the agent's actions ($y = w_{-i}e_{\mathrm{i}}$). We model this based on the theory by Fehr & Schmidt (1999) of inequality aversion (IA) so that an agent suffers from unequal pay. Following them, we assume a linear relationship between inequality and utility. The morale effect that is subtracted in (3) is then given by:

$$M_{IA}(e_i, w_i, w_{-i}) = \alpha \max\{e_i w_{-i} - e_i w_i, 0\} + \beta \max\{e_i w_i - e_i w_{-i}, 0\}. \tag{5}$$

The positive parameters α and β capture two types of disutility from inequality ($\alpha, \beta \geq 0$): Agents suffer more from inequality that is disadvantageous to them than from inequality that is advantageous to them ($\alpha \geq \beta$). Furthermore, agents always like to have more money than less ($\beta < 1$).[5]

Depending on the type of inequality, solving for the optimal effort yields:

$$e_i^* = g'^{-1}(w_i - \alpha(w_{-i} - w_i)) \text{ if } w_{-i} > w_i \text{ and} \tag{6}$$

$$e_i^* = g'^{-1}(w_i - \beta(w_i - w_{-i})) \text{ if } w_i \geq w_{-i}.$$

In our setting, pay inequality cannot be changed by the agents but is exogenously determined through the pay structure. Because $g'^{-1}(\ )$ is monotonically increasing, a change in its input parameters also changes optimal effort accordingly. Based on the theory of Fehr & Schmidt (1999) and their parameterization of α and β we can derive hypotheses about the reactions to pay changes: In our case, we consider disadvantageous inequality where the agent earns less than her stakeholder ($w_i < w_{-i}$). Here, an increase of the agent's pay leads to an increase in effort. However, an increase in the stakeholder's pay has the following effect: the larger the stakeholder's pay the lower the agent's effort.

[5] As Fehr & Schmit (1999) and Breza et al. (2018) point out, other assumptions about α and β can be made. The assumption that people do not like being paid less than their peers captured through a positive α is quite common. However, if people are status seeking, β may also be negative, which would increase the effect.

Appendix D – Raw data

Table C1a: Optimal choices

Group	Treatment	Task A	B	C	D	E	F	G	H	I	J	Age P1	P2	Female P1	P2	Exp. Opp. Cost
1	LOW	0	0	0	0	1	0	1	1	1	1	27	23	0	1	1.72
2	LOW	0	0	0	0	0	1	1	1	1	1	28	26	0	1	0.04
3	LOW	0	0	0	0	0	0	0	0	0	1	24	20	1	0	3.26
4	LOW	1	1	1	1	1	1	1	1	1	1	24	21	0	1	0.00
5	LOW	1	1	1	1	1	1	1	1	1	0	21	25	0	0	0.62
6	LOW	0	1	0	0	0	1	1	0	1	1	27	22	1	0	1.19
7	LOW	0	1	0	1	0	1	1	1	1	1	24	27	1	0	0.02
8	LOW	0	0	0	0	0	1	0	1	0	1	27	22	1	0	1.20
9	LOW	0	0	0	0	0	0	0	0	0	1	23	24	0	1	2.19
10	LOW	1	1	1	1	1	1	1	1	1	0	29	23	0	1	0.06
11	LOW	0	0	0	0	0	1	1	1	1	1	26	24	0	1	0.04
12	LOW	0	0	0	0	1	1	1	1	1	1	25	24	0	1	0.03
13	LOW	1	1	1	1	1	1	1	1	1	1	28	18	1	0	0.00
14	LOW	0	0	0	0	0	0	0	0	0	1	24	22	1	0	1.38
15	LOW	1	1	0	0	1	0	0	1	1	1	26	25	1	1	1.56
16	LOW	0	1	0	0	0	1	1	1	0	1	19	24	1	0	2.22
17	LOW	1	1	1	1	1	1	1	1	1	1	24	22	0	1	0.00
18	LOW	1	0	0	0	1	1	0	0	0	1	25	24	1	1	1.03
19	LOW	1	0	1	0	0	1	1	1	1	1	25	26	0	1	0.03
20	LOW	0	0	1	0	0	1	1	1	1	1	25	22	0	1	0.84
21	LOW	0	0	0	0	1	1	1	1	1	1	24	26	1	0	0.03
22	LOW	0	0	0	0	0	1	1	1	1	1	22	27	0	0	0.04
23	LOW	0	1	1	0	0	1	0	1	1	1	27	24	1	0	3.22
24	LOW	0	0	0	0	0	0	0	1	1	1	22	25	0	0	1.65
25	LOW	0	0	0	0	0	1	1	1	1	1	18	24	0	1	0.04
26	LOW	0	0	0	0	0	1	1	1	1	1	20	27	0	0	0.04
27	LOW	0	1	1	0	0	1	1	1	1	1	27	21	0	1	0.03
28	LOW	1	1	1	0	0	1	1	1	1	1	25	20	0	1	0.02
29	LOW	0	0	0	0	0	0	1	1	1	1	19	17	1	0	1.26
30	LOW	0	0	0	0	0	1	1	0	0	1	19	19	1	0	1.32
31	LOW	0	0	0	0	0	1	1	1	1	1	18	22	1	0	0.04
32	LOW	0	0	0	0	0	1	1	0	1	1	26	25	1	0	0.29
33	LOW	1	1	1	0	1	1	1	1	1	1	23	23	1	0	0.01
34	LOW	0	0	0	0	1	1	1	0	1	1	28	26	1	0	0.28
35	LOW	1	0	1	0	1	1	1	1	1	0	18	23	1	1	0.08
36	LOW	0	0	0	0	0	1	0	1	1	1	22	22	0	1	0.36

Table C1b: Optimal choices

Group	Treatment	Task										Age		Female		Exp. Opp. Cost
		A	B	C	D	E	F	G	H	I	J	P1	P2	P1	P2	
1	HIGH	0	0	1	0	0	1	0	1	1	1	20	24	1	1	1.22
2	HIGH	1	0	0	1	0	1	1	1	1	1	25	25	0	1	0.93
3	HIGH	1	1	1	1	1	1	1	1	1	1	22	20	0	1	0.00
4	HIGH	0	0	0	0	0	1	1	0	1	0	26	25	0	1	1.39
5	HIGH	1	0	1	0	1	1	1	1	0	0	20	26	1	0	2.95
6	HIGH	1	0	1	1	1	1	1	1	1	1	23	21	1	0	0.01
7	HIGH	1	1	1	1	1	1	1	1	1	1	23	25	1	0	0.00
8	HIGH	1	1	1	1	1	1	1	1	1	1	21	28	1	0	0.00
9	HIGH	0	0	0	0	0	1	1	1	1	1	25	20	0	1	0.08
10	HIGH	1	0	1	1	0	1	1	1	1	1	26	26	0	1	1.64
11	HIGH	1	1	1	1	1	1	1	1	1	1	20	20	0	1	0.00
12	HIGH	1	1	0	1	1	1	1	1	1	1	28	23	0	1	0.86
13	HIGH	1	0	1	1	0	1	1	1	1	1	21	24	1	0	0.02
14	HIGH	1	1	1	1	0	1	1	1	1	0	23	31	1	0	1.13
15	HIGH	0	0	0	0	0	0	0	0	0	1	23	27	1	0	1.58
16	HIGH	0	0	0	0	0	1	1	1	1	1	29	25	0	0	0.04
17	HIGH	0	0	0	0	0	1	0	0	0	1	24	31	1	1	2.57
18	HIGH	0	0	0	0	1	1	1	1	1	1	27	29	0	1	0.03
19	HIGH	0	0	0	0	0	0	0	1	1	1	20	22	0	1	0.73
20	HIGH	1	1	0	1	0	1	1	1	1	1	26	27	0	1	0.92
21	HIGH	1	1	1	0	1	1	1	1	1	1	30	26	1	0	0.01
22	HIGH	0	0	1	1	0	1	0	0	0	1	26	22	1	0	2.55
23	HIGH	1	1	0	0	1	1	1	1	1	1	21	27	1	0	0.74
24	HIGH	0	0	0	0	0	1	1	1	1	1	19	21	1	0	0.04
25	HIGH	0	0	0	0	0	1	1	1	1	1	29	22	0	1	0.07
26	HIGH	0	0	0	0	1	1	1	1	1	1	27	20	0	1	0.74
27	HIGH	0	0	0	0	0	1	1	1	1	1	30	21	0	1	0.04
28	HIGH	0	0	0	0	0	1	0	0	0	1	23	29	0	1	1.01
29	HIGH	1	0	0	1	0	1	1	1	1	1	26	19	1	0	0.02
30	HIGH	0	0	1	1	1	0	0	0	1	1	23	22	1	0	2.83
31	HIGH	0	1	1	0	1	1	0	0	1	1	20	28	1	0	1.29
32	HIGH	0	0	0	0	0	1	1	1	1	1	25	20	0	0	0.04
33	HIGH	1	0	1	0	0	1	0	0	1	1	21	22	0	1	0.58
34	HIGH	0	1	1	0	0	1	1	1	1	1	20	25	1	1	0.03
35	HIGH	0	0	0	0	0	1	1	1	1	1	27	29	0	0	0.04
36	HIGH	0	0	0	0	0	1	0	0	0	1	19	26	0	0	1.86

Note: EOC indicates the average expected opportunity cost in ECU.

Table C1c: Optimal choices

Group	Treatment	Task										Age		Female		Exp. Opp.
		A	B	C	D	E	F	G	H	I	J	P1	P2	P1	P2	Cost
1	UNEQUAL	0	0	0	0	0	1	1	1	0	1	25	27	0	1	0.86
2	UNEQUAL	0	0	0	0	0	1	1	1	1	1	24	27	0	1	0.04
3	UNEQUAL	0	1	1	0	0	0	0	1	0	0	23	23	0	0	3.21
4	UNEQUAL	0	0	1	0	0	1	1	1	1	1	27	20	0	1	0.04
5	UNEQUAL	0	0	0	1	0	1	1	1	1	1	23	24	1	0	0.05
6	UNEQUAL	1	1	1	1	0	1	0	1	1	1	20	20	1	0	1.68
7	UNEQUAL	1	1	1	1	1	1	1	1	1	1	21	25	1	0	0.00
8	UNEQUAL	1	1	1	1	1	1	1	1	1	1	24	25	1	0	0.00
9	UNEQUAL	0	1	1	0	0	1	1	1	1	1	24	24	0	1	0.03
10	UNEQUAL	1	1	1	1	1	1	1	1	1	1	27	23	0	0	0.00
11	UNEQUAL	0	1	0	1	1	0	0	0	1	0	28	31	0	1	1.50
12	UNEQUAL	1	1	1	1	1	1	1	1	1	0	20	20	0	1	0.06
13	UNEQUAL	0	0	0	0	0	0	0	0	0	0	18	18	1	0	5.24
14	UNEQUAL	0	0	1	1	1	1	1	1	1	0	20	24	1	0	0.08
15	UNEQUAL	1	1	1	1	0	1	1	1	1	1	35	29	1	0	0.01
16	UNEQUAL	0	0	0	0	0	0	0	0	1	1	27	22	1	0	1.82
17	UNEQUAL	0	0	1	0	0	1	0	1	1	1	21	20	0	1	0.36
18	UNEQUAL	1	1	1	0	1	1	1	1	1	1	20	21	0	1	0.73
19	UNEQUAL	0	0	0	0	0	1	1	1	1	1	19	20	0	0	0.04
20	UNEQUAL	0	0	0	0	0	1	1	1	1	1	22	21	0	1	0.04
21	UNEQUAL	0	0	0	0	0	1	1	0	0	1	21	24	1	0	0.69
22	UNEQUAL	1	1	1	1	0	1	1	1	1	1	25	20	1	0	0.01
23	UNEQUAL	0	0	1	1	1	1	1	1	1	1	26	22	1	0	0.02
24	UNEQUAL	0	0	0	0	0	1	1	1	1	1	27	19	1	0	0.04
25	UNEQUAL	0	0	1	0	0	1	0	0	0	1	32	23	0	0	1.01
26	UNEQUAL	0	1	0	0	1	1	1	1	1	1	22	35	0	1	0.02
27	UNEQUAL	0	1	0	1	1	1	1	1	1	1	23	24	0	1	0.01
28	UNEQUAL	1	1	1	1	1	1	1	1	1	1	20	20	0	1	0.00
29	UNEQUAL	1	1	1	0	1	1	1	1	1	1	23	19	1	0	0.01
30	UNEQUAL	0	0	0	0	0	1	1	0	1	1	24	25	1	0	0.29
31	UNEQUAL	0	0	0	0	1	1	1	1	1	1	21	19	1	0	0.88
32	UNEQUAL	1	1	1	1	0	0	1	0	0	0	20	24	1	0	2.41
33	UNEQUAL	1	1	1	0	1	1	1	1	1	0	26	20	0	1	0.79
34	UNEQUAL	0	0	0	0	1	0	1	1	1	1	28	24	0	1	1.25
35	UNEQUAL	1	1	1	1	1	0	1	1	1	1	21	25	1	1	0.37
36	UNEQUAL	0	1	0	1	0	1	1	1	1	1	24	30	1	0	0.02

Note: EOC indicates the average expected opportunity cost in ECU.

Appendix E – Choices consistent with the directed cognition algorithm (DC) and unclassified choices (other)

Figure E1: Share of DC choices across tasks (all treatments pooled, 95% confidence intervals)

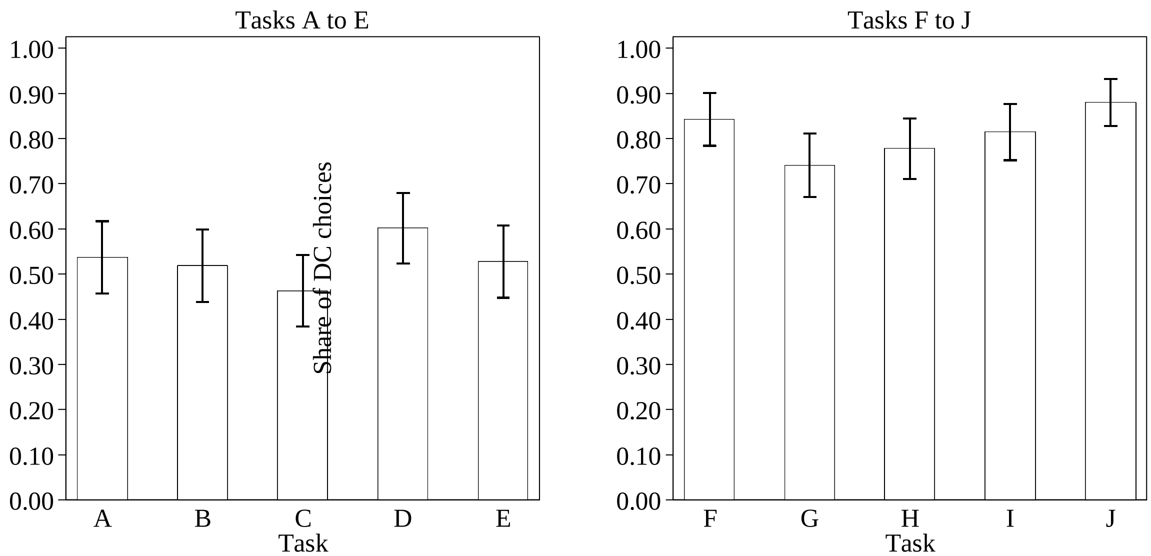

Figure E2: Share of DC choices across tasks by treatment (95% confidence intervals)

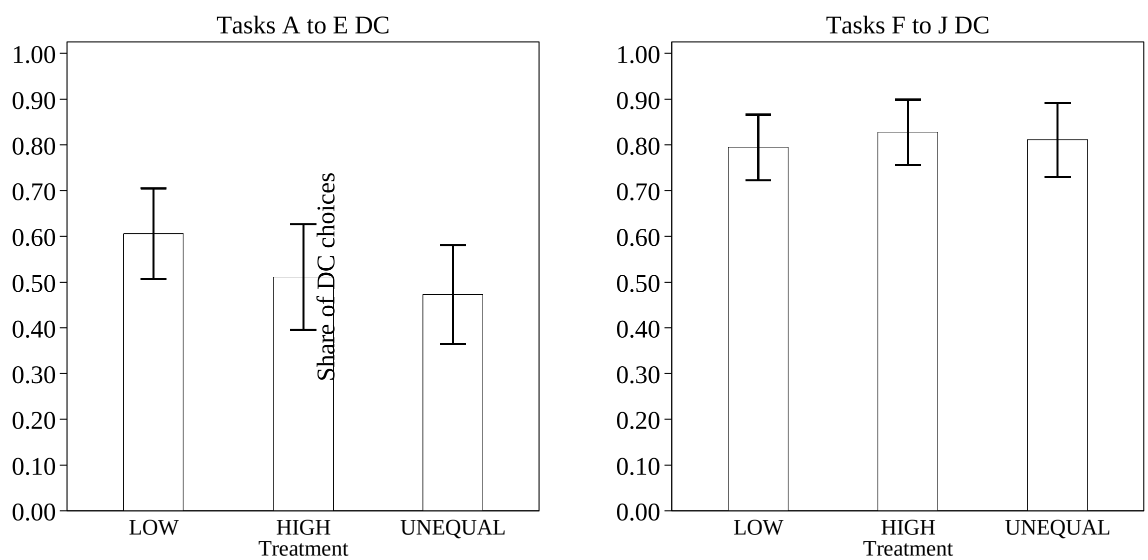

Figure E3: Share of other choices across tasks (all treatments pooled, 95% confidence intervals)

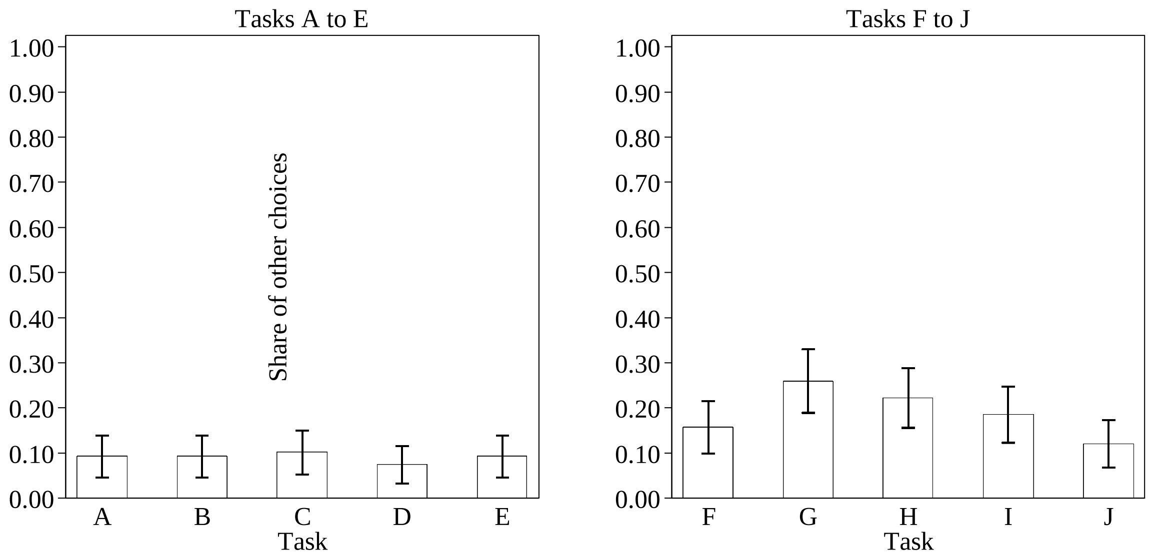

Figure E3: Share of other choices across tasks by treatment (95% confidence intervals)

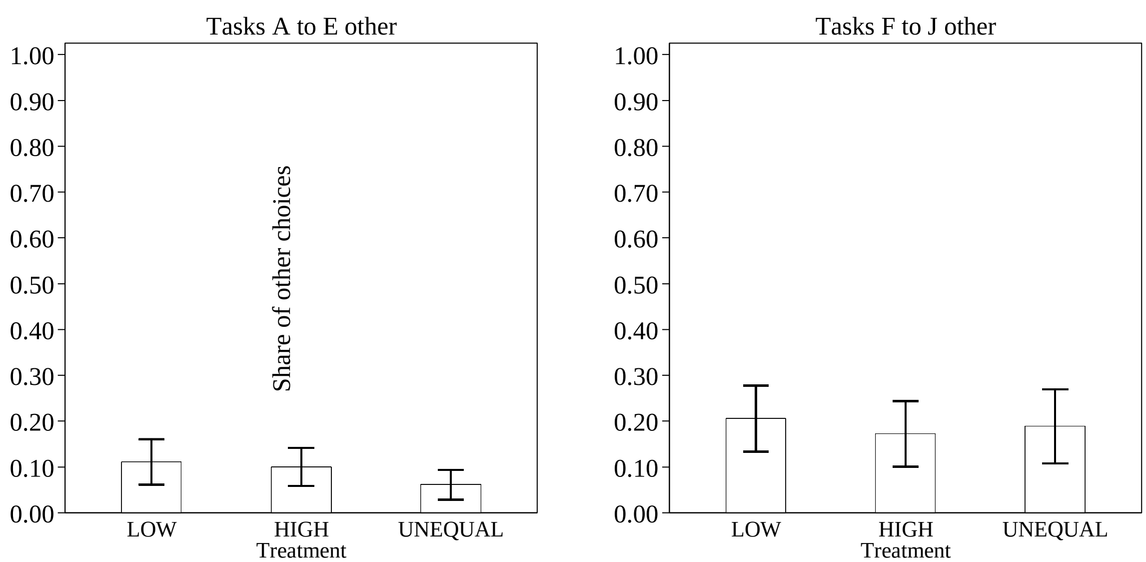

Table E1: OLS regressions, dependent variable: Share of DC choices

	Tasks A to E			Tasks F to J		
	(1)	(2)	(3)	(1)	(2)	(3)
LOW	0.133	0.133	0.134	-0.017	-0.017	-0.018
	(0.087)	(0.088)	(0.088)	(0.064)	(0.063)	(0.064)
HIGH	0.039	0.036	0.036	0.017	0.014	0.013
	(0.094)	(0.092)	(0.093)	(0.064)	(0.064)	(0.064)
Female		-0.122	-0.123		-0.088	-0.081
		(0.073)	(0.076)		(0.051)	(0.051)
Age			-0.001			0.006
			(0.011)			(0.008)
N	108	108	108	108	108	108
LOW vs **HIGH**	0.297	0.274	0.276	0.579	0.600	0.595
R^2	0.021	0.047	0.047	0.003	0.031	0.037

Standard errors in parentheses, * p<0.050, ** p<0.010, *** p<0.001.
LOW vs HIGH represents the *p*-value of a Wald test for equality of coefficients.

Table E2: OLS regressions, dependent variable: Share of other choices

	Tasks A to E			Tasks F to J		
	(1)	(2)	(3)	(1)	(2)	(3)
LOW	0.050	0.050	0.050	0.017	0.017	0.018
	(0.035)	(0.035)	(0.035)	(0.064)	(0.063)	(0.064)
HIGH	0.039	0.040	0.040	-0.017	-0.014	-0.013
	(0.031)	(0.032)	(0.032)	(0.064)	(0.064)	(0.064)
Female		0.037	0.037		0.088	0.081
		(0.028)	(0.029)		(0.051)	(0.051)
Age			0.000			-0.006
			(0.004)			(0.008)
N	108	108	108	108	108	108
LOW vs **HIGH**	0.772	0.790	0.791	0.579	0.600	0.595
R^2	0.021	0.037	0.037	0.003	0.031	0.037

Standard errors in parentheses, * p<0.050, ** p<0.010, *** p<0.001.
LOW vs HIGH represents the *p*-value of a Wald test for equality of coefficients.

Appendix F – Expected opportunity costs

Figure F1: Expected opportunity cost across tasks (all treatments pooled, 95% confidence intervals)

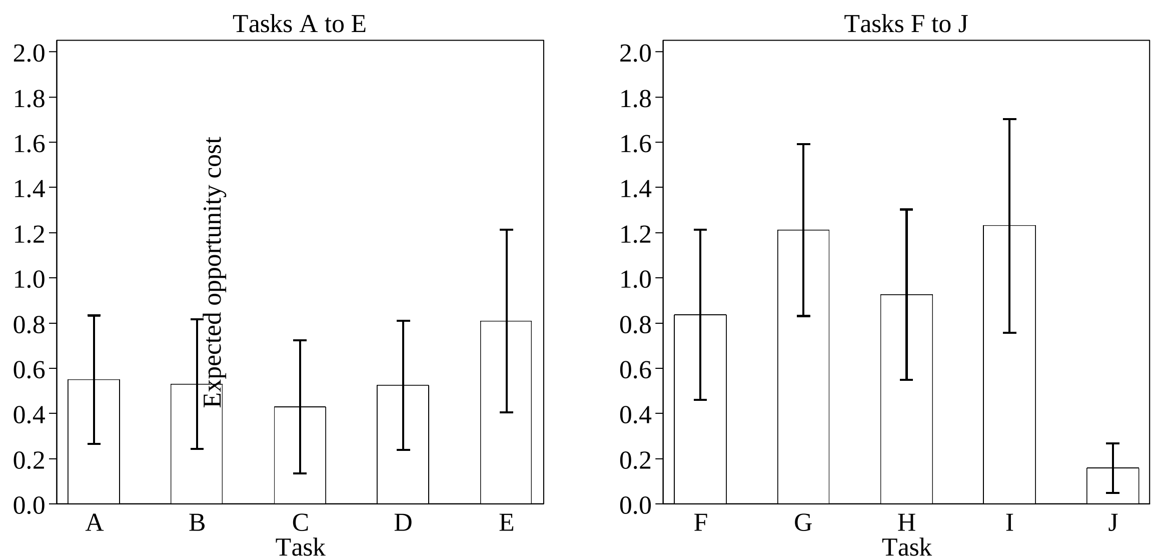

Figure F2: Expected opportunity cost across tasks by treatment (95% confidence intervals)

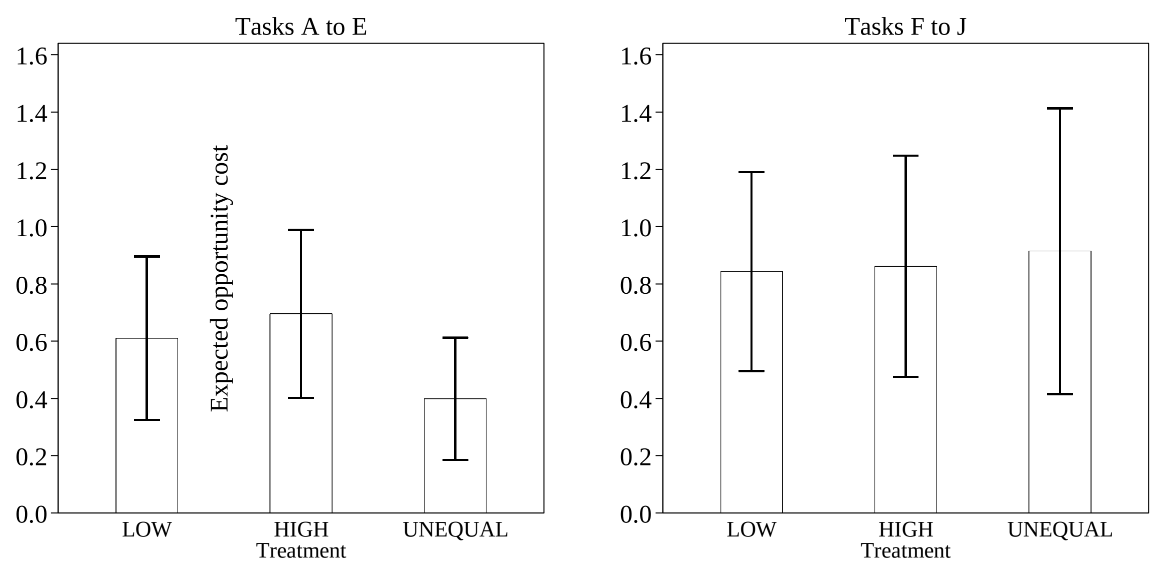

Table F1: OLS regressions, dependent variable: Expected opportunity cost

	Tasks A to E			Tasks F to J		
	(1)	(2)	(3)	(1)	(2)	(3)
LOW	0.212	0.212	0.213	-0.072	-0.072	-0.061
	(0.211)	(0.208)	(0.206)	(0.360)	(0.357)	(0.356)
HIGH	0.297	0.305	0.306	-0.053	-0.039	-0.031
	(0.215)	(0.214)	(0.215)	(0.373)	(0.372)	(0.369)
Female		0.296	0.289		0.513	0.455
		(0.181)	(0.184)		(0.283)	(0.279)
Age			-0.006			-0.050
			(0.027)			(0.043)
N	108	108	108	108	108	108
LOW vs **HIGH**	0.725	0.696	0.699	0.952	0.913	0.921
R^2	0.018	0.042	0.043	0.000	0.031	0.044

Standard errors in parentheses, * $p<0.050$, ** $p<0.010$, *** $p<0.001$.
LOW vs HIGH represents the *p*-value of a Wald test for equality of coefficients.