Nudging innovation: the effect of risk salience

Elisa Matthewes[a], Anis Nassar[b] and Christian Zihlmann[c]

[a]Department of Economics, University of Fribourg, elisa.matthewes@unifr.ch

[b]Department of Economics, University of Fribourg, anis.nassar@unifr.ch

[c]Department of Economics, University of Fribourg, christian.zihlmann@unifr.ch

January 2021*

Abstract

We investigate a nudge designed to foster innovative behavior by changing risk perceptions. In the lab, participants run a virtual lemonade stand. They can either exploit a given strategy or innovate and explore new ones. Their business choices generate respective profits and the subjects receive a performance-based payoff, making innovating risky. To make this risk more salient, we introduce periodic reporting of profits and expect participants to be less innovative. To draw attention away from the risk, we introduce reporting of strategies and expect participants to be more innovative. We find these nudges to affect behavior through the channel of risk attitudes, as only risk-averse participants are affected. However, both treatments inhibited innovation compared to a control with no reporting. We argue this is due to both types of reporting inducing an increased evaluation of risks. We thus encourage further research of risk salience as a tool to foster innovation and recommend practitioners to tailor interventions and their evaluation with care before implementation, as they can backfire.

Keywords: innovation, risk perception, risk aversion, salience, experiment
JEL Classification Codes: C91, D81, D83, D23

*This experiment was pre-registered through the AEA RCT registry (ID: AEARCTR-0005214, and approved by the ethics committee of the University of Fribourg. The authors gratefully acknowledge financial support from the Swiss National Bank fund of the Department of Economics of the University of Fribourg. We are also grateful to Berno Büchel, Holger Herz, Kene Boun My, Christian Zehnder and participants of the CUSO Summer School 2018, and the ESA Round-the-Clock World Meetings 2020 for their valuable comments.

1 Introduction

Innovation is required in many contemporary jobs (Shalley et al., 2000, Unsworth, 2001) and is a driver of firm performance and even survival. Stimulating innovation is thus one of the highest concerns of CEO's (Rudis, 2004).[1] Innovation is inherently a risky endeavour. The literature documents an association between risk attitudes and innovative behavior (Hvide and Panos, 2014, Lazear, 2005, Hall and Woodward, 2010, Kerr et al., 2017, Hudja and Woods, 2019, Koudstaal et al., 2016). Hence, exploiting the possibility to manipulate the *risk* or *risk perceptions* is a promising starting point to tailor interventions designed to promote innovation. However, the literature has so far focused on mechanisms that manipulate the underlying *risk*. In a laboratory experiment, Ederer and Manso (2013) document innovation-inducing effects of contracts that tolerate early failure while rewarding late-stage performance. But changing the incentive schemes within a firm can both be costly and organizationally demanding.

We investigate whether innovation can be induced by channeling *risk perceptions*. Specifically, we ask whether innovative behavior can be nudged by altering the salience of risk, without changing monetary incentives. Adapting the design from Ederer and Manso (2013), we conduct a controlled laboratory experiment: Participants run a virtual lemonade stand over 20 periods. The profit of the virtual stand determines the compensation paid to the participants. Thus, they are incentivized to maximize profits. Since the subjects do not know the profits associated with each of the available choices, they face a trade-off between exploration and exploitation – a core feature of the innovative process (March, 1991): Subjects can either fine-tune well-known strategies (exploitation) or explore untested strategies (exploration) that are thus associated with more risk.

To make the risk of innovating either more or less salient, we introduce two behavioral interventions that are designed to foster innovative behavior. Specifically, we implement a reporting mechanism that happens every three periods. In the profit treatment, participants report their profits for each of the past three periods. The periodical reporting of the profits is ought to increase the salience of the risk. In the strategy treatment, participants report their strategy for the past three periods. This intervention is designed to focus participants on the strategic choice variables and, thereby, intends to make profits, and with it the risk of reduced profits, less salient. In the control group, participants do not report. Hence, we hypothesize that participants in the strategy treatment engage in more innovative behavior while those in the profit treatment are the least innovative. We also hypothesize that the effect of our

[1] To understand which kind of interventions can effectively promote innovative behavior, various determinants of innovation have already been identified and studied. A rich literature investigates, for example, the relationships between innovation and different types of leadership (Shalley and Gilson, 2004, Hughes et al., 2018), personality traits (Amabile, 1996) or incentives (Amabile, 1996, Kohn, 1993). Incentives can improve performance in innovative tasks, depending on the incentive scheme (Ederer and Manso, 2013), the need of hierarchy in team structure (Englmaier et al., 2018), or the type of creative task at hand (Charness and Grieco, 2014).

intervention will be more pronounced among the risk-averse participants.

We find that our interventions affect innovative behaviors through the channel of risk attitudes, as hypothesized: Only risk-averse participants display different patterns of innovation conditional on the treatment. Also in line with our hypothesis, we report that risk-averse participants are more attentive to the profit and explore less when assigned to the profit treatment than when assigned to the control group. However, contrarily to our expectations, the strategy intervention backfired and was detrimental to innovation. Subjects in the strategy treatment do not provide more attention to the strategic variables and participants in this treatment also engage in less exploration than control participants, for a given level of effort. Although we cannot completely exclude that reputational concerns or limited attention matter in our decision environment, these explanatory approaches are unable to explain some key results: First, reputational concerns do not obviously explain why only the more risk-averse participants react to the treatments. Also, maintaining a reputation in our anonymous setting is costly. Hence, we consider it unlikely to be the driving mechanism. Further, we rule out the conjecture that subjects reporting their strategy shift time and effort from a more productive (e.g. information-processing) to a rather unproductive domain, that is writing reports without an instrumental value: We find that the less risk-averse subjects actually decreased their effort and that the more risk-averse, strategy-reporting types innovate less compared to the control group when controlling for the effort level. Thus, we argue that the backfiring intervention is likely a consequence of both reporting treatments inducing an increased perception of risk compared to no reporting.

Our study extends on three main strands of existing research. First, we augment the evidence about fostering innovation through the channel of risk attitudes. Some authors have investigated which incentive structures can successfully induce exploratory behavior. These studies manipulate the underlying *risk* that subjects are facing. Closest to our study, Ederer and Manso (2013) find that incentive schemes that tolerate early failure can induce innovative behavior. Moreover, the authors report that this effect is mainly driven by risk-averse participants.[2] This confirms theoretical predictions by Manso (2011) who shows that the optimal incentive scheme that motivates innovation will tolerate or even reward early failure. This is because innovation involves the exploration of untested approaches that are likely to fail, and thus, standard pay-for-performance schemes that punish failure will be detrimental to innovation. Multiple other studies investigate how different incentive schemes affect creativity. Overall, they find only small differences between standard incentivizations like flat fees, linear payments or tournament incentives (Erat and Gneezy, 2015, Charness and Grieco, 2018).[3]

[2] Such schemes include an exploration contract that is tolerant of early failure but rewards late-stage performance or golden parachutes, that guarantee a certain level of compensation if participants fail in their innovation.

[3] Note that creativity is not an element of the exploitation-exploration trade-off that we study. However, as it shares common aspects, and namely risk-taking, we consider these findings relevant. For perspective on the

However, Knight et al. (2003) show that also non-monetary interventions can reduce risk and foster innovative behavior. They suggest that education leads to more innovation adoption from farmers because it provides them with new skills and reduces uncertainty. We do not manipulate the riskiness of the task. Instead, we design a mechanism that alters the *perception of risk*. We thus propose and investigate a novel approach to stimulate innovative behavior. In line with the previous literature, we find only risk-averse subjects reacting to the induced treatments.

Second, we add to the literature on changing *risk perceptions*. While risk attitudes are commonly considered as stable personality traits and thus as more challenging to manipulate, influencing risk perceptions is suggested to be more effective in changing risky choice behavior (Weber and Milliman, 1997, Sitkin and Weingart, 1995, Pennings and Wansink, 2004). Risk perceptions are influenced by various factors, e.g. stress (Sobkow et al., 2016), graphical visualizations or framing (Lévy-Garboua et al., 2012, Tombu and Mandel, 2015). This can be exploited to (re-)direct behavior (Eppler and Aeschimann, 2009).[4] One wide field for applications is health behavior. The efficacy of various nudges targeting risk perceptions has been extensively investigated in this context (Ferrer and Klein, 2015, Gerrard et al., 1999): For example, Myers (2014) finds that media health messages can change smokers risk perceptions or Banerjee et al. (2021) show that the accuracy of risk perception can be increased by communicating infection risks via raw numbers. Further, how risks are presented is shown to affect risk perceptions in investment settings (Diacon and Hasseldine, 2007, Cohn et al., 2015, Holzmeister et al., 2020). Overall, nudges seem to be effective in influencing risk perceptions, however, this has not yet been exploited to increase innovative behavior. Hence, we contribute to this interdisciplinary field of literature by investigating a new application of a nudge targeting risk perceptions.

Third, we provide a novel perspective on the use of salience as a tool to alter innovative behavior. Considering that people are limited in the attention that they can allocate to information when making decisions (DellaVigna, 2009), varying the salience of certain information affects the decision-making process. Köszegi and Szeidl (2013) and Bordalo et al. (2013) suggest that agents overweight certain parameters of a decision, simply because they are more salient. Salience nudges have thus been shown effective to influence decision-making in various and heterogeneous contexts such as looking before crossing the road, choosing meal options, or selecting the right pod size when farming sea weed (Thaler and Sunstein, 2009, Kurz, 2018, Hanna et al., 2014). However, these contexts have in common that they have a well-defined goal that is known ex-ante: making pedestrians look right before crossing the street in London, influencing students to chose more vegetarian options, or leading farmers to the optimal pod

difference between innovation and creativity, see Hughes et al. (2018).

[4]These approaches affect risk perceptions by either influencing how much people understand the risk (cognitive dimension) or how they feel about the risk (emotional dimension) (Loewenstein et al., 2001).

size. The innovation process in firms, however, implies uncertainty. It follows that it is impossible to make the ex-ante unknown optimal outcome of the innovation process salient. We thus do not nudge our participants to the optimal strategy, but rather use the nudge to alter risk perceptions and encourage innovation. This is novel and closer to what can effectively be implemented in the context of organizations aiming to promote innovative behavior.

Our results confirm the importance of accounting for risk perceptions for practitioners and scholars aiming to promote innovation. We encourage further investigation of salience as a tool for affecting innovative behavior, while warning about a possible backlash effect. This suggests a prudent investigation of behavioral nudges, ideally in controlled environments, before expanding their implementation.

The remainder of the paper is structured as follows. Section 2 presents the experimental design and introduces our hypotheses. Section 3 reports our results. Section 4 provides a discussion and Section 5 concludes.

2 Experimental Design

The experimental task was adapted from Ederer and Manso (2013). Subjects had to solve a task in which they were facing a trade-off between exploration and exploitation: Participants managed a virtual lemonade stand. Over 20 experimental periods, participants had to make decisions on multiple parameters such as the recipe of the lemonade (sugar and lemonade content, color), the location of the lemonade stand and the price of a cup of lemonade. The possible combinations of these choice variables amounts to 23’522’994 combinations. All participants were compensated according to their realized profits such that their aim was to maximize the profit of the fictional lemonade stand, and, thereby, their own earnings. The payoff was determined by a standard pay-for-performance scheme: Participants were paid 50% of the profits they generated during all 20 periods. Participants faced uncertainty since they did not know the profits associated with each of the available choices. However, they did receive a default strategy, i.e. the choices and the associated profit of an imaginary previous manager. The default strategy was not the most profitable strategy.

After each period, participants were informed about the profit of their implemented choices. Additionally, they received a brief customer feedback. The feedback was an informative binary feedback: The computer randomly selected one of the three continuous choice variables (price, lemon or sugar content) and, then, provided a truthful feedback if the value of the selected variable is too high or too low compared to the optimal value (e.g. "Your lemonade is too sweet"). Subjects received no feedback regarding the location and color of the lemonade. Consequently, the feedback was only informative conditional on the chosen location and lemonade color.

The task is characterized by an exploration-exploitation trade-off: subjects can choose to either fine-tune the default strategy – yielding a profit similar to the previous manager (exploitation) – or experiment with new strategies, taking the associated risk of failure but also the chance of success (exploration). The parameters are designed in a way that exploration will increase chances to identify the strategy that leads to the global maximum while exploitation rather leads to local maxima. The parameters to calculate the profits of the lemonade stand in our experiment are one-to-one adapted from Ederer and Manso (2013).[5]

2.1 Treatments

To address our research question, whether innovative behavior can be nudged by a salience intervention, we integrated a reporting stage into the original game: In periods 3, 6, 9 and 12, subjects were requested to submit a report. The focus of these reports was exogenously varied, inducing an attention shift through making a specific aspect of the game salient, and consequently, another aspect of the game less salient. We used a between-subject design with three different groups: First, the control group in which subjects were not asked to report. Second, participants in the profit treatment were requested to report their profits. Third, the strategy treatment requested subjects to report their strategy. Note that the incentive structure between all treatments was identical: The reports were not payment-relevant, i.e. one's reports were not shared with another subject or with the experimenter while playing the game. This was common knowledge for the participants. Hence, the treatment groups only varied in the content of the reports. Controlling for incentive effects ensures that any differences between the treatment groups are causally based on the interventions.

Control Group: Subjects assigned to the control group were not requested to submit a report.

Profit Treatment: After the periods 3, 6, 9 and 12, subjects were requested to report the profits they made within the last three periods. Along with the wording *"Please report your profits of the last three periods."*, subjects faced an entry mask, where they needed to enter the profits of each of the last three periods. The timing and description of the required reports was communicated by the instructions before the start of the business game.

Strategy Treatment: After the periods 3, 6, 9 and 12, subjects were requested to report the strategy they followed within the last three periods. Along with the wording *"Please describe your strategy in the last three periods. Why did you choose this strategy?"*, they faced a free form text field. The timing and description of the required reports was thoroughly communicated by the instructions before the start of the business game.

The treatments served the purpose of shifting attention by making a specific aspect of the game more salient, the strategy or the profit. By manipulating the salient part of the task,

[5] A detailed description of the parametrization can be found in appendix B.2.

we expected to influence the participants judgement and decisions.

2.2 Pre-registered Hypotheses

Ederer and Manso (2013) document that more risk-averse individuals exhibit significantly less innovative behavior in exploration-exploitation tasks than less risk-averse agents. They promote the use of risk-tolerant incentive schemes to increase explorative behavior. While Ederer and Manso (2013) vary the underlying *risk* of the task by using different incentive schemes, we intended to induce more (or less) risk-tolerance by manipulating the *risk perceptions* of the subjects. Based on their findings, we expected to observe more exploratory behavior for agents with lower risk perceptions and the opposite for agents with higher risk perceptions.

Theoretical and experimental literature on salience suggests that shifting attention away from or towards a specific feature, can smoothly guide decision-making processes (e.g. Bordalo et al., 2013, Thaler and Sunstein, 2009). Besides, various studies indicate the effectiveness of framing and nudging in directing risk perceptions (e.g. Banerjee et al., 2021, Holzmeister et al., 2020). Based on this evidence, we expected to influence participants' risk perceptions, and consequent exploratory behaviors, by manipulating the salient part of a task through the content of the reports.

For subjects assigned to the strategy treatment, the salient aspect of the report is the strategy. We thus expected to shift their attention away from the profit. Because decreasing profits is the risk that can inhibit participants' exploration, focusing them away from it is expected to decrease risk perceptions. These participants were thus expected to explore more.

Hypothesis 1a. *Subjects in the strategy treatment explore more than subjects in the control group.*

Conversely, being required to regularly report your previous profits, makes the monetary aspect of the game more salient – and with it the risk to decrease profits. Participants in the profit treatment were thus expected to have an increased perception of risk and to explore less.

Hypothesis 1b. *Subjects in the profit treatment explore less than subjects in the control group.*

The main hypothesis is that our treatment affects behavior and results in higher (lower) exploration behavior for the strategy (profit) treatment. We subsequently investigated the mechanism driving this effect. We posit that a shift of attention to the salient aspect (the strategy, respectively the profit) takes place while the control group serves as a baseline.

Hypothesis 2a. *Subjects in the strategy treatment provide more attention to the strategy than the control group.*

Hypothesis 2b. *Subjects in the profit treatment provide more attention to the profit than the control group.*

We expect that the effect of risk salience on exploratory behavior works through the channel of risk attitudes. An intervention changing the perceived risk should affect risk-averse individuals more strongly than risk-neutral participants. The more risk-averse a participant, the more her innovative behavior should be inhibited by an increased perception of risk. We thus hypothesize heterogeneity in the treatment effect and that the treatment affects specifically risk-averse participants. This line of reasoning goes in the same direction as Ederer and Manso (2013) who find making the incentive structure more risk-tolerant leads to a behavioral change particularly among the risk-averse subjects.

Hypothesis 3a. *Risk-averse subjects assigned to the profit treatment explore less than risk-averse subjects in the control group.*

Hypothesis 3b. *Risk-averse subjects assigned to the strategy treatment explore more than risk-averse subjects in the control group.*

2.3 Sampling

We employed a sequential analysis plan, following the method outlined in Lakens (2014) based on the key outcome variable, the profit obtained in the final period.[6] This plan specified that our data collection will be terminated at $t = 0.5$ – that is half of the required sample size according to our power calculation – when the observed effect size is smaller than our smallest effect size of interest (SESOI), which was set at Cohen's $d = 0.3875$. Since our main effect sizes are lower than the pre-determined SESOI ($d < 0.3875$), we followed the scenario described in our sequential analysis plan and terminated the data collection after collecting 90 observations. The respective results will be described in the following section. A description of the sequential analysis approach and our underlying power analysis can be found in the appendix D. If our expected effect size of $d = 0.5$ came into force, we would require 180 observations to reach statistical significance at the 5%-level (with 80% power). Because the actual effect size is lower than the hypothesized effect size of $d = 0.5$ and since we stopped data collection by adhering to the pre-registered procedure at $t = 0.5$ (90 observations), statistical significance is not attainable.

2.4 Procedures

We collected experimental data in January 2020 at the laboratory of the University of Strasbourg (Laboratoire d'Économie Expérimentale de Strasbourg LEES). The experiment was

[6] In the final period, profit-maximizing subjects should stick to the most profitable strategy that they discovered during all previous periods. Hence, the profit in the final period constitutes a suitable proxy for exploratory behavior in this experiment (see Ederer and Manso, 2013).

programmed with oTree (Chen et al., 2016) and conducted in French. Each of our four sessions lasted approximately 60 minutes. We used experimental currency units called Thalers with an exchange rate of 1:100. All subjects in the laboratory received a fixed show-up fee of 2€, and in addition a performance-based variable payoff. Overall, the average payoff was 15€. Subjects were randomly assigned to the treatment and control groups, constituting the exogenous variation in this study. The random assignment was performed within-session in order to mitigate potential session-specific effects. At the end of each session, we elicited demographics and risk preferences of the subjects (Falk et al., 2018). The sample consists of 90 subjects, i.e. 30 subjects for each treatment group.

3 Results

In the following, we present the results of our study. Each subject is treated as one observation and all standard errors are clustered at the individual level.

3.1 Performance and Explorative Behavior

First, we focus on the effect of our treatments on exploration. We collected several outcome variables that serve as proxies for exploratory behavior, see also Ederer and Manso (2013) for a thorough discussion. In the following, we will focus on i) the profit realized in the final period and ii) the maximum profit over all periods. Other proxy variables are analyzed in the appendix A.1. The results remain qualitatively unchanged.

Figure 1 compares the means of the final and the maximum profit between the treatments. As hypothesized, we find empirical evidence that reporting the profits decreases profits and, with it, exploratory behavior. However, the effect size is small ($d = 0.11$) and lower than the smallest effect size of interest.[7] Thus, this effect is statistically not significant. Further, in contrast to our hypothesis, the data suggest that subjects assigned to the strategy treatment do not realize higher profits than subjects assigned to the control group: the control group earns on average the highest final period profit (146 thalers, profit treatment: 142 thalers, strategy treatment: 133 thalers) and maximum profit of all periods (146 thalers, profit treatment: 142 thalers, strategy treatment: 140 thalers). Thus, against our hypothesis, reporting the strategy does not increase, but decrease profits. The effect size is $d = 0.32$[8] and thus larger than in the profit treatment, but again lower than the SESOI threshold and consequently, not significant.[9]

[7] For this effect to be statistically significant, we would need to collect 1250 observations each in the control group and the profit treatment.

[8] For this effect to be statistically significant, we would need to collect 150 observations each in the control group and the strategy treatment.

[9] Two-sided t-tests: p-values for the final profit are 0.6661 (control and profit treatment), 0.2140 (control and strategy treatment) and 0.4017 (profit and strategy treatment). For the maximum profit p-values are 0.6334, 0.4869 and 0.8246. Similarly, Mann-Whitney U-test, two-sided: p-values of 0.7394, 0.2804 and 0.4964 for final profit; p-values of 0.6843, 0.5249 and 0.8418 for maximum profit.

Figure 1: Means of the final and the maximal profits.

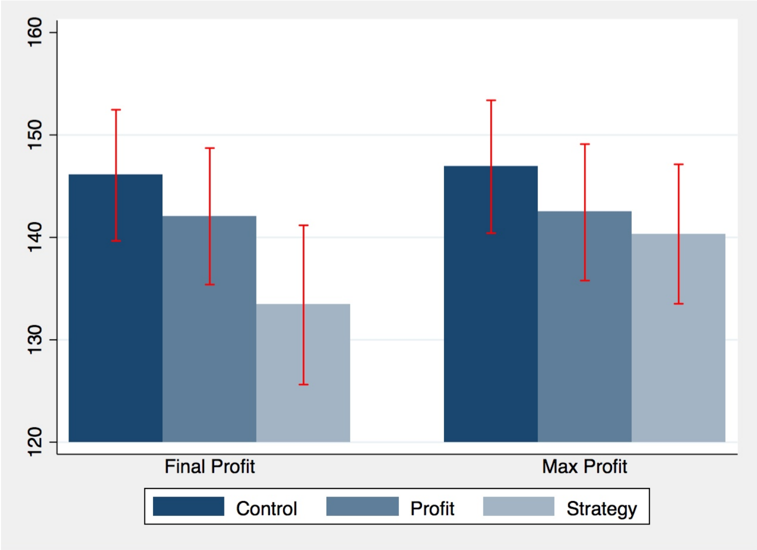

Notes: The figure reports the means of the final and the maximum profit for each of the treatments. Error bars indicate standard errors of the mean.

With the caveat of statistical insignificance due to small effect sizes, we conclude that reporting either the profits or the strategy decreased explorative behavior: we report the expected negative effect for the profit treatment and, in stark contrast to our hypothesis, also find a negative effect for the strategy treatment.

Result 1a. *Subjects in the strategy treatment seem to explore less than subjects in the control group.*

Result 1b. *Subjects in the profit treatment seem to explore less than subjects in the control group.*

We observe that the mean overall profit is the highest in the control group (2323 thalers), followed by the strategy treatment (2151 thalers) and the profit treatment (2143 thalers). Figure 2 shows the evolution of the mean profit in all three treatment groups. The red bars represent periods that precede a reporting screen for the treatment groups. Note that the average profit in the control group is higher than in the reporting treatments in nearly all periods and already before subjects needed to report the first time in period 3. This indicates that subjects in the reporting treatments are already exhibiting a treatment effect in anticipation of future reportings. Indeed, subjects in the profit treatment significantly earn less profit in Period 1 ($p < .01$) than the control group. For the strategy treatment, the difference does not reach statistical significance. We will discuss the implications of a treatment effect already in the first period further in Section 4.

Figure 2: Evolution of profits over time.

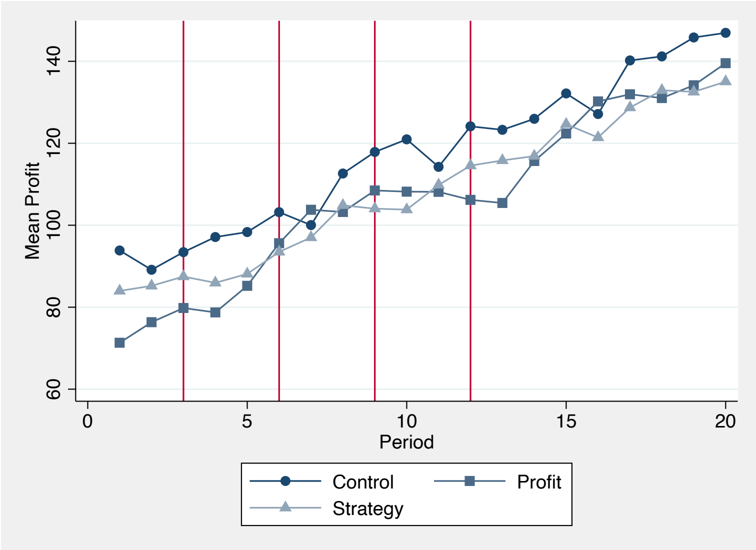

3.2 Attention

According to our hypothesis, the treatment affects behavior because it shifts attention: In the profit treatment, the salience of the profits is increased while in the strategy treatment it is the salience of the strategic choice variables. We construct a measure based on a notes sheet that subjects could fill out voluntarily. We compute the proportion of filled out fields for each subject and, then, derive how many notes each subject took on the strategic variable choices, the periodic profit and the customer feedback relative to the number of total notes taken.

Figure 3 compares the percentage of notes that are taken and illustrates what the subjects across the treatments mainly focused on. It becomes evident that the profit treatment subjects provide less attention to the strategic variables compared to the control group (t-test: $p < 0.01$), but more to the period profits (t-test: $p < 0.01$). Both findings are in line with our hypothesis: Through the reporting mechanism, we successfully shifted the attention of participants in the profit treatment towards the profit. For participants in the strategy treatment, we expect the opposite pattern, namely a shift away from the profit towards the strategy. Contrary to our hypothesis, we observe the same pattern also for the strategy treatment. Yet, the effects do not reach statistical significance at conventional levels.

Result 2a. *Subjects in the profit treatment provide significantly more attention to the profit than the control group.*

Figure 3: Note-taking behavior.

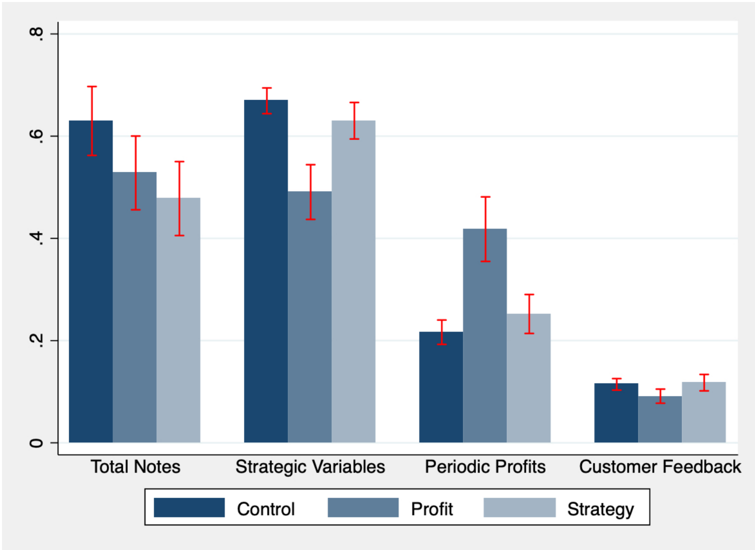

Notes: The figure reports the proportion of total notes taken and, further, the means of subjects' notes of strategic decision variables, the periodic profits and feedback relative to their total notes taken for each of the treatments. Error bars indicate standard errors of the mean.

Result 2b. *Subjects in the strategy treatment do not provide more attention to the strategy than subjects in the control group.*

3.3 Heterogeneity

To test hypotheses 3a and 3b, we elicited the subjects' risk preferences based on the staircase elicitation method by Falk et al. (2018). We split our sample of participants at the median level of risk aversion and classify the participants into less and more risk-averse types.[10] Then, we compare the treatment groups to the control group separately for each type with respect to the final profit and to the maximum profit of all periods.[11]

Interestingly, the results look very different between more and less risk-averse subjects: Figure 4a shows that our treatments do have a negative effect on more risk-averse subjects. That is, requiring more risk-averse subjects to report *reduces* their performance. The effect is most pronounced for the strategy treatment: When more risk-averse subjects are requested to report their strategy, they reduce their profit by nearly 22 thalers, compared to the control group. This is a decrease of approximately 15%, resulting in a relatively large effect size (Cohen's

[10]Median risk aversion, on an index from -3 to +3, amounts to -.03 in our sample, with a mean of .02. Thus, participants in our sample are at median very close to risk-neutrality.

[11]We elicited risk aversion after the experimental game, i.e. the lemonade stand task. Thus, the treatment conditions could potentially have impacted risk preferences. Nevertheless, we do not observe significant differences for risk-preferences among the three groups.

Figure 4: Means of the final and the maximal profits by risk preferences

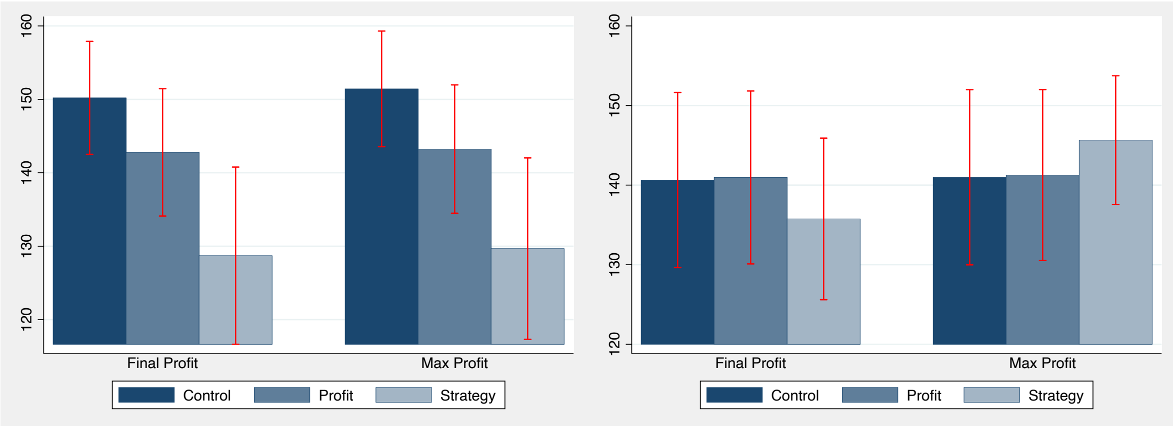

(a) More risk-averse subjects

(b) Less risk-averse subjects

Notes: The figure reports the means of the final and the maximum profit for each of the treatments, for more and less risk-averse subjects separately. Error bars indicate standard errors of the mean.

$d = 0.62$).[12] This effect also holds for the profit treatment, however, the effect as well as its statistical significance are less pronounced.

Figure 4b depicts that the treatment does not lead to a change of behavior for the less risk-averse subjects: They all perform equally well, regardless of the assigned group. The average treatment effect thus seems to be driven by the more risk-averse subjects: only these participants react to the treatment. This is in line with our pre-specified hypothesis that the treatment effect goes through the channel of risk aversion. The detrimental effect for the profit treatment is also in line with our hypothesis, while for the strategy treatment, we find the opposite effect of what we expected: risk-averse subjects decrease their exploratory behavior if they are required to report their strategy.

These findings are corroborated by regression analysis, employing the raw continuous measurement of risk preferences instead of a median split. Table 1 shows that both treatments interacted with risk-aversion lead to a significant and negative coefficient. Thus, compared to subjects in the control group, participants in the treatment groups reduce exploration with higher risk-aversion. The effect is particularly pronounced in the strategy treatment with a significance level below 5%. In short, the higher the risk-aversion, the stronger is the negative reaction to the treatment.

Result 3a. *The more risk-averse participants assigned to the profit treatment explore less than the more risk-averse subjects in the control group.*

Result 3b. *The more risk-averse participants assigned to the strategy treatment explore less*

[12] Which results in low p-values conditional on our sample size: t-test: p=0.11, U-test: p=0.14)

Table 1: Heterogeneous treatment effect

	Final Profit		*Maximal Profit*		*Final Location*		*Exploration Phases*	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Profit	-2.967	0.539	-3.290	-0.570	-0.275	-0.153	-0.055	-0.023
	(9.94)	(9.61)	(9.37)	(9.20)	(0.38)	(0.39)	(0.10)	(0.10)
Strategy	-13.075	-8.708	-7.592	-4.203	-0.506	-0.370	0.010	0.048
	(10.03)	(9.74)	(9.46)	(9.33)	(0.38)	(0.39)	(0.10)	(0.10)
Risk Aversion	10.775	9.521	11.666	10.694	0.796**	0.736*	0.321***	0.320***
	(9.03)	(8.67)	(8.51)	(8.30)	(0.39)	(0.38)	(0.09)	(0.09)
Profit× Risk Aversion	-6.979	-11.193	-7.961	-11.231	-0.775*	-0.878**	-0.338***	-0.383***
	(11.14)	(10.78)	(10.49)	(10.32)	(0.45)	(0.44)	(0.11)	(0.11)
Strategy× Risk Aversion	-22.613*	-22.138*	-28.083**	-27.715**	-1.113**	-1.118**	-0.400***	-0.404***
	(13.30)	(12.76)	(12.54)	(12.21)	(0.52)	(0.51)	(0.13)	(0.13)
Effort (Notes taken)		30.641***		23.771**		0.937**		0.258**
		(10.60)		(10.15)		(0.42)		(0.10)
Constant	144.634***	125.503***	145.356***	130.515***	0.896***	0.341	2.072***	1.905***
	(7.06)	(9.47)	(6.66)	(9.07)	(0.29)	(0.38)	(0.07)	(0.10)
N	90	90	90	90	90	90	90	90
R^2	0.056	0.142	0.066	0.124				
Adjusted R^2	-0.000	0.080	0.011	0.061				

Notes: The dependent variables are different proxies for exploration: the final profit (OLS regressions), the maximal profit (OLS regressions), the location chosen by the subject in the final round (Probit regressions) and the longest duration of an exploration phase (Poisson regressions). Risk Aversion is the subject-specific degree of risk aversion. Effort (Notes taken) proxies effort through the total number of filled out fields on the notes sheet. Level of significance: *p<0.1; **p<0.05; ***p<0.01.

than the more risk-averse subjects in the control group.

Result 3c. *No treatment effect is observed on the less risk-averse participants who perform equally well regardless of the assigned group.*

4 Discussion

Our experimental findings at the interim analysis are only partly in line with the pre-registered hypotheses. As expected, subjects in the profit treatment explore less than subjects in the control group.[13] Further, as hypothesized, we observe a prevalence of the treatment effect on more risk-averse subjects. However, in contrast to our hypotheses, subjects in the strategy treatment do not explore more but *less* than subjects in the control group who face no obligation to report. In this section, we elaborate and discuss potential explanations for this unexpected result. The potential channels are sorted by plausibility in an ascending order.

Hawthorne effect. The Hawthorne effect refers to a change of behavior among the participants of a study in response of their awareness of being observed (Landsberger, 1958). What speaks against this explanation is that only risk-averse subjects reacted to our intervention. Also, recent research that employs sophisticated empirical techniques casts serious doubt about the very existence of the Hawthorne effect (Levitt and List, 2011).

Reputational concerns. Suppose subjects feel reputational concerns when they need to report the strategies or the profits multiple times. Such reputational concerns may arise internally (from themselves) or externally (from the experimenter). The literature hypothesizes

[13] As mentioned previously, the effect is lower than expected and below the SESOI. The reason for the low effect size is that only some subjects react to the treatment induction, namely the rather risk-averse types.

that such concerns lead to conservative behavior, i.e. less innovative and more risk-averse decisions (Holmström, 1999; Prendergast and Stole, 1996; Scharfstein and Stein, 1990; Suurmond et al., 2004). As a result, participants may experience disutility of acting inconsistently: They do not change their strategy too often but consistently apply slight variations of the initial strategy, which results in less exploration. Yet, multiple arguments speak against this explanation: First, it is unclear why in particular risk-averse subjects would experience such reputational concerns. But we observe the treatment effect to happen among risk-averse subjects. We also find that subjects in the strategy treatment are not significantly more risk-averse than subjects in the control group, suggesting that the interventions did not alter the risk preferences of our participants. Second, as Figure 2 depicts, the treatment effect can already be observed in the first period, before subjects even needed to report once. Third, through adhering consistently to a strategy in order to maintain their reputation, subjects will forgo potential profit. Thus, while we cannot fully rule out reputational concerns, we note that if it is the driving force, participants in our experiment would take on substantial costs for maintaining their reputation.

Limited attention. In addition to taking less notes, we also observe that participants in the strategy treatment spend less time on their decision-making and on analyzing their performance (see Figure 11 in the appendix). At first glance, this positive correlation between effort and exploration suggests that subjects in our strategy treatment might have simply exerted less effort for productive subtasks (such as note-taking, information-processing and decision-making), resulting in a hampered performance. This seems to be a reasonable argument since requiring participants to report their strategy is clearly effort-consuming. If subjects have a limited attention span and can only exert a certain fixed level of effort, then participants in the strategy treatment need to trade off time and effort that they would otherwise invest in the decision-making process for writing the reports. While recent literature documents that people exhibit a limited span of effort (Gabaix, 2019), our data do not support this line of reasoning: we find that only the more risk-averse subjects react to the treatment and explore less, however, the less risk-averse participants are the ones that decrease their effort (see Figure 5). Moreover, if limited effort was a reasonable explanation, the exploration level should be the same across treatment groups for a given level of effort. Thus, the production function of explorative behavior – with effort as the input and exploration as the output – should be identical for all subjects, regardless of the treatment condition. Yet, since we observe the more risk-averse type reacting to the treatment but only the less risk-averse type exerting less effort, the production function of the more risk-averse type in the strategy treatment must be impaired. Plotting the production function for more and less risk-averse subjects indeed shows that the production function for the more risk-averse type is altered (see Figure 12 in the appendix): We find that conditional on the effort level, strategy-reporting participants initiate less exploration phases and yield lower profits. Regression analyses underline this finding: In

Table 1 we investigate if the heterogeneous treatment effect persists when we control for the effort level. The interaction term is significant at the 5%-level and remains significant at the same level when controlling for effort. The coefficient of the interaction term remains stable. Hence, the lower exploration level observed in the strategy treatment is not associated with a reduction of effort, but with an impairment of the production function.

Figure 5: Effort exerted by type

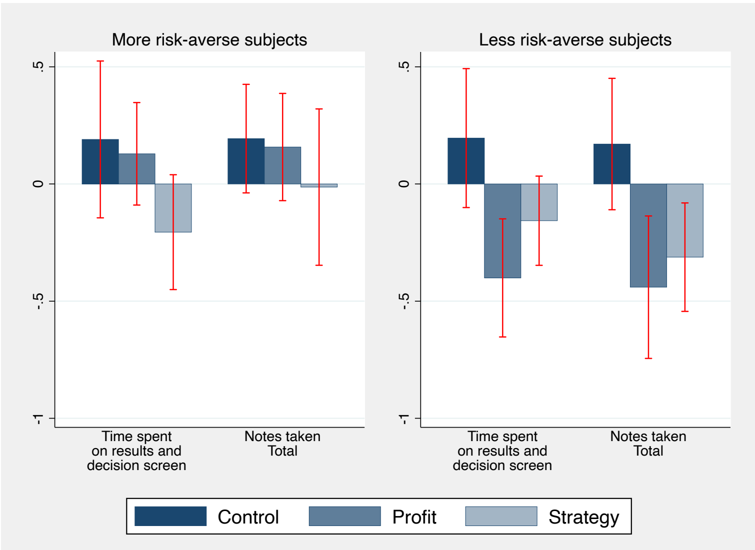

Notes: This figure displays the measured effort through i) time spent on the results and decision screen ii) the total of filled out fields in the notes table sheet. The variables are standardized. Error bars are displayed in red.

Backfiring nudge: Shifting attention *towards* risk. Lastly, the underlying behavioral mechanism might be that our intervention backfired. In the profit treatment, we arguably shifted attention to the profit and with it, to the risky aspect of the game. Consequently, as hypothesized, subjects explore less. In the strategy treatment, we intended to shift attention to the strategic variables, making risk less salient. However, the strategic choice variables come along with and may be non-separable from the profit. Thus, requiring subjects to report their strategy may have focused them on the profit, too: Because subjects in the strategy treatment reported their choices (and with it, their behavior) in an open form field, the risk of losing money may have been more salient than in the control group. Hence, we might have unintentionally nudged subjects in the strategy treatment to focus on the risky aspect of the business game.[14] If risk becomes more salient, the performance of more risk-averse types may be inhibited.[15] Consequently, i) their production function should be impaired and ii) the

[14]As Figure 2 depicts, the treatment effect is already visible before the first reporting took place: From period 1 on, the control group is almost first-order stochastic dominating the strategy and profit treatment. The mere awareness of needing to report may have been enough to increase the salience of risk.

[15]Cognitively, this could work through stress: Making the risk more salient is likely to induce stress for

treatment effect should be more pronounced among risk-averse participants. We find evidence for both these conjectures. Consequently, we deem this channel to be the most plausible.

5 Conclusion

Identifying and evaluating measures that effectively foster innovative behavior is highly relevant for leading decision-makers in organizations (Rudis, 2004). We experimentally investigate whether salience can serve as a tool to nudge exploratory behavior. First, we find that making risk more salient reduces exploration, as hypothesized. Second, contrary to our hypothesis, we observe the same negative effect on exploration for subjects for whom risk was arguably less salient. Therefore, the data reveal that our nudge backfires. We discuss potential reasons and find it most probable that our intervention, aimed at making the risk less salient, likely makes it more salient. Third, we find the effects of the salience nudge to be particularly pronounced among the more risk-averse participants. Subjects that are rather risk tolerant do not react to the intervention. This finding aligns with Ederer and Manso (2013). Thus, we demonstrate that purely behavioral interventions that only change the salience of risk, but do not alter monetary incentives, *are* able to alter explorative behavior.

We apply a state-of-the-art methodological approach by following a pre-registered sequential analysis plan (Lakens, 2014). Since the effects for both groups are lower than the smallest effect size of interest defined in the sequential analysis plan, we adhered to our pre-determined procedure and stopped the data collection after the interim analysis. Consequently, our sample size is not large enough to make reliable inferences given the low effect sizes we observe and our results should be interpreted with prudence.

Still, our results have important implications for practitioners and researchers alike. First, we show that purely behavioral nudges that do not change incentives *are* able to affect innovative behavior by guiding risk perceptions. Second, even nudges that are carefully derived from the existing literature can turn out to be ineffective and may even backfire. Hence, interventions need to be carefully tailored and require meticulous evaluation before implementation, ideally through experimental methods (Banerjee et al., 2017, Sunstein, 2017). This is especially important since nudges have become increasingly popular due to their simplicity and alleged effectiveness (Reisch and Sunstein, 2016, Benartzi et al., 2017, Sunstein et al., 2019). However, DellaVigna and Linos (2020) find that RCTs conducted in so-called nudge units show substantially lower effect sizes than RCTs in published academic studies. The authors show that publication bias can account for the full difference. Also, Camerer et al. (2016) find that replicated studies have a 33% lower effect size than the original studies. By adhering to

the risk-averse subjects. Stress can in turn impair cognitive performance, demonstrated by psychological and neurological research (Lupien et al., 2007, Matthews et al., 2000, Schoofs et al., 2008). Such a channel is well possible and we cannot rule it out. As a matter of fact, rather than a separate explanation, we consider it to be a plausible and potential cognitive mechanism for the channel discussed in this paragraph.

rigorous methods such as pre-registration and the publication of papers unconditional of the results, effect sizes will better reflect reality compared to the past, in which effect sizes and statistical significance are likely inflated (Olken, 2015).

By providing support for the essential role played by risk preferences in innovation, our study contributes to the literature on innovation and entrepreneurship. We are the first to demonstrate that purely behavioral nudges can affect exploratory behavior. This opens up interesting paths for future research since behavioral interventions that cost-effectively foster exploration could be promising approaches to help practitioners in making organizations more innovative. For example, manipulating the customer feedback rather than a reporting mechanism, might be an interesting candidate to test.

References

Amabile, T. M. (1996). *Creativity in context.* Westview Press, Boulder, CO.

Banerjee, A., Banerji, R., Berry, J., Duflo, E., Kannan, H., Mukerji, S., Shotland, M., and Walton, M. (2017). From proof of concept to scalable policies: Challenges and solutions, with an application. *Journal of Economic Perspectives*, 31(4):73–102.

Banerjee, R., Bhattacharya, J., and Majumdar, P. (2021). Exponential-growth prediction bias and compliance with safety measures related to COVID-19. *Social Science & Medicine*, 268:113473.

Benartzi, S., Beshears, J., Milkman, K. L., Sunstein, C. R., Thaler, R. H., Shankar, M., Tucker-Ray, W., Congdon, W. J., and Galing, S. (2017). Should governments invest more in nudging? *Psychological Science*, 28(8):1041–1055.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Camerer, C. F., Dreber, A., Forsell, E., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Almenberg, J., Altmejd, A., Chan, T., Heikensten, E., Holzmeister, F., Imai, T., Isaksson, S., Nave, G., Pfeiffer, T., Razen, M., and Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. 351(6280):1433–1436.

Charness, G. and Grieco, D. (2014). Creativity and financial incentives. *University of California, Santa Barbara, Working Paper.*

Charness, G. and Grieco, D. (2018). Creativity and Incentives. *Journal of the European Economic Association*, 17(2):454–496.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, 105(2):860–85.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–72.

DellaVigna, S. and Linos, E. (2020). Rcts to scale: Comprehensive evidence from two nudge units. Working Paper 27594, National Bureau of Economic Research.

Diacon, S. and Hasseldine, J. (2007). Framing effects and risk perception: The effect of prior performance presentation format on investment fund choice. *Journal of Economic Psychology*, 28(1):31 – 52.

Ederer, F. and Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science*, 59(7):1496–1513.

Englmaier, F., Grimm, S., Schindler, D., and Schudy, S. (2018). The effect of incentives in non-routine analytical teams tasks – evidence from a field experiment.

Eppler, M. J. and Aeschimann, M. (2009). A systematic framework for risk visualization in risk management and communication. *Risk Management*, 11(2):67–89.

Erat, S. and Gneezy, U. (2015). Incentives for creativity. *Experimental Economics*, 19:269–280.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692.

Ferrer, R. A. and Klein, W. M. (2015). Risk perceptions and health behavior. *Current Opinion in Psychology*, 5:85 – 89.

Gabaix, X. (2019). Chapter 4 - behavioral inattention. *Handbook of Behavioral Economics: Applications and Foundations 2.*, 2:261 – 343.

Gerrard, M., Gibbons, F. X., and Reis-Bergan, M. (1999). The effect of risk communication on risk perceptions: the significance of individual differences. *JNCI Monographs*, 1999(25):94 – 100.

Hall, R. E. and Woodward, S. E. (2010). The burden of the nondiversifiable risk of entrepreneurship. *American Economic Review*, 100(3):1163–94.

Hanna, R., Mullainathan, S., and Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies*, 66(1):169–182.

Holzmeister, F., Huber, J., Kirchler, M., Lindner, F., Weitzel, U., and Zeisberger, S. (2020). What drives risk perception? a global survey with financial professionals and laypeople. *Management Science*, 66(9):3977–4002.

Hudja, S. and Woods, D. (2019). Behavioral bandits: Analyzing the exploration versus exploitation trade-off in the lab.

Hughes, D. J., Lee, A., Tian, A. W., Newman, A., and Legood, A. (2018). Leadership, creativity, and innovation: A critical review and practical recommendations. *The Leadership Quarterly*, 29(5):549–569.

Hvide, H. K. and Panos, G. A. (2014). Risk tolerance and entrepreneurship. *Journal of Financial Economics*, 111(1):200–223.

Jennison, C. and Turnbull, B. W. (2001). Group sequential tests with outcome-dependent treatment assignment. *Sequential Analysis*, 20(4):209–234.

Kerr, S. P., Kerr, W. R., and Xu, T. (2017). Personality traits of entrepreneurs: A review of recent literature. Working Paper 24097, National Bureau of Economic Research.

Knight, J., Weir, S., and Woldehanna, T. (2003). The role of education in facilitating risk-taking and innovation in agriculture. *The Journal of Development Studies*, 39(6):1–22.

Kohn, A. (1993). *Punished by Rewards: The Trouble with Gold Stars, Incentive Plans, A's, Praise and Other Bribes.* Houghton Mifflin, Boston.

Köszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1):53–104.

Koudstaal, M., Sloof, R., and van Praag, M. (2016). Risk, uncertainty, and entrepreneurship: Evidence from a lab-in-the-field experiment. *Management Science*, 62(10):2897–2915.

Kurz, V. (2018). Nudging to reduce meat consumption: Immediate and persistent effects of an intervention at a university restaurant. *Journal of Environmental Economics and Management*, 90:317–341.

Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses. *European Journal of Social Psychology*, 44(7):701–710.

Landsberger, H. A. (1958). Hawthorne revisited, cornell studies in industrial and labor relations. vol. ix.

Lazear, E. (2005). Entrepreneurship. *Journal of Labor Economics*, 23(4):649–680.

Levitt, S. D. and List, J. A. (2011). Was there really a hawthorne effect at the hawthorne plant? an analysis of the original illumination experiments. *American Economic Journal: Applied Economics*, 3(1):224–238.

Loewenstein, G., Weber, E., Hsee, C. K., and Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127(2):267–286.

Lupien, S. J., Maheu, F. S., Tu, M. T., Fiocco, A. J., and Schramek, T. E. (2007). The effects of stress and stress hormones on human cognition: Implications for the field of brain and cognition. *Brain and Cognition*, 65(3):209–237.

Lévy-Garboua, L., Maafi, H., Masclet, D., and Terracol, A. (2012). Risk aversion and framing effects. *Experimental Economics*, 15:128–144.

Manso, G. (2011). Motivating innovation. *The Journal of Finance*, 66(5):1823–1860.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87.

Matthews, G., Davies, D., Westerman, S., and & Stammers, R. (2000). *Human performance: Cognition, stress and individual differences*. East Sussex, UK: Psychology Press.

Myers, L. B. (2014). Changing smokers' risk perceptions – for better or worse? *Journal of Health Psychology*, 19(3):325–332.

Olken, B. A. (2015). Promises and perils of pre-analysis plans. *Journal of Economic Perspectives*, 29(3):61–80.

Pennings, J. M. E. and Wansink, B. (2004). Channel contract behavior: The role of risk attitudes, risk perceptions, and channel members' market structures. *The Journal of Business*, 77(4):697–724.

Prendergast, C. and Stole, L. (1996). Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning. *Journal of Political Economy*, 104(6):1105–1134.

Reisch, L. A. and Sunstein, C. R. (2016). Do europeans like nudges? *Judgment and Decision Making*, 11(4):310–325.

Rudis, E. (2004). *CEO Challenge 2004: Perspectives and Analysis*. The Conference Board, New York.

Scharfstein, D. S. and Stein, J. C. (1990). Herd behavior and investment. *The American Economic Review*, 80(3):465–479.

Schoofs, D., Preuss, D., and Wolf, O. T. (2008). Psychosocial stress induces working memory impairments in an n-back paradigm. *Psychoneuroendocrinology*, 33:643–653.

Shalley, C. E. and Gilson, L. L. (2004). What leaders need to know: A review of social and contextual factors that can foster or hinder creativity. *The Leadership Quarterly*, 15:33–53.

Shalley, C. E., Gilson, L. L., and Blum, T. C. (2000). Matching creativity requirements and the work environment: Effects on satisfaction and intentions to leave. *The Academy of Management Journal*, 43(2):215–223.

Sitkin, S. B. and Weingart, L. R. (1995). Determinants of risky decision-making behavior: A test of the mediating role of risk perceptions and propensity. *Academy of Management Journal*, 38:1573–1592.

Sobkow, A., Traczyk, J., and Zaleskiewicz, T. (2016). The affective bases of risk perception: Negative feelings and stress mediate the relationship between mental imagery and risk perception. *Frontiers in Psychology*, page 932.

Sunstein, C. R. (2017). Nudges that fail. *Behavioural Public Policy*, 1(1):4–25.

Sunstein, C. R., Reisch, L. A., and Kaiser, M. (2019). Trusting nudges? lessons from an international survey. *Journal of European Public Policy*, 26(10):1417–1443.

Suurmond, G., Swank, O. H., and Visser, B. (2004). On the bad reputation of reputational concerns. *Journal of Public Economics*, 88(12):2817 – 2838.

Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.

Tombu, M. and Mandel, D. R. (2015). When does framing influence preferences, risk perceptions, and risk attitudes? the explicated valence account. *Journal of Behavioral Decision Making*, 28(5):464–476.

Unsworth, K. (2001). Unpacking creativity. *The Academy of Management Review*, 26(2):289–297.

Weber, E. U. and Milliman, R. A. (1997). Perceived risk attitudes: Relating risk perception to risky choice. *Management Science*, 43(2):123–144.

A Results

A.1 Performance and Explorative Behavior

We investigate different proxies for explorative behavior, in line with Ederer and Manso (2013). Table 2 provides an overview of all measured outcome variables.

Table 2: Overview of proxies for explorative behavior

Variable	Description
final_profit	Profit in final round, continuous variable (min:0,max: 199.1).
max_profit	Highest profit in all rounds, continuous variable (min:0,max: 199.1).
overall_profit	Sum of total profit of all 20 periods, continuous variable (min:0,max: 3982).
final_location	Final location chosen, categorical variable (School, Business, Stadium).
location_non-default	Constructed variable. Count of chosen non-default locations, i.e. non-Business locations. Discrete variable (min:0, max: 20)
max_exploration_phase	Constructed variable. Longest duration of an exploration phase. An exploration phase starts when subjects choose a location other than the default location suggested by the previous manager. An explorative phase is defined as ending when a subject switches back to the default location or when a subject does not change location and lemonade color and also does not change lemon content, sugar content and price by more than 0.25 units. Discrete variable (min:0, max: 20). Adapted 1:1 from EM.
duration_exploration_phase	Constructed variable. Total duration of all exploration phases. Discrete variable (min:0, max: 20)
std_dev_sugar	Constructed variable. Standard deviation for sugar choices over all rounds. Continuous variable.
std_dev_lemon	Constructed variable. Standard deviation for lemon choices over all rounds. Continuous variable.
std_dev_price	Constructed variable. Standard deviation for price choices over all rounds. Continuous variable.
average_std_dev	Constructed variable. The average subject-specific standard deviation of strategy choices for the three continuous variables sugar, price, lemon. Continuous variable.

Some of these variables reflect choices or are constructed based on choices by the subjects and proxy their explorative behavior: the variables based on the location choice indicate whether the subject detected the profit-maximizing location. This is impossible when exclusively following the customer feedback. Further, the higher the standard deviation for the continuous variables (sugar, lemon, price) is, the more explorative the subject behaved. Instead, the outcome variables with respect to profits mirror how these choices are translated into payoffs. Since the business game is designed such that explorative behavior increases the chance of finding a profit-increasing strategy, those outcome variables should be closely correlated. This can be seen by Table 3.

Table 3: Spearman Cross-correlation table of exploration outcome measures

Variables	finalprofit	maxprofit	overallprofit	final_loc_binary	loc_non-def	expl_phase_maxdur	expl_phase_totdur	sd_sugar_1-10	sd_sugar_11-20	sd_lemon_1-10	sd_lemon_11-20	sd_price_1-10	sd_price_11-20	sd_choices_1-10	sd_choices_11-20
finalprofit	1.000														
maxprofit	0.9759	1.000													
	0.000														
overallprofit	0.9057	0.8849	1.000												
	0.000	0.000													
final_loc_binary	0.7661	0.7385	0.7379	1.000											
	0.000	0.000	0.000												
loc_non-def	0.7414	0.7398	0.7976	0.7417	1.000										
	0.000	0.000	0.000	0.000											
expl_phase_maxdur	0.5948	0.6265	0.561	0.53	0.6373	1.000									
	0.000	0.000	0.000	0.000	0.000										
expl_phase_totdur	0.7105	0.7365	0.6668	0.6621	0.78	0.8447	1.000								
	0.000	0.000	0.000	0.000	0.000	0.000									
sd_sugar_1-10	0.3501	0.354	0.3469	0.2708	0.2526	0.42	0.3211	1.000							
	0.0007	0.0006	0.0008	0.0098	0.0163	0.000	0.002								
sd_sugar_11-20	0.1565	0.1908	0.0678	0.1731	0.0898	0.2411	0.2858	0.0744	1.000						
	0.1407	0.0716	0.5255	0.1027	0.4001	0.022	0.0063	0.486							
sd_lemon_1-10	0.3827	0.382	0.3232	0.282	0.1787	0.2807	0.1973	0.4654	0.1184	1.000					
	0.0002	0.0002	0.0019	0.0071	0.0919	0.0074	0.0623	0.000	0.2664						
sd_lemon_11-20	0.3859	0.4099	0.2915	0.3927	0.2923	0.3651	0.4654	0.2855	0.4309	0.2593	1.000				
	0.0002	0.0001	0.0053	0.0001	0.0052	0.0004	0.000	0.0064	0.000	0.0136					
sd_price_1-10	0.3692	0.404	0.2295	0.3056	0.2304	0.3379	0.3536	0.1681	0.2825	0.2455	0.3087	1.000			
	0.0003	0.0001	0.0296	0.0034	0.0289	0.0011	0.0006	0.1133	0.007	0.0197	0.0031				
sd_price_11-20	0.3748	0.4211	0.1951	0.2727	0.1418	0.286	0.3913	0.1525	0.3483	0.2084	0.3267	0.4113	1.000		
	0.0003	0.000	0.0653	0.0093	0.1825	0.0063	0.0001	0.1514	0.0008	0.0487	0.0017	0.0001			
sd_choices_1-10	0.3228	0.3509	0.334	0.1311	0.1557	0.2363	0.2124	0.4962	0.1365	0.4954	0.2001	0.2161	0.1471	1.000	
	0.0019	0.0007	0.0013	0.2183	0.1429	0.0249	0.0444	0.000	0.1997	0.000	0.0586	0.0408	0.1665		
sd_choices_11-20	0.6152	0.6263	0.6248	0.2486	0.392	0.3907	0.442	0.4483	0.0882	0.4197	0.3367	0.1404	0.1712	0.7068	1.000
	0.000	0.000	0.000	0.0182	0.0001	0.0001	0.000	0.000	0.4083	0.000	0.0012	0.187	0.1067	0.000	

Results with respect to the final profit and the maximum profit were already discussed in Section 3.1. In line with these discussed results, the mean overall profit in the control group is higher than in both reporting treatments, as shown by Figure 6.

Figure 6: Mean of overall profit by treatments.

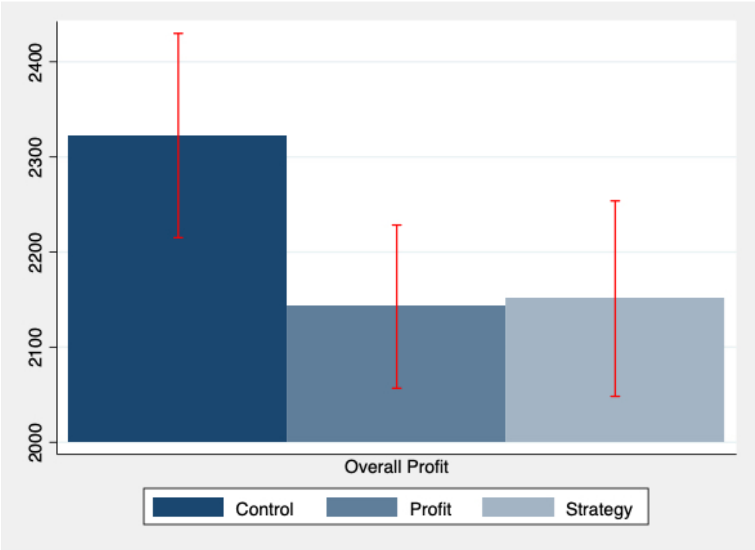

Analyzing the location chosen in the final round of the game shows the same pattern. There are three possible locations, with the school being the optimal one. As Figure 7a depicts, around 80% of subjects in the control group chose the optimal location in the final round. This proportion is higher than in the profit (appr. 73%) and strategy treatment (appr. 66%).[16]

Figure 7: Location measures

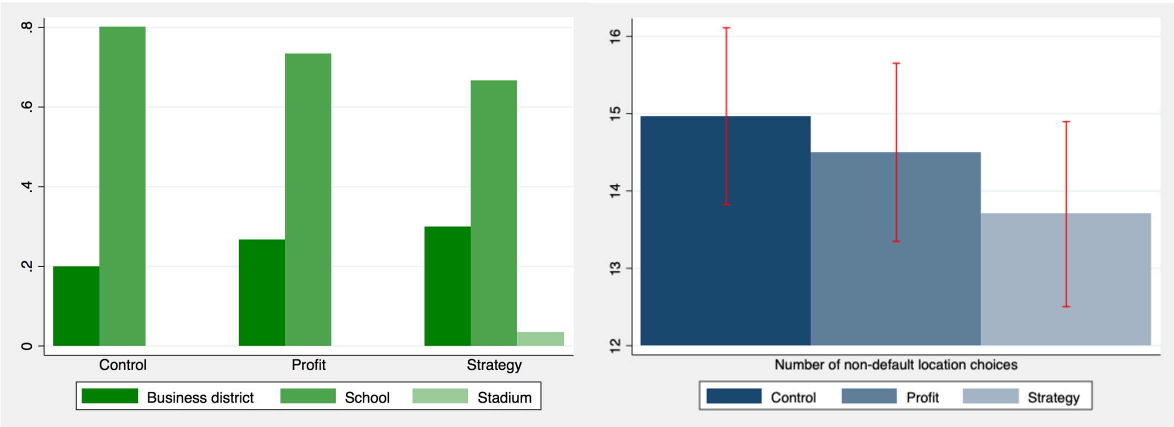

(a) Proportion of subjects by Location in final round.

(b) Number of times the non-default location was chosen.

[16]These difference-in-means are not statistically significant (two-sided Mann-Whitney U-tests: p=0.5449 for control group vs. profit treatment, p=0.2469 for control group vs. strategy treatment).

Results with respect to the exploration phases further support our finding: the control group explores more than the profit and the strategy treatment. However, these differences are, again, not significant.

Figure 8: Exploration phase measures

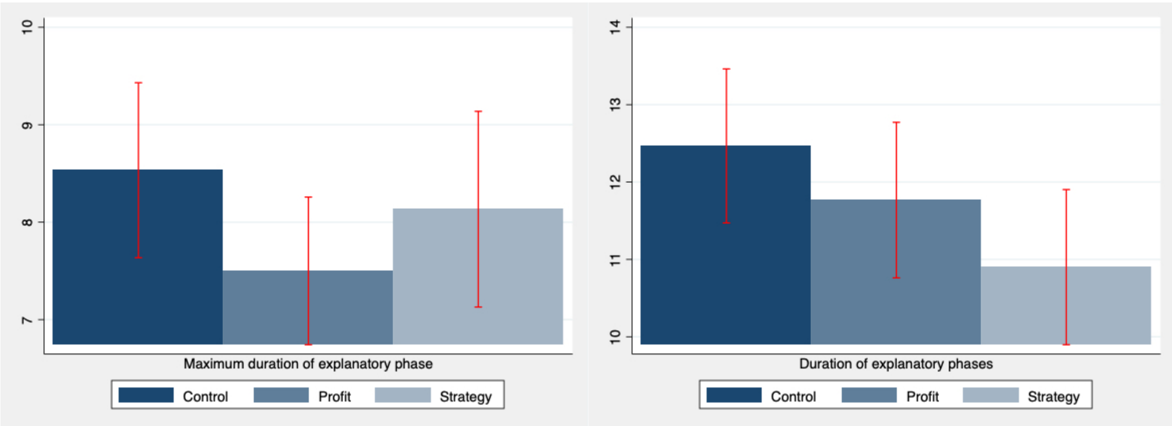

(a) Maximum of all exploration phases

(b) Duration of all exploration phases

Notes: The figure reports the means of the maximum length and of the duration of all exploration phases . Error bars indicate standard errors of the mean.

Lastly, Figure 9 shows that also the results for all outcome variables with respect to the standard deviations are in line with the previous results: the standard deviation in the control group is slightly higher than in the reporting treatments. As expected, the choices within the first ten rounds of the business game vary more than the choices in the last ten rounds. This early exploration is intuitive since the individual can profit from her findings for a longer time horizon than later stage explorative activities.

Figure 9: Standard deviation measures

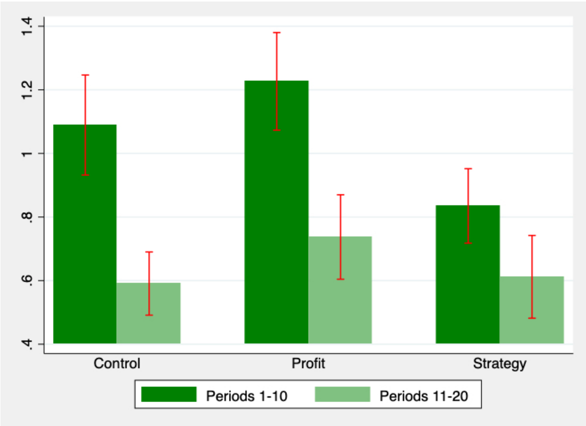

(a) Standard deviation of sugar choices over all rounds.

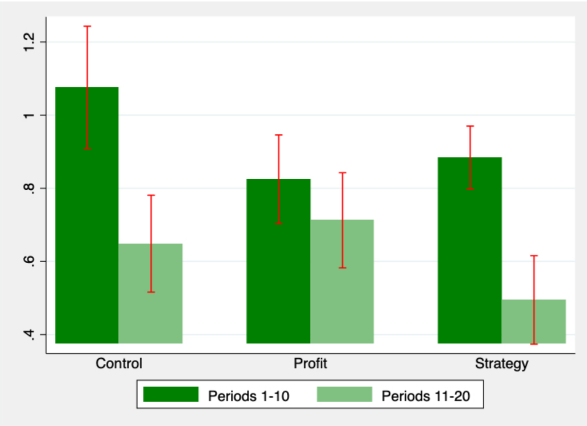

(b) Standard deviation of lemon choices over all rounds.

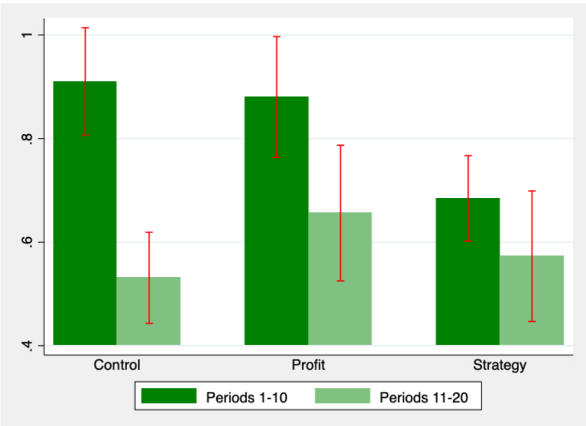

(c) Standard deviation of price choices over all rounds.

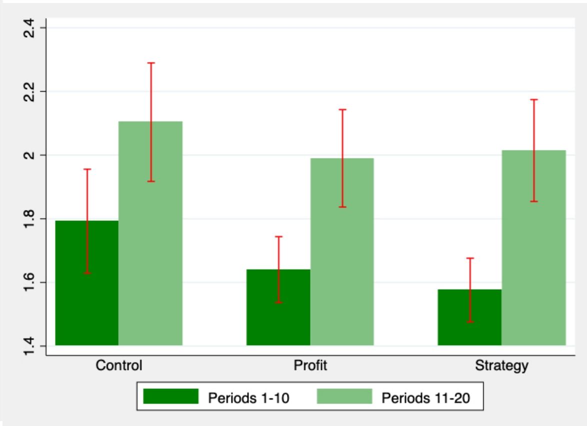

(d) Standard deviation of all choices over all rounds.

To further analyze the explorative behavior of our subjects, we compare the average subject-specific standard deviation of the profits. The variability of profits in all treatments is higher in periods 1-10 (see Figure 10). However, the standard deviation of the profits is not different between the treatments.[17]

[17]Two-sided Mann-Whitney U-tests: p=0.8941 for periods 1-10 and p=1.000 for periods 11-20 between control group and profit treatment. p=0.5444 for periods 1-10 and p=0.1833 for periods 11-20 between control group and strategy treatment and p=0.4598 for periods 1-10 and p=0.3142 for periods 11-20 between profit and strategy treatment.

Figure 10: Standard deviation of realized profits.

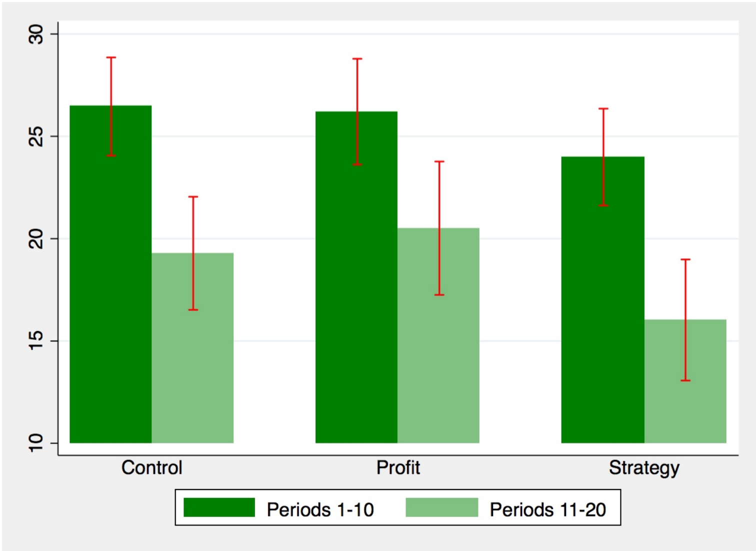

A.2 Attention

We analyze different proxies for exerted effort, again, in line with Ederer and Manso (2013). Table 4 provides an overview of all measured proxies. The respective results are discussed in Section 3.2.

Table 4: Overview of proxies for effort

Variable	Description
total_time_decision	Time elapsed for all 20 Decision-Screens, continuous variable.
focus_time_decision	Focus time elapsed for all 20 Decision-Screens, continuous variable.
total_time_result	Time elapsed for all 20 Result-Screens, continuous variable.
focus_time_result	Focus time elapsed for all 20 Result-Screens, continuous variable.
total_time_reporting	Time elapsed for all 20 Reporting-Screens. Treatment groups only, continuous variable.
focus_time_reporting	Focus time elapsed for all 20 Reporting-Screens. Treatment groups only, continuous variable.
total_notes	Proportion of filled out fields in notes sheet, continuous variable.
notes_strategic	Proportion of notes with respect to strategic variables relative to total_notes, continuous variable.
notes_profit	Proportion of notes with respect to profits relative to total_notes, continuous variable.
notes_feedback	Proportion of notes with respect to customer feedback relative to total_notes, continuous variable.

Figure 11: Effort measured by time elapsed.

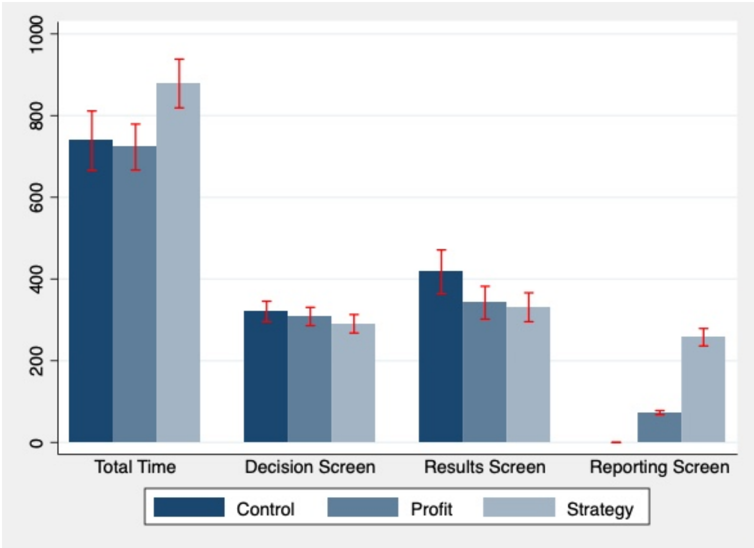

Notes: The figure reports the means of the subjects time elapsed for the total time, the time spent at the decision screen, at the results & feedback screen, and at the reporting screen, respectively.

Table 5: Regressions

	Dependent variables:					
	Overall Time	*Decision Time*	*Results Time*	*Strategy Notes*	*Profit Notes*	*Feedback Notes*
	(1)	(2)	(3)	(4)	(5)	(6)
Profit Treatment	-17.982	-12.933	-77.952	-0.147	0.068	-0.103
	(90.19)	(33.36)	(62.17)	(0.11)	(0.09)	(0.11)
Strategy Treatment	144.199	-27.883	-80.104	-0.167	-0.123	-0.083
	(92.52)	(34.22)	(63.78)	(0.11)	(0.09)	(0.12)
Gender	33.602	-4.345	21.356	0.167*	0.204***	0.104
	(75.42)	(27.90)	(51.99)	(0.09)	(0.08)	(0.09)
Riskaversion	33.238	9.500	39.170	0.092	0.094	0.027
	(77.35)	(28.61)	(53.32)	(0.09)	(0.08)	(0.10)
Constant	704.299***	317.549***	385.747***	0.486***	0.642***	0.471***
	(81.23)	(30.05)	(55.99)	(0.10)	(0.08)	(0.10)
N	90	90	90	90	90	90
R^2	0.046	0.011	0.036	0.092	0.152	0.027
Adjusted R^2	0.001	-0.035	-0.009	0.049	0.112	-0.019

Level of significance: *p<0.1; **p<0.05; ***p<0.01

A.3 Discussion

If limited effort were the valid explanation for our findings, the production function of exploration should still be identical for all subjects. That is, for a given amount of effort, the exploration level should be the same across treatment groups. The reason is that according to the limited effort explanation, participants exert less effort for productive tasks because they need to exert more effort for the (arguably) unproductive reporting.

However, we find that participants in the Strategy treatment explore less conditional on the level of productive effort exerted (measured by time spent in front of the decision and results screen). In other words, the production function of the risk-averse type is impaired. This is best visible in Figure 13 that uses time spent on the reporting and the decision screens as a proxy for effort and the number of times a participant enters into an exploration phase as the outcome variable representing exploration behavior. The figure shows that the production function for the risk-averse type is indeed altered: Conditional on the effort level, treated participants initiate less exploration phases. Further, we observe that the production function of the less risk-averse type is not much different across treatments. Similar patterns are observed when measuring exploration through the final location (Figure 13) and final profit (Figure 14).

Figure 12: Production functions (Exploration regressed on effort invested)

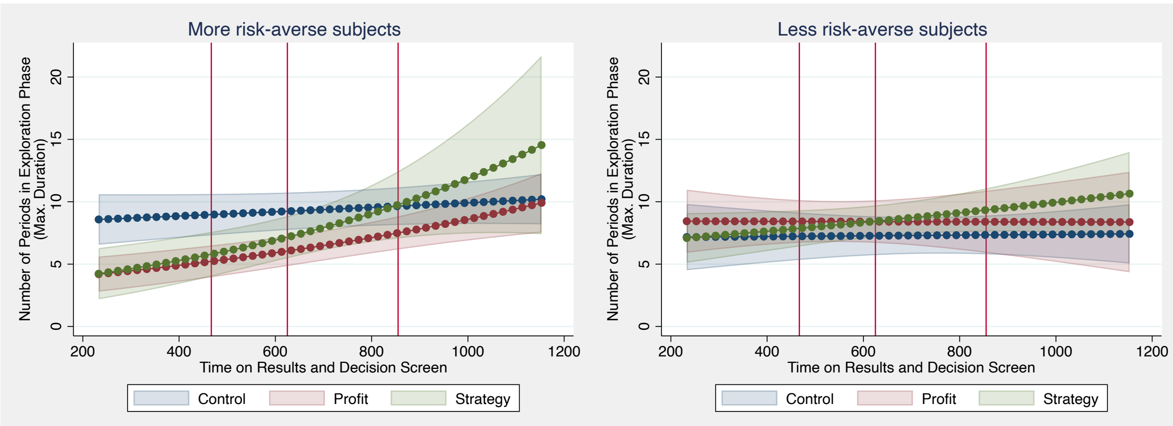

Notes: This figure displays the production function by subject types, split by median risk-aversion into less and more risk-averse. The figure plots the following regression specification: $Explorationphase = Treatment + Timespent + Treatment \times Timespent$. Quartiles of the distribution of $Timespent$ are displayed in red.

Figure 13: Production functions (Final location school regressed on effort invested)

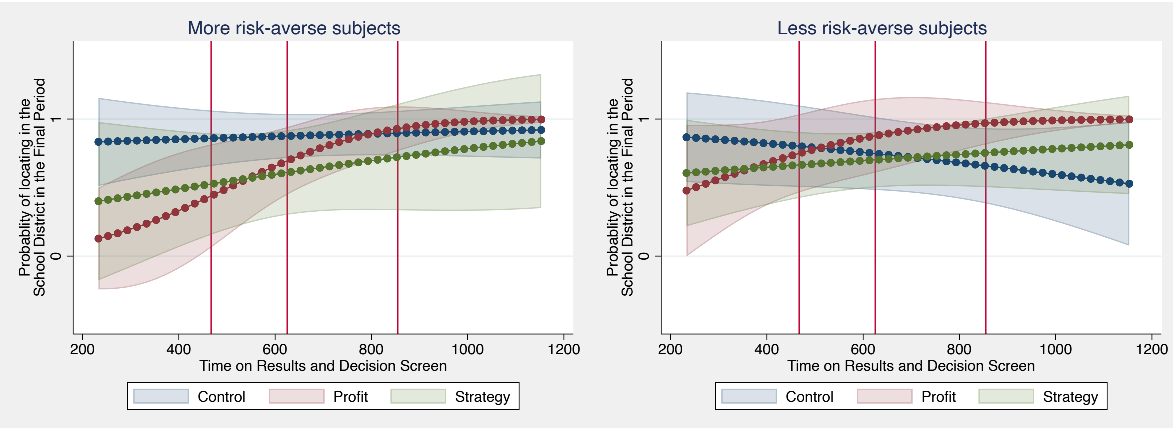

Notes: This figure displays the production function by subject types, split by median risk-aversion into less and more risk-averse. The figure plots the following regression specification: $Finallocation = Treatment + Timespent + Treatment \times Timespent$. Quartiles of the distribution of Time on Results and Decision Screen are displayed in red.

Figure 14: Production functions (Final profit regressed on effort invested)

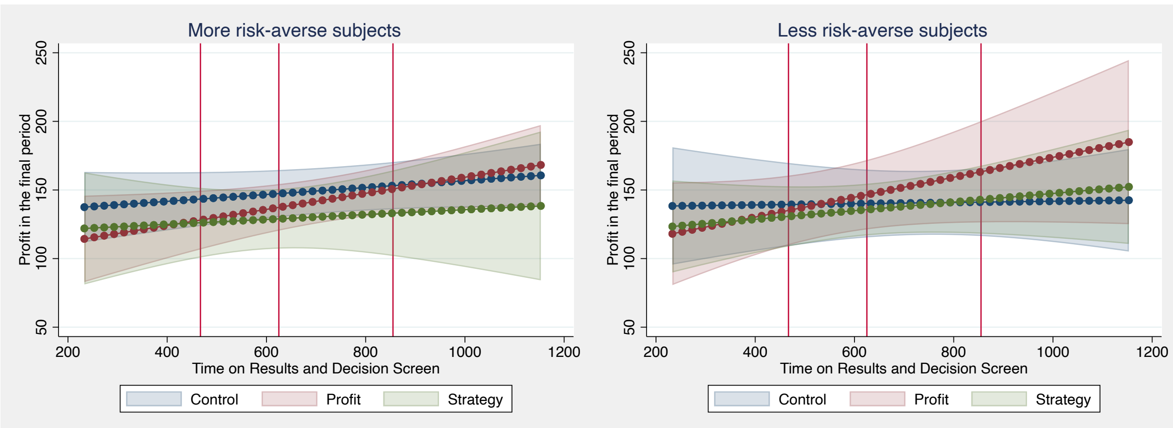

Notes: This figure displays the production function by subject types, split by median risk-aversion into less and more risk-averse. The figure plots the following regression specification: $Finalprofit = Treatment + Timespent + Treatment \times Timespent$. Quartiles of the distribution of Time on Results and Decision Screen are displayed in red.

B Experimental Design

B.1 Instructions

Welcome

You are now taking part in a scientific study. Please read the following instructions carefully. Everything that you need to know in order to participate in this experiment is explained below. Should you have any difficulties in understanding these instructions, please notify us. We will answer your questions at your cubicle. During the course of the experiment you can earn money. The amount that you earn during the experiment depends on your decisions. All the gains that you make during the course of the experiment will be exchanged into cash at the end of the experiment.

The exchange rate will be: 100 thaler = 1 EUR

The experiment is divided into 20 periods. In each period you have to make decisions, which you will enter on a computer screen. The decisions you make and the amount of money you earn will not be made known to the other participants - only you will know them. At the end of the experiment, you will be requested to respond to survey questions. Please note that communication between participants is strictly prohibited during the study. Communication between participants and unnecessary interference with computers will lead to the exclusion from the study. In case you have any questions don't hesitate to ask us.

Procedures

In this experiment, you will take on the role of an individual running a lemonade stand. There will be 20 periods in which you will have to make decisions on how to run the business. These decisions will involve the location of the stand, the sugar and lemon content and the lemonade color and price. The decisions you make in one period, will be the default choices for the next period. At the end of each period, you will learn what profits you made during that period. You will also hear some customer reactions that may help you with your choices in the following periods.

Letter from the Previous Manager

The previous manager of the lemonade stand has left you guidelines on how to run the business. The letter from the previous manager is the following:

Dear X,
I have enclosed the following guidelines that you may find helpful in running your lemonade stand. These guidelines are based on my previous experience running this stand. When running my business, I followed these basic guidelines:

Location: Business District
Sugar Content: 5.2%
Lemon Content: 7.0%

Lemonade Color: Green
Price: 8.2 thaler

With these choices, I was able to make an average profit of about 85 thaler per period. I have experimented with alternative choices of sugar and lemon content, as well as lemonade color and price. The above choices were the ones I found to be the best. I have not experimented with alternative choices of location though. They may require very different strategies.

Regards, Previous Manager

Note that in the first period, these choices will appear as defaults.

Compensation

Your compensation will be based on the profits you make with your lemonade stand. You will get paid 50% of your total lemonade stand profits during the 20 periods of the experiment.

For example, if your total profits during the 20 periods of the experiment were 1700 thaler, you will earn 850 thaler, worth 8.50 EUR.

In addition, you will earn the show-up fee of 2 EUR.

Report

Displayed only if assigned to *profit treatment*: Please report your profits of the last three periods.

Displayed only if assigned to *strategy treatment*: Please describe your strategy in the last three periods. Why did you choose this strategy?

B.2 Parametrization

We adopted the experimental parameters from Ederer and Manso (2013). The participants could make following choices:

- Location = {Business District, School, Stadium}
- Sugar Content = {0,0.1,...,20}
- Lemon Content = {0,0.1,...,20}
- Lemonade Color = {Green, Pink}
- Price = {0,0.1,...,10}

The optimal product mix in each location is shown in table 6.

For the profit calculation in each location, a linear penalty function was implemented. So if the participant did not implement the optimal choices, she was penalized according to the values summarized in table 7. In each location, the penalty factors represented in the table are associated with a deviation of one unit for each of the variables. Note that we implemented a minimum of 0, i.e. participants could not earn negative profit.

Table 6: Optimal Product Mix, by Location

	Business District	School	Stadium
Sugar	1.5%	9.5%	5.5%
Lemon	7.5%	1.5%	5.5%
Lemonade Color	Green	Pink	Green
Price	7.5	2.5	7.5
Maximum Profit	100	200	60

Table 7: Penalty factors, by Location

	Business District	School	Stadium
Sugar	3	6	0.5
Lemon	3	6	0.5
Lemonade Color	20	60	0.5
Price	3	6	0.5

C Elicitation of Control Variables

C.1 Risk preference

For the elicitation of risk preferences, we relied on the staircase elicitation method developed by Falk et al. (2018). We adapted their method without any changes. In the following, the wording for the first decision is summarized, with the subsequent question following the same wording:

Please imagine the following situation: You can choose between a sure payment of a particular amount of money, or a draw, where you would have an equal chance of getting 300 EUR or getting nothing. We will present to you five different situations. The draw is the same in all situations. The sure payment is different in every situation.

What would you prefer: a draw with a 50 percent chance of receiving 300 USD, and the same 50 percent chance of receiving nothing, or the amount of 160 USD as a sure payment?

C.2 Demographics

What is your gender?

What is your age?

What is the highest level of education you have completed or the highest degree you have received?

What is your academic field?

What is your nationality? (If more than one apply, select the one you feel is most representative for you).

In what country did you grow up? (If you grew up in more than one country, please indicate the country you lived the longest while growing up).

In what country do you currently reside?

Do you identify yourself with any of the following religions? ['Christianity', 'Judaism', 'Islam', 'Buddhism', 'Hinduism', 'Other', 'No religion / Atheism']

D Sequential Analysis

Our sequential analysis plan followed the summarized procedure outlined below. It concerned the key outcome variable, the profit in the final round. The assessment on how to proceed after the first stage was based on a hypothesis test conducted with a two sample t-test.

Based on an expected effect size of Cohen's $d = 0.5$, a power analysis indicated that for a two-sided test with an alpha of 0.05, a desired statistical power of 0.8, and two looks using a linear spending function, a total of 180 participants is needed (60 per group). If the expected difference is significant at the first interim analysis (after 90 participants or time = 0.5, with an alpha boundary of 0.025) the data collection will be terminated. The data collection will also be terminated if the observed effect size is smaller than the smallest effect size of interest, which is set at $d = 0.3875$ based on the researcher's willingness to collect at most 300 participants for this study, and the fact that with one interim analysis, 300 participants provide 0.8 power to detect an effect of $d = 0.3875$. If the interim analysis reveals an effect size larger than 0.5, but while $p > 0.025$, the data collection will be continued until 60 participants per group have been collected. If the effect size lies between the smallest effect size of interest ($d = 0.3875$) and the expected effect size ($d = 0.5$), the planned sample size will be increased based on a conditional power analysis to achieve a power of 0.9 (or to a maximum of 100 participants per group, or 300 participants in total). The second analysis is performed at an alpha boundary of 0.0358.

In the original paper of Ederer and Manso (2013), subjects in the pay-for-performance contract yield on average a profit in the last period of 111 thalers. Because our control group is mimics one-to-one the pay-for-performance group in Ederer and Manso (2013), we assume that our control group will yield 111 thalers in the last period too. Furthermore, based on a small test conducted in September 2019, we expect a standard deviation of $SD = 40$. It is our best assumption that the standard deviation is equal for all three groups.

Next, we define the smallest detectable effect size of interest. Based on practical limitations, namely budget restrictions, we are willing to collect at most 300 observations in total. To identify an effect with 80% power when comparing the control with one of the two treatment groups, with one interim look at the data, the effect must be at least 15.5 thalers large, translating in a Cohen's $d = 0.3875$. We deem such an effect also from a real-life perspective as appropriate - implementing a new reporting policy comes along with costs, and thus, the beneficial or detrimental effect should be large enough to be of practical relevance.

Next, we elaborate on the expected effect. In Ederer and Manso (2013), the exploration contract yields in the final period on average a profit of 140. This leads to a profit difference between pay-for-performance and exploration contract of $29 (= 140 - 111)$. With a $SD = 40$, this yields a Cohen's $d = 0.725$. Most probably, our effects will be lower since our treatment interventions are based on a behavioural mechanism but do not change monetary incentives, as this is the case in Ederer and Manso (2013). As a best-estimate, we expect our effect to be 30% (or 9 thalers) lower than in their study. Therefore, we estimate that subjects in the strategy treatment will yield a lower profit in the last period than their exploration contract, and we estimate this to be at 131. This yields an effect size of $20 (= 131 - 111)$, or a Cohen's $d = 0.50$. For the profit treatment, we cannot base our estimates on a previous study due the lack of comparable alternatives. However, we expect the profit treatment to perform similarly as the strategy treatment (but in the opposite direction, of course). Consequently, we adopt the same Cohen's $d = 0.5$ for the profit treatment. Based on the expected effect size of $d = 0.5$, with power 0.8 and an alpha of 0.05, we obtain a sample size $n = 60$ per group or $n = 180$ in total (after both looks). We will have a first look at time $= 0.5$, that is when 90 subjects are collected.

For controlling type 1 error rates, we use a linear spending function (power family function), as outlined in Jennison and Turnbull (2001). The alpha of 0.05 for a single look is adjusted for sequential analyses, namely for two looks using a linear spending function, yielding a nominal alpha of 0.0586. Thus, planning on analyzing the data at two different stages of the experiment, that is with one interim analysis, we formulate the analysis plan outlined in the main body of the text.