CESifo WORKING PAPERS

9122
2021

May 2021

Time Pressure and Regret in Sequential Search

Felix Klimm, Martin G. Kocher, Timm Opitz, Simeon Schudy

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 9122*

Time Pressure and Regret in Sequential Search

Abstract

Perceived urgency and regret are common in many sequential search processes; for example, sellers often pressure buyers in search of the best offer, both time-wise and in terms of potential regret of forgoing unique purchasing opportunities. Theoretically, these strategies result in anticipated and experienced regret, which systematically affect search behavior and thereby distort optimal search. In addition, urgency may alter decision-making processes and thereby the salience of regret. To understand the empirical relevance of these aspects, we study the causal effects of regret, urgency, and their interaction on search behavior in a pre-registered, theory-based, and well-powered experiment. We find that urgency reduces decision times and perceived decision quality but does not alter search length. Only very inexperienced decision-makers buy earlier when pressured. Anticipated regret does not affect search length (neither with nor without time pressure), while experienced regret leads to systematic adjustments in search length. Thus, we recommend that consumer protection policies should particularly focus on markets with inexperienced first-time buyers.

JEL-Codes: C910, D010, D030, D180, D830.

Keywords: sequential search, time pressure, regret, anticipated regret, experienced regret.

Felix Klimm
Department of Economics
University of Munich (LMU) / Germany
felixklimm@gmx.de

Martin G. Kocher
University of Vienna / Austria
martin.kocher@univie.ac.at

Timm Opitz
Max Planck Institute for Innovation and Competition, Munich / Germany
Timm.Opitz@ip.mpg.de

Simeon Schudy
Department of Economics
University of Munich (LMU) / Germany
simeon.schudy@econ.lmu.de

We thank Marvin Deversi, Paul Heidhues, Fabian Herweg, Marieke Huysentruyt, Nicolas Jacquemet, Heiko Karle, Matthias Lang, Robin Mamrak, Christoph Schwaiger, Olga Shurchkov, and seminar audiences at the LMU Munich, the Max Planck Institute for Innovation and Competition Munich, the Competition and Innovation Summer School in Ulcinj, the European ESA Meeting in Dijon and the CRC TRR 190 Workshop 2020 for their helpful thoughts and comments. Financial support by the German Research Foundation (DFG) through CRC TRR 190 (project number 280092119) is gratefully acknowledged.

1 Introduction

Perceived urgency and regret are common in many markets. For instance, in many goods and service markets, sellers pressure buyers searching for the best price with time-limited offers and emphasize potential regret about forgone purchasing opportunities (Sugden, Wang, & Zizzo, 2019). In labor markets, job seekers face deadlines and anticipate (or experience) regret when they reject or accept offers. In financial markets, investors facing rapid price changes may regret forgone selling opportunities when holding onto badly performing assets (Strack & Viefers, 2021).[1] It is thus important to understand to what extent perceived urgency and regret may affect individual choice in dynamic market environments, and whether their combination aggravates or alleviates potential biases in decision making.

Our study investigates the effects of perceived urgency and regret in a pre-registered, theory-based laboratory experiment.[2] Many of the above-mentioned examples for the relevance of urgency and potential regret reflect a search process that can be represented by an optimal stopping problem. In optimal stopping problems, a decision-maker observes a sequence of realizations of some stochastic process and, after observing a realization, decides on whether or not to take an action. For example, buyers may learn about price offers for a flight and then decide on whether to continue searching for a better realization (e.g., by looking at other platforms or waiting another day) or they may stop searching and immediately buy the item for the best available price.[3]

By trading off the best current price with potentially better future prices at higher search costs, decision-makers may experience regret of two types. First, if it turns out that decision-makers could have saved unnecessary search costs, they may regret not having stopped searching earlier (which is often referred to as *inaction regret*). Second, when deciding on whether or not to accept the currently best available price, decision-makers may anticipate that better price realizations can become available after purchase, and thus may anticipate regret from not having searched for longer (i.e., if they observe price realizations after purchase, which is often referred to as anticipated *action regret*).

While an expected utility maximizer is assumed to calculate the optimal search length given her knowledge about the underlying stochastic process and given search costs, perceived urgency may render full optimization unlikely. Time-pressured individuals may rely more on intuitive rather than

[1] In addition, urgency and regret are prevalent in auctions. For instance, in first-price auctions, bidders may anticipate or experience regret when paying too much (relative to the second-highest bid) when winning, or when bidding too little and thus missing an opportunity to win the auction at a favorable price (Engelbrecht-Wiggans & Katok, 2008).

[2] Pre-registration at: AEA RCT Registry; AEARCTR-0004065.

[3] The best available price relates either to the current price offer (optimal stopping with no recall) or the best price among the current and past price offers that the buyer has observed (optimal stopping with recall).

deliberative decision making (Epstein, 1994; Kahneman, 2003, 2011), use heuristics to a greater extent (Gigerenzer & Todd, 1999), or forgo a thorough and in-depth processing of available information (Kruglanski & Freund, 1983).[4] Furthermore, perceived urgency may not only result in lower levels of choice accuracy but may also alleviate anticipated *action regret* because anticipation of regret is less salient when there is (or appears to be) limited time to deliberate.[5]

Our experiment disentangles these channels in a parsimonious dynamic decision-making environment that allows us to identify the role of regret, perceived urgency, and their interaction. Participants in the experiment buy one unit of a product and maximize their payoff by purchasing the item at a low price without searching for too long. They can sequentially request additional price offers and incur a fixed search cost for every offer that they request (see also Cox & Oaxaca, 1989; Hey, 1987; Kogut, 1990; Schotter & Braunstein, 1981; Sonnemans, 1998). In other words, the participants themselves decide to continue the search for another round or to take the best standing offer. They know the distribution from which offers are drawn and that all previously observed offers are attainable (i.e., we employ optimal stopping with recall). Consequently, expected profit maximization is characterized by adherence to a constant reservation price strategy (Lippman & McCall, 1976). Expected payoff-maximizing individuals search until an offer at or below their reservation price is observed and they then buy the item at that price.

Two deviations from the constant reservation price strategy are commonly observed in search environments, in which buyers do not receive post-purchase information on prices: early stopping and the recall of previously rejected prices. Regardless of the context, previous studies show that participants request fewer offers than theoretically predicted (e.g. Cox & Oaxaca, 1989; Einav, 2005; Hey, 1987; Houser & Winter, 2004; Sonnemans, 1998) and they often make use of the recall option (e.g. Hey, 1987; Houser & Winter, 2004; Ibanez, Czermak, & Sutter, 2009; Kogut, 1990; Schotter & Braunstein, 1981; Schunk, 2009; Schunk & Winter, 2009), which is in line with the idea of anticipated *inaction regret*. Indeed, expanding a standard sequential search model (Lippman & McCall, 1976) by regret aversion predicts both of these commonly observed patterns of behavior (see Appendix A.1 for more detail). Consequently, we designed our experiment to ensure that we can empirically assess the relevance of *regret*. By manipulating whether or not information on post-purchase price realizations is available (see also Sugden et al., 2019; Zeelenberg, 1999; Zeelenberg & Beattie, 1997), we exogenously vary whether anticipated *action regret* can prolong search, countervailing the potential effects of *inaction regret*. Further, we employ random variation in feedback to study the

[4]As has been shown, for instance, in the context of risk-taking and loss aversion (see e.g. Ben-Zur & Breznitz, 1981; Kirchler et al., 2017; Kocher, Pahlke, & Trautmann, 2013).

[5]This idea is in line with the finding that, when explaining individuals' behavior with drift-diffusion models, time-pressure reduces barrier height to speed up choices (Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010).

role of experienced *action regret.* To study how perceived urgency alters the role of regret (as well as accuracy in choice), we implement a 2x2 between-subjects design with high or low perceived urgency that avoids potential selection bias due to time pressure. Finally, we vary search costs (within-subjects) to analyze the extent to which participants understand the general logic of the reservation-price strategy.

The results confirm stylized facts from previous experiments. In all treatments, the participants search on average too little (as compared to the expected payoff-maximizing strategy), make use of the recall option, and, with lower search costs and more more experience, participants search longer. Regarding our treatment variations, we find that perceived urgency reduces decision times and perceived decision quality but does not change search length in general. However, in the very first search task, time pressure does affect search length and reduces payoffs substantially. Anticipated *action regret* (i.e., anticipating regret from not stopping early enough) does not increase search length. We observe no significant differences when participants observe post-purchase price realizations. Experienced regret, both *action* and *inaction regret*, leads to systematic adjustments in search length. Learning that one has stopped searching too early, leads to longer search in the subsequent task while searching for too long (and using the recall option) reduces search length. These adjustments nevertheless do not increase payoffs substantially, as some participants over-adjust their search length.

Our experiment allows to identify economically relevant effect sizes (i.e., larger than 0.20 standard deviations), and thus we provide informative results for research and policy. We find strong evidence that regret-avoiding behavior (Bell, 1982; Bikhchandani & Segal, 2014; Buturak & Evren, 2017; Halpern & Leung, 2016; Hayashi, 2008; Loomes & Sugden, 1982; Qin, 2015; Sarver, 2008; Skiadas, 1997), which has been observed in other experimental contexts (Camille et al., 2004; Coricelli et al., 2005; Fioretti, Vostroknutov, & Coricelli, 2020; Strack & Viefers, 2021; Zeelenberg, 1999), seems much less relevant when decision-makers incur salient search costs by actively requesting new price offers.

Our analyses complements and advances earlier experimental findings on active sequential search that excluded post-purchase price information (see e.g. Ibanez et al., 2009; Schunk, 2009; Schunk & Winter, 2009). We study conditions that exclude and conditions that include post-purchase price information and find that under both conditions, time pressure substantially reduces payoffs with inexperienced decision-makers. Our causal experimental findings are also consistent with correlational evidence from the field, which shows that urgency due to being close to a purchasing deadline is associated with decreased search in an environment with price uncertainty (Lemieux & Peterson, 2011).

Further, our study of anticipated *action regret* links to work focusing on the choice of different types of offers (time-limited and non-time-limited, Sugden et al., 2019). In contrast to this work, we focus on how feedback structures and perceived time pressure affect the number of requested (ex-ante identical) offers instead of whether participants choose non-time-limited or time-limited offers. Complementing the findings of Sugden et al. (2019), we provide robust evidence on the limited role of anticipated *action regret* for search length when decision-makers actively incur search cost to receive additional offers.

More generally, our results relate to the literature on anticipated regret within other decisions that involve a sequential revelation of prices. Strack and Viefers (2021) demonstrate regret sensitivity in an asset-selling problem where new offers are automatically updated at no monetary cost and decision-makers have no recall option. To distinguish the behavior of a regret agent from an expected payoff-maximizer, the empirical analysis of Strack and Viefers (2021) relies on random choice behavior. In their analysis, they assess an agent's sensitivity to feelings of *inaction regret* after having continued the search when it was optimal to stop.[6] Our analyses also link to Fioretti et al. (2020), who vary (within-subject) post-purchase information in a setting akin to Strack and Viefers (2021) and find that participants stop later when they may anticipate *action regret.* While these studies focus on situations in which new prices arrive automatically and no recall option exists, our approach involves an active, costly choice for new price requests and allows for recall.

Some of the related experimental literature describes induced learning through experienced regret in sequential decisions (see e.g. Cooke, Meyvis, & Schwartz, 2001; Einav, 2005; Oprea, Friedman, & Anderson, 2009; Sonnemans, 1998). Oprea et al. (2009) provide post-purchase price realizations in all treatments of an investment task and observe that regret associated with stopping decisions in past tasks leads participants to reconsider their strategy in future tasks. Similarly, participants converge faster to an optimal reservation price in a search task with pre-commitment when receiving post-purchase feedback (Einav, 2005). This is in line with findings on the learning-enhancing effect of regret through priming (Reb, 2008; Reb & Connolly, 2009). Our results complement this line of research and show that learning is fast and inefficiencies vanish over time. Even though our setting allows for potential reinforcement of anticipated regret (because our participants have the opportunity to learn about *action regret* over 10 search tasks), we do not find a learning-enhancing effect of experienced regret. While participants in the condition with post-purchase information adjust their behavior after experiencing regret this does not translate into higher levels of efficiency.

[6] Our theoretical predictions are in line with those of Strack and Viefers (2021) for *optimal stopping*. However, their information structure does not allow them to analytically discriminate between a decision-maker with regret aversion and an expected utility decision-maker when analyzing *optimal stopping*.

There are two reasons that may explain why we do not identify strong effects of regret and perceived urgency. First, anticipated regret might not have been very salient for participants because the recall option makes the subjects perceive that good deals are still available, although net benefits from trade are much smaller when searching longer due to search costs. Furthermore, explicit search costs, as well as the fact that a new price requires an active choice, may render the search-prolonging role of anticipated regret less salient. Second, regret might have been salient but the decision environment was too complex to allow for efficiency-enhancing effects. Our results are in line with a combination of both explanations. In the very first task that the participants encounter, anticipated regret plays a minor role (in line with anticipated regret not being very salient); whereas participants who received post-purchase information react to experienced regret in all tasks. However, the participants were not successful in making better decisions in subsequent search tasks with different price realizations and search costs.

Understanding in greater detail how the aversive feelings of regret and urgency connect to actual decision quality in different environments seems a promising route for future research. Previous results point to positive learning effects through experiencing regret in rather simple and repetitive decisions (see e.g. Einav, 2005; Oprea et al., 2009). In our more complex environment, no additional benefits from experiencing regret for learning are observed. Further, urgency has been found to reduce the depth of reasoning and alter information processing (Kocher & Sutter, 2006; Payne, Bettman, & Luce, 1996), which is in line with the subjective perceptions of our participants who judge perceived urgency to be detrimental for task performance and worse performance in the very first search task. However, time pressure does not result in overall lower payoffs. This result points to potentially interesting research questions around the perception of different decision environments, associated emotions, and actual performance.

Finally, in addition to the analysis of the role of regret and perceived urgency in sequential search tasks, our study highlights the need for strategies consumers may employ to protect themselves from searching sub-optimally. One simple strategy that may circumvent inefficient search, is commitment. We analyze whether such pre-commitment to a reservation price strategy can improve optimality of search and find that commitment indeed results in larger payoffs.

The rest of this manuscript is organized as follows. In Section 2, we explain the experimental design. In Section 3, we specify theory-based hypotheses and in Section 4 we present our empirical results. Finally, we discuss our findings in Section 5.

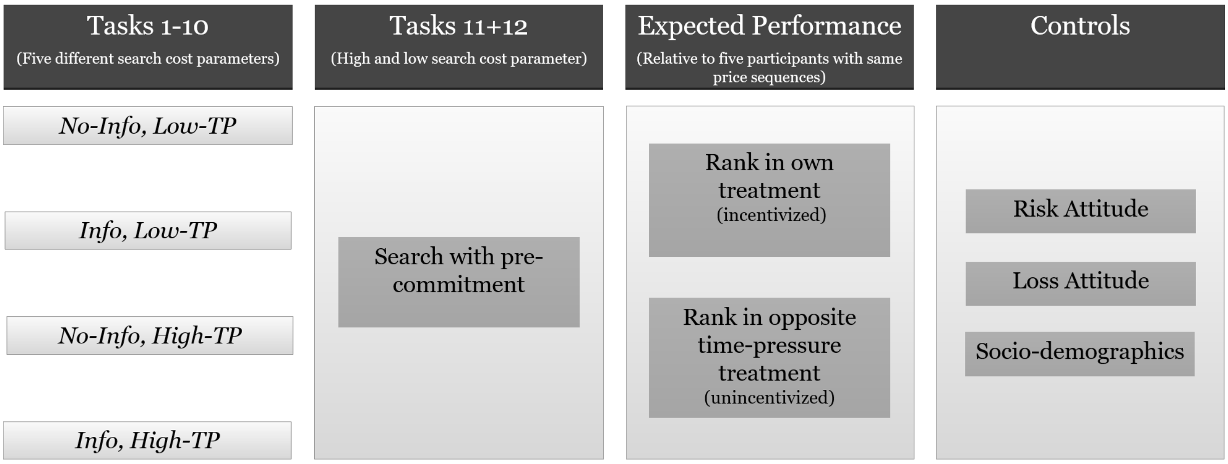

Figure 1: Experimental Design

2 Experimental design

The main part of the preregistered experiment consists of 10 standard sequential search tasks and two additional search tasks with pre-commitment on a reservation price (see also Einav, 2005).[7] For the 10 sequential tasks, we vary perceived urgency by inducing high or low time pressure (*High-TP*, *Low-TP*) and whether participants can anticipate *inaction regret* by providing feedback on post-purchase price offers (*Info*, *No-Info*) in a 2x2 between-subject design, while holding all other aspects of the decision environment constant. After the main part of the experiment, we elicit incentivized measures for the participants' expected relative performance, risk attitudes, and loss attitudes. Furthermore, we elicit a subjective, non-incentivized measure of decision quality relative to participants in the alternative time-pressure condition, and we collect information on socio-demographic characteristics in a short post-experimental questionnaire (see Figure 1). At the end of the experiment, one of the 12 search tasks is randomly drawn to be payoff relevant.[8]

2.1 Sequential search tasks

Participants decide in 10 sequential search tasks whether to buy a fictitious product at the best price observed so far (i.e., optimal stopping with recall).[9] The participants' induced value for the good is $v = 50$. At the beginning of each search task, participants see a first price offer at which they can buy and they then decide whether or not to accept the price or ask for an additional offer. Each additional offer comes at a fixed cost c, which is randomly determined (and altered for each

[7] See AEA RCT Registry; AEARCTR-0004065 for the preregistration.

[8] Negative payoffs in the search task were offset by an additional lump-sum payment for answering the socio-demographic questionnaire at the end of the experiment and additional payments earned in the other incentivized choices.

[9] With perfect recall, previous prices serve as a form of insurance against unsuccessful draws. This reduces the role of risk attitudes on search behavior, allowing us to neatly examine the role of regret.

of the 10 search tasks) but known when deciding upon an additional price request. Price offers are drawn from the known uniform distribution $\{1, 2, ..., 100\}$.[10] We inform the participants that they are free to request new offers as long as there is a possibility to achieve a positive payoff (under any of the search costs used). Given our parameters, this renders the search process finite (because participants can request at most 24 additional offers before making a loss for sure), but the exact number of possible requests is unknown to participants. Only in 0.26 percent of all decisions were 24 additional prices requested (by a total of 4 out of 191 participants). In these cases, the computer automatically bought the product at the best standing price. After the purchase, the current search task was over and the participants proceeded with the next search task.

2.2 Price sequences and search costs

Price sequences were determined randomly in the first two sessions. To keep sequences constant across treatment conditions, the same randomly drawn sequences are used in later sessions. We form within-treatment clusters of six participants who received the same 10 randomly drawn price sequences for the 10 search tasks. Hence, our design allows for a between-subject but within-sequence comparison. Each search task contains eight independent price sequences (because we have 48 participants per treatment and a cluster size of six), and thus the 10 tasks include 80 independently drawn price sequences. We vary the theoretically optimal reservation price strategy by altering search costs between the tasks. We use five different values for the search cost $c \in \{2, 2.5, 3, 3.5, 4\}$. The order in which these parameters appear is randomly determined but held constant for each price sequence and announced for each task as it starts.

2.3 Experimental treatments

2.3.1 Time pressure

We exogenously vary perceived urgency by limiting the amount of time that an individual can spend on each search step (i.e., deciding about buying the product vs. requesting another offer). Instead of resorting to strict time constraints (see, e.g., Ibanez et al., 2009; Sugden et al., 2019), we induce perceived urgency by making longer deliberation more costly. In our high time pressure treatment *High-TP*, participants incur a monetary punishment (1 Taler = 1 unit of the experimental currency) if they fail to accept or ask for a new offer within 4 seconds (and the computer deduces 1 additional Taler every 4 seconds if no decision is made). In our low time pressure treatment *Low-TP*, we set the time limit to reflect on each offer to 60 seconds (i.e., the computer deduces 1 Taler every 60 seconds

[10]We thereby rely on the parametrization of Sonnemans (1998).

if no decision is made). This procedure avoids unwanted selection effects of drop-outs without a deliberate decision (see e.g. Kocher, Schindler, Trautmann, & Xu, 2019), which allows us to impose a time pressure without forcing participants to accept a default (or random) decision after the time ran out and excludes participants from intentionally avoid submitting a choice at all.

2.3.2 Anticipated regret

Orthogonal to the variation in perceived urgency, we vary the feedback after the purchase decision has been made; and thereby, whether decision-makers can anticipate *action regret* from stopping too early. In treatment *Info*, the participants are informed that they will see additional prices after the search, for which they could have bought the product. In *No-Info*, we made the participants aware that they see only those prices that they actively requested until they bought. By varying post-purchase information, we thus exogenously vary whether or not the participants can anticipate *action regret* from buying too early (see also Fioretti et al., 2020; Sugden et al., 2019; Zeelenberg, 1999; Zeelenberg & Beattie, 1997). This anticipation can be reinforced, when experiencing *action regret* in *Info* in previous tasks. To be able to disentangle potential effects of simply seeing additional information (see e.g. Fu, Sefton, & Upward, 2019) as compared to experiencing regret and therefore anticipating regret in later search tasks, we randomly determine the number of displayed offers $k \leq n$ where $n = 25 - \text{OfferNumber}_{accepted}$ such that (for example) a participant who decides to buy after seeing five offers can see between 1 and 20 additional prices.

2.4 Search tasks with pre-commitment

After the 10 sequential search tasks, we confronted all of the participants with two additional search tasks that allow for pre-commitment. In these tasks, the participants pre-specify a price at or below they are willing to buy the good and face no time constraint in that choice. The computer then draws offers until the threshold is reached or undercut. Irrespective of the treatment, the participants have been assigned in the 10 sequential search tasks described earlier, we provide no post-purchase information on additional prices in the tasks with pre-commitment. Thus, the feedback structure rules out anticipated *(action) regret*, and pre-commitment avoids *experiencing (inaction)* regret during the task (as well as the use of the recall option). Search with pre-commitment and without time pressure may therefore counteract potential biases through regret and time pressure. One of the two search tasks involves low search costs ($c_{min} = 2$) and the other involves high search costs ($c_{max} = 4$). This variation allows us to cleanly test for the participants' responsiveness to the search costs.

2.5 Belief elicitation (evaluating own performance)

After the 12 search tasks, the participants have to guess their performance rank (1st to 6th) among those participants who saw the same price offers (i.e., in the within-session price sequence cluster). The subjects are incentivized by a monetary payment if their stated rank matches the actual decision quality (rank) and they receive no payment otherwise. In addition, the participants guess their rank in comparison to the participants who saw the same price sequences and were assigned to the same feedback (*Info* / *No-Info*) condition but to the other time pressure condition. This second, unincentivized measure allows us to study whether participants consider the exogenous increase in perceived urgency to be a less (or more) favorable decision environment.

2.6 Control variables

Given that risk aversion may theoretically shorten search length (empirically, it does not seem to do so, see also Schunk & Winter, 2009; Sonnemans, 1998), we elicit an incentivized proxy for risk attitudes, using the approach by Holt and Laury (2002). We also measure the participants' loss attitudes following the incentive-compatible procedure by Gächter, Johnson, and Herrmann (2007), as suboptimally short search durations may be driven by loss aversion (see e.g. Schunk, 2009). Finally, the participants complete a standard socio-demographic questionnaire (including gender, age as well as their final math grade in high school).

2.7 Procedures

The experiment was conducted at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA) in July and August 2019. In total, 192 participants took part in the experiment.[11] We ran eight sessions (with 24 participants each, two sessions per treatment). The participants were recruited using the online system ORSEE (Greiner, 2015), and we restricted participation to students without experience in sequential search tasks. The experiment was programmed with the software z-Tree (Fischbacher, 2007). On average, participants earned 20 EUR (including a show-up fee of 6 EUR), and the experiment lasted around 60 minutes. Each session was supervised by the same experimenters.

[11] We excluded one participant from the analysis because their search behavior was unresponsive to prices and incentives from task 3 onwards; that is, the participant requested the maximum amount of offers in 8 out of 10 tasks, even when already having encountered extremely favorable offers. Additionally, the decision times of this participant were the fastest across all participants in *Low-TP*. The analyses including this participant are qualitatively the same and can be found in Appendix A.4.1.

3 Predictions

Our main hypotheses concern search behavior; that is, they are directed at differences in the number of requested offers within and across treatment conditions. We also investigate how the number of requested offers corresponds to (ex-ante) efficiency and actual payoffs.

3.1 Regret

Our predictions on the role of regret are based on a theoretical model (see Appendix A.1) which incorporates regret aversion in sequential search building on the formulations of Schunk (2009). This model, reconciles both frequently observed anomalies in empirical search settings without post-purchase information. It predicts that regret-sensitive participants have a higher reservation price (i.e., they request fewer offers) compared to the rational benchmark as they may suffer from *inaction regret* (i.e. from not stopping early enough). The model is also consistent with moderate rates of recall within a task due to *inaction regret.* We specify this prediction in Hypothesis 1:

Hypothesis 1. *In treatment No-Info, regret aversion leads to fewer requested offers when compared to the risk-neutral, regret-free benchmark and it also allows for the use of the recall option.*

The model further predicts that participants request more offers when they know that post-purchase information will be shown (*Info* vs. *No-Info*) because the participants can only regret having stopped too early when learning post-purchase price information. Anticipating this *action regret* prolongs search lengths. We summarize this prediction in Hypothesis 2:

Hypothesis 2. *With anticipated (action) regret, the number of requested offers is lower in treatment No-Info than in treatment Info.*

We additionally hypothesize that experiencing regret reinforces anticipated regret, induces learning, and systematically influences search behavior in subsequent tasks. For Tasks 2 to 10, we specify below one hypothesis for *inaction regret* (i.e., not stopping early enough) that can be present in both information structures and one hypothesis for *action regret* (i.e., having stopped too early) that can only arise under *Info.* We hypothesize that experiencing *inaction regret* leads to a lower number of requested offers in the subsequent search task, whereas we expect experiencing *action regret* to lead to a higher number of requested offers in the subsequent search task.

Hypothesis 3. *The experience of inaction regret (having searched too much) in task t leads to a lower number of requested offers in task $t+1$ in treatments Info and No-Info.*

Hypothesis 4. *The experience of action regret (having searched too little) in task t leads to a higher number of requested offers in tasks $t+1$ in treatment Info.*

Note that empirically testing Hypothesis 2 across all tasks combines the effect of anticipated and experienced regret. In Tasks 2-10, the participants may already have experienced regret in previous tasks, which can directly enhance learning or reinforce the anticipation of regret. To isolate the effect of anticipated regret, we additionally compare search lengths across treatments (*Info* and *No-Info*) in the very first search task participants encounter. Because the participants did not experience regret before this task, the differences between both treatments can be attributed entirely to the anticipation of seeing additional (potentially more favorable) price realizations.

3.2 Time pressure

As alluded to in the introduction, perceived urgency may alter the participants' optimization process and thus result in shorter or longer search length. For example, the participants may tend to accept current offers more frequently when they perceive pressure and thus consider the *High-TP* decision environment to be aversive. Alternatively, the participants may rely on decision heuristics (e.g. Finucane, Alhakami, Slovic, & Johnson, 2000; Gigerenzer & Todd, 1999), which could lead to longer (or shorter) search under time pressure. Because a priori both longer or shorter search is possible and any specific modeling choice seems somewhat arbitrary, the direction of impact remains an empirical question. Consequently, we do not specify a directed hypothesis and instead we formulate the null hypothesis that limiting the time to reflect on an offer does not affect search length.

Hypothesis 5. *The number of requested offers does not differ between treatments High-TP and Low-TP.*

3.3 Potential interaction of time pressure and regret

Building on the idea that time pressure may impair the availability of cognitive resources and thus render the consideration of additional psychological factors less likely (unless they are automatically invoked in the form of heuristics), a potential increase in search length due to the provision of post-purchase price information (i.e., due to the possibility to anticipate regret from requesting too few offers in *Info* and the lack thereof in *No-Info*) should be lower under time pressure. We summarize this prediction in Hypothesis 6, which relies on the assumption that our theory-based prediction for anticipated regret (Hypothesis 2) is also observed empirically:

Hypothesis 6. *Anticipated regret impacts search length to a lesser extent in environments with high levels of perceived urgency.*

4 Results

4.1 Search behavior in general

As outlined above, in this sequential problem, the optimal strategy for a payoff-maximizing regret-free and risk neutral agent is a constant reservation price strategy (see Lippman & McCall, 1976), that is, conditional on search costs, agents derive a cutoff value for the price below which they buy will the good (see also Appendix A.1).[12] Given search costs and realizations of prices in the 10 sequential tasks, this cut-off value translates into an (ex-ante) optimal search length of 4.56 offers in our setting. In the experiment, however, we observe significantly shorter search lengths. Participants stopped on average after seeing 3.78 offers[13] ($p < 0.001$, Wilcoxon signed-ranks test)[14]. Consequently, the participants also earned around 15 percent less than the expected payoff-maximizer would obtain ($p < 0.001$, Wilcoxon signed-ranks test). Furthermore, in a substantial fraction of searches (18.38 percent), the participants make use of the recall option (similar to rates in previous studies between 10-30 percent (e.g. Ibanez et al., 2009; Kogut, 1990; Schotter & Braunstein, 1981), and 74.87 percent of participants do so at least once in the experiment. Hence, we find strong evidence in support of Hypothesis 1:

Result 1. *Participants request significantly fewer offers in No-Info than the risk-neutral and regret-free benchmark predicts and use the recall option.*

4.2 Manipulation of perceived urgency and decision times

Before we present the effects of regret and perceived urgency on search behavior, we briefly establish that our time-pressure intervention indeed resulted in shorter decision times. This is important because our *High-TP* condition deliberately avoids forcing the participants to decide within a strict time limit. Instead of implementing a deadline, the treatment makes slower decisions more costly by deducting 1 point for every 4 seconds that the decision-maker takes to reflect on a price offer. Hence, our treatment variation relies on the assumption that people perceive urgency, and therefore they mostly comply with the time limit.[15]

Our treatment manipulation regarding perceived urgency worked very well. Enforcing a time limit of 4 seconds would be binding in the vast majority of searches under *Low-TP*. Across all tasks,

[12]Depending on the search costs, the reservation price is between 21 and 29 for an expected payoff-maximizer given our parametrization.

[13]This search length corresponds to an average accepted price of 16.88.

[14]All of the reported non-parametric tests in the analysis are two-sided hypothesis tests.

[15]Relative to the average earning in the search task, transgressing the limit once compares to a decrease in earnings of around 4 percent.

participants in *Low-TP* take 5.73s per decision; 44.64 percent of decisions in *Low-TP* take longer than 4 seconds. More importantly, Figure 2 and Table 1 highlight that decision times are substantially and statistically significantly as shorter in *High-TP* than *Low-TP* across all sequential search tasks. Furthermore, the fraction of tasks where all of the decisions were taken within 4 seconds is substantially lower in *Low-TP* when compared to *High-TP* [14.11 percent and 67.19 percent; $p < 0.001$, Mann-Whitney U test (MWU)]. Hence, the participants indeed perceived urgency in *High-TP* and made faster decisions.

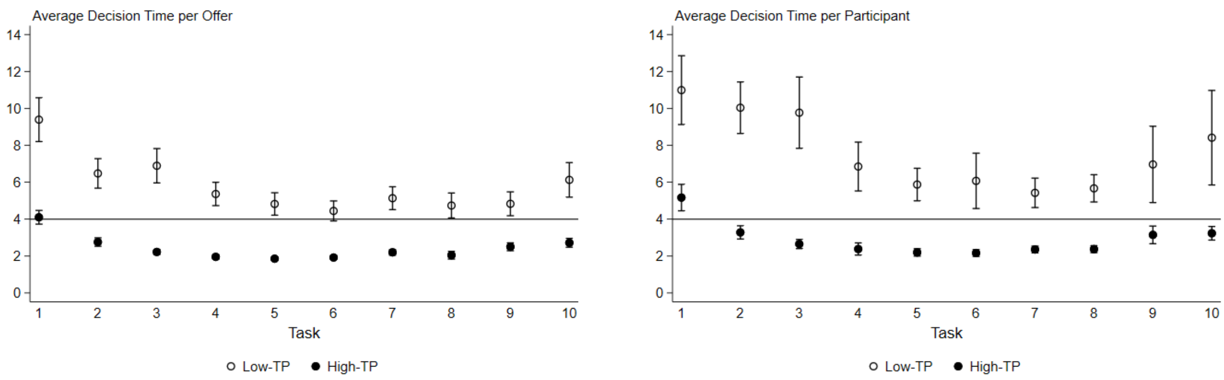

(a) *Average time per offer* (b) *Average time per participant across all offers*

Notes. The error bars indicate 95% confidence intervals.

Figure 2: Decision times across all sequential tasks for *Low-TP* and *High-TP*.

Table 1: Average decision times per task across time pressure conditions

	per Offer		per Subject		
Task	*Low-TP*	*High-TP*	*Low-TP*	*High-TP*	p-value
1	9.39	4.10	10.99	5.17	<0.001
2	6.48	2.76	10.04	3.28	<0.001
3	6.89	2.22	9.77	2.65	<0.001
4	5.36	1.95	6.85	2.38	<0.001
5	4.82	1.86	5.87	2.20	<0.001
6	4.44	1.92	6.07	2.16	<0.001
7	5.13	2.20	5.42	2.35	<0.001
8	4.74	2.05	5.67	2.37	<0.001
9	4.83	2.50	6.97	3.15	<0.001
10	6.13	2.72	8.41	3.23	<0.001

The table shows the average decision times across the time pressure conditions. The p-values are based on non-parametric Mann-Whitney U tests (MWU) on whether the participants' average decision times per task in *Low-TP* and *High-TP* come from the same underlying distribution.

4.3 Search length across treatments

Related to Hypotheses 2 to 5, we compare search behavior across treatments. First, we consider all 10 search tasks jointly and analyze the average effect of time pressure. Then, we consider the joint effect of anticipated and experienced regret on search length. While it may be necessary to experience regret before adjusting behavior in subsequent decisions, a separate analysis of the very first task decision-makers encountered allows us to isolate the effect of anticipated (action) regret. This analysis is presented in section 4.5.[16]

Considering all 10 search tasks, the number of requested offers does not differ significantly across treatments. Neither do we observe a difference between *High-TP* and *Low-TP* ($p = 0.750$, MWU) nor between *No-Info* and *Info* ($p = 0.646$, MWU). The same holds when comparing treatments individually instead of pooling them. Figure 3 illustrates that the average search length is below the (ex-ante) optimal benchmark of 4.56 offers (vertical line) and that the distributions of search lengths across treatments do not differ substantially.

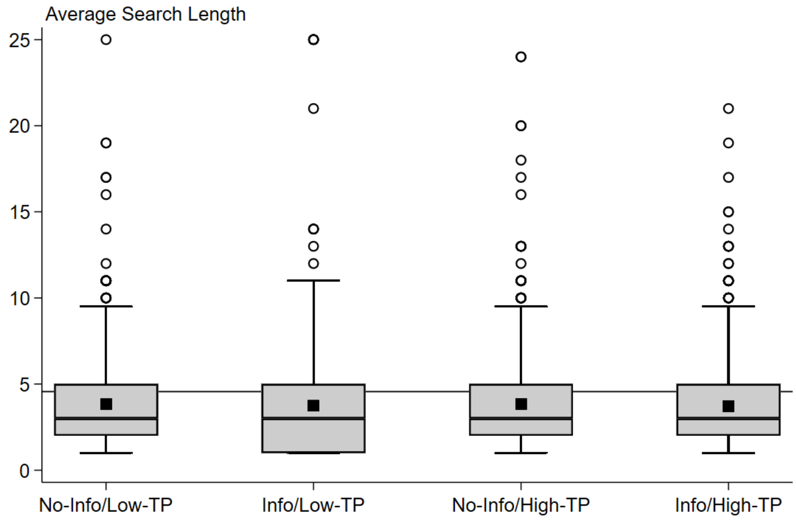

Notes. The figure shows boxplots of search lengths across treatments and a vertical line that indicates the optimal (ex-ante) threshold of a risk-neutral regret-free participant. The length of the whiskers is 1.5 times the interquartile range. The mean search length of each treatment is indicated by a solid square. The vertical line within the box corresponds to the median.

Figure 3: Search length across treatments (Tasks 1-10).

We corroborate these findings in regression analyses (Table 2; columns (1)-(4)). In column (1), we assess the treatment effect, controlling for the number of tasks a decision-maker already completed.

[16]For completeness, we also provide a separate analysis of tasks 2-10. These results mirror the results when considering tasks 1-10 jointly and can be found in Appendix A.3.

In column (2), we add demographic controls, as well as measures of risk and loss attitudes.[17] In column (3), we add fixed effects for the price sequence cluster. In all of the specifications, point estimates for our treatment dummies are consistently close to zero and corroborate the results from the non-parametric analysis—neither perceived urgency nor the variation of the post-purchase information structure affects average search length. In addition to these regression analyses at the search task level, we run Probit regressions for every stopping decision within each search task (see Table A.2, Columns 1 and 2 in the Appendix). This analysis confirms that treatments do not alter search length and shows in addition that decision-makers react systematically to prices. An increase in the current price by one unit approximately leads to a 1 percentage point decrease in the probability of accepting the current price offer. We thus provide robust and consistent evidence that treatments do not affect search length when considering all ten search tasks while, at the same time, decision makers take search costs systematically into account. We thus find no support for Hypotheses 2 but our evidence is in line with Hypotheses 5:

Result 2. *Considering all 10 search tasks, the number of requested offers does neither differ significantly between No-Info and Info nor between High-TP and Low-TP.*

4.4 Efficiency, experiencing regret, and learning over time

Next, we examine how efficient the search behavior is and how it evolves across the 10 search tasks. In total, 57.75 percent of the stopping decisions can be classified as optimal, in 26.60 percent of searches participants should have requested additional offers, and in 16.65 percent of the tasks participants searched too long compared to the reservation price of an expected payoff-maximizer. We observe minor differences across treatments. In *Low-TP*, 62.42 percent of the stopping decisions are optimal; in 24.11 percent of the tasks, too few offers are requested; and in 13.47 percent of the tasks, too many offers are requested. The fraction of optimal decisions in *High-TP* is lower than in *Low-TP* ($p = 0.001$, MWU) and amounts to 53.13 percent. In *High-TP*, the participants request too

[17] Calculating the number of safe choices in the risk elicitation task (Holt & Laury, 2002), participants are on average risk-averse. Meanwhile, 8.38 percent can be classified as risk-loving, 13.61 percent as risk-neutral. In the loss attitude task (Gächter et al., 2007), 4.71 percent of the participants maximize expected payoffs. While the fraction of participants accepting negative expected earnings is negligible (2.09 percent), the vast majority of the participants reject gambles with a positive expected value. The modal response is to accept gambles when the expected value of the gamble is larger than 2 EUR and reject them otherwise. Following the approach of Gächter et al. (2007) we obtain a mean λ of 1.90 (with a standard deviation of 0.57), which is in line with recent literature (Brown, Imai, Vieider, & Camerer, 2021). In the main regressions of Tables 2 and 3, we use a switching point to calculate the measures for risk and loss attitudes. Risk aversion is defined as the row when the participant switches from the safe to the risky lottery. Loss aversion is defined as the (inverse) row when the participant switches from accepting the risky lottery to rejecting it. For example, if a participant does not switch at all, then this is coded as 1. If a participant switches in row 1, then this is coded as 7. The results remain unaffected when we instead control for the number of safe choices (i.e., we take a measure that does not force the participant's responses to comply with monotonicity); see Appendix A.4.3.

Table 2: Search Length

	Number of offers					
	Task 1-10			Task 1		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatments						
High-TP	.022	.072	.071	-.973**	-1.076***	-1.097***
	[-.461,.506]	[-.405,.549]	[-.378,.519]	[-1.737,-.208]	[-1.884,-.268]	[-1.796,-.397]
Info	-.086	-.045	-.059	-.327	-.211	-.214
	[-.571,.399]	[-.515,.425]	[-.474,.357]	[-1.188,.534]	[-1.103,.682]	[-.860,.432]
High-TP X Info	-.033	-.064	-.060	.910	.961	.968
	[-.704,.639]	[-.732,.603]	[-.663,.542]	[-.379,2.199]	[-.344,2.266]	[-.204,2.140]
# Tasks encountered	.079***	.079***	.079***			
	[.032,.125]	[.032,.125]	[.032,.125]			
Risk Aversion		-.036	-.067*		.002	-.080
		[-.117,.044]	[-.145,.011]		[-.176,.181]	[-.256,.096]
Loss Aversion		.017	.017		-.253*	-.233*
		[-.110,.145]	[-.100,.134]		[-.530,.024]	[-.470,.005]
Constant	3.391***	4.295***	4.764***	3.681***	5.747***	4.780***
	[2.988,3.793]	[3.301,5.289]	[3.549,5.979]	[3.081,4.281]	[3.414,8.081]	[2.263,7.297]
Socio-demographic controls	No	Yes	Yes	No	Yes	Yes
Price Sequence Group FE	No	No	Yes	No	No	Yes
Observations	1910	1910	1910	191	191	191

OLS Regressions.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. Values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, which represents the number of offers after which the participant stopped searching. Columns (1)-(3) display search behavior across tasks 1-10, columns (4)-(6) for the first task. Columns (1) and (4) show the effect of the treatments. Columns (2) and (5) add socio-demographic controls (gender, age, cognitive ability) elicited after all search tasks; columns (3) and (6) additionally include price sequence group fixed effects. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

few offers in 27.08 percent of the tasks, and request too many offers in 19.79 percent of the tasks. Hence, behavior is slightly more diverse under *High-TP*. These differences translate into minor payoff differences (*High-TP*: 23.78 vs. *Low-TP*: 25.38; p = 0.080, MWU).

The fractions of optimal stopping decisions under *Info* and *No-Info* are closely aligned (*Info*: 57.37 percent vs. *No-Info*: 58.13 percent; p = 0.879, MWU) and payoffs do not differ substantially across the feedback conditions (*Info*: 24.43 vs. *No-Info*: 24.72; p = 0.727, MWU). Under *No-Info*, in 24.84 percent of the tasks, more offers should have been requested; while in 17.79 percent of the tasks, fewer offers should have been requested. Similarly, in *Info* the fraction of tasks where too few offers were requested is 26.35 percent, and the fraction of tasks where too many offer were requested 15.52 percent.

The closely aligned levels of efficiency across feedback conditions (*No-Info* and *Info*) may result from several reasons. First, participants may not consider the information provided and thus use similar decision processes in both information treatments. Second, participants may process feedback but not react (optimally) to it in subsequent tasks. Third, when participants are confronted with post-purchase information, they may change the overall sensitivity towards their own suboptimal behavior and react differently to similar information in *Info* as compared to *No-Info*. Concerning the first point, we avoided by design that participants simply ignored feedback, as in all treatments participants had to type in the (correct) number of the offer that would have yielded the highest payoff to proceed. Further, we do find evidence that participants spend substantially more time on the feedback screen in *Info* (25.53 seconds) as compared to *No-Info* (14.94 seconds; $p < 0.001$, MWU). It is thus unlikely that participants use similar decision making processes in both information treatments. To investigate the second and third point, we study experienced *inaction regret* (i.e. not having stopped early enough) separately in *Info* and *No-Info* and provide evidence on how experienced *action regret* (i.e., having stopped too early) alters search behavior in *Info* (where participants may learn that they have stopped too early).

Across all conditions, the participants experience *inaction regret* in 22.5 percent of the tasks. *Inaction regret* either arises due to the use of the recall option (79.59 percent of the cases in the data) or when the participants continue the search and encounter a better offer that still does not compensate for the additionally incurred search costs. While (experienced) *inaction regret* does not influence search behavior in general (see Table 3, Column 1), we find evidence that people in *Info* systematically react to the information provided as specified in Hypotheses 3 and 5 (see Table 3, Column 2). Knowing that one should have requested fewer offers in task t results in requesting around 1.14 offers less in task $t+1$ in *Info* compared to participants who did not experience *inaction regret*. In *No-Info*, experiencing *inaction regret*, if at all, slightly increases the number of requested offers (on average they request 0.47 offers more). We summarize this finding in Result 3:

Result 3. *Experiencing inaction regret in task t leads to a lower number of requested offers in task $t+1$ for participants in Info. For participants in No-Info, there is no such effect.*

Next, we assess how *action regret* influences subsequent search behavior. We first compare changes in search behavior in *Info* with changes in search behavior in *No-Info*. That is, we study search in task $t+1$, comparing participants in *Info* who requested too few offers from an ex-ante perspective and were informed by their feedback that they had stopped searching too early in task t with participants in *No-Info* who also requested inefficiently few offers from an ex-ante perspective in task t but did not see post-purchase prices that informed them about their inefficiently short search. For the regression analyses, we simulate the vector of prices participants in *No-Info* would have seen if they had been in the *Info* treatment (i.e., we randomly determine how many post-purchase price realizations they would have observed) and test for the effect of feedback on behavior in task $t+1$. At baseline (*No-Info/Low TP*) in Table 3, Column 3, individuals average search length amounts to 5.02 offers. The average search length of individuals in *Info*, who experience *inaction* *regret* in t is increased by 1.1 offers. In contrast, participants in *No-Info* who also searched too short in task t and thus would have experienced *inaction regret* were they assigned to *Info* instead, continue to search too little (they request around 0.55 offers less in $t+1$). Column 4, which includes experienced *inaction* and *action regret*, confirms these findings.

Table 3: Experienced regret

	Number of offers			
	(1)	(2)	(3)	(4)
Treatments				
High-TP	.203	.192	.211	.203
	[-.295,.700]	[-.280,.664]	[-.265,.687]	[-.257,.663]
Info	-.043	.189	-.335	-.106
	[-.497,.410]	[-.285,.664]	[-.828,.157]	[-.602,.389]
High-TP X Info	-.172	-.129	-.200	-.155
	[-.821,.477]	[-.769,.511]	[-.847,.447]	[-.800,.491]
(Experienced) Inaction Regret	-.082	.473		.420
	[-.497,.332]	[-.127,1.073]		[-.161,1.002]
Inaction Regret X Info		-1.135***		-1.086***
		[-1.885,-.384]		[-1.816,-.356]
(Experienced) Action Regret			-.553*	-.513*
			[-1.111,.006]	[-1.062,.036]
Action Regret X Info			1.095**	1.060**
			[.239,1.951]	[.217,1.903]
# Tasks encountered	.065**	.068**	.061**	.065**
	[.009,.121]	[.012,.123]	[.006,.116]	[.010,.120]
Risk Aversion	-.066	-.065	-.069*	-.068*
	[-.145,.013]	[-.144,.014]	[-.147,.010]	[-.148,.012]
Loss Aversion	.046	.049	.048	.053
	[-.084,.176]	[-.079,.177]	[-.079,.175]	[-.075,.180]
Constant	4.858***	4.666***	5.017***	4.821***
	[3.505,6.210]	[3.297,6.034]	[3.601,6.432]	[3.380,6.262]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1719	1719	1719	1719

OLS Regressions.*** $p<0.01$, ** $p<0.05$, * $p<0.1$. Standard errors clustered at the individual level. Values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching.Columns (1)-(4) display search behavior in tasks 2-10 and investigate the effect of regret experienced in the previous task. All columns include socio-demographic controls (gender, age, cognitive ability) and price sequence group fixed effects. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 2-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

As we randomly determined the number of displayed post-purchase prices within *Info*, we can also compare changes in behavior by participants within *Info* who requested too few offers from an ex-ante perspective and either were informed about having stopped too early and those who did not see more favorable post-purchase price realizations. We find that those who searched too short from an ex-ante perspective and were informed about stopping too early requested on average 0.94 offers more in the subsequent task as compared to those who searched too short but did not see favorable post-purchase price realizations (3.69 vs. 2.75 offers requested in task $t + 1$ after stopping too early in task t; p = 0.058, MWU). Result 4 summarizes these findings:

Result 4. *Experiencing action regret in task t leads to a higher number of requested offers in task $t + 1$.*

Although participants adjust their search behavior, they do not make higher profits after experiencing regret in the previous task (see Appendix Table A.1, Column 1). This is true for both *inaction regret* and *action regret*. Participants who received information that higher earnings were possible had they stopped later (i.e., participants experiencing *action regret*) react by requesting inefficiently many offers in the next task. Table A.1 shows that the likelihood that participants continue to request too few offers remains unaffected (see Column 3), while the likelihood to ask for too many offers increases at the expense of optimal searches (see Columns 2 and 4).[18] Thus, we find evidence that participants react to experienced regret, but do not react optimally and, at the same time, that participants are more sensitive to information about *inaction regret* when experiencing the latter in *Info*.

Finally, we shed light on learning over time in terms of (sub)optimal choice. In the first half of their sequential search (tasks 1-5), the participants request on average around 1.57 fewer offers than ex-ante optimal ($p < 0.001$, Wilcoxon signed-ranks test). That is, suboptimal choice results mainly from stopping too early (participants request too many offers in only 15.39 percent of the first five tasks). Over time, participants request more offers (as shown by the *# Tasks encountered* coefficients in Tables 2 and 3) such that in the second half (tasks 6-10), the difference of the average search length to the optimal search length amounts to only 0.26 fewer offers than ex-ante optimal and does no longer significantly differ from the optimal benchmark (p = 0.352, Wilcoxon signed-ranks test). Overall, the fraction of searches where participants requested too few offers decreases from 36.13 percent in the first half to 15.08 percent in the second half, while the fraction of search tasks in which participants requested too many offers remains fairly constant (15.39 percent to 17.91 percent) across all treatments (see Figure A.2 in the appendix).

4.5 Anticipated regret and inexperienced decision-makers

To isolate the effects of anticipated regret (excluding any experienced regret) and to study the effects of time pressure for inexperienced subjects, we now focus on the first task decision-makers encounter. Similar to our overall finding, participants stop also significantly earlier than optimal

[18] In a robustness check (Table A.7), we show that all results hold in a truncated Poisson specification.

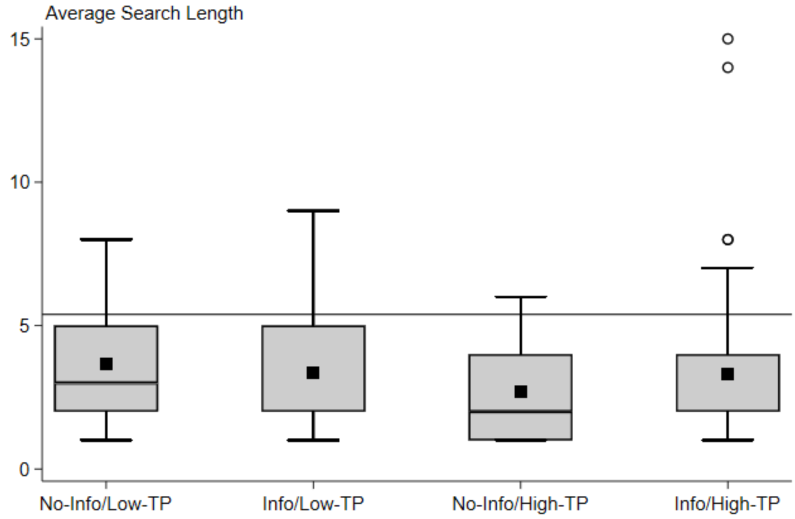

Notes. The figure shows boxplots of search lengths across treatments and a vertical line that indicates the optimal (ex-ante) threshold of a risk-neutral regret-free participant. The length of the whiskers is 1.5 times the interquartile range. The mean search length of each treatment is indicated by a solid square. The vertical line within the box corresponds to the median, which coincides with the lower quartile (lower end of the box) for Info/Low-TP and Info/High-TP.

Figure 4: Search length across treatments (Tasks 1).

in the very first task (in all treatments, see Figure 4). While expected payoff maximizing behavior in the very first task results in stopping after seeing on average 5.39 offers, participants observe on average 3.26 offers. This difference is statistically significant when pooling the treatments and when analyzing them individually ($p < 0.001$ for each individual as well as the pooled test, Wilcoxon signed-ranks test). Search lengths in *Info* and *No-Info* are statistically indistinguishable ($p = 0.805$, MWU), while participants under time pressure search significantly shorter than participants without ($p = 0.019$, MWU). We corroborate the non-parametric analysis by regression analyses (see Table 2; Columns (4)-(6)). The results remain robust when adding demographic controls, and also when using independently elicited preferences as additional controls (column 5) and when including price sequence group fixed effects (column 6). Hence, also for the very first task we find no effects of the feedback environment.

In contrast to our overall result, we do find a strong and statistically significant effect of time pressure on search length in the very first task (see also Table 2; Columns (4)-(6)), which substantially reduces payoffs in *High-TP*. As shown in Figure 5, under *Low-TP*, average payoffs amount to 23.14 Taler whereas in *High-TP*, participants' payoffs are more than 40 percent lower (on average they achieve only 13.33 Taler, $p = 0.004$, MWU).[19]

It is noteworthy, that perceived urgency was detrimental in the sense that subject in *High-TP* would not have fared worse when taking more time (as their counterparts in *Low-TP* did). When

[19]This comparison already excludes the extra cost that participants incurred in *High-TP* when exceeding the time threshold.

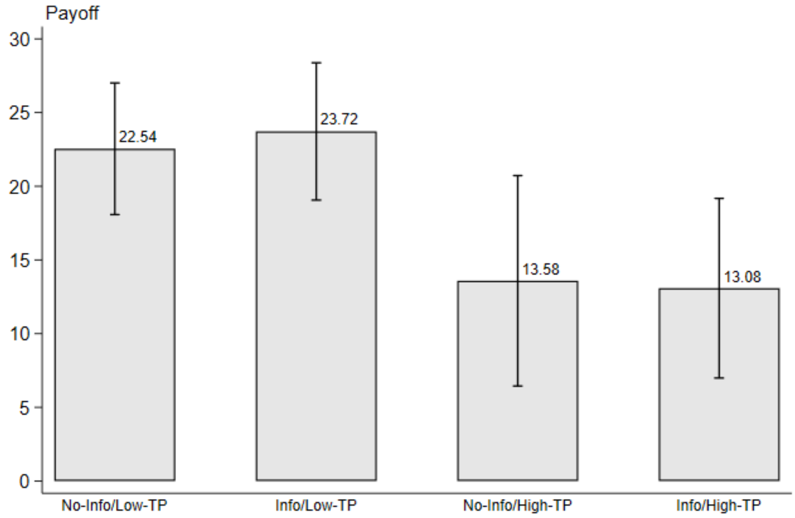

Notes. The figure shows the payoffs (in Taler) from the very first search task, excluding potential deductions for exceeding the time limit in High-TP conditions. The error bars indicate 95% confidence intervals.

Figure 5: Payoffs across treatments in Task 1

taking punishment costs due slower search in *High-TP* into account and applying the same punishment rule hypothetically to participants in *Low-TP*, our data suggests that, if at all, participants could have benefited from making slower choices. Hypothetical payoffs under *Low-TP* (with added costs for exceeding the threshold of 4 seconds) amount to 16.63 whereas those under under *High-TP* amount to 11.75 (when substracting the punishment costs for slow decisions; p = 0.305, MWU). Hence, ignoring the imposed time pressure and acting as if it was absent would have been at least as good in terms of payoffs as the strategies that the participants in *High-TP* resorted to.

Further, we provide additional evidence that participants reacted to pressure in a sub-optimal way in the very first task, by comparing the number of requested offers conditional on the decision times in *Low-TP*. Note that the mere fact of deciding quickly does not imply short search durations in treatment *Low-TP*. Instead, swift decision-making is associated with a larger number of requested offers (Spearman's $rho = -.37$; $p <$ 0.001). In efficiency terms, swift responses do not seem to be related to lower payoffs in *Low-TP*. Participants in *High-TP* who decided within 4 seconds perform substantially, although not significantly, worse (28.09 percent smaller payoffs; p = 0.964, MWU) than those who took more time to reach the decision (including the deduction for violating the time threshold). We interpret this as suggestive evidence that participants who (inefficiently) comply with the time threshold in the *High-TP* treatment by making faster choices than they would without time pressure do so in a systematic way (i.e., by requesting significantly fewer offers).

Summarizing the results for the very first task, we confirm the previously reported Results 1 and 2. Participants request significantly fewer offers in *No-Info* than the risk-neutral and regret-free

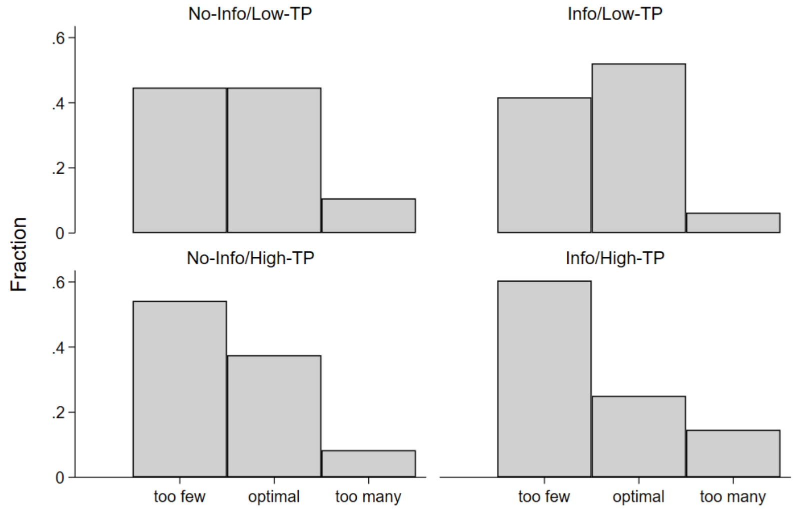

Notes. The figure shows search behavior in Task 1. Behavior is classified as having requested either too few, too many, or the optimal number of offers compared to the (ex-ante) optimal behavior of a risk-neutral regret-free participant.

Figure 6: Efficiency (ex-ante) of search behavior in Task 1

benchmark predicts and there is no (pure) effect of anticipated *action regret* on search behavior. In contrast to the analysis including all tasks, we find a significant effect of time pressure for the first task, which aggravates the existing tendency to request fewer offers than optimal. The latter is also confirmed in additional regression analyses considering every single stopping decision within the first task (see Table A.2, Columns 3 and 4 in the Appendix, which highlight that time pressure makes participants 15 percentage points more likely to stop the search at the current offer).

Result 5 *Participants request significantly fewer offers under High-TP than under Low-TP in the first search task they encounter, forgoing on average more than 40 percent of profits.*

4.6 Perceptions of the decision environment

Sellers often pressure buyers to make quick decisions. Our design allows us not only to study how time pressure alters search behavior and decision quality but also to study how decision-makers perceive their own decision quality. At the end of the experiment, we elicit how decision-makers rank their performance as compared to other buyers. On average, the participants are overconfident in all treatments.[20] Ranking themselves within a group of six (who all observed the same price sequences), they place themselves, on average, around one rank better than they actually are.

Although we do not find strong differences in actual performance across treatment when considering all 10 tasks, in a within-subjects comparison the participants expect to perform worse under

[20]We do not neither observe significant differences between treatments *No-Info* and *Info* (p = 0.165, MWU), nor between *High-TP* and *Low-TP* (p = 0.959, MWU).

time pressure than without ($p < 0.001$, Wilcoxon signed-ranks test). The difference is around 0.38 ranks on average. Although this holds for participants in both urgency conditions, it is stronger for participants in *High-TP* (p = 0.018, MWU test for differences in differences in rankings, comparing those assigned to *High-TP* and *Low-TP*, see also Figure A.1).

4.7 Improving search behavior through commitment

Our study documents inefficient search across all treatment conditions and detrimental effects of time pressure for inexperienced decision-makers. Thereby our findings highlights the need for strategies consumers may employ to protect themselves. One simple strategy that may circumvent suboptimal search is commitment to a reservation price. In two additional search tasks, we explicitly asked participants to commit to a reservation price instead of searching sequentially. We asked for such pre-commitment once with low (c=2) and once with high (c=4) search costs and compare their outcomes to their sequential search behavior.[21] Based on the reservation price stated for low and high search costs and realized prices in the sequential search tasks, we calculate when participants would have stopped the sequential search (if they had adhered to their stated reservation price). Doing so, we compare how the reservation price strategy fares with the same price sequence and with the same search costs as compared to sequentially requesting offers.

We find that commitment improves search efficiency. The percentage of optimal searches is significantly higher with pre-commitment than in the corresponding tasks of the main experiment (70.42 percent vs. 49.74 percent for search costs of c=2 and 80.10 percent vs. 67.02 percent with search costs of c=4; $p < 0.001$ for both search cost parameters, Wilcoxon signed-ranks tests). Hence, average reservation prices with pre-commitment are still above the rational benchmark, but the tendency to systematically request too few offers in early tasks is much less pronounced. Consequently, the participants achieve significantly larger profits with commitment (29.58 vs. 26.61 Taler for search costs of c=2, 21.49 vs. 20.09 Taler for search costs of c=4; $p < 0.001$ and p = 0.014, Wilcoxon signed-ranks tests).[22]

[21]Reassuringly for our analyses of the value of pre-commitment, we find no indication that the treatments in the 10 sequential search tasks had an effect on search behavior in the additional search tasks with pre-commitment. This holds true when comparing the behavior in the two tasks separately (p = 0.529 for *High-TP* vs. *Low-TP* and p = 0.883 for *Info* and *No-Info* for Task 11 (c=2), MWU; p = 0.914 for *High-TP* vs. *Low-TP* and p = 0.167 for *Info* and *No-Info* for Task 12 (c=4), MWU) and jointly (p = 0.61 for *High-TP* vs. *Low-TP* and p = 0.708 for *Info* and *No-Info* for the average reservation price, MWU). In addition, we observe that participants reacted systematically to the incentives that they faced in the tasks with pre-commitment, choosing significantly higher reservation prices with high (as compared to low) search costs ($p < 0.001$, MWU).

[22]This remains unchanged if we only consider treatments without time pressure ($p < 0.001$ and p = 0.007 for search costs of c=2 and c=4, Wilcoxon signed-ranks tests).

Note that this within-subject comparison does not allow us to rigorously disentangle effects of the different decision environment [choice of reservation price (pre-commitment) vs. sequential search] and learning over the experiment (because the tasks with pre-commitment followed after the 10 search tasks). However, we find that efficiency in the two tasks with pre-commitment is higher than in the last two of the 10 sequential tasks (13.48 percentage points more optimal decisions), suggesting that learning alone cannot explain the differences between the sequential search tasks and the tasks with pre-commitment.

To further disentangle learning and the effects of pre-commitment, we replicated the two pre-commitment search tasks in an additional sample, in which participants did not encounter the ten sequential search tasks at all.[23] Again, we find support for the efficiency-enhancing effect of pre-task commitment. Reservation prices in the additional experiment that excluded learning possibilities do not differ significantly from reservation prices in the original experiment (p = 0.405 and p = 0.923 for search costs of c=2 and c=4, MWU). Moreover, we find that reservation price choices in the additional experiment lead to optimal stopping more often than sequential search behavior (with and without time pressure) in the 10 tasks of the main experiment (67.55 percent vs. 57.75 percent jointly; p = 0.007, MWU, with time pressure: mean = 53.13, p = 0.002, without time pressure: mean = 62.42, p = 0.071). Because the participants learned over time in the main experiment (as shown in Section 4.4), the difference is even more pronounced when comparing reservation price choices (which excluded learning possibilities) to the choices made in the first half of the main experiment (64.54 percent vs. 48.48 percent; p = 0.001, MWU). We summarize these findings in Result 6:

Result 6 *Pre-commitment on a reservation price improves profits.*

5 Discussion and conclusion

Perceived urgency and regret may substantially affect individual choice in dynamic market environments and hence aggravate or alleviate any potential biases in decision-making. We used a well-powered experimental study to evaluate the empirical importance of both aspects. The (95 percent) confidence intervals for the treatment effect estimates in our preferred regression specification (Table 2, Column 3) are consistent with differences across treatments of up to 0.66 requested offers, corresponding to 0.17 standard deviations in the number of requested offers. Hence, we can rule out true but undetected effect sizes being larger than 0.17 standard deviations. We obtain very similar results when deriving minimum detectable effect sizes using a simulation-based approach

[23]We recruited 47 subjects from the same pool as in the initial experiment (excluding all participants of the main experiment) and ran the additional sessions at MELESSA in September 2020.

(see Campos-Mercade, 2018). Based on the realized distribution of search lengths, we set the desired level of power to 80 percent and the statistical significance level to 5 percent. We then perform parametric and non-parametric tests, and find that we are able to detect effect sizes of at least 0.15 standard deviations across all tasks. Hence, our study ex-ante allowed us to detect economically meaningful treatment differences.

Our results provide robust evidence that regret and perceived urgency do not generally affect the number of requested offers in sequential search tasks. In particular, we do not find that anticipated regret renders active sequential search. However, we observe that urgency significantly affects search behavior and profits in the very first search task that the participants encounter. Under high time pressure, stopping too early is (even) more prevalent than under low time pressure and profits are substantially reduced. Thus, our results provide a rationale for why sellers often put buyers under time pressure. Our findings emphasize that pressuring buyers by inducing a sense of urgency may be particularly effective when applied to inexperienced customers (i.e., customers who have not encountered the respective search task before). In light of recent policy debates on the effects of sales tactics that "rush consumers into making a decision",[24] these policies should thus especially focus on environments in which decision-makers are not very savvy or in which they search for new products (which they have not searched for before). For example, the British Competition and Markets Authority recently required booking sites to take action against practices of pressure selling (i.e., practices that create perceived urgency) and of displaying potentially misleading unattainable offers [i.e., that give rise to (anticipated) feelings of regret], such as already forgone options. Given that booking flights or hotels is a regular task for many consumers, they may quickly learn to resist perceived urgency and regulation is less immanent. However, other decisions may be more infrequent but substantially more important. Buying a house, taking out life insurance, or making other long-term investment decisions presents most consumers with an unknown decision environment. In these environments, most consumers are inexperienced. Therefore, time pressure may be even more harmful and regulation more important than in the areas that are currently primarily targeted, such as hotel booking or travel sites on the Internet.

Further, our results suggest a simple mechanism that consumers can use to avoid an inefficient search: commitment to a binding reservation price. Commitment turns out to be a simple method that may be applied easily in contexts outside the laboratory. Because pre-commitment increases payoffs, there might be demand for devices that guarantee its provision, even in form of market-

[24]Retrieved from https://www.gov.uk/government/news/cma-launches-enforcement-action-against-hotel-booking-sites on 10/05/2020.

based solutions (as long as enough consumers are sophisticated). Alternatively, public policies may promote or offer such commitment possibilities.

Finally, although our design captures the most essential elements of the trade-off that urgency and the resulting time pressure in real-world settings pose (namely, having to decide on the spot or incurring costs to delay and reflect on one 's decision), decision environments outside the laboratory may confront consumers with additional challenges. For example, the costs from delay are known in the experiment, while they are often unknown or uncertain in environments outside the laboratory, where consumers face uncertainty about the underlying distribution from which prices are drawn and firms may have an incentive to disguise certain pieces of information to create more intransparent decision environments. In addition, outside the laboratory, decision-makers may expect that searching too long may eventually result in a situation in which no purchasing opportunity for the same or an equivalent product or service is left. Consequently, future research should explicitly test for the role of unknown or uncertain search costs and take the disproportional weight of low-probability but potentially detrimental outcomes into account.

References

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, *30*(5), 961–981.

Ben-Zur, H., & Breznitz, S. J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica*, *47*(2), 89–104.

Bikhchandani, S., & Segal, U. (2014). Transitive regret over statistically independent lotteries. *Journal of Economic Theory*, *152*, 237–248.

Bleichrodt, H., Cillo, A., & Diecidue, E. (2010). A quantitative measurement of regret theory. *Management Science*, *56*(1), 161–175.

Braun, M., & Muermann, A. (2004). The impact of regret on the demand for insurance. *Journal of Risk and Insurance*, *71*(4), 737–767.

Brown, A. L., Imai, T., Vieider, F., & Camerer, C. (2021). Meta-analysis of empirical estimates of loss-aversion. *CESifo Working Paper*.

Buturak, G., & Evren, Ö. (2017). Choice overload and asymmetric regret. *Theoretical Economics*, *12*(3), 1029–1056.

Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J.-R., & Sirigu, A. (2004). The involvement of the orbitofrontal cortex in the experience of regret. *Science*, *304*(5674), 1167–1170.

Campos-Mercade, P. (2018). *Power analysis through simulations in stata: a guide for dummies.* Unpublished manuscript.

Cooke, A. D., Meyvis, T., & Schwartz, A. (2001). Avoiding future regret in purchase-timing decisions. *Journal of Consumer Research*, *27*(4), 447–459.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: a neuroimaging study of choice behavior. *Nature Neuroscience*, *8*(9), 1255–1262.

Cox, J. C., & Oaxaca, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, *2*(3), 301–329.

Einav, L. (2005). Informational asymmetries and observational learning in search. *Journal of Risk and Uncertainty*, *30*(3), 241–259.

Engelbrecht-Wiggans, R., & Katok, E. (2008). Regret and feedback information in first-price sealed-bid auctions. *Management Science*, *54*(4), 808–819.
Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist*, *49*(8), 709.
Finucane, M. L., Alhakami, A., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making*, *13*(1), 1–17.
Fioretti, M., Vostroknutov, A., & Coricelli, G. (2020). Dynamic regret avoidance. *USC-INET Research Paper*(17-07).
Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.
Fu, J., Sefton, M., & Upward, R. (2019). Social comparisons in job search. *Journal of Economic Behavior & Organization*, *168*, 338–361.
Gächter, S., Johnson, E. J., & Herrmann, A. (2007). Individual-level loss aversion in riskless and risky choices. *IZA Discussion Paper 2961.*
Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart.* Oxford University Press, USA.
Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, *1*(1), 114–125.
Halpern, J. Y., & Leung, S. (2016). Minimizing regret in dynamic decision problems. *Theory and Decision*, *81*(1), 123–151.
Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, *139*(1), 242–268.
Hey, J. D. (1987). Still searching. *Journal of Economic Behavior & Organization*, *8*(1), 137–144.
Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655.
Houser, D., & Winter, J. (2004). How do behavioral assumptions affect structural inference? Evidence from a laboratory experiment. *Journal of Business & Economic Statistics*, *22*(1), 64–79.
Ibanez, M., Czermak, S., & Sutter, M. (2009). Searching for a better deal–on the influence of group decision making, time pressure and gender on search behavior. *Journal of Economic Psychology*, *30*(1), 1–10.
Irons, B., & Hepburn, C. (2007). Regret theory and the tyranny of choice. *Economic Record*, *83*(261), 191–203.
Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58*(9), 697.
Kahneman, D. (2011). *Thinking, fast and slow.* Macmillan.
Kirchler, M., Andersson, D., Bonn, C., Johannesson, M., Sørensen, E. Ø., Stefan, M., … Västfjäll, D. (2017). The effect of fast and slow decisions on risk taking. *Journal of Risk and Uncertainty*, *54*(1), 37–59.
Kocher, M. G., Pahlke, J., & Trautmann, S. T. (2013). Tempus fugit: Time pressure in risky decisions. *Management Science*, *59*(10), 2380–2391.
Kocher, M. G., Schindler, D., Trautmann, S. T., & Xu, Y. (2019). Risk, time pressure, and selection effects. *Experimental Economics*, *22*(1), 216–246.
Kocher, M. G., & Sutter, M. (2006). Time is money—time pressure, incentives, and the quality of decision-making. *Journal of Economic Behavior & Organization*, *61*(3), 375–392.
Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, *14*(3), 381–392.
Krähmer, D., & Stone, R. (2005). Regret in dynamic decision problems. *SFB/TR 15 Discussion Paper*.
Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay-inferences: Effects on impressional primacy, ethnic stereotyping, and numerical anchoring. *Journal of Experimental Social Psychology*, *19*(5), 448–468.

Lemieux, J., & Peterson, R. A. (2011). Purchase deadline as a moderator of the effects of price uncertainty on search duration. *Journal of Economic Psychology*, *32*(1), 33–44.
Lippman, S. A., & McCall, J. (1976). The economics of job search: A survey. *Economic Inquiry*, *14*(2), 155–189.
Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, *92*(368), 805–824.
Michenaud, S., & Solnik, B. (2008). Applying regret theory to investment choices: Currency hedging decisions. *Journal of International Money and Finance*, *27*(5), 677–694.
Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making*, *5*(6), 437.
Muermann, A., & Volkman, J. M. (2007). Regret, pride, and the disposition effect. *Working Paper*.
Oprea, R., Friedman, D., & Anderson, S. T. (2009). Learning to wait: A laboratory investigation. *Review of Economic Studies*, *76*(3), 1103–1124.
Payne, J. W., Bettman, J. R., & Luce, M. F. (1996). When time is money: Decision behavior under opportunity-cost time pressure. *Organizational Behavior and Human Decision Processes*, *66*(2), 131–152.
Qin, J. (2015). A model of regret, investor behavior, and market turbulence. *Journal of Economic Theory*, *160*, 150–174.
Reb, J. (2008). Regret aversion and decision process quality: Effects of regret salience on decision process carefulness. *Organizational Behavior and Human Decision Processes*, *105*(2), 169–182.
Reb, J., & Connolly, T. (2009). Myopic regret avoidance: Feedback avoidance and learning in repeated decision making. *Organizational Behavior and Human Decision Processes*, *109*(2), 182–189.
Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, *76*(2), 263–305.
Schotter, A., & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, *19*(1), 1–25.
Schunk, D. (2009). Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence. *Journal of Economic Dynamics and Control*, *33*(9), 1719–1738.
Schunk, D., & Winter, J. (2009). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization*, *71*(2), 347–360.
Skiadas, C. (1997). Conditioning and aggregation of preferences. *Econometrica*, 347–367.
Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, *35*(3), 309–332.
Strack, P., & Viefers, P. (2021). Too proud to stop: Regret in dynamic decisions. *Journal of the European Economic Association*, *19*(1), 165–199.
Sugden, R., Wang, M., & Zizzo, D. J. (2019). Take it or leave it: Experimental evidence on the effect of time-limited offers on consumer behaviour. *Journal of Economic Behavior & Organization*, *168*, 1–23.
Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, *12*(2), 93.
Zeelenberg, M., & Beattie, J. (1997). Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Organizational Behavior and Human Decision Processes*, *72*(1), 63–78.

A Online appendix

A.1 Theoretical search model

Standard information environment

To derive testable behavioral hypotheses for the experimental design, we incorporate regret aversion into one of the most classic and simple search models, building on the formulation of Schunk (2009).[25] In the model, agents have an inelastic demand for one unit of a good, receive offers sequentially, and they incur a (fixed) search cost for every offer that they request. We allow for perfect recall, such that agents can always take the lowest price encountered so far. There is no limit on the number of offers that can be requested and the prices are randomly drawn from a previously known discrete uniform distribution. The distribution function from which the offers are drawn is $F(.)$ with range $[l, h]$. The search costs for each requested offer are denoted as c. The participants maximize profits (π), which are calculated as the difference between induced valuation (v) for the good and the costs for the purchase. This cost consists of the total search cost plus the final price to be paid (p). The best price observed so far is denoted by (m_t). Intuitively, to request another offer, the sure loss of c must be outweighed by the possibility of finding a better price in $t + 1$.[26]

Payoff-maximizing agent. The optimal behavior for a risk-neutral agent is a constant reservation price strategy (Lippman & McCall, 1976). To calculate this reservation price, it is sufficient for the agent to compare the benefits from stopping the search now and the benefits requesting one additional offer and stopping afterward. This is displayed in Equation 1.

$$
\begin{aligned}
\pi(v - m_t) &= [1 - F(m_t)]\pi(v - m_t - c) + \int_l^{m_t} \pi(v - x - c)\,\mathrm{d}F(x) \qquad (1)\\
\Leftrightarrow \pi(v - m_t) &= [1 - F(m_t)]\pi(v - m_t - c) + \int_l^{m_t - c} \pi(v - x - c)\,\mathrm{d}F(x) + \int_{m_t - c}^{m_t} \pi(v - x - c)\,\mathrm{d}F(x)
\end{aligned}
$$

The left-hand side represents the value from stopping the search. The right-hand side is the value from requesting another offer. The first term on the right side corresponds to the cases where no better price is found. The second term in the first line corresponds to prices that are below the current best price (m_t) and weights the resulting profits by their probability. In the second row, we

[25] This relates to other theoretical models that incorporate regret in static frameworks like currency hedging (Michenaud & Solnik, 2008), insurance choices (Braun & Muermann, 2004) or the expansion of the choice set (Irons & Hepburn, 2007). In sequential decisions, general approaches to model dynamic choices under regret (e.g. Krähmer & Stone, 2005) have been applied to investment decisions (Muermann & Volkman, 2007) and asset-selling problems (Strack & Viefers, 2021).

[26] We refer to every decision between stopping at offer t or requesting offer $t + 1$ as a round, meaning that every search task consists of up to 25 rounds.

distinguish between the cases where better prices outweigh the search costs ($m_t - m_{t+1} > c$) and the cases where they do not. This allows us to draw a comparison with the optimization problem of a regret-sensitive agent.

Regret-sensitive agent. We also derive predictions for a regret-sensitive agent (Bell, 1982; Loomes & Sugden, 1982). We make the simplifying assumption that regret is a function of the difference between the payoffs of the chosen and the unchosen option. Accordingly, the utility from choosing option i over k under the state of the world j is defined as: $m^k_{ij} = \pi(x_{ij}) - R[\pi(x_{kj}) - \pi(x_{ij})]$. The agent both derives utility from the material benefits from the choice of i, but also from the comparison of the chosen and the unchosen option. The regret/rejoice-function R specifies how much the comparison of actual and counterfactual outcomes affects the individual's utility. As common (e.g Michenaud & Solnik, 2008; Muermann & Volkman, 2007; Zeelenberg, 1999), we build on the observation that regret is felt more intensely than rejoice (Bleichrodt, Cillo, & Diecidue, 2010). For simplicity, we assume that the agent does not experience (and anticipate) any rejoice. The agent experiences negative utility if the unchosen option had led to higher payoffs. Conversely, the agent does not experience positive utility if the chosen alternative led to higher profits. We assume regret aversion; that is, a convex R in the positive domain of regret.

The experience and anticipation of *inaction regret* induce the two commonly observed anomalies in standard search tasks: early stopping and the recall of previously rejected offers. The utility from stopping at a lowest price m_t in round t becomes $u(m_t) = \pi(v - m_t) - R(\pi_{max_t} - \pi_t)$. Regret is defined as a function of the foregone profits by not having stopped at the payoff-maximizing offers up to t. π_{max_t} denotes the payoffs at the ex-post optimal stopping point. This maximum serves as a reference point for the feelings of regret. π_t denotes the payoff from stopping in round t.

We incorporate *inaction regret* into Equation 1. Equation 2 models optimal decision making for regret-sensitive agents using one-step forward-induction. Current feelings of regret enter on the left-hand side, anticipated feelings on the right-hand side. On the right-hand side, the first term captures the case where the next draw does not yield a better price than m_t. The second term describes the situations in which a payoff-increasing price was drawn. The third term corresponds to prices that are better than m_t, but do not outweigh the search costs (c).

$$\begin{aligned}\pi(v - m_t) - R(\pi_{max_t} - \pi_{m_t}) =& [1 - F(m_t)][\pi(v - m_t - c) - R(\pi_{max_t} - \pi_t - c)] \\ &+ \int_l^{m_t - c} \pi(v - x - c)\,\mathrm{d}F(x) \\ &+ \int_{m_t - c}^{m_t} [\pi(v - x - c) - R(\pi_x - \pi_{max_t} - c)]\,\mathrm{d}F(x)\end{aligned} \tag{2}$$

Why would a regret-averse agent search shorter than an expected profit-maximizing individual? In the standard information environment, no feedback about foregone options after stopping is revealed. You only feel regret if you have searched for too long (*inaction regret*). At each decision node, the experience of (additional) regret can occur only by continuing, not by stopping. Accordingly, regret-averse agents have a higher reservation price and therefore request fewer offers. For simplicity, we assume that the current price is the best offer so far. Given $\pi_{max_t} = \pi_t$, the left hand sides of Equations 1 and 2 are the same. Nevertheless, the expected value from continuing the search is strictly lower for regret-averse agents. If no better price is found, then not only does the material loss of c reduce utility but so does the regret of not stopping in the previous round. As the continuation value is lower, a regret-averse agent stops searching at a higher price than a pure payoff-maximizer due to the anticipation of (potential) *inaction regret.*

Why would regret-averse agents sometimes exercise recall? A regret-averse agent may use the recall option to avoid additional *inaction regret.* Suppose that a regret-averse agent rationally chose to continue searching in round t and does not find a better price in the subsequent round. Now they experience regret $R(c)$ and anticipate that not finding a better price in the next round leads to $R(2c)$. Because the regret function is convex, the (potential) increase in aversive feelings of regret is higher in this decision than in the previous decision. This may translate into a higher reservation price and a reversal of the choice to continue the search.

Post-purchase information environment

While seeing subsequent prices does not alter the utility function of pure payoff-maximizers, regret-sensitive agents are affected by this variation. Seeing subsequent prices may lead to *action regret.* Participants may blame themselves for having stopped too early when continuing the search would have yielded a higher payoff.[27] Thus, seeing subsequent prices directly affects the utility from stopping and enters the left-hand side of Equation 2. For simplification, we assume that the agent encountered the best draw in round t. We also ignore *inaction regret* because it is constant across conditions and enters the utility function independently.

The (expected) utility from stopping the search in round t while anticipating to see the next draw in case of stopping becomes $\pi(v - m_t) - \int_l^{m_t - c} R(m_t - c - x)\, \mathrm{d}F(x)$. The second term captures that regret is experienced when the price of the next draw (x_{t+1}) is lower than the previously best

[27] This entails the implicit assumption that the agent needs to see the price realization to experience *action regret* (or not), instead of incorporating expectation-based regret (without ever knowing the realization) into every decision.

price m_t and also compensates for the search cost. If one anticipates seeing all of the draws, then the feelings of regret add up to $\sum_{n=1}^{\infty}\int_l^{m_t-nc}R(m_t-nc-x)\mathrm{d}F(x)$, n denoting the (future) draws.[28]

For a regret-averse agent, the expected utility from stopping the search in t is strictly lower when additional draws are revealed after the end of the search. An agent who solves the problem based on one-step forward-induction anticipates that the same holds when stopping the search after requesting another offer ($t+1$). To avoid additional subscripts, the next offer x_{t+1} is denoted as z in the following optimization problem with *action regret*.

$$\begin{aligned}\pi(v-m_t)-\sum_{n=1}^{\infty}\int_l^{m_t-nc}R(m_t-nc-x)\mathrm{d}F(x) = &\\ [1-F(m_t)][\pi(v-m_t-c)-\sum_{n=2}^{\infty}\int_l^{m_t-nc}R(m_t-nc-x)] &\\ +\int_l^{m_t}[\pi(v-z-c)-\sum_{n=2}^{\infty}\int_l^{z-(n-1)c}R(z-(n-1)c-x)]\,\mathrm{d}F(x) &\end{aligned} \tag{3}$$

If the next draw does not yield a better price, then the probability of experiencing *action regret* when stopping the search in $t+1$ is lower than in t. This happens because future offers must also compensate for the additional search costs incurred to be advantageous. If a better offer is found in $t+1$, then the expectation of regretting the purchase at the new price is lower because it becomes less likely that future draws will yield a better payoff. Therefore, the variation in the information structure increases the (relative) attractiveness of requesting another offer and induces longer search durations for regret-sensitive agents.

[28]The upper limit of the integral changes because the likelihood of finding a more favorable offer decreases in each round as it has to compensate for all additional search costs. This is not necessary when defining R only in the positive domain. To allow for a more general definition of R, we maintain this notation. An alternative approach would be to define regret only with respect to the best forgone option. While possible, calculating the probabilities of each regret level conditional on being the highest would have been more complicated.

A.2 Additional figures and tables

Table A.1: Experienced Regret: Optimality of Search

	Forgone Profits	Optimal	Too few offers	Too many offer
	(1)	(2)	(3)	(4)
Treatments				
High-TP	.971	-.087**	.006	.081**
	[-.271,2.213]	[-.158,-.017]	[-.051,.064]	[.011,.152]
Info	.173	.016	.006	-.022
	[-1.261,1.607]	[-.056,.087]	[-.054,.066]	[-.089,.045]
High-TP X Info	-.606	.008	.009	-.017
	[-2.874,1.662]	[-.094,.109]	[-.072,.091]	[-.110,.075]
(Experienced) Inaction Regret	1.915**	-.124***	-.013	.137***
	[.399,3.430]	[-.215,-.033]	[-.080,.054]	[.057,.217]
Inaction Regret X Info	-1.505	.070	.019	-.089*
	[-3.502,.492]	[-.048,.188]	[-.077,.115]	[-.187,.010]
(Experienced) Action Regret	-1.136	.078**	.002	-.081***
	[-2.669,.397]	[.006,.151]	[-.068,.072]	[-.131,-.030]
Action Regret X Info	3.526**	-.108**	.006	.102***
	[.851,6.200]	[-.215,-.000]	[-.089,.101]	[.026,.178]
# Tasks encountered	-.200**	.023***	-.031***	.008**
	[-.371,-.029]	[.015,.031]	[-.038,-.024]	[.001,.015]
Risk Aversion	-.160	.009	.002	-.011*
	[-.569,.250]	[-.005,.022]	[-.009,.013]	[-.022,.000]
Loss Aversion	-.167	-.010	-.005	.015
	[-.640,.307]	[-.030,.011]	[-.020,.010]	[-.005,.035]
Constant	6.334**	.624***	.413***	-.037
	[.324,12.344]	[.409,.840]	[.214,.612]	[-.214,.140]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1719	1719	1719	1719

OLS Regressions.*** $p < 0.01$, **$p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. Columns (1) shows an OLS regression, estimating the forgone profits compared to the ex-ante optimal benchmark. Column (2) estimates the likelihood that search behavior was optimal (compared to the ex-ante optimal benchmark) with a (binary) OLS regression. The (binary) dependent variable takes the value 1 if the participant requested the optimal number of offers in the task and 0 otherwise. Column (3) shows the corresponding analysis with the dependent variable taking the value 1 if too few offers were requested and 0 otherwise. In Column (4), the dependent variable takes the value 1 if too many offers were requested and 0 otherwise. All columns refer to search behavior in tasks 2-10. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

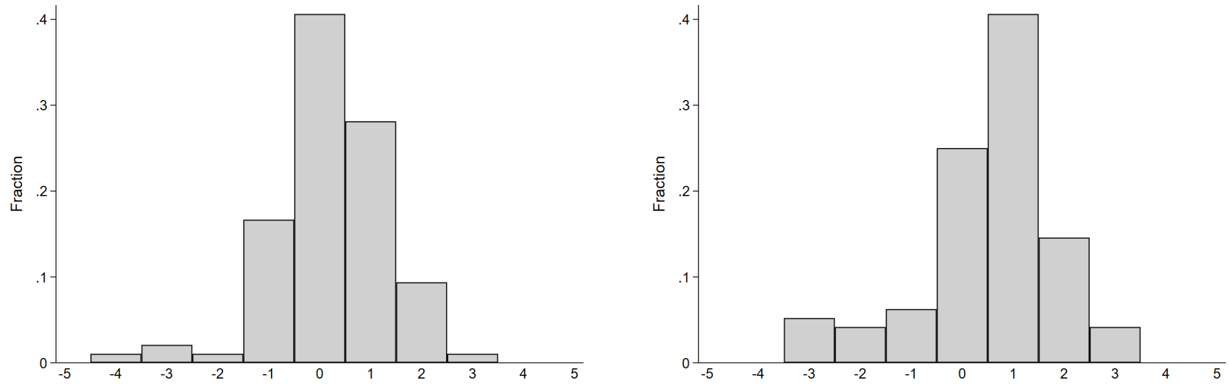

(a) *LowTP* (b) *High-TP*

Notes. The figure shows the perceived advantage of having 60 sec for each decision. Positive values indicate that the participant expected to perform better with 60 seconds than with 4 seconds. For example, a value of 1 in the left-hand panel (*Low-TP*) means that a participant expects to have scored one rank lower in the group of six if they had only had 4 seconds. In the right-hand panel (*High-TP*), a value of 1 means that a participant expects to have scored one rank higher in the group of six if they had had 60 seconds.

Figure A.1: Perceived Advantage of having 60 seconds for the decision (in ranks), by treatment assignment

Table A.2: Probit Regression: Stopping the search

	1[Stopped Search]			
	(1)	(2)	(3)	(4)
Treatments				
High-TP	-.011	-.007	.148***	.156***
	[-.049,.028]	[-.043,.028]	[.048,.249]	[.057,.256]
Info	-.001	.001	.041	.046
	[-.037,.035]	[-.032,.034]	[-.033,.115]	[-.029,.122]
High-TP X Info	.016	.012	-.132*	-.137*
	[-.039,.071]	[-.039,.063]	[-.277,.013]	[-.274,.000]
# Tasks encountered	-.006***	-.005***		
	[-.009,-.003]	[-.009,-.002]		
Price	-.009***	-.009***	-.010***	-.010***
	[-.010,-.008]	[-.010,-.008]	[-.011,-.008]	[-.012,-.008]
Risk Aversion	.003	.004	.010	.009
	[-.005,.011]	[-.003,.012]	[-.012,.032]	[-.012,.031]
Loss Aversion	-.003	.001	.027*	.032**
	[-.014,.008]	[-.009,.011]	[-.004,.058]	[.001,.063]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	No	Yes	No	Yes
Observations (# of choices)	7226	7226	622	622

Probit Regression.*** $p<0.01$, **$p<0.05$, * $p<0.1$. Standard errors clustered at the individual level. The table shows marginal effects at the mean from a probit regression. Columns (1) & (2) display search behavior across tasks 1-10, columns (3) & (4) in Task 1. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Price* is the price of the current offer [1,100] the participant faces. *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

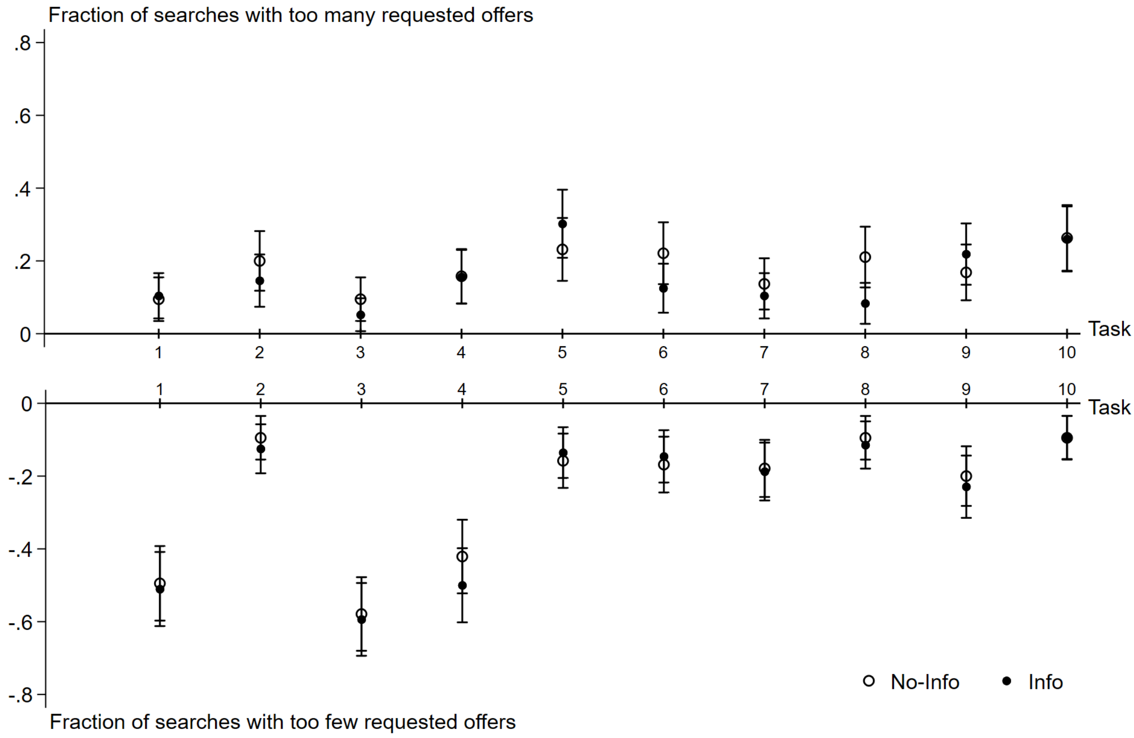

Notes. The upper panel displays the fraction of searches per task in which too many offers were requested. The lower panel shows the fraction of searches, where too few offers were requested. Larger (absolute) values correspond to higher deviations from optimal search behavior.

Figure A.2: Deviation from optimal behavior across tasks, split by *Info* condition.

A.3 Tasks 2-10

In tasks 2-10, the participants stop on average after seeing 3.84 offers, which are significantly fewer offers compared to the (ex-ante) optimal strategy of an expected payoff-maximizer, requesting 4.47 offers on average ($p < 0.001$, Wilcoxon signed-ranks test). The number of requested offers is very similar across treatments. Search length neither differs between *High-TP* and *Low-TP* ($p = 0.589$; MWU) nor between *No-Info* and *Info* ($p = 0.714$; MWU). This holds equally true when comparing treatments individually and when re-calculating the main regression outcomes for the tasks 2-10 (see Table A.3).

Table A.3: OLS Regression Search Length (Task 2-10)

	Number of offers (Task 2-10)		
	(1)	(2)	(3)
Treatments			
High-TP	.133	.200	.200
	[-.403,.669]	[-.327,.726]	[-.292,.693]
Info	-.059	-.026	-.041
	[-.573,.455]	[-.523,.471]	[-.494,.411]
High-TP X Info	-.138	-.178	-.175
	[-.854,.579]	[-.888,.531]	[-.819,.470]
# Tasks encountered	.064**	.064**	.064**
	[.009,.120]	[.009,.120]	[.009,.120]
Risk Aversion		-.041	-.066*
		[-.122,.040]	[-.144,.012]
Loss Aversion		.048	.045
		[-.095,.190]	[-.084,.174]
Constant	3.453***	4.229***	4.857***
	[2.962,3.945]	[3.139,5.319]	[3.512,6.202]
Socio-demographic controls	No	Yes	Yes
Price Sequence Group FE	No	No	Yes
Observations	1719	1719	1719

OLS Regressions.*** $p < 0.01$, **$p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching. Columns (1)-(3) display search behavior across tasks 2-10. Column (2) adds socio-demographic controls (gender, age, cognitive ability) elicited after all search tasks; column (3) and (6) additionally include price sequence group fixed effects. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 2-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

A.4 Robustness checks

A.4.1 Inclusion of unresponsive participant

In this section, we show that our main regression analyses (Table 2 and 3) are robust to the inclusion of one participant who was unresponsive to the price offers from Task 3 onward.

Table A.4: OLS Regression Search Length

	Number of offers					
	Task 1-10			Task 1		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatments						
High-TP	-.321	-.290	-.275	-.937**	-1.052***	-1.087***
	[-1.146,.505]	[-1.135,.555]	[-1.052,.502]	[-1.695,-.180]	[-1.849,-.255]	[-1.779,-.395]
Feedback	-.429	-.430	-.433	-.292	-.185	-.203
	[-1.255,.397]	[-1.311,.451]	[-1.249,.382]	[-1.147,.563]	[-1.065,.695]	[-.838,.432]
High-TP X Feedback	.310	.284	.271	.875	.938	.958
	[-.638,1.258]	[-.666,1.235]	[-.600,1.143]	[-.410,2.160]	[-.361,2.236]	[-.214,2.131]
# Tasks encountered	.090***	.090***	.090***			
	[.039,.141]	[.039,.141]	[.039,.141]			
Risk Aversion		-.059	-.088*		.004	-.079
		[-.151,.033]	[-.177,.001]		[-.174,.182]	[-.255,.097]
Loss Aversion		.034	.008		-.254*	-.232*
		[-.099,.167]	[-.114,.130]		[-.530,.022]	[-.470,.005]
Constant	3.671***	4.616***	6.054***	3.646***	5.721***	4.741***
	[2.995,4.347]	[3.452,5.781]	[3.312,8.796]	[3.054,4.237]	[3.396,8.047]	[2.271,7.211]
Socio-demographic controls	No	Yes	Yes	No	Yes	Yes
Price Sequence Group FE	No	No	Yes	No	No	Yes
Observations	1920	1920	1920	192	192	192

OLS Regressions.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching. Columns (1)-(3) display search behavior across tasks 1-10, columns (4)-(6) for the first task. Columns (1) and (4) show the effect of the treatments. Columns (2) and (5) add socio-demographic controls (gender, age, cognitive ability) elicited after all search tasks; columns (3) and (6) additionally include price sequence group fixed effects. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

Table A.5: Experienced Regret

	Number of offers			
	(1)	(2)	(3)	(4)
Treatments				
High-TP	-.190	-.196	-.174	-.185
	[-1.045,.666]	[-1.018,.626]	[-1.010,.662]	[-.988,.618]
Feedback	-.447	-.068	-.833*	-.436
	[-1.313,.419]	[-.745,.608]	[-1.812,.145]	[-1.152,.280]
High-TP X Feedback	.180	.238	.171	.214
	[-.732,1.092]	[-.681,1.157]	[-.769,1.111]	[-.699,1.128]
(Experienced) Inaction Regret	.316	1.178*		1.099
	[-.519,1.151]	[-.224,2.581]		[-.264,2.461]
Inaction Regret X Info		-1.790**		-1.715**
		[-3.189,-.390]		[-3.075,-.356]
(Experienced) Action Regret			-.913**	-.816**
			[-1.609,-.218]	[-1.442,-.189]
Action Regret X Info			1.397***	1.305***
			[.477,2.317]	[.426,2.183]
# Tasks encountered	.071**	.075***	.067**	.070**
	[.015,.127]	[.019,.131]	[.011,.122]	[.016,.124]
Risk Aversion	-.087*	-.084*	-.092**	-.087*
	[-.175,.002]	[-.172,.004]	[-.183,-.001]	[-.175,.000]
Loss Aversion	.030	.035	.037	.038
	[-.102,.162]	[-.093,.164]	[-.093,.167]	[-.088,.164]
Constant	6.299***	5.958***	6.572***	6.202***
	[3.272,9.327]	[3.162,8.754]	[3.373,9.771]	[3.319,9.085]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1728	1728	1728	1728

OLS Regressions.*** $p<0.01$, **$p<0.05$, * $p<0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching. Columns (1)-(4) display search behavior in tasks 2-10 and investigate the effect of regret experienced in the previous task. All columns include socio-demographic controls (gender, age, cognitive ability) and price sequence group fixed effects. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 2-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

A.4.2 Truncated Poisson regressions

In this section, we show that our main regression analyses (Table 2 and 3) are robust to a truncated Poisson specification.

Table A.6: Search Length

	Number of offers					
	Task 1-10			Task 1		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatments						
High-TP	.006	.021	.021	-.366**	-.400***	-.425***
	[-.131,.144]	[-.114,.156]	[-.106,.149]	[-.652,-.080]	[-.691,-.109]	[-.678,-.172]
Info	-.025	-.013	-.014	-.105	-.062	-.100
	[-.166,.116]	[-.149,.123]	[-.135,.106]	[-.381,.170]	[-.345,.222]	[-.298,.098]
High-TP X Info	-.010	-.019	-.020	.344	.367	.401*
	[-.204,.185]	[-.212,.175]	[-.195,.155]	[-.115,.802]	[-.086,.821]	[-.007,.810]
# Tasks encountered	.023***	.023***	.023***			
	[.010,.036]	[.010,.036]	[.010,.037]			
Risk Aversion		-.011	-.020*		.006	-.025
		[-.033,.012]	[-.043,.004]		[-.056,.068]	[-.088,.039]
Loss Aversion		.005	.005		-.088*	-.089**
		[-.032,.042]	[-.030,.040]		[-.177,.000]	[-.164,-.013]
Constant	1.188***	1.475***	1.580***	1.275***	1.940***	1.587***
	[1.067,1.308]	[1.157,1.794]	[1.232,1.928]	[1.096,1.454]	[1.126,2.754]	[.430,2.743]
Socio-demographic controls	No	Yes	Yes	No	Yes	Yes
Price Sequence Group FE	No	No	Yes	No	No	Yes
Observations	1910	1910	1910	191	191	191

Truncated Poisson Regressions.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching. Columns (1)-(3) display search behavior across tasks 1-10, columns (4)-(6) for the first task. Columns (1) and (4) show the effect of the treatments. Columns (2) and (5) add socio-demographic controls (gender, age, cognitive ability) elicited after all search tasks; columns (3) and (6) additionally include price sequence group fixed effects. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

Table A.7: Experienced Regret

	Number of offers			
	(1)	(2)	(3)	(4)
Treatments				
High-TP	.058	.054	.060	.057
	[-.080,.196]	[-.077,.185]	[-.071,.192]	[-.070,.184]
Info	-.009	.057	-.093	-.028
	[-.140,.122]	[-.080,.193]	[-.233,.047]	[-.170,.113]
High-TP X Info	-.049	-.035	-.057	-.043
	[-.233,.135]	[-.218,.147]	[-.241,.127]	[-.226,.141]
(Experienced) Inaction Regret	-.024	.126		.114
	[-.143,.095]	[-.029,.282]		[-.037,.265]
Inaction Regret X Info		-.328***		-.317***
		[-.537,-.118]		[-.521,-.113]
(Experienced) Action Regret			-.167**	-.158*
			[-.331,-.003]	[-.320,.005]
Action Regret X Info			.317***	.310**
			[.077,.558]	[.073,.547]
# Tasks encountered	.019**	.020**	.018**	.019**
	[.003,.035]	[.004,.036]	[.002,.033]	[.003,.035]
Risk Aversion	-.019	-.018	-.020*	-.019
	[-.042,.004]	[-.041,.005]	[-.043,.003]	[-.042,.004]
Loss Aversion	.013	.014	.014	.015
	[-.025,.051]	[-.023,.051]	[-.023,.051]	[-.022,.052]
Constant	1.592***	1.535***	1.635***	1.576***
	[1.219,1.965]	[1.158,1.911]	[1.250,2.021]	[1.183,1.968]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1719	1719	1719	1719

Truncated Poisson Regressions.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching. Columns (1)-(4) display search behavior in tasks 2-10 and investigate the effect of regret experienced in the previous task. All columns include socio-demographic controls (gender, age, cognitive ability) and price sequence group fixed effects. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 2-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

A.4.3 No switchpoint

In this section, we show that our main regression analyses (Table 2 and 3) are robust to controlling for risk attitudes and loss attitudes without by calculating the number of safe choices instead of a switchpoint.

Table A.8: Search Length

	Number of offers					
	Task 1-10			Task 1		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatments						
High-TP	.022	.098	.108	-.973**	-1.090***	-1.066***
	[-.461,.506]	[-.378,.575]	[-.340,.555]	[-1.737,-.208]	[-1.906,-.274]	[-1.764,-.368]
Info	-.086	-.041	-.056	-.327	-.188	-.190
	[-.571,.399]	[-.512,.430]	[-.471,.360]	[-1.188,.534]	[-1.079,.702]	[-.838,.457]
High-TP X Info	-.033	-.075	-.073	.910	.988	.978
	[-.704,.639]	[-.741,.590]	[-.672,.526]	[-.379,2.199]	[-.329,2.306]	[-.201,2.158]
# Tasks encountered	.079***	.079***	.079***			
	[.032,.125]	[.032,.125]	[.032,.125]			
Risk Aversion		-.044	-.077		-.018	-.105
		[-.142,.054]	[-.169,.016]		[-.226,.189]	[-.314,.104]
Loss Aversion		.044	.047		-.158	-.104
		[-.104,.192]	[-.078,.171]		[-.462,.146]	[-.358,.150]
Constant	3.391***	4.251***	4.691***	3.681***	5.324***	4.157***
	[2.988,3.793]	[3.280,5.223]	[3.500,5.883]	[3.081,4.281]	[3.190,7.457]	[1.770,6.545]
Socio-demographic controls	No	Yes	Yes	No	Yes	Yes
Price Sequence Group FE	No	No	Yes	No	No	Yes
Observations	1910	1910	1910	191	191	191

OLS Regressions.*** p <0.01, **p <0.05, * p <0.1. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching.Columns (1)-(3) display search behavior across tasks 1-10, columns (4)-(6) for the first task. Columns (1) and (4) show the effect of the treatments. Columns (2) and (5) add socio-demographic controls (gender, age, cognitive ability) elicited after all search tasks; columns (3) and (6) additionally include price sequence group fixed effects. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as the number of safe choices, as described in Footnote 17.

Table A.9: Experienced Regret

	Number of offers			
	(1)	(2)	(3)	(4)
Treatments				
High-TP	.241	.230	.252	.245
	[-.256,.738]	[-.239,.700]	[-.224,.727]	[-.213,.703]
Info	-.043	.192	-.335	-.105
	[-.495,.410]	[-.282,.666]	[-.827,.156]	[-.600,.390]
High-TP X Info	-.188	-.146	-.217	-.173
	[-.833,.457]	[-.781,.490]	[-.860,.426]	[-.814,.469]
(Experienced) Inaction Regret	-.080	.478		.425
	[-.494,.334]	[-.121,1.077]		[-.156,1.006]
Inaction Regret X Info		-1.140***		-1.092***
		[-1.890,-.390]		[-1.822,-.362]
(Experienced) Action Regret			-.556*	-.516*
			[-1.115,.003]	[-1.066,.034]
Action Regret X Info			1.101**	1.066**
			[.246,1.957]	[.224,1.909]
# Tasks encountered	.065**	.068**	.061**	.065**
	[.009,.121]	[.012,.123]	[.006,.116]	[.010,.120]
Risk Aversion	-.074	-.073	-.078*	-.078*
	[-.166,.018]	[-.166,.020]	[-.169,.013]	[-.171,.015]
Loss Aversion	.064	.069	.070	.077
	[-.072,.200]	[-.066,.204]	[-.062,.203]	[-.057,.210]
Constant	4.847***	4.656***	5.004***	4.810***
	[3.527,6.167]	[3.322,5.990]	[3.622,6.386]	[3.405,6.215]
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1719	1719	1719	1719

OLS Regressions.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. The dependent variable is a count variable, representing the number of offers after which the participant stopped searching.Columns (1)-(4) display search behavior in tasks 2-10 and investigate the effect of regret experienced in the previous task. All columns include socio-demographic controls (gender, age, cognitive ability) and price sequence group fixed effects. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 2-10). *Risk Aversion* and *Loss Aversion* are defined as the number of safe choices, as described in Footnote 17.

A.4.4 Probit regression: Optimality after experienced regret

This section shows that Table A.1 is robust to a probit specification in Columns (2)-(4).

Table A.10: Probit Regression: Stopping the search

	Forgone Profits	Optimal	Too few offers	Too many offers
	(1)	(2)	(3)	(4)
Treatments				
High-TP	.971	-.090**	.009	.075**
	[-.271,2.213]	[-.162,-.018]	[-.051,.069]	[.010,.139]
Info	.173	.016	.006	-.026
	[-1.261,1.607]	[-.060,.091]	[-.057,.069]	[-.100,.047]
High-TP X Info	-.606	.009	.005	-.009
	[-2.874,1.662]	[-.095,.113]	[-.080,.090]	[-.098,.080]
(Experienced) Inaction Regret	1.915**	-.125***	-.019	.111***
	[.399,3.430]	[-.216,-.034]	[-.090,.051]	[.050,.172]
Inaction Regret X Info	-1.505	.069	.028	-.062
	[-3.502,.492]	[-.049,.187]	[-.071,.126]	[-.141,.017]
(Experienced) Action Regret	-1.136	.080**	.007	-.087***
	[-2.669,.397]	[.005,.155]	[-.058,.073]	[-.146,-.028]
Action Regret X Info	3.526**	-.110**	.005	.110***
	[.851,6.200]	[-.219,-.001]	[-.085,.095]	[.029,.191]
# Tasks encountered	-.200**	.023***	-.031***	.009**
	[-.371,-.029]	[.015,.031]	[-.039,-.024]	[.002,.015]
Risk Aversion	-.160	.009	.003	-.010*
	[-.569,.250]	[-.005,.023]	[-.009,.015]	[-.020,.001]
Loss Aversion	-.167	-.010	-.007	.014
	[-.640,.307]	[-.030,.011]	[-.023,.009]	[-.004,.033]
Constant	6.334**			
	[.324,12.344]			
Socio-demographic controls	Yes	Yes	Yes	Yes
Price Sequence Group FE	Yes	Yes	Yes	Yes
Observations	1719	1719	1719	1719

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the individual level. The values in square brackets represent the 95% confidence intervals. Columns (1) shows an OLS regression, estimating the forgone profits compared to the ex-ante optimal benchmark. Column (2) estimates the likelihood that search behavior was optimal (compared to the ex-ante optimal benchmark) with a probit regression. The (binary) dependent variable takes the value 1 if the participant requested the optimal number of offers in the task and 0 otherwise. Column (3) shows the corresponding analysis with the dependent variable taking the value 1 if too few offers were requested and 0 otherwise. In Column (4), the dependent variable takes the value 1 if too many offers were requested and 0 otherwise. Columns (2)-(4) show marginal effects at the mean. All columns refer to search behavior in tasks 2-10. *(Experienced) Inaction Regret* is an indicator variable, taking a value of 1 if the participant experienced *inaction regret* in the previous task. *Inaction Regret X Info* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *Info*. *(Experienced) Action Regret* and *Action Regret X Info* are defined accordingly. *# Tasks encountered* is a count variable, indicating the number of the current task (Task 1-10). *Risk Aversion* and *Loss Aversion* are defined as switching points, as described in Footnote 17.

A.5 Instructions

Appendix A.5 includes the translated instructions (from German). The participants received the instructions for the experiment in print. Additional short instructions and control questions were later displayed on the computer screen. Treatment specific parts are shown in *italics* and the corresponding treatment clearly indicated.

Welcome to the experiment and thank you for your participation!

Please do not speak from now on with any other participant

General Procedures

In this experiment, we study economic decision-making. You can earn money by participating. The money you earn will be paid to you privately and in cash after the experiment. The experiment lasts for around 60 minutes and consists of multiple parts (the exact number of parts is unknown to all participants). At the beginning of every part, you receive detailed instructions. If you have questions after reading the instructions or during the experiment, please raise your hand or press the red button on your keyboard. One of the experimenters will then come to you and answer your question(s) privately.

Important: Depending on the decision, you will see an expiring clock at two different places on the screen. If you see the clock with the tag "Remaining time" in the center of the screen it indicates how much time you have for the decision. Further information will be provided in the instructions. During other decisions, you will see a (small) expiring clock at the right-upper part of the screen. This time only gives you an indication, how long the current decision should take. You can also take more time if you need it. Entering a decision is also possible before time expires.

Anonymity

The analysis of the experiment is anonymous; that is, we will never link your name with the data generated in the experiment. At the end of the experiment, you will be asked to sign a receipt to confirm the payments you received. This receipt will only be used for accounting purposes. No further personal data will be passed on.

Tools

You find a pen at your desk. Please leave the pen and the instructions on the table after the experiment.

Payment

In addition to the income that you earn during the experiment, you will receive 6 € for showing up on time. During the experiment, we do not talk about Euro, but about Taler. We convert the Taler into Euros at the end of the experiment and pay those in addition to the 6 € for your punctual appearance in cash.

Procedure

This experiment consists of **multiple decisions on the purchase of a fictitious product**. In the following, the rules that determine the payoff from your decisions, are explained in detail. At the end of the experiment, one of the buying decisions will be randomly chosen and you receive the corresponding payoff. Every purchase decision is equally likely to be randomly chosen.

After the purchase decisions, you can earn additional money through correct assessments and further decisions.

Following this, we will ask you to respond to a few questions conscientiously. After that, the experiment ends. You will then receive the money that you earned through your decisions, as well as 6€ in cash for your punctual appearance.

Exchange rate in the purchasing decisions

In some parts of the experiment, we do not task about Euros, instead we refer to Taler. These will be converted into Euros at the end of the experiment. Please note the following exchange rate:

100 Taler = 12 €

Your task

The experiment has several tasks. In every task, the objective is to obtain as many Taler as possible through the purchase of a fictitious product. In general, a task proceeds as follows.

In every task, the number of Taler you receive from a purchase decision is calculated as the difference between the value of the product and the costs that you incur through making the purchase.

Taler from the purchasing decision = Value of the product – Price – Cost for price offers

Value of the product

The **product is worth 50 Taler for you.**

When you buy the product, **you receive 50 Taler**. At the same time, you have to pay a price for the purchase of the product.

Price of the product and cost for the price offers

The computer offers the product to you by displaying a purchase price, at which you can buy the product. You can then decide whether you want to request another offer in the form of a new purchase price or whether you want to buy the product for the lowest purchase price offered so far. You can request as many offers as you want (as long as there is a possibility to achieve a positive payoff under any search cost). However, every offer you request is associated with a cost for you:

Every offer you request costs a fixed amount of Taler.

In the following, **these costs will be called search costs**. The search cost can vary across tasks. You will know the exact cost level before each purchase decision.

You **can always buy the product at the lowest standing offer** (even if you have requested additional offers that might have been higher). Therefore, amount of Taler you receive from a purchase decision is

50 – (lowest price received) – search cost*(number of offers you requested).

Accordingly, the amount of Taler you receive is higher when the price at which you purchase the product is lower. The amount of Taler decreases by the amount of search cost with every offer you request. (For the first, automatically displayed offer, you do **not** incur any costs.)

Time for the decision

You only have limited time to make your decision. After every offer you have *60 seconds [Low-TP/No-Info and Low-TP/Info]/ 4 seconds [High-TP/No-Info and High-TP/Info]* to decide whether you want to buy for the best price observed so far or whether you want to request another offer. If you neither decide to buy the product nor request an additional offer, we will deduct 1 Taler from your payoff in this task. Afterward, you have an additional *60 seconds [Low-TP/No-Info and Low-TP/Info]/ 4 seconds [High-TP/No-Info and High-TP/Info]* to make the decision (purchasing vs. requesting another offer). If you do not decide within that time once again, you will be again deducted 1 Taler in this task. This procedure is repeated until you make a decision.

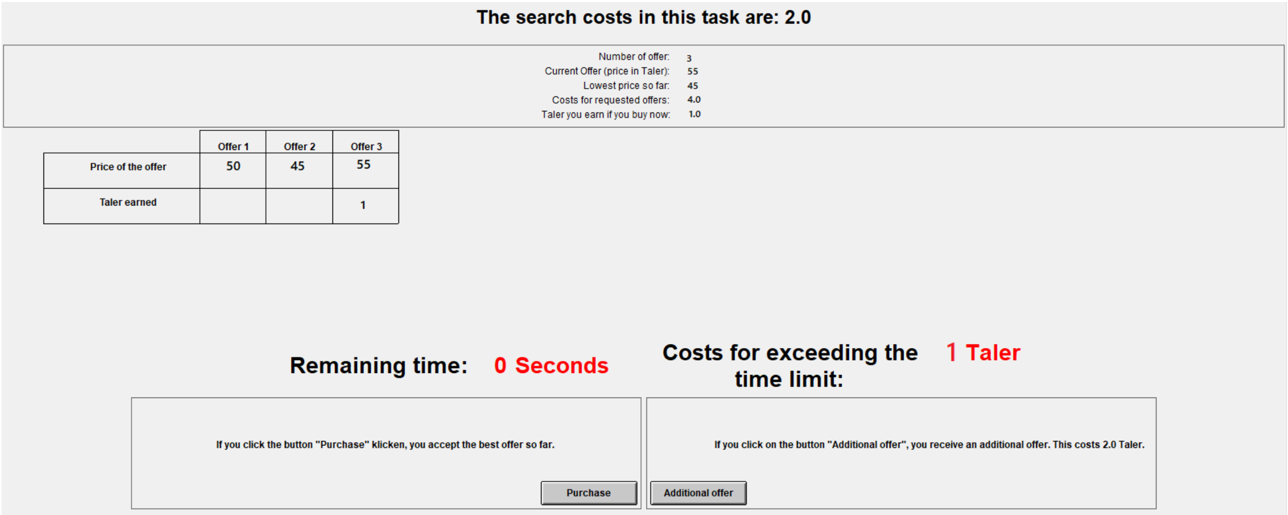

Information on the offers of the computer

The price offers of the computer are integers and can take the values 1, 2, 3… to 100 Taler. The computer draws each price independently and randomly with the same probability of 1% (draws with replacement). You can imagine the procedure like this: an urn contains 100 balls, which are numbered from 1 to 100. At each offer, the computer draws one of those balls, displays the number on the ball as a price offer, and puts the ball back into the urn, such that each ball in the next draw will be again drawn with a probability of 1%.

On-screen procedure

To illustrate the decision screen, below you can see an example of a task, where—in addition to the first offer of the computer (price of 50)—two more offers were requested:
In the upper part, you see the search cost for this task. Below you see how many offers are already displayed, as well as which offer is the current offer and which is the best one. Additionally, you see the costs that have to be paid for the offers requested so far.

In the lower part you make your purchase decision. To accept the best offer so far, you click on the button: "Buy". To request another offer and incur the above-displayed search cost, you click on the button: "Additional offer". In the central part, you see an overview of the offers received so far, as well as your current payoff for the task if you click "Buy."

In the displayed example, the first offer was equal to 50 Taler. Because the product is worth 50 Taler, buying the product at this price would have resulted in a payoff of 0 Taler in this task. In the example, we assumed, that another offer was requested at the (search) cost of 2 Taler.

The second price offered to you, was 45 Taler in the example. Deciding to buy at this offer would have led to receiving the product at the lowest price so far observed (i.e., 45 Taler). Hence, your payoffs would have been determined as follows:

Received Taler	= value of the product	− lowest price	− search cost (2 Taler for each requested offer)	
.	= 50	− 45	− 2	= **3**

In the example, we assumed that another offer at the cost of 2 Taler was requested. This time, the randomly drawn price was 55 Taler. If you decided to purchase the product at this point within the remaining time, then you would receive 1 Taler for this task (as you can always purchase the product for the lowest price seen so far):

Received Taler	= value of the product	− lowest price	− search cost (2 Taler for each requested offer)	
.	= 50	− 45	− 2*2	= **3**

If you instead requested another offer, then you would incur the cost of 2 Taler again and the computer would display an additional randomly drawn price.

Beneath the offers seen so far, you see the "Remaining Time" for the decision. This shows how much time you have remaining to decide between "Buy" and "Additional offer". On the right-hand side, you see how many Taler were already deducted from your payoff due to exceeding the time limit in this task.

In the example, we assumed that the decision time has just expired, such that an additional cost of 1 Taler through exceeding the time limit has to be paid. After the expiration of the decision time, the "Remaining time" further runs down. Should you decide to buy the product after offer 3 in the next 60 seconds, you receive 0 Taler in this task. Should you request another offer within this time, then you pay the search cost of 2 Taler and the computer displays an additional randomly drawn price. Should you neither buy the product nor request another offer within the next 60 seconds, you incur a cost of 1 Taler again. This procedure is repeated until you make a decision.

Note

In every task it is possible, that you receive a negative payoff. If this task is drawn as payoff relevant, this loss will be offset by your payoff from the other parts of the experiment.

Procedure

After every purchase decision, you will see all the offers until your purchase decision once again. *Furthermore, you see additional offers, which would have been displayed to you later, if you had not*

made a purchase decision at that point. This means, you will see whether requesting an or multiple additional offers would have yielded more (or less) Taler. [only in Low-TP/Info and High-TP/Info]
To conclude the task, please type in the number of the offer, at which you would have received the highest payoff. After the purchasing decisions, you will be additionally asked for assessments of your own behavior and you will be asked to make additional decisions, with which you can earn or lose money. At the end of the experiment, you see your payoff on a separate screen. You will also be shown, which of the purchasing decisions has been randomly drawn to be relevant for your payoff.

Comprehension questions

To verify your understanding of the task and the payoff scheme, you will be confronted with some control questions before the purchasing decisions start. The first purchasing decision starts when all participants have answered the questions correctly. Important: Your answers to the comprehension questions do not affect your payoff.