
Time Pressure and Regret in Sequential Search

Felix Klimm (LMU Munich and CESifo)
Martin G. Kocher (University of Vienna, CESifo and University of Gothenburg)
Timm Opitz (MPI and LMU Munich)
Simeon Schudy (LMU Munich and CESifo)

Discussion Paper No. 359

December 27, 2022

Time Pressure and Regret in Sequential Search

Felix Klimm¹ Martin G. Kocher² Timm Opitz^{1,3} Simeon Schudy¹

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. Empirically, we find that anticipated regret does not affect search behavior either with or without time pressure, while experienced regret leads to systematic adjustments in search length. Urgency reduces decision times and perceived decision quality, but does not generally alter search length. Only very inexperienced decision-makers buy earlier when pressured. Thus, consumer protection measures against pressure selling tactics can help inexperienced consumers in particular.

JEL codes: C91, D01, D03, D18, D83

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

¹ Department of Economics, University of Munich (LMU), Geschwister-Scholl-Platz 1, 80539 Munich, Germany [felix.klimm@econ.lmu.de; simeon.schudy@econ.lmu.de]

² Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1080 Vienna, Austria; University of Gothenburg, 405 30 Gothenburg, Sweden [martin.kocher@univie.ac.at]

³ Max Planck Institute for Innovation and Competition, Marstallplatz 1, 80539 Munich, Germany [timm.opitz@ip.mpg.de]

We thank Marvin Deversi, Paul Heidhues, Fabian Herweg, Marieke Huysentruyt, Nicolas Jacquemet, Heiko Karle, Matthias Lang, Robin Mamrak, Franz Ostrizek, 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, the CRC TRR 190 Workshop 2020, and the Vfs Annual Meeting 2021 for their helpful thoughts and comments. Financial support by the German Research Foundation (DFG) through CRC TRR 190 (project number 280092119) is gratefully acknowledged. Declarations of interest: none.

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

¹In 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).

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

³The 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, 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 role of experienced *action regret*. As buyers are also often pressured time-

⁴As 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).

⁵This 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).

wise, we further study how perceived urgency alters search behavior and the role of regret. We implement a 2x2 between-subjects design with high or low perceived urgency that avoids potential selection bias due to time pressure, and vary search costs (within-subjects) to analyze the extent to which participants understand the general logic of the reservation-price strategy.

Our empirical results confirm stylized facts from previous experiments, as in all treatments, participants search on average too little (as compared to the expected payoff-maximizing strategy), make use of the recall option, and search longer with lower search costs and more experience. In our main analyses, we study the causal role of perceived urgency, regret, and their interaction for search behavior. 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 stopping too early) does not increase search length, while 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 reduces search length. These adjustments do not increase payoffs substantially, as some participants over-adjust their search length. Finally, perceived urgency does not substantially alter the observed role of regret.⁶ In addition to our main analysis, our study highlights the need for strategies consumers may employ to protect themselves from searching sub-optimally. Thus, we also discuss commitment to reservation prices as a simple strategy that may circumvent inefficient search and provide empirical evidence showing that such commitment can indeed improve the optimality of search and results in larger payoffs.

The rest of this manuscript is organized as follows. Section 2 discusses the three-fold contribution of our approach (i.e., understanding the role of time-pressure, regret, and their potential interaction in sequential search tasks) relative to the existing literature. In Section 3, we explain the experimental design. In Section 4, we specify theory-based hypotheses, which we test in our main empirical analyses in Section 5. In Section 6, we discuss our findings and their robustness. Section 7 concludes.

2 Related Literature

Our search design builds on classical search experiments (e.g., Houser & Winter, 2004; Schotter & Braunstein, 1981; Sonnemans, 1998) which revealed two commonly observed anomalies in sequential

⁶Importantly, our experiment allows to identify economically relevant effect sizes (i.e., larger than 0.20 standard deviations, for more details, see also Section 7.

search problems: early stopping and recall. Our experimental treatment variations complement and advance earlier experimental findings on active sequential search under conditions with or without perceived urgency as well as with or without post-purchase price information.

Our analyses on perceived urgency extends earlier experimental findings on time pressure in sequential search environments that excluded post-purchase price information. Ibanez et al. (2009) document inefficiently short search patterns for inexperienced decision-makers under mild time constraints (without post-purchase feedback). Through making deliberation more costly in our design, we confirm that time pressure substantially reduces payoffs with inexperienced decision-makers both with and without post-purchase price information. These 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).

Regarding the study of anticipated *action regret*, our approach links to that has studied the effects of post-purchase information on search behavior. Sugden et al. (2019) study whether time-limited offers are chosen more often without post-purchase information, finding no evidence of regret effects. In contrast, we focus on how feedback structures and perceived time pressure affect the number of requested (ex-ante identical) offers. In line with 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. Our findings further complement important recent evidence on the search-enhancing effect of anticipated *action regret* when decision makers search through repeatedly stating reservation prices and post-purchase information only includes (potentially) better offers (Jhunhunwala, 2021). Relating to this work, we provide evidence from additional experimental treatments (see Section 6.3) which underscores the critical role of the nature of post-purchase information which may generate behavioral changes through anticipated regret.

More generally, our results regarding anticipated and experienced regret relate to the broader literature on optimal stopping problems. Strack and Viefers (2021) demonstrate regret sensitivity in an asset-selling task 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.⁷ Our analyses also link to work by Fioretti, Vostroknutov,

⁷Our 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*.

and Coricelli (2022), who vary (within-subject) post-purchase information in a setting akin to Strack and Viefers (2021) and find –consistent with our theoretical predictions– that participants stop later when they may anticipate *action regret*.⁸ 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. These changes may render the role of regret less salient in our setting. On the other hand, avoiding *action regret* may be perceived as less costly in Fioretti et al. (2022). The stochastic mean-reverting process that determines the prices in Fioretti et al. (2022) leads to a multimodal distribution of prices (and payoffs) over time. Thus, it becomes likely that participants encounter similar payoffs in the future, even when not selling early on. As the cost of delaying the purchase in early periods becomes less costly, participants may stop later and at the same time achieve similar payoffs while reducing the probability of *action regret*.

Finally, our setup allows us to study the role of experienced regret which may induce learning across time (see e.g., Cooke, Meyvis, & Schwartz, 2001; Einav, 2005; Oprea, Friedman, & Anderson, 2009; Sonnemans, 1998). Sonnemans (1998) (Experiment 2) shows that participants change their reservation prices after learning that they searched too long. Similarly, participants converge faster to an optimal reservation price in a search task with pre-commitment when receiving post-purchase feedback (Einav, 2005). 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. 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. In general, we find that the fraction of searches that are too long remains constant across time while the fraction of searches that are too short decreases within the first half of the experiment, thereby reducing inefficiencies to some extent. Experienced (action and inaction) regret alters search length in our setting systematically. In particular, participants in the treatment condition with post-purchase information increase (decrease) search length after experiencing action (inaction) regret. However, such learning from experienced regret does not translate into higher levels of efficiency, presumably because participants face different search costs and prices across search tasks, rendering profitable adjustments more complex. Finally, experiencing *action regret* from searching to little does not reinforce anticipated regret. That is, differences in search lengths across feedback conditions do not substantially change across the 10 search tasks.

⁸Note that contrary to classical experimental search tasks, the environment of Fioretti et al. (2022) already leads to longer search than theoretically predicted in the condition without post-purchase feedback, while the classical anomaly in search tasks goes in the opposite direction compared to the rational benchmark.

3 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). 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. At the end of the experiment, one of the 12 search tasks is randomly drawn to be payoff relevant. Figure 1 summarizes the experimental procedures, showing the different parts of the experiment. To avoid unwanted effects of anticipating the content of subsequent parts, we inform participants only at the beginning of each part about its content. Further, participants of the subject pool are aware that they receive a flat payment of 6 Euro and that they can make losses during some parts of the experiment which will be compensated by the 6 Euro flat payment and potential earnings from other experimental parts. For example, given the nature of the search task in our experiment, participants could encounter losses in Tasks 1-12, if they decided to pay a price higher than their valuation or when searching too long and thus incurring search costs larger than the gains from trade.⁹

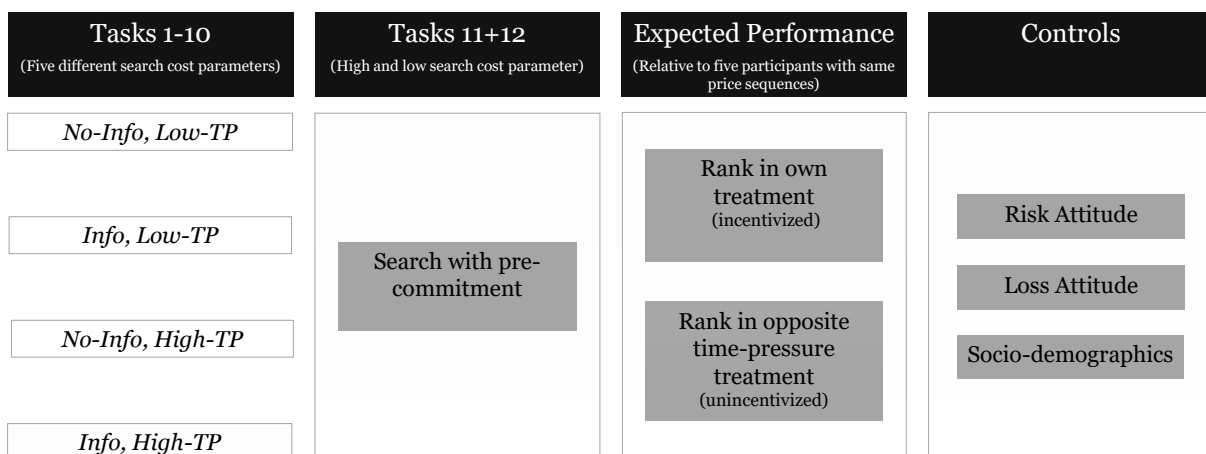


Figure 1: Experimental Design

⁹All participants received positive payoffs and this procedure does not alter our theoretical predictions.

3.1 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).¹⁰ The participants' induced value for the good is $v = 50$ and stays the same across all search tasks. 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 stays constant within each of the ten search tasks (but varies across tasks) such that participants are aware about their search cost when deciding upon an additional price request. Price offers are drawn from the known uniform distribution $\{1, 2, \dots, 100\}$.¹¹ We inform the participants that they are free to request new offers as long as there is a possibility to achieve a positive payoff given search costs. This renders the search process finite (because participants can request at most 24 additional offers before making a loss for sure given our parameter values, although we do not state the exact number of possible requests to participants).¹² After purchasing the product the current search task ends and participants proceeded with the next search task.

3.2 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. To ensure that perceived urgency can affect search behavior also in later tasks, we vary the theoretically optimal reservation price strategy by altering search costs between the tasks. We use five different values for the search cost $c \in \{2, 2.5, 3, 3.5, 4\}$. Each parameter value occurs twice and 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.

¹⁰With 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.

¹¹We thereby rely on the parametrization of Sonnemans (1998).

¹²Only 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.

3.3 Experimental treatments

3.3.1 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 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 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.

3.3.2 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 *No-Info* participants are informed that they see only those prices that they actively requested until they purchase the product. In treatment *Info*, the participants are informed that they will see additional price offers for which they could have bought the product, after purchasing it. We randomly determine the number of displayed offers $k \leq n$ where $n = 25 - \text{OfferNumber}_{\text{accepted}}$, such that (for example) a participant who decides to buy after seeing five offers can see between 1 and 20 additional prices. This design feature renders learning about the maximum possible search length similar in both treatments. By varying the availability of 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., 2022; Jhunjhunwala, 2021; Sugden et al., 2019; Zeelenberg, 1999). This anticipation can be reinforced, when experiencing *action regret* in *Info* in previous tasks. Because we randomize the number of additional prizes displayed, we vary whether participants experience regret given the same search behavior and price sequence. This allows us to analyze the effect of experienced action regret both within the *Info* treatment and across treatments, and disentangles potential effects of simply seeing additional prizes (e.g., by

familiarizing oneself with the random process of price draws) as compared to experiencing regret due observing particularly attractive prices.

3.4 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 also allows us to cleanly test for the participants' responsiveness to the search costs.

3.5 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.

3.6 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 (2022), 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).

3.7 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.¹³ 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.

4 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.

4.1 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.*

¹³We 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.

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* theoretically 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 directional learning, and systematically influences search behavior in subsequent tasks. Decisions in repeated search tasks may reflect the experience of regret in the previous task, translating into higher awareness and sensitivity to anticipated regret. 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 k leads to a lower number of requested offers in task $k + 1$ in treatments Info and No-Info.*

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

4.2 Time pressure

Perceived urgency has been found to reduce the depth of reasoning and alter information processing (Kocher & Sutter, 2006; Payne, Bettman, & Luce, 1996). Altering participants' optimization process, perceived urgency may thus result in shorter or longer search length. The observation that sellers

use practices that create a sense of urgency suggests a reduction in search length as higher accepted prices benefit sellers. Participants may also tend to accept current offers more frequently when they perceive pressure and thus consider the *High-TP* decision environment to be aversive. At the same time, time pressure may impair the availability of cognitive resources and thus render the consideration of additional psychological factors less likely. If these are the reason for (inefficiently) short search, time pressure may increase search length. Further, if participants rely increasingly on decision heuristics under time pressure (e.g., Finucane, Alhakami, Slovic, & Johnson, 2000; Gigerenzer & Todd, 1999), search length may increase or decrease (depending on the decision heuristic). 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.*

4.3 Potential interaction of time pressure and regret

Building on the idea that time pressure renders 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. The lower availability of cognitive resources leads to regret being less relevant for the decision. 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.*

5 Main results

5.1 Search behavior without feedback

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

will buy the good (see also Appendix A.1).¹⁴ 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 substantially shorter search lengths (see Table 1). Relating to earlier literature, we first focus on the standard environment without information about future prices (and discuss potential treatment differences in Section 5.3). Participants stopped on average after seeing 3.83 offers when receiving no information on post-purchase prices (pooled across both time pressure conditions; $p < 0.001$, Wilcoxon signed-ranks test).¹⁵ This result also holds when analyzing each time pressure condition individually (*No-Info/Low-TP*: $p < 0.001$, *No-Info/High-TP*: $p < 0.001$, Wilcoxon signed-ranks tests). The search length corresponds to an average accepted price of 16.59. Consequently, the participants also earned around 11 percent less than the expected payoff-maximizer would obtain ($p < 0.001$, Wilcoxon signed-ranks test). Furthermore, in a substantial fraction of searches (18.84 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 78.95 percent of participants do so at least once in the experiment. Recall rates do not differ statistically significantly across time pressure conditions in the standard search environment (17.02 percent in *No-Info/Low-TP*, 20.63 percent in *No-Info/High-TP*; $p = 0.435$, MWU). 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.*

Table 1: Descriptive statistics on search behavior

| | Search Length | | | | Accepted Price | n |
|------------------------|---------------|-----|-----|------|----------------|-----|
| | Mean | Min | Max | SD | Mean | |
| <i>No-Info/Low-TP</i> | 3.82 | 1 | 25 | 3.11 | 15.75 | 470 |
| <i>No-Info/High-TP</i> | 3.85 | 1 | 24 | 3.24 | 17.42 | 480 |
| <i>Info/Low-TP</i> | 3.74 | 1 | 25 | 3.40 | 16.44 | 480 |
| <i>Info/High-TP</i> | 3.73 | 1 | 21 | 2.97 | 17.91 | 480 |

This table shows descriptive statistics on search behavior across the four treatments. *Mean*, *Min*, *Max*, *SD* denote the mean, the minimum, the maximum, and the standard deviation, respectively. *n* denotes the number of observations.


¹⁴Depending on the search costs, the reservation price is between 20 and 29 for an expected payoff-maximizer given our parametrization.

¹⁵For the non-parametric tests, we form within-subject averages across the respective tasks so that we consider one data point per individual. All of the reported non-parametric tests in the analysis are two-sided hypothesis tests.

5.2 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.¹⁶

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, 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 2 highlight that decision times are substantially and statistically significantly shorter in *High-TP* than *Low-TP* in all sequential search tasks (pooled across both feedback conditions)¹⁷ 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.



Notes. The error bars indicate 95% confidence intervals.

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

¹⁶Relative to the average earning in the search task, transgressing the limit once compares to a decrease in earnings of around 4 percent.

¹⁷Table A.1 corroborates that the decision times significantly decrease in both feedback conditions.

Table 2: Average decision times per task across time pressure conditions

| Task | per Offer | | per Subject | | p-value |
|------|---------------|----------------|---------------|----------------|---------|
| | <i>Low-TP</i> | <i>High-TP</i> | <i>Low-TP</i> | <i>High-TP</i> | |
| 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.

5.3 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 (see Section 5.5).¹⁸

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* (pooling in terms of Info, $p = 0.750$, MWU) nor between *No-Info* and *Info* (pooling in terms of time-pressure, $p = 0.646$, MWU). The same holds when comparing treatments individually instead of pooling them (see Table 1). Time pressure neither changes the number of requested offers without ($p = 0.941$, MWU) nor with feedback ($p = 0.575$, MWU); the feedback structure neither affects average search length without ($p = 0.451$, MWU) nor with time pressure ($p = 0.967$, MWU). 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.

¹⁸For 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.



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 3; Columns (1)-(3)). In Column (1), we assess the treatment effect, controlling for the number of tasks a decision-maker already completed. In Column (2), we add demographic controls, as well as measures of risk and loss attitudes.¹⁹ 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 Appendix Table A.4, Columns 1 and 2). 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

¹⁹Calculating 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., 2022), 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. (2022) 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, in press). In the main regressions of Tables 3 and 4, 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 3: Search Length

| | Number of offers | | | | | |
|----------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-----------------------------|-----------------------------|
| | Task 1-10 | | | Task 1 | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| High-TP | .022 [-.461,.506] | .072 [-.405,.549] | .071 [-.378,.519] | -.973** [-1.737,-.208] | -1.076*** [-1.884,-.268] | -1.097*** [-1.796,-.397] |
| Info | -.086 [-.571,.399] | -.045 [-.515,.425] | -.059 [-.474,.357] | -.327 [-1.188,.534] | -.211 [-1.103,.682] | -.214 [-.860,.432] |
| High-TP X Info | -.033 [-.704,.639] | -.064 [-.732,.603] | -.060 [-.663,.542] | .910 [-.379,2.199] | .961 [-.344,2.266] | .968 [-.204,2.140] |
| # Tasks encountered | .079*** [.032,.125] | .079*** [.032,.125] | .079*** [.032,.125] | | | |
| Risk Aversion | | -.036 [-.117,.044] | -.067* [-.145,.011] | | .002 [-.176,.181] | -.080 [-.256,.096] |
| Loss Aversion | | .017 [-.110,.145] | .017 [-.100,.134] | | -.253* [-.530,.024] | -.233* [-.470,.005] |
| Constant | 3.391*** [2.988,3.793] | 4.295*** [3.301,5.289] | 4.764*** [3.549,5.979] | 3.681*** [3.081,4.281] | 5.747*** [3.414,8.081] | 4.780*** [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 19.

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.*

5.4 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 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).²⁰ 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 4, Column 1), we find evidence that people in *Info* systematically react to the information provided as specified in Hypotheses 3 and 5 (see Table 4, Column 2). Knowing that one should have requested fewer offers in task k results in requesting

²⁰Because of this very similar efficiency across both feedback conditions, the lower efficiency under *High-TP* holds both without feedback ($p = 0.028$) and with feedback ($p = 0.016$).

around 1.14 offers less in task $k + 1$ 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). That is, *inaction regret* (although possible in both treatments) affects subsequent behavior only in *Info*. This finding appears surprising but is consistent with an increased awareness towards regret feelings in general due to feedback provision in *Info*. In line with this idea, we observe that participants spend around 30% more time on the feedback screen in *Info* than in *No-Info* when experiencing *inaction regret* (*Info*: 22.37 seconds, *No-Info*: 17.17 seconds). Further, seeing additional prices in *Info* may reinforce *inaction regret* when the additional prices shown are inferior to the accepted price. We summarize this finding in Result 3:

Result 3. *Experiencing inaction regret in task k leads to a lower number of requested offers in task $k + 1$ 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 $k + 1$, 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

Table 4: Experienced regret

| | Number of offers | | | | |
|-------------------------------|---------------------------|-----------------------------|---------------------------|-----------------------------|-----------------------------|
| | (1) | (2) | (3) | (4) | (5) |
| Treatments | | | | | |
| High-TP | .203 [-.295,.700] | .192 [-.280,.664] | .211 [-.265,.687] | .203 [-.257,.663] | .432 [-.103,.967] |
| Info | -.043 [-.497,.410] | .189 [-.285,.664] | -.335 [-.828,.157] | -.106 [-.602,.389] | -.115 [-.610,.380] |
| High-TP X Info | -.172 [-.821,.477] | -.129 [-.769,.511] | -.200 [-.847,.447] | -.155 [-.800,.491] | -.135 [-.771,.500] |
| (Experienced) Inaction Regret | -.082 [-.497,.332] | .473 [-.127,1.073] | | .420 [-.161,1.002] | .470 [-.228,1.168] |
| Inaction Regret X Info | | -1.135*** [-1.885,-.384] | | -1.086*** [-1.816,-.356] | -1.082*** [-1.815,-.349] |
| Inaction Regret X High-TP | | | | | -.104 [-.858,.650] |
| (Experienced) Action Regret | | | -.553* [-1.111,.006] | -.513* [-1.062,.036] | -.124 [-.806,.558] |
| Action Regret X Info | | | 1.095** [.239,1.951] | 1.060** [.217,1.903] | 1.058** [.225,1.891] |
| Action Regret X High-TP | | | | | -.757* [-1.598,.084] |
| # Tasks encountered | .065** [.009,.121] | .068** [.012,.123] | .061** [.006,.116] | .065** [.010,.120] | .064** [.008,.119] |
| Risk Aversion | -.066 [-.145,.013] | -.065 [-.144,.014] | -.069* [-.147,.010] | -.068* [-.148,.012] | -.065 [-.146,.015] |
| Loss Aversion | .046 [-.084,.176] | .049 [-.079,.177] | .048 [-.079,.175] | .053 [-.075,.180] | .054 [-.075,.183] |
| Constant | 4.858*** [3.505,6.210] | 4.666*** [3.297,6.034] | 5.017*** [3.601,6.432] | 4.821*** [3.380,6.262] | 4.667*** [3.229,6.106] |
| Socio-demographic controls | Yes | Yes | Yes | Yes | Yes |
| Price Sequence Group FE | Yes | Yes | Yes | Yes | Yes |
| Observations | 1719 | 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. *Inaction Regret X High-TP* is an indicator, taking a value of 1 if the participant experienced *inaction regret* in the previous task and was randomly assigned to treatments *High-TP*. *Action Regret X High-TP* is 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 19.

k with participants in *No-Info* who also requested inefficiently few offers from an ex-ante perspective in task k 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 $k + 1$. At baseline (*No-Info/Low-TP*) in Table 4, Column 3, individuals average search length amounts to 5.02 offers. The average search length of individuals in *Info*, who experience *action regret* in t is increased by 1.1 offers. In contrast, participants in *No-Info* who also searched too short in task k and thus would have experienced *action regret* were they assigned to *Info* instead, continue to search too little (they request around 0.55 offers less in $k + 1$). Column 4, which includes experienced *inaction* and *action regret*, and Column 5 which additionally includes interactions of both types of regret and time pressure, confirm these findings.²¹

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 $k + 1$ after stopping too early in task k ; $p = 0.058$, MWU). Result 4 summarizes these findings:

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

Although participants adjust their search behavior directionally, they do not make higher profits after experiencing regret in the previous task (see Appendix Table A.3, 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.3 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).²² Thus, we find

²¹The interaction between (experienced) *inaction regret* and time pressure in Column 5 implies that previous feelings of regret do not influence search behavior differentially when there is less time for deliberation. The constant coefficient for the interaction between *inaction regret* and *Info* also implies that the effect of time pressure after experiencing *inaction regret* is orthogonal to the *Info* treatment. The same holds true when adding the interaction between (experienced) *action regret* and *High-TP*. Here, the interaction term between *action regret* and *High-TP* suggests that participants who are under time pressure are somewhat more likely to search too short again (i.e., less likely to adjust their behavior). In Table A.2 in the Appendix, we show that the effect of experienced regret and time pressure is similar in both feedback structures.

²²In a robustness check (Table A.9), we show that all results hold in a truncated Poisson specification.


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 3 and 4) 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).

5.5 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 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 3; 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 3; Columns (4)-(6)), which 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, 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 (Task 1).

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).²³

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 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 *High-TP* amount to 11.75 (when subtracting 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 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

²³This 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

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 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.4, 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.*



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

6 Discussion

Our main findings documented limited differences across decision environments that may (or may not) involve anticipated regret and perceived urgency. In this section, we i) provide further insights into more subtle changes in search behaviors across decision environments and discuss how participants perceive their own decision quality (Section 6.1), ii) study whether commitment can serve as a simple tool to improve search efficiency (Section 6.2), and iii) provide evidence on the robustness of our result regarding the insensitivity to anticipated regret (Section 6.3).

6.1 Search heuristics and perceptions of search environments

Our main analyses focused on search behavior in the experiment compared to the risk neutral optimal benchmark (i.e., a constant reservation price). At the same time, previous literature (e.g., Hey, 1982; Houser & Winter, 2004; Moon & Martin, 1990; Schunk & Winter, 2009) highlights the importance of heuristics given the complexity to derive the optimal stopping rule in search tasks. Although our experimental design does not allow us to study all candidate heuristics discussed in this important previous work, we shed more light on individual search behavior related to i) salient stopping prices, ii) bounce-heuristics, and iii) streak-based heuristics across treatments (see Appendix Section A.6 for the details). We find that perceived urgency reduces the probability of acceptance of salient favorable prices ($p \leq 10$) and increases the probability of acceptance of salient unfavorable prices ($p > 50$) whereas our information conditions do not alter these probabilities.

Further, we show that the use of bounce-heuristics does not strongly differ across information and time pressure conditions. For example, when analyzing the one-bounce heuristics following Houser and Winter (2004) and Schunk and Winter (2009) (i.e., “Have at least 2 searches and stop if a price quote larger than the previous quote is received.”), we find that overall about 11 percent of decisions are consistent this heuristic, but this fraction does not strongly differ across treatments. Similarly, our additional analyses on streak-based heuristics (see Appendix Table A.14) results in minor treatment differences.

While overall treatment differences in search behavior appear minor, inexperienced participants do suffer from urgency. In turn, it is important to ask whether participants are aware of the influence of time pressure on decision quality. To study these perceptions in more detail, we elicit how decision-makers rank their performance as compared to other buyers at the end of the experiment. We find that on average, participants are overconfident in all treatments. 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.²⁴ 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 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 who actually experienced time pressure 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).

6.2 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

²⁴We do not neither observe significant differences in between treatments *No-Info* and *Info* ($p = 0.165$, MWU), nor between *High-TP* and *Low-TP* ($p = 0.959$, MWU).

outcomes to their sequential search behavior.²⁵ 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).²⁶

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.²⁷ 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

²⁵Reassuringly 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).

²⁶This 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).

²⁷We 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.

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 5.4), 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).

6.3 Robustness: Non-binding reservation prices

Active sequential search takes place in different search environments. In many tasks, consumers are faced with the decision of whether to buy the product at a certain price or continue the search. In other tasks, the decision is characterized by setting a maximum acceptable price for the product, and continuing the search if the price was above this threshold (then with a potentially different new reservation value). Theoretically, both decisions are equivalent. Consumers *should* buy the product as soon the price is below their reservation value, irrespective of whether they first see the offer and then decide about buying or not, or whether they first specify their reservation value, and only then learn the value of the next offer. However, from a behavioral economics perspective, these decisions may be perceived differently. Asking repeatedly which future prices are acceptable (reservation value elicitation) may render behavior more future-oriented and thus alter the importance of anticipated regret. For example, recent evidence from Jhunjhunwala (2021) indicates that regret may play a more important role when repeated reservation value elicitation is used. In an additional pre-registered experiment (see *AsPredicted*; #80046), we investigate whether our findings are robust to such repeated reservation value elicitation. This experiment includes our treatments *No-Info* and *Info* without time pressure as baseline, as well as two treatments in which participants set an initial reservation price before the first offer is drawn (*Reservation/Info* and *Reservation/No-Info*). If the first offer drawn is below the initial reservation price, the product is bought at the offered price. If the first offer is higher than the initial reservation value, participants can adjust their reservation value, and another offer is drawn for which search costs are incurred.

Interestingly, we find that information about post-purchase price realizations does not alter search behavior in an environment with non-binding reservation prices, either. Average search lengths are indistinguishable (*Reservation/Info*: 2.88 vs. *Reservation/No-Info*: 2.93; $p = 0.719$, MWU). Also the fraction of optimal stopping decisions under *Reservation/Info* and *Reservation/No-Info* is nearly identical (*Reservation/Info*: 61.88 percent vs. *Reservation/No-Info*: 60.00 percent), translating

into very similar payoffs (*Reservation/Info*: 23.27 vs. *Reservation/No-Info*: 23.48; $p = 0.877$, MWU). In Appendix A.5, we discuss these findings in more detail and compare search behavior under both elicitation procedures. Further, we show that the exact replication of the baseline treatments (*Info* and *No-Info* without time pressure) confirms the negligible role of anticipated post-purchase information on search length and payoffs observed in the main experiment.

7 Conclusion

From a theoretical perspective, 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 and their interaction. The 95 percent confidence intervals for the treatment effect estimates in our preferred regression specification (Table 3, 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 (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 and thereby allows us to study the role of time-pressure, regret, and their potential interaction in sequential search tasks.

First, we find 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 one rationale for why sellers often put buyers under time pressure. Clearly, short-lived discounts can deter search, since they limit consumers' ability to consider alternative offers before the discount expires. Search deterrence can be even more pronounced if sellers can discriminate against buyers who do not purchase at the first opportunity (Armstrong & Zhou, 2016). Our findings additionally emphasize a channel of bounded rationality. 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). With experience, participants in our experiment were not significantly affected by time pressure. Consumer protection policies against

sales tactics that "rush consumers into making a decision",²⁸ can thus be especially helpful for inexperienced consumers. These are, for example, customers who are in an environment where they are not very savvy, or who are searching for products that they have not previously looked for. 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 the sense of urgency and make better decisions. However, other purchase 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. As we find that perceived urgency particularly harms decision quality of inexperienced participants, regulation may be more important in such 'unknown' environments than in areas that are currently primarily targeted (e.g., hotel booking or travel websites).

Second, our results provide robust empirical evidence that anticipated regret does 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. While avoiding anticipated regret (Bell, 1982; Bikhchandani & Segal, 2014; Buturak & Evren, 2017; Halpern & Leung, 2016; Hayashi, 2008; Loomes & Sugden, 1982; Qin, 2015; Sarver, 2008; Skiadas, 1997) has been observed in other experimental contexts (Camille et al., 2004; Coricelli et al., 2005; Fioretti et al., 2022; Strack & Viefers, 2021; Zeelenberg, 1999), such regret seems to play a minor role when decision-makers incur salient search costs by actively requesting new price offers and learn advantageous and disadvantageous post-purchase prices. We replicate this result in an additional experiment (see Section 6.3) and show that the observed insensitivity towards anticipated regret in our setting does not hinge on whether participants directly chose to buy the product or repeatedly specify reservation prices. Recent evidence by Jhunjhunwala (2021) suggests that such search behavior can be affected when the feedback structure only highlights potentially better offers. In contrast, in our setting participants see a random subset of actual future price realizations (and associated payoffs). Thus, our results show the tight boundaries of changes in search behavior through post-purchase information: in a search environment where consumers may learn (a subset of) all competitors' prices post purchase, changes in search behavior due to anticipated regret appear unlikely while in environments, where consumers anticipate to only see prices that provide a better deals (e.g., because competitors may

²⁸Retrieved from <https://www.gov.uk/government/news/cma-launches-enforcement-action-against-hotel-booking-sites> on 10/05/2020.

be more likely to advertise such prices), anticipated regret may result in searching longer. There are two reasons that may explain why we do not identify strong efficiency effects of regret. First, anticipated *action 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 still reacted to experienced *action* and *inaction regret*. However, participants were not successful in making better decisions in subsequent search tasks with different price realizations and search costs.

Third, we do not find a substantial interaction between anticipated regret and perceived urgency. Independent of the decision environment, our results indicate that individuals search too little. However, our results also hint at a simple mechanism that consumers may use to avoid such inefficient search: commitment to a binding reservation price. In the experiment, commitment increases search length and payoffs. As such, commitment may also be applied as a potential solution outside the laboratory, and sophisticated consumers may demand commitment devices in the form of public policies or market-based solutions.

Finally, although our design captures many important elements of the trade-off that urgency, resulting time pressure, and regret in real-world settings may pose, decision environments outside the laboratory may both confront consumers with additional challenges or relieve them of some that exist in our setting. On the one hand, repeated search tasks in which search costs stay constant may render learning from past experiences and regret easier (see e.g., Einav, 2005; Oprea et al., 2009) while more complex environments with varying search costs (as in our experiment) may render learning harder. On the other hand, in many search environments outside the laboratory, 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, thereby complicating optimal search. Hence, 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.

References

- Armstrong, M., & Zhou, J. (2016). Search deterrence. *Review of Economic Studies*, 83(1), 26–57.
- 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. (in press). Meta-analysis of empirical estimates of loss-aversion. *Journal of Economic Literature*.
- 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. (2022). Dynamic regret avoidance. *American Economic Journal: Microeconomics*, 14(1), 70–93.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Gächter, S., Johnson, E. J., & Herrmann, A. (2022). Individual-level loss aversion in riskless and risky choices. *Theory and Decision*, 92(3), 599–624.
- 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. (1982). Search for rules for search. *Journal of Economic Behavior & Organization*, 3(1),

- 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.
- Jhunjhunwala, T. (2021). Searching to avoid regret: An experimental evidence. *Journal of Economic Behavior & Organization*, 189, 298–319.
- 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. *Working Paper*.
- Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay-inferences: Effects on impression 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.
- Moon, P., & Martin, A. (1990). Better heuristics for economic search—experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3(3), 175–193.
- 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.

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).²⁹ 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 $F(\cdot)$ with range $[l, h]$. The search costs for each requested offer are denoted as c . Both the distribution function and the search costs are known to the agent. The agent maximizes 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 $t + 1$.³⁰

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 A.1.

$$\begin{aligned} \pi(v - m_t) &= [1 - F(m_t)]\pi(v - m_t - c) + \int_l^{m_t} \pi(v - x - c) dF(x) & (A.1) \\ \Leftrightarrow \pi(v - m_t) &= [1 - F(m_t)]\pi(v - m_t - c) + \int_l^{m_t - c} \pi(v - x - c) dF(x) + \int_{m_t - c}^{m_t} \pi(v - x - c) dF(x) \end{aligned}$$

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

²⁹This 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).

³⁰We refer to every decision between stopping at offer t or requesting offer $t + 1$ as a round, meaning that every search task k consists of up to 25 rounds.

best price (m_t) and weights the resulting profits by their probability. Given the parametrization in the experiment ($v=50$; discrete uniform distribution with range $[1, 100]$), we solve

$$\begin{aligned} 50 - m_t &= \frac{100 + c - m_t}{100} \cdot (50 - c - m_t) + \sum_{i=1}^{10} (50 - c - x) \cdot \frac{1}{100} \\ &= \frac{100 + c - m_t}{100} \cdot (50 - c - m_t) + \frac{m_t - c}{100} \cdot \frac{99 - m_t - c}{2} \end{aligned}$$

for the reservation price m_t which equals the benefits from stopping the search now with the benefits of requesting one additional offer. For example, in the case of search cost $c = 2$, this results in a reservation price of 20.56.

In the second row of Equation A.1, we 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: $m_{ij}^k = \pi(x_{ij}) - R[\pi(x_{kj}) - \pi(x_{ij})]$. 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, an increasing 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 A.1. Equation A.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).

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

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 A.1 and A.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*.

We illustrate the higher reservation price of regret-sensitive agents with the parameters of our experimental design. We assume that the decision-maker receives an initial offer of $m_1 = 22$ and faces search cost of $c = 2$. As illustrated above, a payoff-maximizing agent would continue the search as the offer is above the reservation price. The regret-sensitive decision maker also anticipates aversive feelings of size $\frac{R(1)}{100} + \frac{R(2)*78}{100}$ if they continue the search without encountering a more favorable offer. The first term corresponds to the case in which the next offer m_2 is equal to 21, the second term to cases where the next offer is weakly higher than the current one ($m_2 \geq 22$). Whether to stop the search at $m_1 = 22$ depends on the relative importance of anticipated regret. Assuming the regret function takes the following functional form $R(\pi_{max_t} - \pi_t) = \rho(\pi_{max_t} - \pi_t)^2$, an agent would only continue the search for $\rho < 0.604$. If the sensitivity to feelings of regret is larger, the decision maker stops the search at $m_1 = 22$.

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.³¹ Thus, seeing subsequent prices directly affects the utility from stopping and enters the left-hand side of Equation A.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 $\pi(v - m_t) - \int_l^{m_t - c} R(m_t - c - x) dF(x)$. The second term captures that regret is experienced when the price of the next draw (x_{t+1}) is lower than the previously best 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 $\sum_{n=1}^{\infty} \int_l^{m_t - nc} R(m_t - nc - x) dF(x)$, n denoting the (future) draws.³²

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*.

$$\begin{aligned} \pi(v - m_t) - \sum_{n=1}^{\infty} \int_l^{m_t - nc} R(m_t - nc - x) dF(x) = \\ [1 - F(m_t)] [\pi(v - m_t - c) - \sum_{n=2}^{\infty} \int_l^{m_t - nc} R(m_t - nc - x)] \\ + \int_l^{m_t} [\pi(v - z - c) - \sum_{n=2}^{\infty} \int_l^{z - (n-1)c} R(z - (n-1)c - x)] dF(x) \end{aligned} \quad (\text{A.3})$$

³¹This 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.

³²The 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.

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.

Previous regret experience and urgency: Linking experimental design and theoretical model

In the experimental design, we go beyond the stylized one-period model outlined so far. We allow for the experience of regret in a previous task as participants face multiple search tasks. We hypothesize that the experience of regret in task k intensifies the anticipation of regret in task $k + 1$. As a consequence, experiencing (inaction) regret due to requesting too many offers in task k translates into shorter search in the next task. Experiencing (action) regret due to requesting too few offers in task t translates into longer search in task $t + 1$. In our model, this is both consistent with a payoff-maximizing agent becoming regret-sensitive through experience (extensive margin effect) and with the functional form of the regret function R being subject to regret experiences (intensive margin effect).

Our design also takes into account urgency, which our model does not. We can formally link urgency and regret through the introduction of cognitive capacities if we assume that the anticipation of regret depends on the amount of available cognitive resources. A straightforward approach would be to think about cognitive capacities ($\lambda \in [0, 1]$) as a scaling factor for regret. The perceived utility from stopping at a lowest price m_t in round t became $u(m_t) = \pi(v - m_t) - \lambda R(\pi_{max_t} - \pi_t)$. For example, if the agent does not have any cognitive resources available ($\lambda = 0$), there is no anticipation of regret. Hence, the reduction of cognitive resources during the decision-making process through time pressure makes the agent less sensitive to feelings of regret. The impact of the regret manipulation on search length is therefore expected to be smaller in treatments with high levels of perceived urgency.³³

³³This modeling approach would yield a directed hypothesis on the effect of urgency in environments without post-purchase information. With urgency, we should observe longer (and more efficient) search. At the same time, we acknowledge the multiple channels through which urgency may impact search behavior and do not specify a directed hypothesis in the main text.

A.2 Additional figures and tables

Table A.1: Average decision times per task across time pressure conditions by feedback condition

| Task | <i>No-Info</i> | | | <i>Info</i> | | |
|------|----------------|----------------|---------|---------------|----------------|---------|
| | <i>Low-TP</i> | <i>High-TP</i> | p-value | <i>Low-TP</i> | <i>High-TP</i> | p-value |
| 1 | 10.86 | 5.87 | <0.001 | 11.12 | 4.46 | <0.001 |
| 2 | 9.24 | 3.50 | <0.001 | 10.82 | 3.06 | <0.001 |
| 3 | 8.54 | 2.74 | <0.001 | 10.97 | 2.56 | <0.001 |
| 4 | 5.84 | 2.28 | <0.001 | 7.84 | 2.49 | <0.001 |
| 5 | 5.74 | 2.20 | <0.001 | 6.01 | 2.19 | <0.001 |
| 6 | 5.63 | 2.17 | <0.001 | 6.50 | 2.16 | <0.001 |
| 7 | 5.26 | 2.38 | <0.001 | 5.59 | 2.33 | <0.001 |
| 8 | 6.27 | 2.38 | <0.001 | 5.08 | 2.37 | <0.001 |
| 9 | 8.60 | 3.36 | <0.001 | 5.37 | 2.93 | <0.001 |
| 10 | 10.03 | 3.33 | <0.001 | 6.83 | 3.14 | <0.001 |

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

Table A.2: Experienced regret


| | Number of offers | | |
|-------------------------------|-----------------------------|---------------------------|---------------------------|
| | Full Sample | No-Info | Info |
| | (1) | (2) | (3) |
| Treatments | | | |
| High-TP | .432 [-.103,.967] | .415 [-.152,.981] | .301 [-.269,.872] |
| Info | -.115 [-.610,.380] | | |
| High-TP X Info | -.135 [-.771,.500] | | |
| (Experienced) Inaction Regret | .470 [-.228,1.168] | .413 [-.392,1.218] | -.695* [-1.408,.019] |
| Inaction Regret X Info | -1.082*** [-1.815,-.349] | | |
| Inaction Regret X High-TP | -.104 [-.858,.650] | .071 [-1.046,1.188] | -.077 [-1.042,.888] |
| (Experienced) Action Regret | -.124 [-.806,.558] | -.121 [-.892,.651] | 1.029* [-.063,2.121] |
| Action Regret X Info | 1.058** [.225,1.891] | | |
| Action Regret X High-TP | -.757* [-1.598,.084] | -.798 [-1.889,.293] | -.772 [-2.062,.518] |
| # Tasks encountered | .064** [.008,.119] | .049 [-.033,.131] | .079** [.002,.155] |
| Risk Aversion | -.065 [-.146,.015] | .021 [-.087,.129] | -.123** [-.234,-.012] |
| Loss Aversion | .054 [-.075,.183] | .011 [-.151,.174] | .159 [-.033,.351] |
| Constant | 4.667*** [3.229,6.106] | 5.684*** [3.791,7.576] | 3.469*** [1.836,5.101] |
| Socio-demographic controls | Yes | Yes | Yes |
| Price Sequence Group FE | Yes | Yes | Yes |
| Observations | 1719 | 855 | 864 |

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 19.

Table A.3: Experienced Regret: Optimality of Search

| | Forgone Profits | Optimal | Too few offers | Too many offer |
|-------------------------------|--------------------------|---------------------------|---------------------------|---------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | .971 [-.271,2.213] | -.087** [-.158,-.017] | .006 [-.051,.064] | .081** [.011,.152] |
| Info | .173 [-1.261,1.607] | .016 [-.056,.087] | .006 [-.054,.066] | -.022 [-.089,.045] |
| High-TP X Info | -.606 [-2.874,1.662] | .008 [-.094,.109] | .009 [-.072,.091] | -.017 [-.110,.075] |
| (Experienced) Inaction Regret | 1.915** [.399,3.430] | -.124*** [-.215,-.033] | -.013 [-.080,.054] | .137*** [.057,.217] |
| Inaction Regret X Info | -1.505 [-3.502,.492] | .070 [-.048,.188] | .019 [-.077,.115] | -.089* [-.187,.010] |
| (Experienced) Action Regret | -1.136 [-2.669,.397] | .078** [.006,.151] | .002 [-.068,.072] | -.081*** [-.131,-.030] |
| Action Regret X Info | 3.526** [.851,6.200] | -.108** [-.215,-.000] | .006 [-.089,.101] | .102*** [.026,.178] |
| # Tasks encountered | -.200** [-.371,-.029] | .023*** [.015,.031] | -.031*** [-.038,-.024] | .008** [.001,.015] |
| Risk Aversion | -.160 [-.569,.250] | .009 [-.005,.022] | .002 [-.009,.013] | -.011* [-.022,.000] |
| Loss Aversion | -.167 [-.640,.307] | -.010 [-.030,.011] | -.005 [-.020,.010] | .015 [-.005,.035] |
| Constant | 6.334** [.324,12.344] | .624*** [.409,.840] | .413*** [.214,.612] | -.037 [-.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. Column (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 19.




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.4: Probit Regression: Stopping the search

| | 1[Stopped Search] | | | |
|-----------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | -.011 [-.049,.028] | -.007 [-.043,.028] | .148*** [.048,.249] | .156*** [.057,.256] |
| Info | -.001 [-.037,.035] | .001 [-.032,.034] | .041 [-.033,.115] | .046 [-.029,.122] |
| High-TP X Info | .016 [-.039,.071] | .012 [-.039,.063] | -.132* [-.277,.013] | -.137* [-.274,.000] |
| # Tasks encountered | -.006*** [-.009,-.003] | -.005*** [-.009,-.002] | | |
| Price | -.009*** [-.010,-.008] | -.009*** [-.010,-.008] | -.010*** [-.011,-.008] | -.010*** [-.012,-.008] |
| Risk Aversion | .003 [-.005,.011] | .004 [-.003,.012] | .010 [-.012,.032] | .009 [-.012,.031] |
| Loss Aversion | -.003 [-.014,.008] | .001 [-.009,.011] | .027* [-.004,.058] | .032** [.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 19.



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.5).

Table A.5: OLS Regression Search Length (Task 2-10)

| | Number of offers (Task 2-10) | | |
|----------------------------|------------------------------|---------------------------|---------------------------|
| | (1) | (2) | (3) |
| Treatments | | | |
| High-TP | .133 [-.403,.669] | .200 [-.327,.726] | .200 [-.292,.693] |
| Info | -.059 [-.573,.455] | -.026 [-.523,.471] | -.041 [-.494,.411] |
| High-TP X Info | -.138 [-.854,.579] | -.178 [-.888,.531] | -.175 [-.819,.470] |
| # Tasks encountered | .064** [.009,.120] | .064** [.009,.120] | .064** [.009,.120] |
| Risk Aversion | | -.041 [-.122,.040] | -.066* [-.144,.012] |
| Loss Aversion | | .048 [-.095,.190] | .045 [-.084,.174] |
| Constant | 3.453*** [2.962,3.945] | 4.229*** [3.139,5.319] | 4.857*** [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 19.

A.4 Robustness checks

A.4.1 Inclusion of unresponsive participant

In this section, we show that our main regression analyses (Table 3 and 4) are robust to the inclusion of one participant who was unresponsive to the price offers from Task 3 onward.

Table A.6: OLS Regression Search Length

| | Number of offers | | | | | |
|----------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-----------------------------|-----------------------------|
| | Task 1-10 | | | Task 1 | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| High-TP | -.321 [-1.146,.505] | -.290 [-1.135,.555] | -.275 [-1.052,.502] | -.937** [-1.695,-.180] | -1.052*** [-1.849,-.255] | -1.087*** [-1.779,-.395] |
| Info | -.429 [-1.255,.397] | -.430 [-1.311,.451] | -.433 [-1.249,.382] | -.292 [-1.147,.563] | -.185 [-1.065,.695] | -.203 [-.838,.432] |
| High-TP X Info | .310 [-.638,1.258] | .284 [-.666,1.235] | .271 [-.600,1.143] | .875 [-.410,2.160] | .938 [-.361,2.236] | .958 [-.214,2.131] |
| # Tasks encountered | .090*** [.039,.141] | .090*** [.039,.141] | .090*** [.039,.141] | | | |
| Risk Aversion | | -.059 [-.151,.033] | -.088* [-.177,.001] | | .004 [-.174,.182] | -.079 [-.255,.097] |
| Loss Aversion | | .034 [-.099,.167] | .008 [-.114,.130] | | -.254* [-.530,.022] | -.232* [-.470,.005] |
| Constant | 3.671*** [2.995,4.347] | 4.616*** [3.452,5.781] | 6.054*** [3.312,8.796] | 3.646*** [3.054,4.237] | 5.721*** [3.396,8.047] | 4.741*** [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 19.

Table A.7: Experienced Regret

| | Number of offers | | | |
|-------------------------------|---------------------------|----------------------------|---------------------------|----------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | -190 [-1.045,.666] | -196 [-1.018,.626] | -174 [-1.010,.662] | -185 [-.988,.618] |
| Info | -447 [-1.313,.419] | -068 [-.745,.608] | -833* [-1.812,.145] | -436 [-1.152,.280] |
| High-TP X Info | .180 [-.732,1.092] | .238 [-.681,1.157] | .171 [-.769,1.111] | .214 [-.699,1.128] |
| (Experienced) Inaction Regret | .316 [-.519,1.151] | 1.178* [-.224,2.581] | | 1.099 [-.264,2.461] |
| Inaction Regret X Info | | -1.790** [-3.189,-.390] | | -1.715** [-3.075,-.356] |
| (Experienced) Action Regret | | | -913** [-1.609,-.218] | -816** [-1.442,-.189] |
| Action Regret X Info | | | 1.397*** [.477,2.317] | 1.305*** [.426,2.183] |
| # Tasks encountered | .071** [.015,.127] | .075*** [.019,.131] | .067** [.011,.122] | .070** [.016,.124] |
| Risk Aversion | -.087* [-.175,.002] | -.084* [-.172,.004] | -.092** [-.183,-.001] | -.087* [-.175,.000] |
| Loss Aversion | .030 [-.102,.162] | .035 [-.093,.164] | .037 [-.093,.167] | .038 [-.088,.164] |
| Constant | 6.299*** [3.272,9.327] | 5.958*** [3.162,8.754] | 6.572*** [3.373,9.771] | 6.202*** [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 19.

A.4.2 Truncated Poisson regressions

In this section, we show that our main regression analyses (Table 3 and 4) are robust to a truncated Poisson specification.

Table A.8: Search Length

| | Number of offers | | | | | |
|----------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| | Task 1-10 | | | Task 1 | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| High-TP | .006 [-.131,.144] | .021 [-.114,.156] | .021 [-.106,.149] | -.366** [-.652,-.080] | -.400*** [-.691,-.109] | -.425*** [-.678,-.172] |
| Info | -.025 [-.166,.116] | -.013 [-.149,.123] | -.014 [-.135,.106] | -.105 [-.381,.170] | -.062 [-.345,.222] | -.100 [-.298,.098] |
| High-TP X Info | -.010 [-.204,.185] | -.019 [-.212,.175] | -.020 [-.195,.155] | .344 [-.115,.802] | .367 [-.086,.821] | .401* [-.007,.810] |
| # Tasks encountered | .023*** [.010,.036] | .023*** [.010,.036] | .023*** [.010,.037] | | | |
| Risk Aversion | | -.011 [-.033,.012] | -.020* [-.043,.004] | | .006 [-.056,.068] | -.025 [-.088,.039] |
| Loss Aversion | | .005 [-.032,.042] | .005 [-.030,.040] | | -.088* [-.177,.000] | -.089** [-.164,-.013] |
| Constant | 1.188*** [1.067,1.308] | 1.475*** [1.157,1.794] | 1.580*** [1.232,1.928] | 1.275*** [1.096,1.454] | 1.940*** [1.126,2.754] | 1.587*** [.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 19.

Table A.9: Experienced Regret

| | Number of offers | | | |
|-------------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | .058 [-.080,.196] | .054 [-.077,.185] | .060 [-.071,.192] | .057 [-.070,.184] |
| Info | -.009 [-.140,.122] | .057 [-.080,.193] | -.093 [-.233,.047] | -.028 [-.170,.113] |
| High-TP X Info | -.049 [-.233,.135] | -.035 [-.218,.147] | -.057 [-.241,.127] | -.043 [-.226,.141] |
| (Experienced) Inaction Regret | -.024 [-.143,.095] | .126 [-.029,.282] | | .114 [-.037,.265] |
| Inaction Regret X Info | | -.328*** [-.537,-.118] | | -.317*** [-.521,-.113] |
| (Experienced) Action Regret | | | -.167** [-.331,-.003] | -.158* [-.320,.005] |
| Action Regret X Info | | | .317*** [.077,.558] | .310** [.073,.547] |
| # Tasks encountered | .019** [.003,.035] | .020** [.004,.036] | .018** [.002,.033] | .019** [.003,.035] |
| Risk Aversion | -.019 [-.042,.004] | -.018 [-.041,.005] | -.020* [-.043,.003] | -.019 [-.042,.004] |
| Loss Aversion | .013 [-.025,.051] | .014 [-.023,.051] | .014 [-.023,.051] | .015 [-.022,.052] |
| Constant | 1.592*** [1.219,1.965] | 1.535*** [1.158,1.911] | 1.635*** [1.250,2.021] | 1.576*** [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 19.

A.4.3 No switchpoint

In this section, we show that our main regression analyses (Table 3 and 4) are robust to controlling for risk attitudes and loss attitudes without by calculating the number of safe choices instead of a switchpoint.

Table A.10: Search Length

| | Number of offers | | | | | |
|----------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-----------------------------|-----------------------------|
| | Task 1-10 | | | Task 1 | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| High-TP | .022 [-.461,.506] | .098 [-.378,.575] | .108 [-.340,.555] | -.973** [-1.737,-.208] | -1.090*** [-1.906,-.274] | -1.066*** [-1.764,-.368] |
| Info | -.086 [-.571,.399] | -.041 [-.512,.430] | -.056 [-.471,.360] | -.327 [-1.188,.534] | -.188 [-1.079,.702] | -.190 [-.838,.457] |
| High-TP X Info | -.033 [-.704,.639] | -.075 [-.741,.590] | -.073 [-.672,.526] | .910 [-.379,2.199] | .988 [-.329,2.306] | .978 [-.201,2.158] |
| # Tasks encountered | .079*** [.032,.125] | .079*** [.032,.125] | .079*** [.032,.125] | | | |
| Risk Aversion | | -.044 [-.142,.054] | -.077 [-.169,.016] | | -.018 [-.226,.189] | -.105 [-.314,.104] |
| Loss Aversion | | .044 [-.104,.192] | .047 [-.078,.171] | | -.158 [-.462,.146] | -.104 [-.358,.150] |
| Constant | 3.391*** [2.988,3.793] | 4.251*** [3.280,5.223] | 4.691*** [3.500,5.883] | 3.681*** [3.081,4.281] | 5.324*** [3.190,7.457] | 4.157*** [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 19.

Table A.11: Experienced Regret

| | Number of offers | | | |
|-------------------------------|---------------------------|-----------------------------|---------------------------|-----------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | .241 [-.256,.738] | .230 [-.239,.700] | .252 [-.224,.727] | .245 [-.213,.703] |
| Info | -.043 [-.495,.410] | .192 [-.282,.666] | -.335 [-.827,.156] | -.105 [-.600,.390] |
| High-TP X Info | -.188 [-.833,.457] | -.146 [-.781,.490] | -.217 [-.860,.426] | -.173 [-.814,.469] |
| (Experienced) Inaction Regret | -.080 [-.494,.334] | .478 [-.121,1.077] | | .425 [-.156,1.006] |
| Inaction Regret X Info | | -1.140*** [-1.890,-.390] | | -1.092*** [-1.822,-.362] |
| (Experienced) Action Regret | | | -.556* [-1.115,.003] | -.516* [-1.066,.034] |
| Action Regret X Info | | | 1.101** [.246,1.957] | 1.066** [.224,1.909] |
| # Tasks encountered | .065** [.009,.121] | .068** [.012,.123] | .061** [.006,.116] | .065** [.010,.120] |
| Risk Aversion | -.074 [-.166,.018] | -.073 [-.166,.020] | -.078* [-.169,.013] | -.078* [-.171,.015] |
| Loss Aversion | .064 [-.072,.200] | .069 [-.066,.204] | .070 [-.062,.203] | .077 [-.057,.210] |
| Constant | 4.847*** [3.527,6.167] | 4.656*** [3.322,5.990] | 5.004*** [3.622,6.386] | 4.810*** [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 19.

A.4.4 Probit regression: Optimality after experienced regret

This section shows that Table A.3 is robust to a probit specification in Columns (2)-(4).

Table A.12: Probit Regression: Stopping the search


| | Forgone Profits | Optimal | Too few offers | Too many offers |
|-------------------------------|--------------------------|---------------------------|---------------------------|---------------------------|
| | (1) | (2) | (3) | (4) |
| Treatments | | | | |
| High-TP | .971 [-.271,2.213] | -.090** [-.162,-.018] | .009 [-.051,.069] | .075** [.010,.139] |
| Info | .173 [-1.261,1.607] | .016 [-.060,.091] | .006 [-.057,.069] | -.026 [-.100,.047] |
| High-TP X Info | -.606 [-2.874,1.662] | .009 [-.095,.113] | .005 [-.080,.090] | -.009 [-.098,.080] |
| (Experienced) Inaction Regret | 1.915** [.399,3.430] | -.125*** [-.216,-.034] | -.019 [-.090,.051] | .111*** [.050,.172] |
| Inaction Regret X Info | -1.505 [-3.502,.492] | .069 [-.049,.187] | .028 [-.071,.126] | -.062 [-.141,.017] |
| (Experienced) Action Regret | -1.136 [-2.669,.397] | .080** [.005,.155] | .007 [-.058,.073] | -.087*** [-.146,-.028] |
| Action Regret X Info | 3.526** [.851,6.200] | -.110** [-.219,-.001] | .005 [-.085,.095] | .110*** [.029,.191] |
| # Tasks encountered | -.200** [-.371,-.029] | .023*** [.015,.031] | -.031*** [-.039,-.024] | .009** [.002,.015] |
| Risk Aversion | -.160 [-.569,.250] | .009 [-.005,.023] | .003 [-.009,.015] | -.010* [-.020,.001] |
| Loss Aversion | -.167 [-.640,.307] | -.010 [-.030,.011] | -.007 [-.023,.009] | .014 [-.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. Column (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 19.

A.5 Non-binding reservation prices

In our main experiment, we show that future price realizations do not alter search behavior when participants choose whether to buy the product or to continue the search after every offer. To study the robustness of our results, we introduce an additional pre-registered experiment in which

participants repeatedly choose their reservation price before every offer in Section 6.3. Figure A.3 summarizes the main findings from this experiment. Most importantly, anticipating post-purchase information does not increase search length if participants repeatedly choose their reservation price before every offer (*Reservation/No-Info*: 2.93 vs. *Reservation/Info*: 2.88; $p = 0.719$, MWU). Interestingly, a direct comparison of both elicitation procedures shows that when participants make their choices through a reservation price, they request fewer offers. This holds both true with information (*Info*: 3.77 vs. *Reservation/Info*: 2.88; $p < 0.001$, MWU) and without information (*No-Info*: 3.56 vs. *Reservation/No-Info*: 2.93; $p = 0.001$, MWU) about future price realizations. Regression analyses (Table A.13) show that this effect is less pronounced in the first search task (Task 1).



Notes. The figure shows boxplots of search lengths across treatments in the additional experiment. The 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 A.3: Search length across information structures and elicitation procedures (Tasks 1-10).

Finally, the replication of our baseline treatments (*Info* and *No-Info* without time pressure) show that search behavior is unaffected by the provision of post-purchase information. Average search lengths are similar (*Info*: 3.77 vs. *No-Info*: 3.56; $p = 0.418$, MWU), and payoffs closely aligned (*Info*: 25.20 vs. *No-Info*: 24.96; $p = 0.739$, MWU) across the two information structures. Regression analyses (Table A.13) corroborate these non-parametric results.

Table A.13: Search Length

| | Number of offers | | | | | |
|----------------------------|---------------------------|----------------------------|---------------------------|---------------------------|-------------------------|-------------------------|
| | Task 1-10 | | | Task 1 | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| Reservation Price | -.627** [-1.104,-.150] | -.619*** [-1.084,-.155] | -.553** [-.995,-.112] | -.313 [-1.059,.434] | -.205 [-.901,.491] | -.363 [-1.131,.404] |
| Info | .212 [-.296,.721] | .213 [-.278,.705] | .156 [-.345,.658] | -.021 [-.761,.719] | -.013 [-.756,.729] | .166 [-.638,.970] |
| Reservation X Info | -.267 [-.929,.396] | -.253 [-.885,.378] | -.267 [-.838,.303] | .000 [-1.054,1.054] | -.016 [-1.082,1.051] | -.053 [-1.073,.968] |
| # Tasks encountered | .070*** [.032,.107] | .070*** [.032,.107] | .070*** [.032,.107] | | | |
| Risk Aversion | | .055 [-.033,.143] | .072 [-.017,.160] | | .227*** [.091,.364] | .234*** [.104,.365] |
| Loss Aversion | | -.097 [-.276,.083] | -.095 [-.266,.076] | | .053 [-.177,.284] | .103 [-.130,.336] |
| Constant | 3.175*** [2.807,3.543] | 3.153*** [2.019,4.286] | 3.343*** [2.133,4.553] | 3.292*** [2.750,3.833] | 1.198 [-.734,3.130] | 1.026 [-1.065,3.118] |
| 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. 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.

A.6 Search heuristics


In this section, we shed more light on individual search behavior related to i) salient stopping prices, ii) bounce-heuristics, and iii) streak-based heuristics across treatments.

A.6.1 Salient stopping prices

First, we look at whether stopping behavior around salient reservation prices differs across experimental treatments. To do so, we define salient unfavorable prices as prices that always leads to a negative payoff irrespective of search costs and search length (i.e., prices larger than 50) and salient favorable prices as prices at or below 10. Overall, the probability to stop searching with all received price offers larger than 50 is very low (2.25 percent). In *Low-TP*, only 0.8 percent of stopping decisions happen with salient unfavorable prices (*Low-TP/No-Info*: 0.6 percent and *Low-TP/Info*: 1.0 percent). in *High-TP* this fraction amounts to 3.6 percent (*High-TP/No-Info*: 3.1 percent, *High-TP/Info*: 4.2 percent). Hence, with time pressure, salient unfavorable prices are more likely to be accepted ($p > 0.001$, MWU). Across information conditions, we do not find significant differences with respect to

the acceptance of salient unfavorable prices (*No-Info*: 1.9 percent vs. *Info*: 2.6 percent; $p = 0.678$, MWU). Salient favorable prices are accepted with a much higher probability across all treatments (in 89 percent of the decisions) but mistakes are again more likely to occur with time pressure. Participants in *Low-TP* accept prices at or below 10 in 90.5 percent of the cases, participants in *High-TP* in 87.5 percent, with the difference being marginally statistically significant ($p < 0.060$). Again, we do not find strong differences across information conditions (*No-Info*: 87.3 percent vs. *Info*: 90.6 percent; $p = 0.254$, MWU).

An alternative way of studying whether particular salient prices influence stopping is to compare accepted prices across treatments. Figure A.4 shows histograms of accepted prices and highlights that the distribution is very similar across feedback (*No-Info* vs. *Info*; $p = 0.837$ Kolmogorov-Smirnov) and time pressure conditions (*Low-TP* vs. *High-TP*; $p = 0.388$ Kolmogorov-Smirnov test). Hence, we do not observe more frequent stopping at some salient cutoffs (e.g., 20/30/40) in particular treatments.



Notes. The figure shows the histograms of accepted prices across Tasks 1-10 in each of the four treatment conditions.

Figure A.4: Accepted prices across treatments (Tasks 1-10).

Lastly, we expand our regression analysis on individual stopping behavior in Table A.14. In Column (1), we corroborate that participants in *High-TP* are somewhat less responsive to the current price. This is consistent with the finding that participants in *Low-TP* make fewer mistakes in the sense of accepting salient unfavorable prices or rejecting salient favorable prices.

A.6.2 Bounce heuristics

“Bounce” heuristics describe individual search behaviors where the search was continued beyond the ultimately accepted price (e.g., a once-bounce heuristic could refer to “Have at least 2 searches and stop if a price quote larger than the previous quote is received”, see Houser and Winter (2004) and Schunk and Winter (2009)). In the main part of the paper, we discussed the use of the recall option, which may reflect such bounce heuristics (i.e., “Stop at a price which is higher than the best price you have encountered so far.”). We find that recall rates do not differ substantially across treatments. In *No-Info*, participants exercise the recall option in 18.8 percent of decisions, in *Info* in 17.9 ($p = 0.998$; MWU). In *High-TP*, rates are somewhat higher (20.5 percent) than in *Low-TP* with 16.2 percent ($p = 0.094$; MWU). Similarly, when focusing on other bounce heuristics, treatment differences are small. Analyzing the one-bounce heuristics following Houser and Winter (2004) and Schunk and Winter (2009), overall 10.9 percent of decisions are consistent with the one-bounce strategy: “Have at least 2 searches and stop if a price quote larger than the previous quote is received.”, but we do not find treatment differences across feedback conditions (*No-Info*: 10.9 percent vs. *Info*: 10.8 percent; $p = 0.936$, MWU), and only small differences across time pressure conditions (*Low-TP*: 9.7 percent vs. *High-TP*: 12.1 percent; $p = 0.093$, MWU). We also analyze a modified one-bounce rule: “Have at least 2 searches and stop if a price quote larger than the previous quote less the search cost is received.” Also here, we find no differences across feedback conditions (*No-Info*: 11.7 percent vs. *Info*: 11.4 percent; $p = 0.821$, MWU) and minor differences between *High-TP* and *Low-TP* (*Low-TP*: 10.1 vs. *High-TP*: 12.9; $p = 0.054$, MWU).

A.6.3 Streak-based heuristics

In addition, we investigate how streaks in unfavorable past prices impact stopping behavior, akin to the idea of (losing) streak-based heuristics proposed in the literature (e.g., Houser & Winter, 2004; Schunk & Winter, 2009). Table A.14 Column (2) shows that, across treatments, participants are equally likely to stop after they encountered two times an unfavorable price in a row. Columns (3)-(6) confirm that this holds across treatments. Based on our analyses, we do not find convincing evidence that participants resort to different heuristics across treatments.

Table A.14: Probit Regression: Stopping the search

| | 1[Stopped Search] | | | | | |
|----------------------------------|-------------------|---------------|---------------|---------------|---------------|---------------|
| | Full Sample | | No-Info | Info | Low-TP | High-TP |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Treatments | | | | | | |
| High-TP | -.077* | -.007 | -.010 | .004 | | |
| | [-.158,.005] | [-.043,.028] | [-.042,.022] | [-.030,.039] | | |
| Info | .040 | .001 | | | .001 | .014 |
| | [-.039,.119] | [-.032,.034] | | | [-.029,.032] | [-.023,.051] |
| High-TP X Info | .014 | .012 | | | | |
| | [-.037,.065] | [-.039,.063] | | | | |
| Price | -.009*** | -.009*** | -.009*** | -.009*** | -.009*** | -.008*** |
| | [-.011,-.008] | [-.010,-.008] | [-.010,-.007] | [-.010,-.008] | [-.010,-.008] | [-.009,-.007] |
| Price X High-TP | .002** | | | | | |
| | [.000,.004] | | | | | |
| Price X Info | -.001 | | | | | |
| | [-.003,.001] | | | | | |
| Previous Two Prices[≥ 50] | | -.004 | -.007 | -.002 | -.014 | .007 |
| | | [-.027,.018] | [-.037,.023] | [-.036,.032] | [-.043,.014] | [-.027,.041] |
| # Task encountered | -.005*** | -.005*** | -.005** | -.005** | -.004* | -.007*** |
| | [-.008,-.002] | [-.009,-.002] | [-.010,-.001] | [-.010,-.000] | [-.008,.000] | [-.011,-.002] |
| Risk Aversion | .004 | .004 | -.004 | .011** | .000 | .005 |
| | [-.003,.012] | [-.003,.012] | [-.013,.005] | [.000,.021] | [-.011,.012] | [-.004,.013] |
| Loss Aversion | .000 | .001 | .006 | -.010 | .001 | -.000 |
| | [-.010,.011] | [-.009,.011] | [-.006,.018] | [-.026,.006] | [-.018,.020] | [-.012,.011] |
| # of choices | 7226 | 7226 | 3643 | 3583 | 3591 | 3635 |
| Price FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Pseudo R-sq | .3 | .3 | .28 | .33 | .35 | .26 |

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 all treatments, columns (3) & (4) in the respective feedback environment. *Price* is the price of the current offer [1,100] the participant faces. *Previous Two Prices*[≥ 50] is an indicator variable taking the value 1 if the previous two prices were ≥ 50 . *# 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 19.

A.7 Instructions

A.7.1 Main Experiment

Appendix A.7.1 includes the translated instructions of the main experiment (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.

During the experiment, you and the other participants will be asked to make decisions. These can affect the payoffs for you, and potentially for other participants. How your decisions relate to the payoffs will be explained in more detail in the instructions (or later on the screen).

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 talk 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:

$$100 \text{ Taler} = 12 \text{ €}$$

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.

$$\text{Taler from the purchasing decision} = \text{Value of the product} - \text{Price} - \text{Cost for price offers}$$

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

$$50 - (\text{lowest price received}) - \text{search cost} * (\text{number of offers you requested}).$$

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:

The search costs in this task are: 2.0

| | |
|---------------------------------|-----|
| Number of offer: | 3 |
| Current Offer (price in Taler): | 55 |
| Lowest price so far: | 45 |
| Costs for requested offers: | 4.0 |
| Taler you earn if you buy now: | 1.0 |

| | Offer 1 | Offer 2 | Offer 3 |
|--------------------|---------|---------|---------|
| Price of the offer | 50 | 45 | 55 |
| Taler earned | | | 1 |

Remaining time: 0 Seconds **Costs for exceeding the time limit: 1 Taler**

If you click the button "Purchase" klicken, you accept the best offer so far.

If you click on the button "Additional offer", you receive an additional offer. This costs 2.0 Taler.

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 |
| | | | | | = 1 |

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.

Additional on-screen instructions throughout the experiment

Expected Performance; rank in own treatment

You made several purchase decisions in the first part of the experiment.

Please think back to the **first 10 purchase decisions**, where you could decide after each offer whether to accept it or not.

5 other people have seen the same price sequences as you in this part.

Below we ask you to rate how successful you were in this part compared to the other people.

For this, we have calculated the **average payout of all 10 rounds**.

Below we ask you to rate how successful you were in this part compared to the other people.

For a correct estimation, you will receive 2 EUR. Otherwise, you will receive 0 EUR.

Estimate your rank based on the **average payout**:

1 2 3 4 5 6

Expected Performance; rank in opposite time-pressure treatment

Please think back again to the **first 10 purchase decisions**, where you could decide after each offer whether to accept it or not.

We ask you again to compare yourself with 5 other people.

These have also seen the same price sequences.

However, these participants each had **60 seconds** [*High-TP/No-Info and High-TP/Info*]/ **4 seconds** [*Low-TP/No-Info and Low-TP/Info*] to make a decision.

As a reminder, you had *4 seconds* [*High-TP/No-Info and High-TP/Info*]/ *60 seconds* [*Low-TP/No-Info and Low-TP/Info*].

Below we ask you to rate how successful you were in this part compared to the other people.

Unlike the previous decision, this question is hypothetical and you will not receive a payout based on your answer.

Nevertheless, your answer to this question is of great interest.

Estimate your rank based on the **average payout**:

1 2 3 4 5 6

Loss attitudes [Task A] (Gächter et al., 2022)

Task A consists of 6 decisions where you can accept up to 6 offers.

The offers consist of a lottery through which you can lose or win money. You have to decide for each of the 6 offers whether to accept it or not.

For each accepted offer, the computer loses or wins an amount of money.

At the end of the experiment, your decision is implemented for one of the 6 offers. The computer randomly selects (with equal probability) which offer will be implemented.

Decide for each offer whether you want to accept it.

| | | |
|---|--|---|
| 1 | With 50% probability you lose 2 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |
| 2 | With 50% probability you lose 3 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |
| 3 | With 50% probability you lose 4 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |
| 4 | With 50% probability you lose 5 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |
| 5 | With 50% probability you lose 6 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |
| 6 | With 50% probability you lose 7 euros; with 50% probability you win 6 euros. | <input type="radio"/> accept <input type="radio"/> reject |

Risk attitudes [Task B] (Holt & Laury, 2002)

Task B consists of 10 decisions, each of which allows you to choose between 2 offers.

The offers consist of a lottery through which you win money. You must choose lottery X or Y for each of the 10 choices.

For each lottery you choose, the computer will draw the amount of money you win.

At the end of the experiment, one of the 10 decisions is implemented. The computer randomly

selects (with equal probability) which decision will be implemented.

Decide in each case whether you want to accept X or Y.

| | Option X | Option Y | |
|----|--|--|---|
| 1 | With 10% probability you win 2.00 Euro; with 90% probability you win 1.60 Euro. | With 10% probability you win 3.85 Euro; with 90% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 2 | With 20% probability you win 2.00 Euro; with 80% probability you win 1.60 Euro. | With 20% probability you win 3.85 Euro; with 80% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 3 | With 30% probability you win 2.00 Euro; with 70% probability you win 1.60 Euro. | With 30% probability you win 3.85 Euro; with 70% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 4 | With 40% probability you win 2.00 Euro; with 60% probability you win 1.60 Euro. | With 40% probability you win 3.85 Euro; with 60% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 5 | With 50% probability you win 2.00 Euro; with 50% probability you win 1.60 Euro. | With 50% probability you win 3.85 Euro; with 50% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 6 | With 60% probability you win 2.00 Euro; with 40% probability you win 1.60 Euro. | With 60% probability you win 3.85 Euro; with 40% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 7 | With 70% probability you win 2.00 Euro; with 30% probability you win 1.60 Euro. | With 70% probability you win 3.85 Euro; with 30% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 8 | With 80% probability you win 2.00 Euro; with 20% probability you win 1.60 Euro. | With 80% probability you win 3.85 Euro; with 20% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 9 | With 90% probability you win 2.00 Euro; with 10% probability you win 1.60 Euro. | With 90% probability you win 3.85 Euro; with 10% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |
| 10 | With 100% probability you win 2.00 Euro; with 0% probability you win 1.60 Euro. | With 100% probability you win 3.85 Euro; with 0% probability you win 0.10 Euro. | <input type="radio"/> X <input type="radio"/> Y |

Socio-demographics

Please provide the following statistical information.

- Gender [male; female]
- Age [integer]
- Field of study (faculty/major)
 - 1=Humanities
 - 2=Engineering
 - 3=Medicine
 - 4=Natural Science
 - 5=Law
 - 6=Economics
 - 7=Social Science
 - 8=Other
- What is your high school graduation grade in mathematics? [integer; 1-6]
- What language(s) is (are) your native language(s)? [string]

- How many times have you participated in an economic laboratory study (including outside of this laboratory)? [integer]
- How many participants from the experiment do you know personally? [integer]
- If there is anything else you would like to tell us regarding the experiment, please enter it here: [string]

A.7.2 Additional Experiment: Non-binding reservation price

Appendix A.7.2 includes the translated instructions for the treatments with a repeated reservation price elicitation (*Reservation/Info* and *Reservation/No-Info*) of the additional experiment (from German).

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.

During the experiment, you and the other participants will be asked to make decisions. These can affect the payoffs for you, and potentially for other participants. How your decisions relate to the payoffs will be explained in more detail in the instructions (or later on the screen).

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 talk 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:

$$100 \text{ Taler} = 12 \text{ €}$$

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.

$$\text{Taler from the purchasing decision} = \text{Value of the product} - \text{Price} - \text{Cost for price offers}$$

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 you the product. It makes you offers in the form of purchase prices at which you can buy the product. In the process, the computer makes one offer after another. Offers made remain valid, so you can always purchase the product at the lowest purchase price offered so far. However, every offer is associated with a cost for you. This means that you can receive as many offers as you want (as long as there is a possibility of achieving a positive payout amount in one round), but you pay for each offer:

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

$$50 - (\text{lowest price received}) - \text{search cost}^*(\text{number of offers you requested}).$$

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.)

Your purchase decision

Before each offer you receive from the computer, you specify your *maximum purchase price*. Your *maximum purchase price* determines the price up to which you would buy the product. If the computer's next offer is lower than (or equal to) your *maximum purchase price*, you buy the product at the offered purchase price. If the offer is higher than your *maximum purchase price*, you will not buy the product. In this case, you will be asked again to enter a *maximum purchase price*. This can be different from your last entry, but it does not have to be. After that, you will receive another offer (for which you will pay the search cost displayed on the screen).

Time for the decision

You only have limited time to make your decision. After every offer you have 60 seconds 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 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)—another offer has already been made by the computer:

The search cost in this round are: 2.0

Number Offer: 2
 Current Offer (price in Taler): 45
 Lowest Price so far: 45
 Cost for received offers: 2.0

| | Offer 1 | Offer 2 |
|--------------------|---------|---------|
| Price of the offer | 50 | 45 |
| Taler earned | | 3 |

Remaining Time: 0 Sekunden **Cost for exceeding the time limit: 1 Taler**

The randomly drawn price of this round is higher than your maximum purchase price. Before the next offer is drawn, we ask you to enter your maximum purchase price again. This may or may not be different from your last entry.

Please enter your maximum purchase price here:

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 received so far.

In the lower part you make your decision by entering your *maximum purchase price* in the free field. This can be between 1 and 100 Taler. To confirm it and receive the next offer, click "Confirm".

In the central part, you see an overview of the offers received so far, as well as the payoff you would have received if you had purchased at the current lowest price.

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 your first *maximum purchase price* was lower than 50 Taler and therefore you did not buy the product. Instead, you were asked again for your *maximum purchase price*, entered it, and then received another offer at the cost of 2 Taler.

The second price offered to you, was 45 Taler in the example. If you had specified a *maximum purchase price* of 45 Taler or higher after the first offer, you would have received the product for the lowest price so far: 45 Taler. Hence, the achieved Taler 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 your *maximum purchase price* after the first offer was lower than 45. Therefore, you now enter your maximum purchase price again and then receive another offer at a cost of 2 Taler. This is the current situation that you see on the screenshot.

If you now enter a *maximum purchase price* of at least 45 Taler within the remaining time, you will receive the product regardless of the next offer (as you can always buy the product at the lowest price offered so far). Therefore, if the next price is higher than 45 Taler (for example, price = 55 Taler), you received 1 Taler for the purchase in this round:

| | | | | | |
|---|----|---|----|---|-----|
| Received Taler = value of the product – lowest price – search cost (2 Taler for each requested offer) | | | | | |
| = | 50 | – | 45 | – | 2*2 |
| | | | | | = 1 |

If instead you enter a *maximum purchase price* of less than 45 Taler and the next randomly drawn price is 55 Taler, you would not buy the product and then be asked for your *maximum purchase price*. By entering it again, you would incur additional cost of 2 Taler, and the computer would show you another randomly drawn price.

If offer 3 is less than 45 Taler you buy the product, provided that your *maximum purchase price* after the second offer was greater than this.

Beneath the offers seen so far, you see the “Remaining Time” for the decision. This shows how much time you have remaining to decide about your *maximum purchase price*. 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. If you then decide to make an entry within the next 60 seconds and it leads to a purchase at the price of 45 Taler, you will receive 0 Taler in this round (since the product is worth 50 Taler and you incurred a total search cost of 4 Taler, plus 1 Taler for exceeding the time limit). Should not make an entry again within 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 Reservation/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.