



Sequential search with a price freeze option: theory and experimental evidence

Emanuel Marcu¹ · Charles N. Noussair² 

Received: 30 May 2022 / Revised: 3 July 2024 / Accepted: 10 July 2024
© The Author(s), under exclusive licence to Economic Science Association 2024

Abstract

We introduce price freeze options into a model of sequential search. The model's predictions are tested in a laboratory experiment. The experiment varies (1) whether freezing is possible or not, (2) the cost of freezing, and (3) the time horizon. Overall, the observed treatment effects are consistent with the predictions of our model. Assuming that individuals experience regret, fail to ignore sunk search costs, misperceive the number of periods remaining, or are risk-averse, does not improve upon the performance of the model. Our results support the use of the assumption of optimal search behavior in theoretical and empirical studies.

Keywords Sequential search · Price freezing · Experiment · Under-searching · Imperfect recall

JEL Classification Numbers: D83 · C91

We especially thank Asaf Plan, Amnon Rapoport, the editor and three anonymous referees, for providing us with detailed comments. We thank Jaap Abbring, Maria Arbatskaya, Jordi Brandts, Bart Bronnenberg, Jeff Campbell, David Cesarini, Guillaume Frechette, Ernan Haruvy, Ron Oaxaca, Andrew Schotter, Tiemen Woutersen, participants at the 2017 North American Economic Science Association Meetings, as well as seminar participants at the Universities of Arizona, Alicante, Maryland, Melbourne, Technology-Sydney, and Texas-Dallas, as well as Emory, New York and Tilburg Universities for their helpful comments. We thank Abdolkarim Sadrieh for providing helpful software code. We thank the Economic Science Laboratory and the Eller College of Management at the University of Arizona for funding. The replication material for the study is available at DOI 10.17605/OSF.IO/TRJQZ.

✉ Charles N. Noussair
cnoussair@email.arizona.edu
Emanuel Marcu
e.y.marcu@uvt.nl

¹ Department of Econometrics and Operations Research, Tilburg University, Tilburg, Netherlands

² Department of Economics and the Economic Science Laboratory, University of Arizona, Tucson, USA

1 Introduction

Purchases of many goods are characterized by a tradeoff between accepting the best price currently available and waiting to see if a better offer appears at a later date. For example, a seller of a home must decide whether to accept an offer from a buyer or to turn it down and wait for a higher price from another prospective purchaser. This decision may be made a number of times before a transaction occurs. Markets with this feature are called *search* markets. These markets innovate over time as information technology progresses. Recent years have seen the advent of additional services that are offered as part of search processes, changing the decision problem that searchers face. This paper focuses on a particular innovation that has become prevalent in such markets, the *Price Freeze Option*, or PFO. A PFO is a service offered by firms, which guarantees the availability of an observed offer for potential future acceptance. Purchasing a PFO gives a searcher the possibility of accepting an offer at a later stage, while not committing herself to that offer. That is, once an offer is frozen, the searcher can go back and accept the frozen offer at a later date.

PFOs are making inroads into various important markets. As an example, consider commercial airlines, a large number of whom have introduced the option to freeze the price of airline tickets. For example, at the time of this writing, United Airlines offers an option, called *Farelock*, to freeze a price of an airline ticket for one week, for a fee of between 5 and 10 US dollars.¹ Lufthansa, Air France, and other global and regional carriers offer similar possibilities.² The proliferation of such offers suggests that there is some demand for the option to lock in prices, and that the airlines find the practice profitable, or at least necessary to stay competitive. Similar features exist in the market for mortgages, in which it is often possible to lock in an interest rate for a limited period of time, and for hotel and rental car reservations that can be canceled for some time after they are made.³

To our knowledge, the institution of price freezing has not been studied by economists.⁴ Thus, it is not well-understood how the existence of the option to freeze

¹ In the US, airlines are legally required to offer the option to cancel within 24 hours free of charge. Therefore, the PFOs offered by US airlines are always for a period that is longer than 24 h. Guidelines regarding this regulation can be found on the website of the Department of Transportation at the following link: https://www.transportation.gov/sites/dot.gov/files/docs/Notice_24hour_hold_final20130530.pdf.

² See Table 7 in Appendix A for examples of the fee and duration of PFOs being offered by various airlines in 2017.

³ In the mortgage market, when one applies for a loan, the lender offering the loan guarantees an interest rate for some time period, typically 7 days. This is essentially a price freeze offer at a price of 0. It is an option to take out a loan that is good for a prespecified period, just like the price freeze option in our experiment. Booking hotel reservations that are cancellable has the same feature. There is typically a date provided that is the latest date at which one can cancel with no penalty. This arrangement is essentially the purchase of a price freeze option with a fee of 0 for the period of time that free cancellation is permitted. Booking rental cars with a temporary option to cancel are similar. For these settings, our theoretical model also applies with $f = 0$.

⁴ A rich and well-developed literature investigates investors' decisions to purchase financial options. There are many differences between our environment and those characteristic of financial markets, two of which are fundamental. The first is that a basic assumption of the asset pricing literature is that the

prices affects the behavior of consumers. It has not been established under which, if any, conditions the availability of a PFO lengthens or shortens searches, benefits the searcher, or is profitable for the party offering the option. Furthermore, it is unknown how the response of consumers to the opportunity to purchase such a price guarantee depends on the length of time they have to make the purchase and the cost of the freezing option.

We use both theoretical and experimental methods to study the effect of PFOs on search behavior and outcomes. We study the searcher's decision only, and take the price setting process and freeze fees as exogenous. While this emphasis on one side of the market does not allow us to evaluate equilibrium predictions, it does permit us to focus on the searcher's ability to solve a dynamic search problem without having to consider strategic uncertainty about the behavior of agents on the other side of the market. The theoretical analysis identifies benchmark decisions and outcomes that result from optimal behavior of a risk neutral agent. The experiment is used to consider which aspects of the model are likely to find empirical support, and where its predictions might exhibit inaccuracies.

Our purpose is to study price freezing as an institution, and not to simulate or investigate the airline, mortgage, or any other specific market. Our focus is, rather, on the PFO itself and its implications on searchers' behavior, and our goal is to obtain some general insights regarding the properties of this institution. While we do find it striking that numerous firms have recently adopted price freezing with such vigor, we do not address the forces behind the decision to adopt a policy of offering PFOs. Rather, in our experiment, individuals are randomly assigned in different phases of the sessions to different markets that may or may not permit price freezing, and the implications of this institution on search behavior are studied *ceteris paribus*.

Our model builds on a classical homogeneous good sequential search model with a finite horizon and a risk neutral agent, which we extend to include the presence of PFOs. We consider how behavior and outcomes respond to changes in (i) the price of the freezing option, and (ii) the length of the time horizon available to the decision maker. We do so both in the absence and the presence of the possibility of recalling and accepting a prior offer other than the one that was frozen. We consider a setting in which there is a particular good available for purchase that has no close substitutes, such as an airline ticket to travel to a meeting on the only flight available on the day that one must travel.

We characterize the optimal decision rule, which is a non-stationary policy that we call a *reservation / double reservation policy*. This policy dictates that in the stages just before the terminal stage, there are two price thresholds, and in

Footnote 4 (continued)

value of the underlying asset follows a Brownian motion or a discrete random walk. In our environment, the price of the underlying good itself, rather than the change in its price, is independently and identically distributed at each stage. This means that in our environment, the price at which the good can be purchased at time $t - 1$ has no relationship to its price at time t . The second difference is that financial options can be traded so that part of their value results from the ability to resell them. In our setting, PFOs cannot be transferred or resold. This means that techniques used to solve for option values cannot be applied to our environment.

sufficiently early periods, there is one threshold. When there are two thresholds, offers more favorable than a cutoff level are accepted, those in an intermediate range are frozen, and those that are less favorable than a second cutoff are rejected. If there is one threshold, only acceptances and rejections occur. The threshold price levels depend on the offer distribution, the cost of search, the fee for freezing an offer, and the number of periods available to continue the search.

We then report a laboratory experiment, in which we study whether some of the conclusions of the model are borne out in the data. To evaluate the comparative statics of the model, we vary, in different treatments, the fee charged for the freeze option and the length of the time horizon. We test whether higher freeze fees decrease the average length of searches and lower the incidence of the use of the freeze option. We also consider whether lengthening the time horizon increases search length, leads to less use of the freeze option, and increases profits. We evaluate these predictions in an environment in which recall of prior offers is not possible, as well as one in which recall is possible, but imperfect.⁵ Finally, we investigate individual decisions and consider how well these conform to the optimal strategy that the model predicts.

The data show that the model predicts the general patterns in the experiment very well. The differences between experimental treatments with regard to search length, the usage of freezing, and the payoff to the searcher, are consistent with the comparative statics of the model. The treatment effects are strong despite a relatively small number of observations. There is strong evidence for the use of reservation and double reservation price rules in the predicted manner. A logistic functional form describes the probability of accepting an offer at a given price very accurately. Our model also outperforms four alternatives, with different underlying mechanisms, that have been applied to sequential search in previous literature. We adapt these mechanisms to our setting, where freezing is allowed, and compute the resulting predictions. The alternative mechanisms are: (1) risk aversion, (2) cognitive information acquisition costs, (3) anticipated regret, and (4) inclusion of sunk costs in the payoff calculation. Models based on these mechanisms have been applied to experimental data on search without freezing by Cox and Oaxaca (1989), Gabaix et al. (2006),⁶ Weng (2009) and Kogut (1990), respectively. We consider whether these mechanisms explain our data better than the model we propose in Sect. 3.

The main departure from the model is a modest, though statistically significant, tendency to end searches too early in most treatments. As discussed in Sect. 2, this pattern has also been documented in a number of prior studies in which price freezing is not possible. Beyond documenting that under-searching is common, we

⁵ The experiment is not designed to measure the effect of allowing recall on freezing, search length, earnings, or other variables. Our model's point predictions under no recall are very similar to those obtained under imperfect recall, implying that tests for treatment effects of recall would be underpowered. Rather, the experiment is designed to test the effect of changing the freeze fee and time horizon on behavior in two distinct environments, one with no recall and one with imperfect recall. Varying the extent to which recall is possible can be thought of as a robustness exercise.

⁶ We note that in contrast to the other three mechanisms, the model in Gabaix et al. (2006) was intended to describe search among heterogeneous goods, as in Weitzman (1979). This is an essential difference from our homogeneous goods sequential search setting.

observe two other patterns. Firstly, we show that lowering the freeze fee magnifies the extent to which searchers' exploration is below the optimal level. Secondly, we show that while under-searching is prevalent, its adverse effects on individuals' profits are small. Thus, in the short run, one-sided model we consider, where firms do not respond to consumer behavior, searchers' welfare loss is not substantial.⁷

As we describe in Sect. 4, in some of our treatments, offers that have been rejected cannot be recalled later. These conditions are referred to as the *No Recall* treatments. In other treatments, if an offer is rejected and search continues, there is a positive probability less than one of being able to recall the offer and accept it later. We refer to these as *Imperfect Recall* treatments.⁸ Imperfect recall can typically arise in two ways in the field. Firstly, the good may become sold out. For example, popular concert tickets may run out in minutes. On retailing websites on Cyber Monday, it may be a matter of seconds. In such cases, there is limited opportunity to recall earlier offers. Secondly, prices may exhibit high volatility, so that even if revisiting previous vendors is possible, the price is likely to have changed. In such cases, recall of earlier offers may be possible, but is far from guaranteed. It is evident that the value of a PFO decreases in the recall probability. While a PFO and the ability to recall earlier offers share some similar features, they also have important differences. In particular, a PFO pertains to a specific offer that was chosen by the searcher to be frozen, whereas recall applies to the best offer seen at any prior stage. Our results regarding the effects of PFOs on behavior are robust to both settings with no recall and imperfect recall.

The paper proceeds as follows. In Sect. 2 we discuss related literature. We briefly review the experimental literature on search, and discuss several explanations that have been proposed for the tendency for search to be terminated too early. We then briefly mention the varying assumptions on the recall of prior offers in prior literature. Section 3 develops a theoretical model of sequential search featuring a PFO and characterizes the optimal solution for a risk neutral agent when recall of offers is not possible. The first part of the section assumes that the PFO is exogenous and shows that if it is accepted, the acceptance must occur in either the first period of search or in the last period in which it is possible. The second part extends this analysis to allow the searcher to purchase the PFO at any time. The results in this section serve as the source of hypotheses for our experiment. Section 4 describes our experimental design, indicates the model's specific predictions, and states the hypotheses for our experiment. Section 5 presents the results from the No Recall treatments. We first present results regarding treatment effects and then move on to individual decision patterns. Later in the section, we evaluate the four alternative models mentioned above against the model we derive in Sect. 3. Section 6 summarizes the findings and

⁷ It is conceivable that in the long run, accounting for firms' responses, under-searching can lead to less favorable price distributions, which can have economically significant welfare effects.

⁸ The search literature has mostly focused on the polar cases of perfect recall (usually in the consumer choice context) and no recall (usually in the labor market context). Perfect recall is common in goods markets in which there is no shortage and price volatility is low, whereas no recall (each offer being in a take-it-or-leave-it format) has been considered to be a reasonable description of job search when labor supply is sufficiently high in relation to the number of available jobs.

offers concluding thoughts. The results from treatments with Imperfect Recall are reported in Appendix D.

2 Background

The economic literature on search dates back to the seminal analysis of Stigler (1961). He develops a simultaneous search model, where a consumer decides ex-ante how many price offers to sample. Sequential search, in which a decision maker takes a sequence of decisions about whether or not to accept new offers, was introduced by McCall (1970), DeGroot (1970), Kohn and Shavell (1974) and others.⁹ The basic sequential search model is very simple. Every period, a risk-neutral agent draws an offer from a fixed and known distribution, and then chooses between accepting and rejecting it. Acceptance terminates the decision problem, while rejection moves the task on to the next period, where a new offer is received, and so on. A constant search cost is paid for every offer drawn. Under these simple assumptions, the optimal strategy is a reservation rule (i.e. a cutoff strategy), where the cutoff is chosen so that the expected marginal benefit of continuation to the next stage equals the per-period search cost. It is well-known that this reservation rule is stationary when the time horizon is infinite. Under a finite horizon assumption, the reservation rule has the property that the cutoff is monotonically increasing over time, as individuals become less selective when the expected benefit from continuation decreases.¹⁰

The experimental literature on testing predictions of search models dates back to Kahan et al. (1967) and Rapoport and Tversky (1970). Kahan et al. (1967) vary the distribution of offers and do not observe a strong effect on the quality of decisions. They observe that subjects tested in groups search longer than those participating individually. Rapoport and Tversky (1970) observe decisions close to optimal with some early stopping. Schotter and Braunstein (1981) evaluate various theoretical implications. For instance, they study how search behavior is affected by exogenously-induced risk aversion, changes in the offer distribution and in the information the searcher holds, as well as the degree of recall. They also test whether individuals are following an optimal threshold rule directly by eliciting the payment that subjects require as compensation for not engaging in search. A large part of the subsequent experimental search literature focuses on whether individuals apply reservation rules in optimal stopping problems. While the elegant, simple, and intuitive optimal reservation price rule prediction is one of the most appealing attributes of these models, the typical empirical finding is that individuals tend to stop earlier

⁹ Weitzman (1979) extended the sequential analysis to the case of differentiated goods, where searchers choose which offer to examine before deciding on whether to accept or continue.

¹⁰ Armstrong and Zhou (2016) characterize the combinations of search costs and prices that make accepting a current offer, not accepting it, and taking the outside option, optimal.

than is predicted by this rule (see, for example, the experiments of Cox and Oaxaca (1989), Sonnemans (1998), Kogut (1990), Einav (2005) and Schorvitz (1998)).¹¹

Various explanations have been proposed to rationalize early stopping. We consider how well models based on these explanations predict decisions in our task compared to our model. A commonly discussed explanation is risk aversion, under which searchers value the action of stopping at a premium because it provides a deterministic payoff, avoiding the variability involved in continuation. Thus, risk averse buyers have higher reservation prices. While Cox and Oaxaca (1989) attribute under-searching to risk aversion, Sonnemans (1998) argues that only a small fraction of decisions to stop early can be rationalized by reasonably risk averse preferences.

A second mechanism, based on cognitive information acquisition costs, is proposed by Gabaix et al. (2006). Their behavioral assumption is that at any stage, individuals treat the search problem as having fewer remaining future stages than there actually are. This immediately implies under-searching because reservation prices are increasing over time as the end of the horizon approaches. Their model was intended for application to an environment with heterogeneous goods, in which individuals direct their search. The model successfully predicts experimental results for that setting. Here, we have a homogeneous good setting, no directed search and an additional action, freezing, that can be taken. Thus, applying their model to our data is not a test of their theory, but rather an exploration of whether a similar mechanism might be at work in our environment.

The third mechanism we consider, anticipated regret preferences, is based on the anticipated regret theory of decision making under uncertainty proposed by Loomes and Sugden (1986). In their model, agents anticipate that after uncertainty is resolved, they will compare the realized payoff of the chosen alternative with the payoff that would have been obtained by choosing a different action (and having the uncertainty resolved in the same manner). Weng (2009) shows that in a model with perfect recall, incorporating anticipated regret implies a standard reservation rule strategy, but with higher buyer reservation prices than in the benchmark model, leading to under-searching.

The fourth mechanism that we consider is one proposed by Kogut (1990). Under this account, individuals do not treat the accumulated search costs incurred before the current period as sunk. Instead, though the actual search costs are constant over time, individuals act as if they are increasing. This leads them to stop their search earlier than they would otherwise. We consider a version of this mechanism, with a failure to treat both search costs and freeze fees as sunk.¹²

¹¹ See also Zwick et al., (2003) for a discussion of the motives for under-searching and the possibility of over-searching when searchers use particular heuristic decision rules.

¹² There have been other rationalizations of under-searching in previous literature which we do not consider. For example, under a mechanism proposed by Sonnemans (2000) and formally tested by Einav (2005), under-searching is a consequence of the asymmetric information structure in the feedback searchers receive between search episodes. In particular, subjects cannot observe what they would have been offered had they stopped later. Thus, they can only form downward regret between search problems, where they wish they would have been less picky and not skipped a lucrative offer. This causes under-searching because of the one-sidedness of the implied directional learning: negative feedback is only obtained when searching too long and not when searching too little.

The impact of different assumptions on the capacity to recall previous offers has also been studied. Landsberger and Peled (1977) and Karni and Schwartz (1977) formalize sequential search with imperfect recall, and Janssen and Parakhonyak (2014) analyze the implications of costly recall. Landsberger and Peled (1977) augment the homogeneous good sequential search model by allowing for any recall probability (encompassing perfect, imperfect and no recall). In their model, the recall probability is independent of the time elapsed since the best prior offer, and consumers know the price distribution. They interpret imperfect recall as an indicator of market conditions - prices are less likely to be available for recall in the future when demand is greater or supply is more constrained. In the model of Karni and Schwartz (1977), the recall probability decreases as time elapses since the best offer has been seen, and consumers do not know the price distribution *ex-ante*. They characterize a family of learning processes for which a reservation price strategy is optimal. The way we model imperfect recall is similar to the approach taken in Landsberger and Peled (1977).

3 Theory

This section is organized in the following manner. In Subsection 3.1, we develop a model of sequential search with an exogenously given outside option. In Subsection 3.2, we endogenize the outside option by introducing a PFO, which is in effect an opportunity to purchase the availability of an outside option. We assume a finite horizon throughout the entire section. In both subsections, we assume that recall of prior offers is not possible. We consider the case of imperfect recall in Appendix B, solving it using numerical methods. The model is formulated, in line with the theoretical literature on consumer search, as the decision problem of a buyer facing a sequence of offers from potential sellers and who searches for the lowest price. The theory can be readily translated in a symmetric manner into the environment of the experiment, in which subjects are sellers confronting a sequence of offers to buy, and we perform this translation when we evaluate its predictions. The proofs of the propositions in this section are given in Appendix C.

3.1 Search with an outside option

We begin by analyzing a sequential search model with an outside option, which can also be interpreted as an offer that was previously frozen. Consider a potential buyer of a good, who can receive a price offer in each of a sequence of T stages, indexed by $t \in \{1, \dots, T\}$. Assume that price offers $(p_t)_{t=1}^T$ are independently and identically distributed on $[0, \bar{p}]$, according to a continuous distribution F and are drawn sequentially. Each offer in a stage $t > 1$ is drawn at a cost c (the first offer, at $t = 1$, is free). The searcher is risk neutral and knows the price distribution. The price under the outside option available by withdrawing from the search is denoted by k .

In each stage t , the searcher chooses between (1) accepting p_t , (2) rejecting p_t , and continuing to the next stage after paying the search cost of c , or (3) taking the

outside option k . If the player accepts an offer in stage t , she purchases the item and pays the current offer price p_t . Denote by $\tilde{R}_t^T(k)$ the expected payment that the searcher would make if she rejects at stage t and continues optimally thereafter, when the horizon is of length T . We call $\tilde{R}_t^T(k)$ the *post-freeze rejection payment*. By “payment” we mean the expected price to be paid for the item and all expected search costs to be expended in subsequent stages, provided that the agent proceeds optimally. We suppress the dependence of $\tilde{R}_t(k)$ on T to simplify the notation.

The *post-freeze expected payment* of the individual in stage t is denoted by $\tilde{V}_t(p_t, k)$. This is what the individual expects to pay, in terms of both the price for the item and in search costs, if she makes optimal decisions from stage t onward. It is the minimum of the current price, the outside option, and the post-freeze rejection payment. Thus, when $t < T$, we have

$$\tilde{V}_t(p_t, k) = \min\{p_t, k, \tilde{R}_t(k)\} \tag{1}$$

$$\tilde{R}_t(k) = \int_0^{\bar{p}} \tilde{V}_{t+1}(p_{t+1}, k) dF(p_{t+1}) + c. \tag{2}$$

Rejection in the terminal period T yields a payoff of 0, so that $\tilde{V}_T(p_T, k) = \min\{p_T, k\}$.

The following functions $h(x)$ and $g(x)$ are useful in deriving some of our results.¹³

$$h(x) = \int_0^{\bar{p}} \min\{p', x\} dF(p') = \int_0^x p' dF(p') + (1 - F(x))x \tag{3}$$

$$g(x) = x - h(x) = \int_0^x F(p') dp'. \tag{4}$$

A threshold strategy is one in which there exists a cutoff price for every stage. The buyer accepts all offers below the cutoff, and rejects all prices above it. Consider the function $h(x)$, as formulated in equation (3). The expression denotes the expected price paid in the last search period when an outside option is available. The first term describes the case of an offer lower than the outside option, in which case the offer is accepted. The second term corresponds to the case in which the offer is greater than the outside option, in which event the outside option is taken. Notice

¹³ $h(x)$ is the post-freeze rejection payment in stage $T - 1$, net of search costs, when holding x as an outside option. $g(x)$ is the difference between accepting the outside option and rejecting, net of search costs, in stage $T - 1$.

that $h(\bar{p}) = E(p)$, so that if the buyer adopts the strategy of accepting any price offered, she accepts all offers and pays in expectation the average offer.¹⁴

Assumption 1 $c < g(\bar{p})$.

Let $p^* = g^{-1}(c)$. It is well known that p^* is the stationary optimal reservation price in a sequential search model with no recall and an infinite horizon.¹⁵ This price is the basis of our first proposition, which describes the optimal decision rule for the buyer in the presence of an outside option. The proposition states that if the outside option price is lower than p^* , the buyer accepts the outside option in the first stage. If not, the buyer uses a reservation price strategy with a dynamic threshold that increases in each stage.

Proposition 1 *When $k \leq p^*$, search ends immediately by accepting either p_1 or k . When $k > p^*$, an increasing reservation price strategy is optimal and k is either never chosen or chosen at the terminal stage.*

Figure 1 shows the optimal strategy in p -space. The bold line indicates the optimal decision rule, the ranges of offers for which it is optimal to accept p_t , settle for k , or reject p_t . The left panel shows that the post-freeze rejection payment is always higher than the outside option if $k \leq p^*$. The right panel illustrates how, in the case where $k > p^*$, post-freeze rejection payments increase over time as the end of the horizon approaches.

We end this subsection by noting a limit result, which we use in Subsection 3.2. Namely, provided that $k > p^*$, for any current stage t , as the number of future stages $T - t$ becomes large, the post-freeze rejection payment approaches p^* .

Proposition 2 *The sequence $(\tilde{R}_t^T(k))_{T=2}^\infty$ converges uniformly to $\tilde{R}_t^\infty(k) = p^*$ for $k > p^*$.*

3.2 Endogenous freezing of offers

Now suppose that the searcher can, at any stage, purchase an option to freeze the current offer, in effect buying an outside option of the type described in the previous

¹⁴ We make the standard assumption that the search cost c is sufficiently low for search to be initiated in the first place. Otherwise, for any T , the decision in period $T - 1$ is to accept any offer. Thus, period $T - 1$ can essentially be viewed as the terminal period. Proceeding with this logic and applying backward reasoning, it is clear that if assumption 1 is violated, search terminates in the first period.

¹⁵ It is also known (e.g. Lippman and McCall (1976) and Landsberger and Peled (1977)) that this price is the optimal reservation price in a sequential search model with perfect recall and a finite horizon. There, p^* can be interpreted as the price that makes the individual indifferent between (a) recalling and accepting the current best offer, and (b) keeping the current best offer as an outside option and searching for one more period.

subsection. Assume that the searcher can freeze one offer by paying a fixed fee $f > 0$, which we will refer to as the *freeze fee*. Once one offer is frozen, a second offer may not be frozen.¹⁶ Offers may not be unfrozen. There is no other outside option available. To analyze this situation, we introduce several functions that are analogous to those used in Subsect. 3.1. $V_t(p_t)$ is the pre-freeze expected payment, R_t is the pre-freeze rejection payment, and $K_t(p_t)$, which we shall refer to as the *freeze payment*, is the expected payment when freezing p_t and continuing to the next stage with p_t as an outside option. Formally,

$$V_t(p_t) = \min\{p_t, R_t, K_t(p_t)\} \tag{5}$$

$$K_t(p_t) = \tilde{R}_t(p_t) + f \tag{6}$$

$$R_t = \int_0^{\bar{p}} V_{t+1}(p_{t+1})dF(p_{t+1}) + c. \tag{7}$$

We begin with a lemma stating the straightforward fact that the pre-freeze rejection payment is weakly lower than the worst possible post-freeze rejection payment.

Lemma 1 $R_t \leq \tilde{R}_t(\bar{p})$ for all $t < T$

Next, we make the assumption that the freeze fee is sufficiently low to guarantee the existence of prices that are frozen. To assure this, we assume that the fee is low enough so that there exists some price that would be frozen in the next-to-last stage:

Assumption 2 The freeze fee f satisfies

$$f < \int_{h(\bar{p})+c}^{\bar{p}} [p - (h(\bar{p}) + c)]dF(p) \tag{8}$$

$$=h(\bar{p}) - h(h(\bar{p}) + c) \tag{9}$$

To gain intuition for this condition, notice that it can be written as $K_{T-1}(R_{T-1}) < R_{T-1}$. This means that at $T - 1$, it is better to freeze an offer equal to the pre-freeze rejection payment than rejecting. We now turn to the optimal strategy of the consumer when a PFO is available. We use the following terminology to describe the optimal rule.

Definition 1 Under a reservation rule (RR), in stage t ,

$$V_t(p_t) = \begin{cases} p_t & \text{if } p_t < R_t \\ R_t & \text{otherwise.} \end{cases} \tag{10}$$

¹⁶ If a person can buy a second option, it simplifies the problem. It introduces a source of stationarity in the environment that we do not have here. In our environment, the decision problem changes when one has already frozen an offer in a way that would not occur if one could freeze more than one offer.

Under a double reservation rule (DRR), in stage t ,

$$V_t(p_t) = \begin{cases} p_t & \text{if } p_t < a_t \\ K_t(p_t) & \text{if } p_t \in [a_t, b_t) \\ R_t & \text{otherwise} \end{cases} \quad (11)$$

where $0 < a_t < b_t < \bar{p}$. A subset of stages in which a RR is used is denoted by $T^{RR} \subseteq T$ and a subset of stages in which a DRR is used is denoted by $T^{DRR} \subseteq T$.

A RR specifies a stage-specific threshold price below which offers are accepted and above which they are rejected. A DRR specifies a stage-specific threshold below which offers are accepted, another one above which offers are rejected, and an intermediate range between the two thresholds, in which offers are frozen. We define a strictly increasing RR sub-policy to be one in which $R_t < R_{t'}$ for $t < t'$, $\{t, t'\} \subseteq T^{RR}$. Similarly, we define a strictly increasing DRR sub-policy to be one in which $a_t < a_{t'}$ and $b_t < b_{t'}$ for $t < t'$, $\{t, t'\} \subseteq T^{DRR}$. That is, a DRR is strictly increasing if the acceptance region strictly increases and the rejection region strictly decreases over time.

The following lemma allows us to restrict ourselves to particular strategies when solving for the optimal policy. The lemma shows that the consumer will always use either a RR or a DRR at every stage.

Lemma 2 *Under the optimal solution, $t \in T^{RR} \cup T^{DRR}$ for all $t < T$.*

The following lemma is a monotonicity result which is useful for characterizing the optimal solution. It shows that the range of offers that is rejected decreases over time.

Lemma 3 *Any offer that is rejected in some period t is also rejected in period $t - 1$. That is,*

$$V_t(p) = R_t \implies V_{t-1}(p) = R_{t-1} \quad (12)$$

In Proposition 3, we describe the structure of the optimal policy for the consumer to follow.

Proposition 3 *In a given stage t , as long as an offer was not frozen before, the optimal policy consists of a strictly increasing RR for $t \leq t^*$ and a strictly increasing DRR for $t > t^*$, where $0 \leq t^* < T - 1$. If an offer was frozen before stage t , then the optimal policy follows an increasing RR thereafter. In period T it is optimal to accept p_T or the previously frozen offer in the event that such an offer exists, whichever is lower.*

Thus, the solution can take two forms. When $t^* = 0$, the optimal policy is an increasing DRR. When $t^* > 0$, an increasing RR is optimal in stages 1 through t^* followed by an increasing DRR thereafter. We now present a final result, which

states that when the horizon is sufficiently long, there will be at least one stage at the beginning of the search sequence where it is optimal not to freeze any offer.

Proposition 4 *There exists a $T^* < \infty$ such that for $T > T^* \implies t^* > 0$.*

The left panel of Fig. 2 illustrates a stage in which a DRR is optimal ($t \in T^{DRR}$). The pre-freeze expected payment for the stage is given by the lower envelope of the functions p_t , $K_t(p_t)$, and R_t . The figure illustrates how a_t is a cutoff value, above which it is optimal to freeze an offer, and below which it is optimal to accept. b_t defines a similar threshold between freezing and rejection as optimal actions. The right panel plots the solution at $t = 1$ when $T \rightarrow \infty$, where we must have $1 \in T^{RR}$.

4 Experimental design

4.1 General structure

The sessions were conducted in the Economic Science Laboratory at the Eller College of Management of the University of Arizona in late 2016. All subjects were undergraduate students at the university. A total of 177 subjects participated in the experiment and the number of individuals present varied across sessions. The experiment consisted exclusively of individual choice tasks. Upon arrival, a first set of instructions, which pertained to the risk elicitation protocols, was read aloud by the experimenter. Subjects then performed the protocols. Upon completing these two tasks, a second set of instructions, describing the 180 search problems participants were about to face, were read aloud. Afterward, the main part of the experiment began.¹⁷ These were followed by the main part of the experiment, which consisted of 180 search problems that had the possibility of counting toward earnings.¹⁸ The session concluded with a brief questionnaire. Subjects could complete the sequence of tasks at their own pace. The entire sequence of tasks took between 60

¹⁷ The results from these risk elicitation protocols are given in Appendix H. The measures were uncorrelated with each other, and thus were not amenable to constructing a convincing overall measure of risk aversion. Therefore, we do not use them in our analysis. Schunk and Winter (2009) examine whether risk and loss aversion correlate with behavior in a search task. Their task differs from ours in a number of important ways, including in that the horizon is indefinite and some offers are greater than the willingness-to-pay of the searcher. They find that their independent measures of risk aversion do not correlate with search length, as we do here. However, they do find that measures of loss aversion correlate with the length of the search, with more loss averse individuals terminating their search earlier.

¹⁸ One of the 182 tasks counted toward participant earnings, and the task that counted was determined by a random draw at the end of the session. With only one of the tasks counting, the expected payoff of each task becomes relatively small. In principle, this may serve to weaken the incentive to make good decisions. However, a number of studies have shown that within-subject randomization protocols do not lead to different decisions than if each decision counted with probability one (Starmar and Sugden (1991); Cubitt et al. (1998); Laury (2005); Hey and Lee (2005); Baltussen et al. (2011)) despite the lower expected stakes.

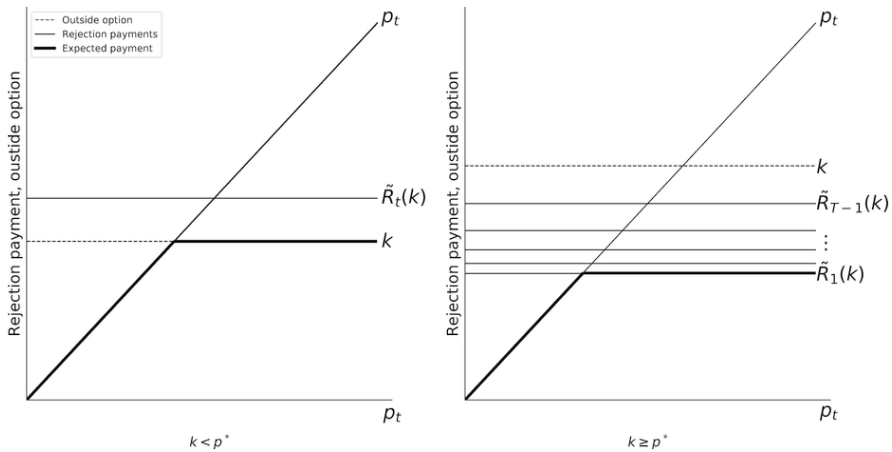


Fig. 1 Post-freeze expected payment in p -space

and 100 minutes to complete. The experiment was programmed using Z-tree (Fischbacher (2007)).

The 180 search problems that could count toward subjects' earnings were divided into three equal blocks of 60 search problems, as described in Subsection 4.2. There was a mandatory two-minute pause between each block of 60 trials to allow participants to rest. There was also a requirement that subjects stay in the laboratory for at least one hour, to prevent them from completing the task as rapidly as they could in order to leave the session early.¹⁹ Subjects were paid for one randomly selected task (either one of the two risk measurement tasks or one of the 180 search problems), plus a \$5 show up fee.²⁰ Each of the 182 tasks was equally likely to be selected to count toward participants' earnings.

There are two notions of time in the experiment. We will use the term *stage* to refer to each time the subject must make a decision on an offer (as this term was used in Sect. 3), whereas the term *round* refers to an entire search problem, which consists of a sequence of stages. We use the terms *round* and *sequence* interchangeably. Our experiment, therefore, includes 180 rounds, and each round consists of multiple stages.

The subjects in the experiment are potential sellers of a fictitious item. In each stage of a round, subjects receive an offer, drawn from a discrete uniform distribution on $\{0, \dots, 1000\}$, in which each of the 1001 integers in the range is equally likely. A search cost of $c = 10$ is paid for every offer (except for the offer in the first stage of each round). The offers in each stage are independent of those in preceding

¹⁹ The use of cellphones or any other electronic devices was forbidden for the entire session, even for subjects who had completed the task. This rule made it more salient to subjects that there was nothing to gain by rushing their decisions.

²⁰ Appendix E includes the instructions and screen shots of the interface.

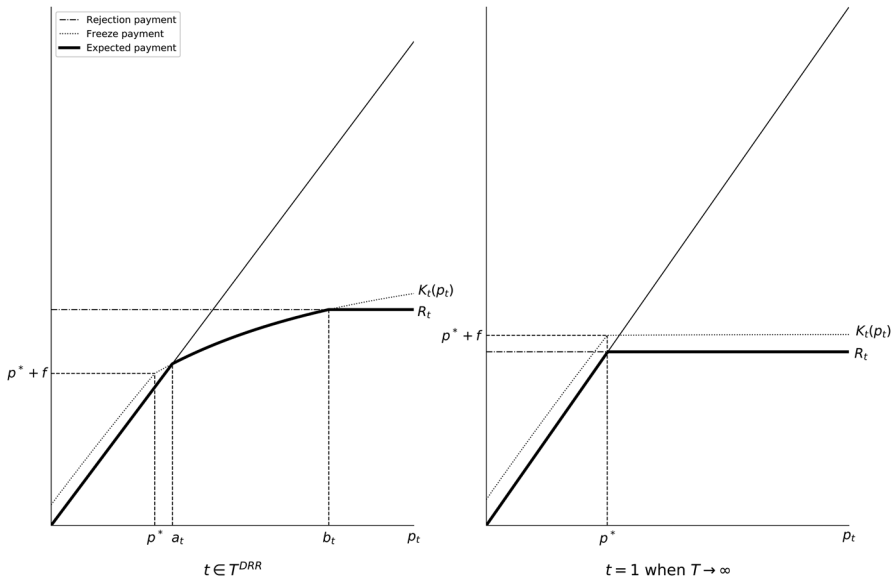


Fig. 2 Example of a DRR, and of behavior as $T \rightarrow \infty$

or subsequent stages. A player may accept an offer at any stage. If she accepts an offer, she receives the offer price minus the accumulated search costs within the round. The round ends when an offer is accepted. If she rejects the offer in a given stage, the round continues to the next stage. Rejection is not possible in the terminal stage of a round. Offers and costs are denominated in terms of an experimental currency, which is convertible to US dollars at a rate of 70 to 1.

4.2 The three 60-period blocks

Every subject faced the exact same 180 offer sequences. These 180 sequences consisted of 60 sequences that were drawn in advance. The sequences were repeated three times, with a modification that we describe below. Thus, there were three blocks of 60 rounds. The environment was identical for all individuals, except for the exogenous variations in the freeze fee, horizon length and the recall probability (denoted as f , T and q respectively), described in the next subsection. The number of stages T and the recall probability q were varied between subjects, while f was varied within subjects across blocks. Each 60-period block was preceded by four practice periods.

The first block of sequences, those employed in rounds 1–60, were random independent draws from $U\{0, \dots, 1000\}$. In the second and third blocks, subjects face

perturbed offers. These are created by adding a random, relatively small, integer to the corresponding offers from the first block.²¹

4.3 The treatments

In different treatments, we vary the freeze fee f , the time horizon T , and the recall probability q . The freeze fee is varied within-subject across blocks. Under the *Low Freeze Fee (Lo)* condition, individuals must pay 10 to freeze an offer, and under *High Freeze Fee (Hi)* they must pay 40. In a third condition, *No Freezing (No)* freezing is not possible, so that the freeze fee can be thought of as infinite. In some sessions, the freeze fee is varied in ascending order across the three blocks; *Lo* in the first 60 rounds of the search task, *Hi* in rounds 61–120, and *No* in the last 60 rounds. In other sessions, the reverse sequence is in effect, and the fees appear in descending order across the three blocks. The cost c of generating an offer in the next stage is always 10.

T and q are varied between subjects. The time horizon is fixed at $T = 4$ in some sessions and $T = 10$ in others. The recall probability is set to $q = 0$ in some sessions (No Recall or NR conditions) and $q = .5$ (referred to as Imperfect Recall or IR) in others. When $q = .5$, recall during a given stage of a sequence is possible with probability .5. When recall is possible in stage t , the highest offer from stages 1 to $t - 1$ may be recalled and accepted. Whether recall is available within a given stage of a sequence is independently drawn in each stage. For example, recall of prior offers from the first two stages may not be possible when facing the third offer in a sequence, but may be possible when facing the fourth.²²

Thus, the experiment has a $2 \times 2 \times 3$ structure. We refer to the treatments in an abbreviated form by the time horizon (4 or 10 periods), whether recall was not possible or imperfect (IR or NR), and the level of the freeze fee (Lo, Hi or No). For example, in *Hi10IR*, subjects could sample up to 10 offers in a round, the high freeze fee of 40 was in effect, and there was a recall probability of .5. Our hypotheses

²¹ If the last two digits of the offer in the first block are not in $\{98, 99, 00, 01\}$, then the random number which is added to the corresponding offer from the first block is in $\{-2, -1, 1, 2\}$, each value with equal probability. If the last two digits are in $\{98, 99, 00, 01\}$ then the support for the random draw is such that perturbed offers do not change the hundreds digit, and each value in the support is drawn with equal probability. For example, if an offer in the first block was 298, then the support of the perturbations consists of -2, -1 or 1, but not 2, because that would imply changing the first digit from 2 to 3. The idea behind this distinction is to avoid biases similar to those that arise from “.99 cent pricing”, as in the theory proposed by Basu (2006) and documented experimentally by Ruffle and Shtudiner (2006), in which the last two digits in the price are ignored and thus the third-to-last digit has great prominence. By perturbing the offers in the second and third blocks, in contrast to presenting subjects with the exact same offers, we presumably guarantee that even if subjects remember a particular sequence or sub-sequence, they would not see the exact same sequence again. This is intended to make them, in practice, unable to remember offer sequences or sub-sequences while allowing us to keep the sequences nearly identical across treatments.

²² This situation arises, for example, if you would like to purchase a coat and find that it is currently out of stock. If you return to the shop later and a new shipment of coats has come in, the coat becomes available again after being unavailable.

concern comparisons between different levels of f and T , which are evaluated under No Recall and Imperfect Recall separately.

Table 1 provides some information about the participants in each treatment, in terms of gender distribution, number of years of university study, the number of previous economic experiments and the percentage of participants who studies economics or business. Our sample is relatively experienced in participating in experiments. This is reassuring as this increases our confidence that subjects are aware of the direct positive relationship between understanding the instructions and expected earnings. Participant earnings averaged \$15.30 with a standard deviation of \$3.66. The sessions averaged approximately 1 hour and 25 minutes in duration.

4.4 Optimal decisions

The optimal strategy, using the parameters described in Subsection 4.3, in each of the No Recall and Imperfect Recall treatments, are illustrated in the panels on the left and right sides of Fig. 3, respectively. In each of these panels, we plot the solution for one level of f , when no offer has yet been frozen, for periods 1 to $T-1$.²³ The black region denotes offers that are accepted, gray stands for offers that are frozen, and offers in the white area are rejected. For the IR treatments, we plot the solutions for the case in which the highest offer seen so far is zero, implying that recall is not available (so that the value of recalling enters through the freeze and pre-freeze rejection payments only).

The figure shows that the optimal thresholds under No Recall are monotonic and concave. The acceptance and freezing regions are increasing over time. More freezing occurs when the freeze fee is low. For example, in the *Lo* conditions ($f = 10$), the freezing region is already larger than 10% of the range of possible offers in stage $T - 9$, whereas under *Hi* the freezing region surpasses 10% only in stage $T - 5$. The horizon $T = 10$ is long enough so that no freezing is predicted to occur in the initial stage. In the Imperfect Recall case, optimal thresholds are also monotonic, and are concave once the DRR sub-policy is in effect.

4.5 Hypotheses

The hypotheses guiding the design of our experiment are derived from the theoretical results presented in the Sect. 3 for the No Recall treatments, and from the computation of the optimal solution for the Imperfect Recall treatments, as described in Appendix B. We solve the models for each of our treatments using the offer and recall realizations faced by subjects in the experiment. Table 2 presents the resulting mean search lengths, the percentage of rounds in which it is optimal to freeze an offer, the average earnings of the searcher (which equal 1000 minus the price

²³ In the $T = 4$ treatments, the decisions from stages $T - 3$ to $T - 1$ constitute the model's prediction for stages 1–3, while for $T = 10$, stages $T - 9$ to $T - 1$ correspond to stages 1–9. In the last stage T , all offers are accepted.

Table 1 Subject numbers and characteristics in the different treatments

| T | No recall | | Imperfect recall | |
|-----------------------------|-----------|-------|------------------|-------|
| | 4 | 10 | 4 | 10 |
| Num. participants | 42 | 42 | 51 | 40 |
| Perc. female | 33.33 | 40.48 | 56.86 | 55.00 |
| Education | 2.62 | 2.74 | 2.90 | 2.94 |
| No. prior experiments | 9.75 | 8.88 | 6.90 | 8.11 |
| Perc. economics or business | 34.09 | 59.52 | 60.78 | 57.50 |

Averages of subject characteristics in each (T, q) treatment. The variables are the number of participants in each treatment, percentage of females, average years of university education, average number of economic experiments participated in previously, and percentage of subjects who are studying economics or business

paid minus total search costs minus the freeze fee if an offer is frozen) and the other party's surplus (the earnings of a hypothetical individual on the other side of the market, whose surplus equals 1000 minus any offer accepted, plus any freeze fee paid). The numbers in the table reveal some of the comparative statics of the model. On average, a higher freeze fee results in shorter search and less freezing. It can also be seen from the table that increasing the horizon T results in longer search and higher earnings. The effect on search length of changes in the freeze fee and horizon length constitute our Hypothesis 1. Hypothesis 2 concerns use of the freeze option, and asserts that it is more common when the freeze fee is lower and when the horizon is shorter.²⁴

Within Hypothesis 2, our model makes a prediction regarding the effect of freezing that at first glance may seem counterintuitive. One's intuition might be that, just like a financial option, freezing is more valuable when there is more time remaining during which one can accept the frozen offer, and thus for a given freeze fee, one might be more likely to pay to freeze an offer when there are more stages remaining. However, the optimal policy has the property that freezing is more likely when the horizon is relatively short (4 periods) than when it is long (10 periods). We interpret support for this prediction as strong evidence in favor of our model. Our third hypothesis is that individuals employ the optimal RR-DRR policy. That is, they make acceptance, freezing, and rejection decisions that are consistent with the model. They employ the thresholds depicted in Fig. 3.

Summarizing, the experiment is designed to test the following hypotheses. If all three hypotheses are supported, we would conclude that our model is strongly supported.

²⁴ Participants had strong incentives to make good decisions. A simple rule of accepting the first offer one gets, which gives an expected payoff of 500 ECU (\$7.14) under either $T = 4$ or $T = 10$. A simple rule of accepting the first offer that is above 500 ECU and never freezing, gives an expected payoff of 652.5 ECU (\$9.32) under $T = 4$ and 739.1 ECU (\$10.55) under $T = 10$. These rules yield expected payoffs well below the optima and well below the averages realized in the experiment.

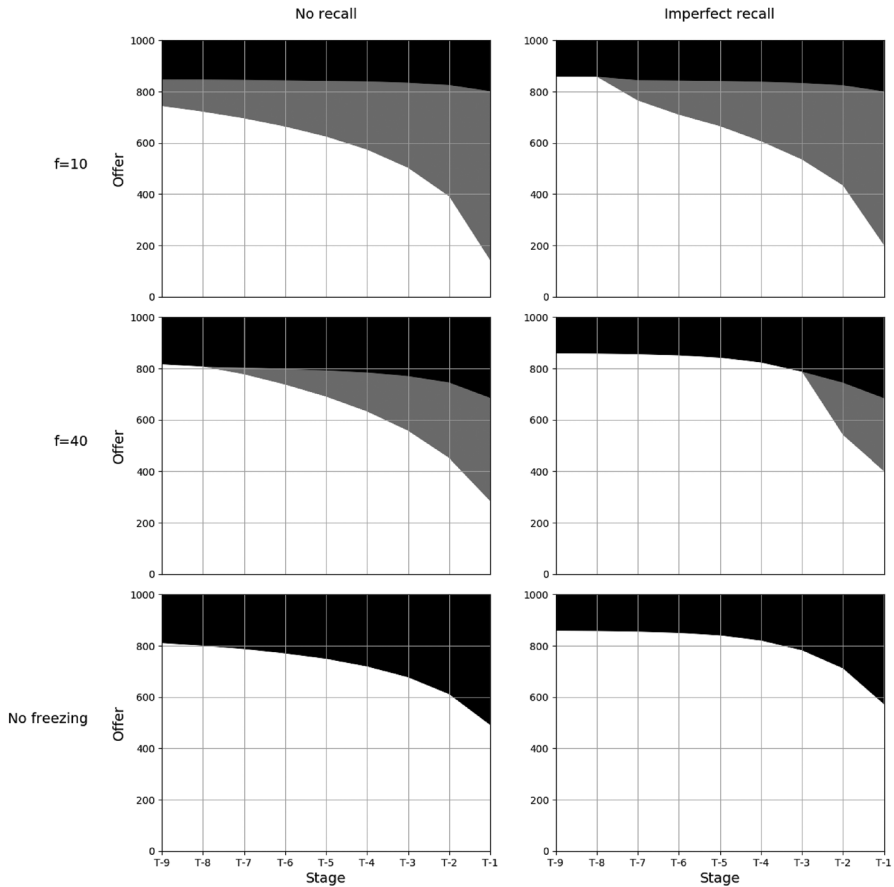


Fig. 3 Optimal policy with $p \sim U[0, 1000]$. Optimal behavior, according to the model, for the different treatments. The left panels correspond to the case of No Recall. The right panels are for Imperfect Recall, assuming that in each stage the highest offer observed so far is zero, so that the impact of recall is through the freeze and pre-freeze rejection payments when freezing or rejecting. Each row plots the solution for a different freeze fee condition. Offers are accepted in the black area, frozen in the gray area, and rejected in the white area. In treatments where $T = 10$ ($T = 4$), the data from $T - 9$ ($T - 3$) onward are applicable

Hypothesis 1: Search length is longer when the freeze fee f is smaller and the time horizon T is greater.

Hypothesis 2: Freezing is less frequent when the freeze fee f is greater and the time horizon T is greater.

Hypothesis 3: Individuals employ the optimal RR-DRR policy.

Table 2 Model predictions, by treatment

| | T = 4 | | | T = 10 | | |
|----------------|------------------|--------|--------|--------|--------|--------|
| | f = 10 | f = 40 | f = ∞ | f = 10 | f = 40 | f = ∞ |
| | No recall | | | | | |
| Search length | 2.98 | 2.58 | 2.23 | 4.62 | 4.05 | 3.92 |
| Freezing usage | 0.72 | 0.05 | | 0.38 | 0.12 | |
| Earnings | 791.20 | 772.82 | 764.58 | 852.82 | 853.85 | 850.07 |
| | Imperfect recall | | | | | |
| Search length | 2.98 | 2.63 | 2.57 | 4.63 | 4.37 | 4.35 |
| Freezing usage | 0.68 | 0.15 | | 0.18 | 0.03 | |
| Earnings | 791.53 | 782.32 | 781.32 | 857.93 | 859.47 | 860.97 |

This table includes the model's predictions for key variables, computed for the actual offers and recall realization sequences drawn in the experiment, and aggregated across rounds. Search length is the mean number of stages. Freezing usage is the percentage of rounds where it is optimal to freeze an offer in any stage. Earnings are the payoff of the decision maker

5 Results

5.1 Summary statistics

Figure 4 illustrates some general patterns in the data from each of the twelve treatments. This figure shows the means and percentages of key variables in the different treatments, aggregated over subjects and rounds, compared to the theoretical predictions (represented by circles and *x*s, respectively). The data displayed are the average search length (panel 4a), the percentage of rounds in which an offer has been frozen (panel 4b) and the earnings per round (panel 4c)). In panel 4b, the different shades indicate the eventual fate of frozen offers, whether they are accepted in the last possible round, taken in a round other than the last, or not accepted at all.²⁵

Panel (a) shows that search length tends to be greater for smaller f and larger T under both No Recall and Imperfect Recall. Search length tends to be shorter than predicted. Panel (b) shows that there are more offers frozen under $T = 4$ than $T = 10$. These patterns are consistent with our model. While under $T = 4$, the vast majority of acceptances of frozen offers are in the terminal stage, as predicted, this is not the case for the $T = 10$ treatments. This raises some doubt about the validity of the model's prediction that frozen offers are only accepted in the terminal round. Panel (c) reveals that earnings are greater when the time horizon is longer, as the model predicts.

²⁵ Figure 11 in Appendix F contains the same information as Fig. 4, but showing only the second half of each block (30 rounds instead of 60). In these 30 rounds, the average decision takes place with greater prior experience with the particular freeze fee in effect, and can thus be presumed to reflect more informed decision making. The results are qualitatively similar with regard to the comparative statics between treatments.

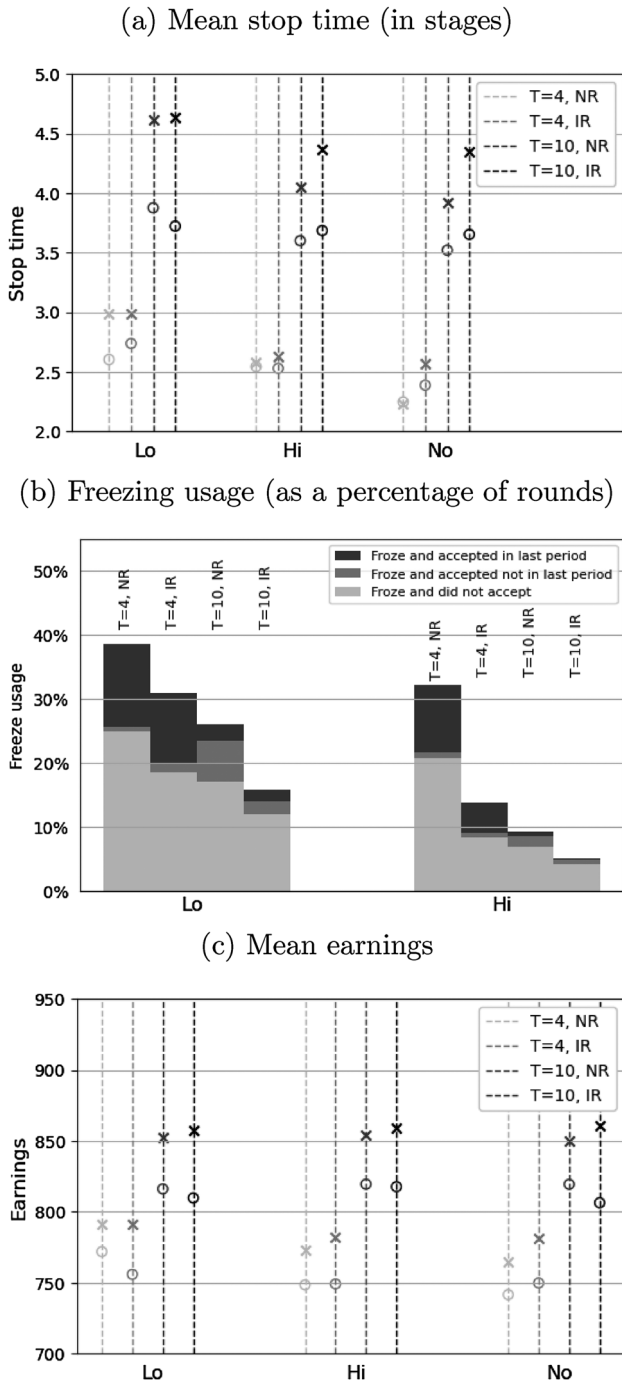


Fig. 4 Within-Treatment Means and Frequencies. This figure contains the observed within-treatment means and frequencies for key variables (as circles) and the theoretical predictions (as xs). Panel (b) also indicates whether frozen offers were accepted later and whether the acceptances occur in the last stage

Table 3 The effect of treatment on earnings, search length, and usage of the freezing option

| | (1) Earnings | (2) Search length | (3) Freeze usage |
|--------------------|--------------------|----------------------|---------------------|
| Lo4 (β_1) | 771.973 (1.702) | 2.603 (0.048) | 0.385 (0.031) |
| | 791.200 0.000 | 2.983 0.000 | 0.7167 0.000 |
| Hi4 (β_2) | 748.526 (3.720) | 2.544 (0.052) | 0.321 (0.031) |
| | 772.817 0.000 | 2.583 0.464 | 0.400 0.013 |
| No4 (β_3) | 741.624 (4.027) | 2.248 (0.031) | |
| | 764.583 0.000 | 2.233 0.626 | |
| Lo10 (β_4) | 816.255 (9.128) | 3.875 (0.135) | 0.259 (0.026) |
| | 852.817 0.000 | 4.617 0.000 | 0.3833 0.000 |
| Hi10 (β_5) | 819.587 (7.956) | 3.600 (0.128) | 0.092 (0.019) |
| | 853.850 0.000 | 4.050 0.001 | 0.1167 0.186 |
| No10 (β_6) | 819.502 (7.278) | 3.521 (0.115) | |
| | 850.067 0.000 | 3.917 0.001 | |
| R^2 | 0.035 | 0.082 | 0.162 |

Notes: Each of the dependent variables is regressed on six treatment dummy variables, corresponding to each of the no recall treatments. Standard errors are clustered by subject. In each row, there are four entries: (1) the estimated coefficient, (2) the standard error of the coefficient estimate, (3) the value predicted by the theoretical model (as given in Table 2), and (4) the p -value from a t -test of the equality of the estimated coefficient to its value as predicted in the model

5.2 Treatment effects

In this section, we consider the effect of treatments in the No Recall treatments. The estimated treatment effects on search length, freezing usage and earnings are reported in Tables 3 and 4. In Table 3, we regress each of these dependent variables on six treatment dummy variables, corresponding to each of the No Recall treatments. The regression has no constant and standard errors are clustered by subject. Therefore the coefficients can be interpreted as treatment averages. In each row, there are four entries: (1) the estimated coefficient, (2) the standard error, (3) the

Table 4 Hypothesis tests of treatment differences

| Null hypothesis | | <i>p</i> -value |
|---|--------------------------|-----------------|
| Earnings decrease in <i>T</i> | (1) $\beta_1 > \beta_4$ | 0.000* |
| | (2) $\beta_2 > \beta_5$ | 0.000* |
| | (3) $\beta_3 > \beta_6$ | 0.000* |
| Earnings increase in <i>f</i> when <i>T</i> = 4 | (4) $\beta_1 < \beta_2$ | 0.000* |
| | (5) $\beta_1 < \beta_3$ | 0.000* |
| | (6) $\beta_2 < \beta_3$ | 0.004 |
| Earnings unchanged by <i>f</i> when <i>T</i> = 10 | (7) $\beta_1 = \beta_2$ | 0.474 |
| | (8) $\beta_1 = \beta_3$ | 0.613 |
| | (9) $\beta_2 = \beta_3$ | 0.977 |
| Search length decreases in <i>T</i> | (10) $\beta_1 > \beta_4$ | 0.000* |
| | (11) $\beta_2 > \beta_5$ | 0.000* |
| | (12) $\beta_3 > \beta_6$ | 0.000* |
| Search length increases in <i>f</i> | (13) $\beta_1 > \beta_2$ | 0.093 |
| | (14) $\beta_1 > \beta_3$ | 0.000* |
| | (15) $\beta_2 > \beta_3$ | 0.000* |
| | (16) $\beta_4 > \beta_5$ | 0.000* |
| | (17) $\beta_4 > \beta_6$ | 0.000* |
| | (18) $\beta_5 > \beta_6$ | 0.091 |
| Freeze usage decreases in <i>T</i> | (19) $\beta_1 < \beta_4$ | 0.001* |
| | (20) $\beta_2 < \beta_5$ | 0.000* |
| Freeze usage decreases in <i>f</i> | (21) $\beta_1 < \beta_2$ | 0.033 |
| | (22) $\beta_4 < \beta_5$ | 0.000* |

Notes: Tests based on estimates in Table 3. * denotes significant treatment difference at $p < .05$ after correcting for multiple hypothesis testing

value predicted by the model (as given in Table 2), and (4) the *p*-value from a *t*-test for the equality of the coefficient to its predicted value.

In Table 4 we use the results of Table 3 to conduct hypothesis tests for each of our directional hypotheses. The first column contains the null hypothesis, the second column is its interpretation in terms of the coefficients from the regression in Table 3 and the third column presents the *p*-value of the hypothesis test. To correct for multiple hypothesis testing, we require a Bonferroni correction factor of 19 for the 19 directional hypotheses in the table so that we consider a result as significant at $p < .05/19 = .0026$. To correct for testing the three hypotheses of no treatment differences, we consider $p < .05 * 3 = .15$ as a conservative threshold to reject the hypothesis of no difference. Table 4, along with the patterns shown in Fig. 4, provide the basis for our first four results.

Result 1 The length of search is (1) decreasing in the freeze fee and (2) increasing in the horizon length, as predicted by the model. Search length is lower than predicted in most treatments.

Support for Result 1: The tests reported in Table 4 show that search length exhibits the predicted treatment differences as T changes in all three treatment comparisons (tests 10–12). It moves in the predicted direction as f changes in 4 of 6 cases (tests 13–18). Table 3 shows that search length is significantly less than the model prediction in all treatments except for *Hi4* and *No4*. \square

Result 2 The usage of the freezing option decreases in the freeze fee and the horizon length, as predicted by the model. There is less freezing than predicted by the model. When $T = 4$, the large majority of frozen offers are accepted in the last stage, as predicted, whereas this is not the case when $T = 10$.

Support for Result 2: The last four tests in Table 4 indicate that the incidence of freezing is significantly lower when $T = 10$ than when $T = 4$. It is also decreasing in the freeze fee f and the effect is significant when $T = 10$. The last column of Table 3 shows that there is significantly less freezing than predicted in every treatment. Figure 4 shows that under $T = 4$, the majority of the acceptances of frozen offers do occur in the last stage, as the model predicts. \square

Result 3 Earnings are affected by freeze fees when $T = 4$, but not when $T = 10$, as predicted by the model. Earnings are 96.6% of the optimal level.

Support for Result 3: The first nine tests of Table 4 show that eight out of nine treatment effects are in the same direction as those that the model predicts. Tests (4)–(9) reveal that earnings are affected by freeze fees when $T = 4$, but not when $T = 10$. Tests (1)–(3) in the table show that earnings are greater in the $T = 10$ than in the $T = 4$ treatments, as the model anticipates. The estimates in the first column of Table 3 indicate the average earnings compared to the model prediction. Averaging the data across treatments reveals that earnings average 96.6% of the predicted level. \square

As indicated earlier, some previous studies have found under-searching (searching for fewer stages than is optimal) relative to the optimal length of search when no freezing is possible. As summarized in Result 4, this general tendency, though modest, also appears when freezing is possible. Our next result is that low freeze fees tend increase the tendency to under-search.

Result 4 The extent to which individuals under-search decreases in f . That is, the availability of an affordable freezing option amplifies under-searching.

Support for Result 4: The left half of Table 5 presents the average number of stages that individuals under-search. In all treatments, except for *No4*, there is under-searching, albeit by negligible amounts in some treatments. In the $T = 10$ treatments, the under-searching ranges between 0.39 and 0.74 stages, with the extent of under-searching monotonically decreasing in f .

To shed light on heterogeneity in under-searching, we classified whether each of our subjects was an under-searcher. We ran one-tailed Wilcoxon signed-rank tests, one for each subject, where the null hypothesis is that the observed search length is

equal to or higher than the predicted search length. The right panel of Table 5 presents, for each treatment, the percentage of subjects for whom the null hypothesis is rejected at a 0.05 significance level. Thus, the table can be interpreted as the share of under-searchers in each treatment. The share of under-searchers is greater, the lower the freeze fee, and the longer the time horizon T .²⁶ □

5.3 Individual responses to changes in freeze fees

Our experimental design allows us to examine the effects of a higher freezing fee, within-subject. Tables 10–12 in Appendix F present augmented confusion matrices, portraying the predicted effects of freeze fee variation on search length, PFO usage and earnings, respectively. Since we have paired observations, the comparisons are for the same round and thus involve the same sequence of offers (modulo the minor perturbation) and the same person, and only the freeze fee varies. For each freeze-cost variation (given in the first column), we report in the third column the overall percentage of rounds in which this variation was predicted to yield a positive, negative or no effect on the outcome variable. Then, in the next three columns we report the percentage of rounds where the outcome variable increased, decreased, or stayed the same in the instances in which the model made each possible prediction. The last column reports, out of the observations where either a positive or negative change was observed, the percentage where the change was as predicted. This last column aims to measure the extent to which subjects' behavioral adjustments are in the same direction as the theory predicts. The unit of observation is the behavior of an individual subject in one of the 60 matched pairs of sequences.

Table 10 shows the percentage of changes in search length that are consistent with the model predictions. The percentage of all changes that are in the predicted direction range between 51.7% and 90.2%. The most common deviation is a failure to change one's search length when a change is predicted. Thus, there is under-responsiveness to treatment variation, in that there are more zero changes observed than are predicted. However, when there is responsiveness it tends to be in the direction predicted by our model as only 8.15% of all decisions involve changes that are in the opposite direction of the model. Table 11 shows the data for freezing usage. It shows that over 90% of the changes between treatments are consistent with model predictions, and that the principal deviation is a failure to adjust at all when a change is predicted. The data for earnings are shown in Table 12. They show a similar pattern in that most decisions involve no change in earnings even when a change is predicted. Most changes are in the correct direction, with cells indicating a majority in the incorrect direction involving relatively few observations.

²⁶ We also tested for gender differences in behavior, particularly in search length. For each treatment, we performed a two-tailed Mann–Whitney test for differences between genders in search length. None of the treatments exhibit a gender difference in search length at $p < .1$. The results of this analysis are given in Table 13 in Appendix F.

Table 5 Average number of stages under-searched and percentage of individuals who are under-searchers

| <i>T</i> | Avg. stages of under-search | | | Percentage of participants under-searching | | |
|----------|--------------------------------|-----|------|---|-------|-------|
| | <i>f</i> | | | | | |
| | Lo | Hi | No | Lo | Hi | No |
| 4 | .38 | .03 | -.02 | 66.7% | 35.7% | 14.3% |
| 10 | .74 | .44 | .39 | 66.7% | 50.0% | 38.1% |

The left part of the table indicates the average number of stages by which subjects under-search, compared to the model's predictions (positive numbers denote that search is terminated earlier than predicted on average). The right part of the table shows the share of under-searchers according to a Wilcoxon signed-rank test, using a 0.05 significance level. Observations in which the difference between the optimal and observed search lengths is zero (i.e. the subjects behaved optimally) are included in the tests

5.4 Do individuals use optimal reservation price strategies?

In this section, we evaluate Hypothesis 3, which concerns the use of the optimal RR-DRR policy. Figure 5 illustrates the fraction of offers of different magnitudes that have been accepted and frozen in the first stage of each round under each treatment. Behavior in the first stage provides the most stringent test of our model, since the backward induction task involved in optimizing in the first stage is the most demanding among all of the stages.²⁷ Each panel corresponds to one treatment condition. The graphs show the pooled data from all individuals. Each dot denotes the percentage of offers at the value indicated on the horizontal axis that were accepted. The x symbol indicates the percentage of offers at different levels that were frozen.

Our model predicts that all offers that are below a certain threshold are rejected. In treatments with no freezing, shown in the bottom two panels, it is predicted that all offers at or above this threshold are accepted. In some of the conditions in which freezing is possible, given in the other four panels, there is one threshold distinguishing offers that are rejected and those that are frozen (the reject-freeze threshold), and another cutoff dividing those offers that are frozen and those that are accepted (the freeze-accept threshold). In each panel, the predicted thresholds are denoted with vertical lines.

First consider acceptance decisions. We use nonlinear least squares minimization to fit the plotted acceptance frequencies to the logistic function

$$l(x; m, r) = [1 + \exp(-r(x - m))]^{-1} \quad (13)$$

²⁷ This analysis uses data from the first stage only. In later stages, selection and small sample issues begin to appear. Similar figures and estimates are shown for stages 2 and 3 in Appendix F (Figs. 7–10). Another complication for stages after the first is that the threshold in the IR case is history-dependent. Figures 8 and 10 in Appendix F assume that no past prices could be recalled.

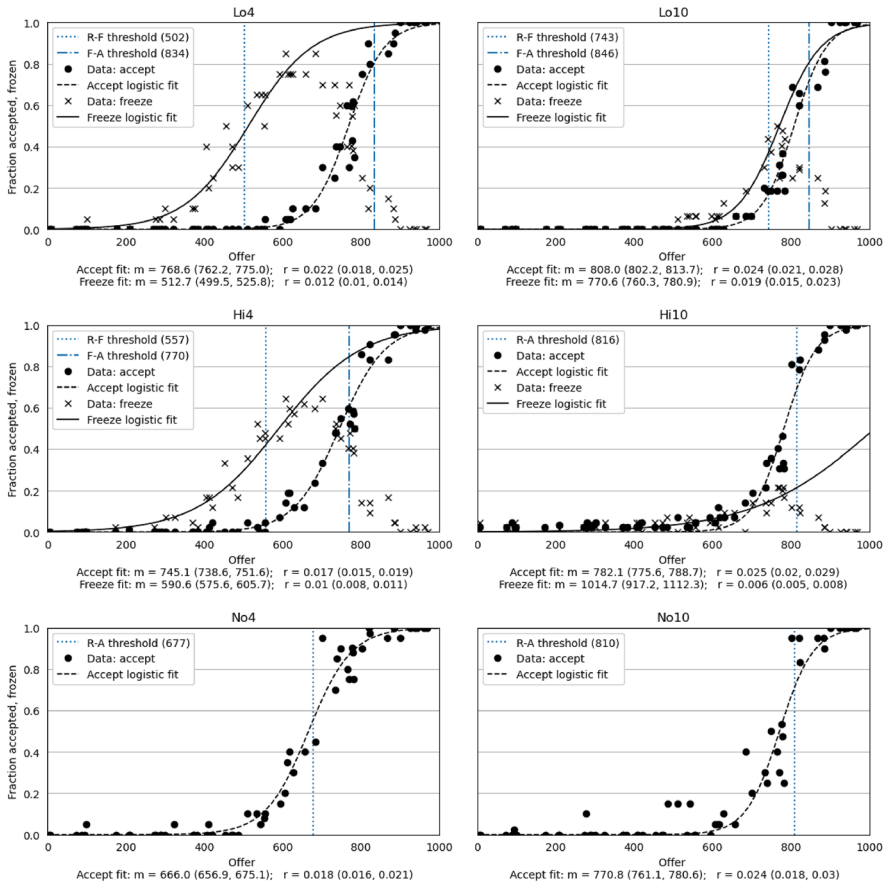


Fig. 5 Empirical frequencies of decisions in the first stage. The plots contain empirical frequencies of accepting and freezing in the first stage of each round for the six No Recall treatments. We fit a logistic function to the acceptance frequencies by nonlinear least squares minimization. We also fit a logistic function to the freezing frequencies for offers that are at or below the offer for which the freezing frequency is highest. Vertical lines represent the decision thresholds implied by the model

where m and r are parameters representing the midpoint and the curvature, respectively. Under each panel in Fig. 5, we report the fitted values of m and r , and a 95% confidence interval in parentheses. Large curvature is evidence for the use of a threshold strategy in the acceptance decision, because it indicates that the acceptance probability increases very rapidly at or near a particular value. We interpret large curvature, along with a midpoint m close to the optimal threshold, as evidence supporting the reservation strategy indicated by the model. The logistic functional form has the feature that it permits more deviations from the optimal strategy for offers close to the threshold. This means that errors are less likely, the more costly that they are. Note that as observations in this figure are aggregated across individuals, the results we report pertain to a representative consumer, potentially masking heterogeneity in underlying individual behavior.

In the treatments where there exist offers that are predicted to be frozen in the first stage, we apply a similar strategy to evaluate the threshold between the regions of rejection and freezing. For fitting the logistic function to the freezing fractions, we consider all offers below the offer for which the freezing frequency is highest.

The observed patterns can be summarized as the following result.

Result 5 Aggregated acceptance and freezing decisions in stage 1 are well approximated by a logistic function. However, the estimated acceptance thresholds are typically lower than predicted by the model. The freezing thresholds are close to predicted levels.

Support for Result 5: The figures reveal some strong and consistent patterns. The first is that the percentage of offers in excess of the theoretical acceptance thresholds that are indeed accepted is very high. In other words, in almost all instances in which acceptance of the current offer is predicted, acceptance occurs. The second is that the acceptance probabilities are described very well by a logistic function. The third is that the estimated acceptance thresholds are modestly, though significantly, lower than the predicted level in all treatments. This means that there is a greater tendency to accept offers below the theoretical cutoff for acceptance than to reject offers that are above the cutoff. This corresponds to the early termination of search on average. Fourth, the probability of an offer being frozen is much higher in the range in which it is predicted than when it is not predicted, though it does not reach a level greater than .8 for any range of offers in any treatment. Fifth, a logistic specification fits the relationship between the probability of freezing and that of rejecting quite well in the treatments in which freezing is predicted to occur. Sixth, the estimated thresholds between rejection and freezing are very close to the models' predictions, which lie within the 95% confidence interval in *Lo4*, and just outside of it in the other two treatments in which freezing is predicted. \square

5.5 Alternative models

We now consider whether adaptations of four models of decision making that have been proposed in previous literature outperform our model, presented in Sect. 3, which assumes optimal decision making under risk neutrality. We recognize that the models were proposed for environments without freezing, and in some cases, for settings that differed in other ways. Thus, any failure of one of the models to explain patterns in our data does not suggest that they are not appropriate in other environments. The four models are expressed for the situation in which the searcher is a seller, as in our experiment.

The first model we consider is in the spirit of the Regret Model developed by Loomes and Sugden (1986) and applied to sequential search by Weng (2009). Under this model, the individual incurs a disutility cost when she accepts an offer that was less favorable than the best offer that she has previously rejected. Specifically, the individual incurs a disutility $\beta(x - y)$ if she chooses an alternative that yields

pecuniary payoff x rather than another which would have resulted in a greater payoff y . In other words,

$$u(x, y) = x - \beta \mathbf{1}\{y > x\}(y - x) \tag{14}$$

Applied to our setting, in which x is the offer the searcher accepts and y is the best offer foregone, the rejection and freezing payments are:

$$R_t(\check{p}_t) = \int_0^{\bar{p}} \max\{u(p_{t+1}, \check{p}_t), K_{t+1}(p_{t+1}, \check{p}_t), R_{t+1}(p_{t+1})\} dF(p_{t+1}) - c, \tag{15}$$

$$\tilde{R}_t(k, \check{p}_t) = \int_0^{\bar{p}} \max\{u(p_{t+1}, \max\{\check{p}_t, k\}), u(k, \check{p}_t), \tilde{R}_{t+1}(k, \check{p}_t)\} dF(p_{t+1}) - c, \tag{16}$$

$$K_t(p_t, \check{p}_t) = \tilde{R}_t(p_t, \check{p}_t) - f, \tag{17}$$

where $\check{p}_t = \max\{p_1, \dots, p_t\}$. Under this specification, the possibility of regret lowers the value of accepting, freezing, and rejecting compared to the risk neutral model. However, compared with the risk neutral optimum, it lowers the value of accepting the most and rejecting the least. Thus, it can be shown that regret of this form leads to longer searches than under our model. One intuition for this is the following. In our model, individuals become less picky over time, so that they will sometimes be in a situation in which they are accepting an offer that they turned down at an earlier stage. However, an individual who feels regret incurs an additional cost when accepting such an offer, making her less likely to do so. The tendency to reject these offers serves to lengthen the searches of those who experience regret. This behavior is at odds with our data, and we do not discuss this model further.

The second model we consider is what we shall term the Cognitive Acquisition Cost (CAC) model. This model is inspired by the work of Gabaix et al. (2006), who modeled under-search over a finite horizon as a consequence of applying backward reasoning for an insufficient number of stages. We consider three versions of this model. The first version assumes that the choice made in stage $t < T$ is the optimal decision for stage $\min\{t + z, T - 1\}$, with $z = 1$. That is, decisions taken at stage $t < T$ are those that would be optimal at stage $t + 1$, as long as $t + 1 < T$. The second version of the Cognitive Acquisition Cost model is similar, except that the searcher behaves optimally under the assumption that there are $T - 2$ stages. Thus, the choice made in stage $t < T$ is the optimal decision for stage $\min\{t + z, T - 1\}$, with $z = 2$. The third version assumes that the individual always behaves as if it were stage $T - 1$, which is the assumption in Gabaix et al. (2006). Each of the three versions assumes a different type of failure of backward reasoning. The first two are consistent with applying an insufficient number of steps of backward reasoning, by 1 and 2 periods, respectively. The third version is consistent with the capacity to reason backward for only one step. The three versions of the model are evaluated here

in this subsection, with some additional detail provided in Appendix G. The conclusion from the analysis is that none of the adaptations of the Cognitive Acquisition Cost Model predict the decisions made by our participants as effectively as our model assuming risk-neutral and optimal decisions. The model predicts more freezing than the risk-neutral model, since the probability of freezing increases in later stages under the optimal policy, while we observe the opposite pattern in our data.

The third model is that of Kogut (1990), who proposes that agents under-search because they are susceptible to a type of sunk cost fallacy. We refer to this model as the Sunk Cost Fallacy (SCF) model. Under the SCF model, searchers use the total costs incurred to date rather than the marginal cost of searching for an additional stage when they make their decisions, as if they do not realize that previous costs are sunk. Specifically, they behave as if their freezing and rejection payments are given by:

$$\tilde{R}_t(k) = \int_0^{\bar{p}} \tilde{V}_{t+1}(p_{t+1}, k) dF(p_{t+1}) + tc + f, \quad (18)$$

$$\tilde{K}_t(p) = \tilde{R}_t(p), \quad (19)$$

$$R_t = \int_0^{\bar{p}} V_{t+1}(p_{t+1}) dF(p_{t+1}) + tc. \quad (20)$$

Thus, the expected costs of rejecting the current offer and obtaining a new one are always perceived as the accumulated search costs of all stages up to the present, tc . Because the individual perceives the cost of continuing the search as greater than under optimal decision making, searches tend to terminate earlier.

The fourth model we examine assumes optimal behavior under risk aversion. This has been previously proposed as an explanation for under-searching (see for example Cox and Oaxaca (1989)). We model risk aversion using the CRRA utility function $u(x) = \frac{x^{1-\alpha}}{1-\alpha}$ and assume that subjects integrate their earnings and costs in each stage in the following manner. Accepting an offer of p yields $u(p)$. The other actions, namely rejecting, freezing and rejecting while having an offer k frozen, respectively, yield:

$$R_t = \int_0^{\bar{p}} \max\{u(p_{t+1}), K_{t+1}(p_{t+1}), R_{t+1}\} dF(p_{t+1}) - u(c), \quad (21)$$

$$K_t(p_t) = \tilde{R}_t(p_t) + u(c) - u(c + f), \quad (22)$$

and

$$\tilde{R}_i(k) = \int_0^{\bar{p}} \max\{u(p_{t+1}), u(k), \tilde{R}_{t+1}(k)\} dF(p_{t+1}) - u(c). \quad (23)$$

We evaluate the alternative models, in comparison to the risk neutral model, by considering only those rounds in which there are differing predictions for the first stage. For each model and every treatment, Table 6 reports the proportion of decisions in these instances that were correctly predicted by our risk neutral model, the remaining proportion that was predicted by the alternative model, the p -value from testing for a difference in proportions between the two models, the number of rounds for which first round predictions differ between the two models and the corresponding number of observations (the number of participants in the treatment times the number of rounds in which the two models make competing predictions). We report only those model-treatment combinations for which there are at least five rounds of competing predictions, limiting the extent to which our conclusions are driven by the particular sequences observed by subjects (the treatment-pairs with fewer than five rounds differing are included Table 14, provided in Appendix F). For the model with risk aversion, we take the value of $\alpha = .3$, which is close to typical estimates reported in the literature (Holt and Laury (2002); Noussair et al. (2014)). In all treatments, the risk neutral model outperforms the model assuming risk aversion and the three versions of the CAC model. It also does better than the SCF model in two of three comparisons. Overall, none of the alternative models presented here outperforms the model we have proposed in Sect. 3.

6 Discussion

We have analyzed the effect of an option to freeze price offers on the behavior of agents engaged in sequential search. Our model generates a number of predictions, which serve as hypotheses for our experiment. The model predicts that the existence of the freeze option and the length of the time horizon available increase search length. Lower freeze fees also increase search length. These patterns are observed in the data. Furthermore, as predicted, the usage of freezing decreases as it becomes more expensive and as the time horizon increases. Our results regarding freezing are robust to having imperfect recall, though the impact of the presence of a freeze option is more pronounced when recall is impossible. Individuals have a strong tendency to behave as if they employ threshold strategies with regard to the range of offers that they accept, freeze and reject. The frequency of accepting offers increases sharply at specific threshold levels. These levels tend to be at or somewhat below those predicted by our model. Freezing is most common in the range of offers for which it is predicted. Our overall interpretation is that the model is quite successful in predicting behavior and outcomes. In our view, this result provides a behavioral foundation for the assumption of optimality in search behavior in structural modeling of demand in markets featuring consumer search.

The main deviation from optimal behavior is a modest tendency toward under-search, the termination of the search earlier than predicted. Under-searching relative

Table 6 Comparative performance of the optimal risk neutral model against alternative models

| Other model | T | f | Perc. RN | Perc. other | p -value | n | N |
|---------------------------------|-----|----------|----------|-------------|------------|-----|------|
| Risk aversion ($\alpha = .3$) | 4 | 10 | 48.3 | 47.1 | 0.693 | 15 | 630 |
| | | 40 | 54.4 | 29.2 | 0.000 | 12 | 504 |
| | | ∞ | 67.9 | 32.1 | 0.000 | 6 | 252 |
| | 10 | 10 | 43.2 | 29.4 | 0.000 | 14 | 588 |
| | | 40 | 48.4 | 36.1 | 0.000 | 12 | 504 |
| | | ∞ | 56.7 | 43.3 | 0.000 | 12 | 504 |
| Sunk costs fallacy | 4 | 40 | 53.9 | 38.1 | 0.000 | 11 | 462 |
| | 10 | 40 | 36.5 | 44.8 | 0.057 | 6 | 252 |
| | | ∞ | 50.3 | 49.7 | 0.877 | 8 | 336 |
| CAC | 4 | 10 | 78.9 | 20.9 | 0.000 | 22 | 924 |
| | | 40 | 72.3 | 26.1 | 0.000 | 25 | 1050 |
| | | ∞ | 77.1 | 22.9 | 0.000 | 11 | 462 |
| | 10 | 10 | 82.3 | 12.9 | 0.000 | 37 | 1554 |
| | | 40 | 77.5 | 14.6 | 0.000 | 37 | 1554 |
| | | ∞ | 72.9 | 27.1 | 0.000 | 23 | 966 |
| CAC ₁ | 4 | 10 | 72.8 | 26.5 | 0.000 | 7 | 294 |
| | | 40 | 59.5 | 39.4 | 0.000 | 11 | 462 |
| CAC ₂ | 4 | 10 | 78.9 | 20.9 | 0.000 | 22 | 924 |
| | | 40 | 72.3 | 26.1 | 0.000 | 25 | 1050 |
| | 10 | 40 | 34.3 | 31.9 | 0.604 | 5 | 210 |

The table compares each of the models discussed in this section to our risk neutral model of Sect. 3. For each model-treatment combination, the columns in the table contain the following. The first column is the mean of the indicator variable for subjects' actions adhering to the risk neutral model's prediction, that is, the percentage of instances in which the data are consistent with the model among those decisions where the two models make competing predictions. The second column is analogous, using the alternative model that we compare to our benchmark. The third value is the p -value from a statistical test in which the null hypothesis is that the proportion of correct predictions in the alternative model equals the analogous proportion in the benchmark model. The remaining columns are the number of sequences for which a given model-treatment combination has differing predictions in the first stage, and the total number of observations of decisions taken in which the two models make competing predictions in our data set

to the risk-neutral optimum has been a prominent finding in the experimental search literature. Beyond documenting that it continues to appear when freezing is possible, we document that affordable PFOs do increase the extent of under-searching. Nevertheless, this under-searching leads only to small losses in earnings.²⁸

Several other avenues for future research come to mind. On the empirical side, PFO data, which is increasingly available online, along with data on behavior and/

²⁸ It is plausible that such losses may be magnified under equilibrium considerations, such as in a Burdett and Judd (1983) model, because firms' pricing responses will result in a less favorable price distribution. Thus, while under-searching may be relatively harmless in the short run (i.e. earnings are unaffected), it can lead to welfare loss on the part of the searchers in the long run.

or prices, can be used to structurally estimate demand, assuming decision making that is as in the model presented in the current paper. To that end, our model can be of use even though that it allows only for one offer to be frozen, provided that one can show evidence that freezing multiple offers is rare. This is particularly likely to be the case when typical consumers search very little, as in, for example, Moraga-Gonzalez et al. (2018), or in settings where search or freezing carries a high cost.

Our theoretical model could be extended to allow the freezing of multiple offers, or the possibility of freezing offers that expire before date T . The price of the PFOs could also be modeled as proportional to the underlying price of the good. Similar experiments could be conducted, but augmented to include sellers. Studying both sides of the market would allow for the evaluation of long run welfare effects. Finally, policymakers and firms may want to experiment with introducing PFOs for wage offers in labor markets.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-024-09839-9>.

References

- Armstrong, M., & Zhou, J. (2016). Search deterrence. *Review of Economic Studies*, *83*, 26–57.
- Baltussen, G., Post, T., van den Assem, M., & Wakker, P. (2011). Random incentive systems in a dynamic choice experiment. *Experimental Economics*, *15*, 418–443.
- Basu, K. (2006). Consumer cognition and pricing in the nines in oligopolistic markets. *Journal of Economics and Management Strategy*, *15*, 125–141.
- Burdett, K., & Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, *55*, 955–969.
- 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.
- Cubitt, R., Starmer, C., & Sugden, R. (1998). On the validity of the random lottery incentive system. *Experimental Economics*, *1*, 115–131.
- DeGroot, M. H. (1970). *Optimal statistical decisions* (Vol. 82). McGraw-Hill.
- Einav, L. (2005). Informational asymmetries and observational learning in search. *Journal of Risk and Uncertainty*, *30*(3), 241–259.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.
- Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, *96*(4), 1043–1068.
- Hey, J., & Lee, J. (2005). Do subjects separate (or are they sophisticated). *Experimental Economics*, *8*(3), 233–265.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655.
- Janssen, M. C. W., & Parakhonyak, A. (2014). Consumer search markets with costly revisits. *Economic Theory*, *55*(2), 481–514.
- Kahan, J. P., Rapoport, A., & Jones, L. V. (1967). Decision making in a sequential search task. *Perception & Psychophysics*, *2*(8), 374–376.
- Karni, E., & Schwartz, A. (1977). Search theory: The case of search with uncertain recall. *Journal of Economic Theory*, *16*(1), 38–52.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, *14*(3), 381–392.
- Kohn, M., & Shavell, S. (1974). The theory of search. *Journal of Economic Theory*, *9*(2), 93–123.
- Landsberger, M., & Peled, D. (1977). Duration of offers, price structure, and the gain from search. *Journal of Economic Theory*, *16*(1), 17–37.

- Laury, S. (2005). Pay one or pay all: Random selection of one choice for payment. *Andrew Young School of Policy Studies working paper*, pp. 6–13.
- Lippman, S. A., & McCall, J. J. (1976). Job search in a dynamic economy. *Journal of Economic Theory*, 12(3), 365–390.
- Loomes, G., & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2), 271–282.
- McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84(1), 113–126.
- Moraga-Gonzalez, J., Sándor, Z., & Wildenbeest, M. (2018). Consumer search and prices in the automobile market. *Unpublished working paper*.
- Noussair, C. N., Trautmann, S. T., & van de Kuilen, G. (2014). Higher order risk attitudes, demographics, and financial decisions. *The Review of Economic Studies*, 81(1), 325–355.
- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5(2), 105–120.
- Ruffle, B. J., & Shtudiner, Z. (2006). 99: Are retailers best responding to rational consumers? Experimental evidence. *Managerial and Decision Economics*, 27(6), 459–475.
- Schorvitz, E. B. (1998). *Experimental tests of fundamental economic theories*. Ph.D. thesis, University of Arizona.
- Schotter, A., & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, 19(1), 1–25.
- Shunk, 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.
- Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior and Organization*, 35(3), 309–332.
- Sonnemans, J. (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, 21(1), 91–102.
- Starmer, C., & Sugden, R. (1991). Does the random lottery incentive system elicit true preferences? An experimental investigation. *American Economic Review*, 81(4), 971–979.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213–225.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 47(3), 641–654.
- Weng, D. (2009). Does regret explain why people search too little? A model of sequential search with anticipated regret and rejoicing. *Unpublished working paper*.
- Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, 22(4), 503–519.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.