Tilburg University

Sequential Search with a Price Freeze Option - Theory and Experimental Evidence

Marcu, Emanuel; Noussair, Charles

Publication date:
2018

Document Version
Early version, also known as pre-print

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Marcu, E., & Noussair, C. (2018). *Sequential Search with a Price Freeze Option - Theory and Experimental Evidence*. (CentER Discussion Paper; Vol. 2018-029). CentER, Center for Economic Research.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Sequential Search with a Price Freeze Option: Theory and Experimental Evidence

Emanuel Marcu* Charles N. Noussair[†]

December 19, 2018

Abstract

We study the economics of price freeze options (PFOs), by introducing them into a model of sequential search. The model makes a number of predictions, which we test in a laboratory experiment. The experiment varies (1) whether freezing is possible or not, (2) the cost of freezing, and (3) the length of the time horizon. We find that the observed treatment effects are consistent with the predictions of our model, though on average searches are terminated earlier than predicted. Interestingly, the extent of under-searching is magnified by the presence of an affordable PFO. We also show that while under-searching is relatively harmless at the individual level, it can lead to welfare losses in the long run.

Keywords: Sequential Search, Price Freezing, Experiment, Under-searching, Imperfect Recall.

JEL Classification Numbers: D83, C91.

*Department of Econometrics and Operations Research, Tilburg University. e.y.marcu@uvt.nl.

[†]Department of Economics and the Economic Science Laboratory, University of Arizona. cnoussair@email.arizona.edu. We thank Jaap Abbring, Jordi Brandts, Bart Bronnenberg, Jeff Campbell, Ernan Haruvy, Dan Houser, Ron Oaxaca, Asaf Plan, Amnon Rapoport, Tiemen Woutersen, as well as participants at the 2017 North American Economic Science Association Meetings and the 2018 Latin American Meeting of the Econometric Society, seminar participants at the University of Arizona, Tilburg University, The University of Alicante and the University of Texas-Dallas for their helpful comments. We thank the Economic Science Laboratory and the Eller College of Management at the University of Arizona for funding.

1 Introduction

Search markets are innovating over time as information technology progresses. Recent years have seen an advent of services that are offered as part of search processes, changing the decision problem that searchers face. This paper focuses on a particular innovation that is becoming prevalent in such markets. Namely, the PFO, a costly service offered by firms which guarantees the availability of an observed offer for future solicitation. 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 in subsequent periods.

To our knowledge, the institution of price freezing, however, has not been studied by economists.¹ Thus, it is not well-understood how the existence of the option to freeze prices affects the behavior of consumers. In particular, it is unknown how the response of consumers to the opportunity to purchase such a price guarantee interacts with the length of time they have to make the purchase, the cost of the freezing option and their ability to go back and accept other previously offered prices that were not frozen.

We use both theoretical and experimental methods to study the effect of PFOs on searching behavior and outcomes. We focus only on the consumer's decision, and do not endogenize the price setting process. The theoretical analysis identifies benchmark decisions and outcomes that result from rational 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.

PFOs are making inroads into various online 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,

¹A 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 value of the underlying asset follows a Brownian motion or a discrete random walk. In our environment, the price of the underlying good, rather than the change in its price, is independently and identically distributed at each stage. This means that the price at which the good can be purchased at stage $t - 1$ has no relationship to its value in 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, price freeze options cannot be transferred or resold. This means that techniques used to solve for option values cannot be readily applied to our environment.

for a fee of between 5 and 10 US dollars.² Lufthansa, Air France, and other global and regional carriers offer similar possibilities.³ An example of how the option is offered can be seen in the screen shot shown in Figure 1, taken from the website of Air France. This screen is shown to consumers immediately after they receive a price quote for a proposed itinerary. It indicates the duration of the PFO and the fee the consumer is asked to pay for the option (€5). The fee and duration of the PFO can vary depending on flight characteristics. 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.

Our purpose is to study price freezing as an institution, and not to simulate or investigate the airline industry, or any other industry in particular.⁴ Our focus is, rather, on the PFO itself and its implications on searchers' behavior. While we do find it striking that the 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. In our experiment, individuals are exogenously assigned in different phases of the experiment to different markets that may or may not permit price freezing, and the implications of the institution are studied *ceteris paribus*.

Our model builds on a classical homogeneous good sequential search model with a finite horizon, that we extend to include the presence of costly PFO. We consider how behavior and outcomes respond to changes in (i) the price of the freezing option, (ii) the length of the time horizon available to the decision maker, and (iii) the possibility of recalling and accepting some prior offer other than the one that was frozen.

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 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

²In the US, airlines must follow the “24-hour rule”, overseen by the Department of Transportation, stating that airlines must offer the option to cancel within 24 hours, or hold the ticket for that time frame, free of charge. Therefore, the PFOs offered by US airlines are for a period that is longer than 24 hours. Guidelines regarding this regulation can be found on the website of the Department of Transportation in the following link: https://www.transportation.gov/sites/dot.gov/files/docs/Notice_24hour_hold_final20130530.pdf.

³See Table 7 in Appendix A for an overview of the fee and duration of PFOs being offered by various airlines at the time of this writing.

⁴Pricing for airline travel is characterized by a number of features that have been explored in many interesting studies by other authors (see for example Borenstein (1992); Berry and Jia (2010) and Sengupta and Wiggins (2014)).

that are less favorable than a second cutoff are rejected. If there is one threshold, only acceptances and rejections occur. We show that the solution is monotonic in the sense that a stage with two thresholds can follow a stage with one threshold, but not vice versa. The threshold price levels depend on the offer distribution, the cost of freezing offers 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 costs 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. There is strong evidence of the use of reservation and double reservation price rules in the predicted manner, at least in early stages.

The main departure from the model is a modest, though significant, tendency to end searches too early in most treatments. We pay particular attention to this previously documented issue of under-searching. That is, we evaluate whether average search length is as predicted by our model or whether we observe the tendency toward shorter than optimal. More importantly, beyond documenting that under-searching is common, we relate it to two economically meaningful phenomena. Firstly, we show that lowering the freeze fee magnifies the extent to which searchers' exploration is insufficient. This new result can shed light on the mechanisms that cause under-searching. Secondly, We show that while under-searching is prevalent,

⁵The experiment is not designed to measure the effect of allowing recall on freezing, search length, earnings, or other variables. The reason is that the 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 are underpowered. Thus, the experiment is designed to test the effect of changing the freezing fee and time horizon on behavior in two distinct environments, one with no recall and one with imperfect recall. This can be thought of as a robustness exercise.

its adverse effects on individuals' profits are negligible. Thus, in the short run model we consider, where firms do not respond, searchers' welfare loss is not substantial. However, in the long run, accounting for firms' responses, under-searching can lead to less favorable price distributions, which can have economically significant welfare effects.

The relationship between price freezing and the ability to recall prior offers merits some comment, as they share some commonalities. The inability to recall prior offers can arise in markets mainly for two reasons. 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 no 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. We refer to this situation as one of imperfect recall. The main difference between a PFO and the ability to recall is that the PFO pertains to a specific offer that was chosen by the searcher to be frozen, whereas recall pertains to the best offer seen so far. Therefore, the value of a PFO decreases in the recall probability.

In the search literature, imperfect recall is typically modeled with a probability of being able to accept a prior offer specified to be strictly between zero and one. That is, if an offer is rejected and search continues, then there is a risk of not having the possibility to return to that offer and accepting it later.⁶ In contrast, if a price-freezing option is available, searchers can pay a one-time fixed fee, which increases the recall probability of that particular frozen offer to one. That is, the frozen offer will be readily available at any time in the future with certainty. Our results regarding the effects of PFOs on behavior are robust to both settings, though, as predicted, the impact of the presence of a freeze option is more pronounced when recall is impossible.

The paper proceeds as follows. In section 2 we discuss related literature. Section 3 develops a theoretical model of sequential search featuring a PFO and characterizes the optimal solution for a risk neutral agent. Section 4 describes our experimental

⁶The 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). A potential explanation for this pattern in the literature is that 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.

design, section 5 presents the results and discusses them and section 6 concludes.

Figure 1: PFO at Air France

NEED MORE TIME TO THINK?

Lock in this fare and pay for your ticket later on our website.
Reservation for your selected dates and fares guaranteed until Monday 15 August 2016 at 22:00 (local time).

By purchasing the Time to Think option now, you can guarantee your flights and fares. Head to the "Your Reservations" area to confirm your purchase. This option is nonrefundable and does not allow you to change your original reservation.

5 €

per passenger

☐ **ADD THE OPTION**

SUMMARY OF YOUR TRIP

YOUR DEPARTURE FLIGHT ECONOMY MINI Amsterdam > Paris Wed 14 Sep 2016 at 06h45	YOUR RETURN FLIGHT ECONOMY MINI Paris > Amsterdam Fri 23 Sep 2016 at 09h30	Passengers 1 Duration 10 d	TOTAL INCLUDING TAX 108,13 € Fare conditions (for the selected fares)
--	---	-------------------------------------	---

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 \(1970\)](#) and others.⁷ 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.

The experimental literature on testing predictions of search models dates back to [Braunstein and Schotter \(1981\)](#) who evaluated various theoretical implications. For

⁷[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.

instance, they test how search behavior is affected by 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. Indeed, 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 than is predicted by this rule (see, for example, the experiments by [Cox and Oaxaca \(1989\)](#), [Sonnemans \(1998\)](#), [Kogut \(1990\)](#), [Einav \(2005\)](#) and [Schorvitz \(1998\)](#)).⁸

Various explanations have been proposed for rationalizing early stopping. We evaluate three of these. 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 agents have higher reservation prices. While [Cox and Oaxaca \(1989\)](#) attribute under-searching to risk aversion, [Sonnemans \(1998\)](#) shows 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 horizon approaches.

The last mechanism we consider, anticipated regret/rejoice preferences, is based on the regret/rejoice 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 yielded by choosing a different action (and having uncertainty resolve in the same way). [Weng \(2009\)](#) showed that in a model with perfect recall, incorporating such regret/rejoice attitudes and assuming that the regret disutility outweighs rejoice utility implies a standard reservation rule strategy, but with higher reservation prices than in the benchmark model, leading to under-searching. Regret disutility being stronger rejoice utility is a “stylized fact” from psychological

⁸See 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.

research. This mechanism is the one that has been explored the least in previous literature.⁹

[Landsberger and Peled \(1977\)](#) and [Karni and Schwartz \(1977\)](#) formalize sequential search with imperfect 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 \(1977\)](#).

More recently, search models have been used by scholars in industrial organization and quantitative marketing to estimate demand, and to quantify the consequences of economic frictions in settings where incomplete information is important. In coarse terms, this literature can be divided into two strands. The first strand is structural demand estimation under incomplete information, typically attempting to recover search costs using various types of data and the structure implied by a search model. For example, [Hong and Shum \(2006\)](#) and [Moraga-González and Wildenbeest \(2008\)](#) use price data, coupled with the structure implied by a model similar to [Burdett and Judd \(1983\)](#), to estimate consumer search costs in equilibrium. [Kim et al. \(2010\)](#) estimate demand and recover search costs using view-rank data, along with data on prices and product characteristics, employing a structural model based on [Weitzman \(1979\)](#). A second strand in the literature, which is more closely related to this paper, consists of studies providing descriptive evidence on how individuals search

⁹There have been other rationalizations of under-searching in previous literature which we do not consider. For example, under mental accounting as in [Thaler \(1985\)](#), even though past search costs are sunk, individuals “do not close” these accounts, but rather make decisions based on cumulative search costs. Under a mechanism proposed by [Sonnemans \(2000\)](#) and formally tested by [Einav \(2005\)](#), under-searching is a consequence of the asymmetric information structure in terms of the feedback searchers receive between search problems. 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 more picky and not skipped a lucrative offer. This causes under-searching because of the one-sidedness of the implied directional learning: feedback is only obtained when searching too much and not when searching too little.

in different settings. For example, [Bronnenberg *et al.* \(2016\)](#) describe behavioral patterns in consumers' online search for a differentiated good, whereas [De los Santos *et al.* \(2012\)](#) compare the performance of a simultaneous search model to its sequential counterpart in explaining online search for a homogeneous good. In our paper, rather than attempting to recover search costs, we fix these and evaluate the predictions of a theoretical model that we develop.

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.¹⁰ The model is formulated, in line with the theoretical literature on 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.¹¹

The logic of this section is as follows. We first show that, when an outside option is exogenous, its effect depends on whether it is greater than or equal to the price at which an individual is indifferent between accepting the current offer or searching for exactly one more period. If the outside option lies below this value, then in the next-to-last stage, there is no benefit of continuing the search, as the consumer is better off taking the outside option, or any better offer, rather than searching. However, in that case, the decision problem never reaches the last stage, and with at most one stage remaining, it becomes better to accept the outside option at the second-to-last stage, rather than to continue. The problem unravels in this manner and the consumer

¹⁰The manner in which we introduce imperfect recall, along with the outside option, complicates the analysis substantially, as we discuss in the appendix. As a result, we do not solve for the optimal policy analytically, but rather use numerical methods.

¹¹While the theoretical model operates symmetrically regardless of whether a consumer or a seller is considered, it is an open behavioral question whether consumers facing a sequence of offers to sell approach the problem similarly to sellers or workers facing a sequence of job offers.

terminates her search immediately, accepting either the first offer she receives or the outside option. If the outside option is at a price higher than that which makes the individual indifferent between accepting the current price and continuing to search for exactly one more period, one would always rather reject than take the outside option, except in the last period if the outside option price is lower than the final price offered. Thus, an outside option will be accepted either immediately, in the final stage, or not at all, depending on its magnitude and the initial and final offers.

In the second subsection, we show that when the searcher can either accept, reject, or freeze offers, there can be an intermediate range of offers, which are optimal to freeze. In this range, freezing an offer has the benefit that it lowers the expected eventual price paid when a subsequent offer is rejected. Some offers have the property that it is more profitable to pay a small fee to gain this benefit than to either accept or reject them. The range of offers that fall within this range increases as the final deadline to accept an offer approaches. Generically, the optimal strategy is characterized in the following manner. When the final deadline is sufficiently far in the future, it is optimal to set one threshold above which offers are rejected and below which they are accepted. As the deadline approaches, it becomes optimal to use two thresholds and accept low offers, freeze intermediate ones, and reject high ones. In the terminal stage, it is trivially optimal to accept the lowest among the current offer or a frozen offer, if one is available. All proofs are 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 *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 *payment* in stage t from the individual 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 point t onward. It is the minimum of the current price, the outside option, and the rejection payment. Thus, when $t < T$, we have

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

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

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

The functions $h(x)$ and $g(x)$ are useful in deriving some of our results.¹²

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

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

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

¹² $h(x)$ is the rejection payment in stage $T - 1$, net of search costs, when holding x as an outside option. Note that this interpretation holds also for a model with perfect recall and without an outside option, with x being the lowest price observed so far, regardless of t . $g(x)$ is the difference between accepting the outside option and rejecting, net of search costs, in stage $T - 1$.

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 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.

We make the standard assumption that the search cost c is sufficiently low for search to be initiated in the first place. ¹³

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. ¹⁴ 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.*

The relationship of this result to the effect of the introduction of the PFO is that a searcher may freeze a price for the purpose of having it serve as an outside option in subsequent stages. Having the frozen price available decreases future expected payments by decreasing all future rejection payments. This is the case even though the offer itself may only be accepted in the terminal stage, if at all. A sufficiently low price will never be frozen because it is better to accept it. Thus, there is a lower bound on the prices that may be frozen. An upper bound on the offers that may be frozen will exist if, as we will assume, there is a positive freezing fee.

Figure 2 plots the rejection payments in k -space. Rejection payments are increasing in k . That is, the less attractive the outside option, the higher the rejection

¹³Otherwise, 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, it is clear that if assumption 1 is violated, search terminates in the first period.

¹⁴It 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.

payment. When $k \leq p^*$, all rejection payments are the same, because rejection is always optimally followed by accepting whichever is lower among the next offer and the outside option k . When $k > p^*$, the rejection payments shift upward over time. Figure 3 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 rejection payment is always higher than the outside option if $k \leq p^*$. The right panel illustrates how, in the case where $k > p^*$, rejection payments increase over time as the end of the horizon approaches.

Figure 2: Search with an outside option: rejection payment in k -space

Figure 3: Search with an outside option: rejection payment in p -space

We end this subsection by noting a limit result, which we use in subsection 3.2. Namely, as the number of stages tends to infinity, the rejection payment approaches p^* when $k > p^*$. This means that the outside option is never used, except for potentially in the first stage.

Lemma 1. *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 subsection. Assume that the searcher can freeze one offer by paying a fixed fee $f > 0$, which we will refer to as the *freezing fee*. Once one offer is frozen, a second offer may not be frozen. 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 section 3.1. $V_t(p_t)$ is the payment when there is no frozen alternative, R_t is the rejection payment, and $K_t(p_t)$, which we shall refer to as the *freezing payment*, is the expected payment when freezing p_t and continuing to the next stage with p_t as an outside option. Formally,

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

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

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

We begin with a lemma stating that the rejection payment is weakly lower when freezing is available than when it is not.

Lemma 2. $R_t \leq \tilde{R}_t(\bar{p})$ for all $t < T$

Next, we make an assumption on the freezing fee, which is that it 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*

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

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

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 rejection payment than rejecting. We now turn to the optimal strategy of the consumer under a freeze option. We use the following terminology to describe the optimal rule.

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

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

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 increasing if the acceptance region strictly increases and the rejection region strictly decreases over time.

The following lemma lets us 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 strategy at every stage.

Lemma 3. *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. Offers greater than a threshold level are rejected and the threshold is strictly increasing over time, as the consumer becomes less picky.

Lemma 4. *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 2, we describe the structure of the optimal policy for the consumer to follow.

Proposition 2. *The optimal policy consists of a strictly increasing RR for $t \leq t^*$ and a strictly increasing DRR for $t > t^*$, where $0 \leq t^* < T - 1$.*

Thus, the solution can take two forms. When $t^* = 0$, the optimal policy is an increasing DRR, whereas 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 3. *There exists a $T^* < \infty$ such that for $T > T^* \implies t^* > 0$.*

The left panel Figure 4 illustrates a stage in which a DRR is optimal ($t \in T^{DRR}$). The 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 $t \in T^{RR}$.

Figure 4: Example of a DRR, and of behavior as $T \rightarrow \infty$

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. Each session began with two risk elicitation protocols. These were followed by the main part of the experiment, which consisted of 180 search problems that had the possibility of counting toward earnings. The session concluded with a brief questionnaire.¹⁵ Subjects could complete the sequence of tasks at their own pace. The entire sequence of tasks took between 60 and 100 minutes to complete.

The 180 search problems that counted 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.¹⁶ 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.¹⁷ 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 section 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

¹⁵Upon 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.

¹⁶The 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.

¹⁷Appendix D includes the instructions and screen shots of the interface.

stage of a round, subjects receive an offer, drawn from a discrete uniform distribution on $[0, 1000]$, 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 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.

Since the freeze fee is varied within subjects across blocks, its effect on behavior is only isolated if we use the same sequences of offers in each block, or if participants perceive these sequences as being the same. However, if the 60 sequences were repeated precisely, individuals might notice that they were experiencing the exact same sequences when encountering them for the second or third time. In this case, they would have information on future draws within a sequence, which is of course not desirable. While it is cognitively difficult to remember 60 sequences of offers, subjects might remember several specific sequences, or may remember which offer came after some particular offers. We devised a solution to this problem. The first block of sequences, those employed in rounds 1 - 60, were simply random independent draws from $U[0, 1000]$.¹⁸ In the second and third block we amended these sequences

¹⁸Specifically, they are pseudo-random integers drawn from $U\{0, \dots, 1000\}$ after setting the seed in Python's Numpy package to 0.

in a way that allows us to compare the three 60-period blocks, while controlling for the sequences of offers individuals faced, and preventing subjects from remembering useful information from previous sequences. This allows us to identify the effect of different levels of f with a within-subject comparison.

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. 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 are not changing the "hundreds" digit, and each value in the support is drawn with equal probability.¹⁹ 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 makes them, in practice, unable to remember offer sequences or sub-sequences.

4.3 The Treatments

In different treatments, we vary the recall probability q , the time horizon T , and the freeze cost f . 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 (referred to as No Recall or NR treatments) and $q = .5$ (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 , any 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.

The freeze fee is varied within-subject across blocks. There are three conditions.

¹⁹For 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. A related idea is the "price points theory" (see e.g. Kashyap (1995)) whereby there is a kink in consumers' demand around particular price endings.

Under the *Low Freeze Fee* (LF) condition, individuals must pay 10 to freeze an offer, and under *High Freeze Fee* (HF) they must pay 40. In a third condition, *No Freeze* (NF) 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; Low Freeze Fee in the first 60 rounds of the search task, High Freeze Fee in rounds 61 - 120, and No Freezing allowed 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 of generating an offer in the next stage is always 10.

Thus, the experiment has a 2 x 2 x 3 structure. We refer to the treatments in an abbreviated form by the time horizon (T4 or T10), whether recall was not possible or imperfect (NR or IR), and the level of the freeze cost (LF, HF or NF). For example, T10IRHF denotes the treatment in which individuals could sample up to 10 offers in a round, there was a recall probability of .5, and the high freeze cost of 40 was in effect.

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.

Table 1: Subject Numbers and Characteristics in the Different Treatments

q T	0.0		0.5	
	4	10	4	10
Count	44.00	42.00	51.00	40.00
Perc. female	33.33	40.00	56.86	55.26
Education	2.62	2.74	2.90	2.94
No. experiments	9.75	8.88	6.90	8.11
Perc. Economics / business	34.09	59.52	60.78	57.50

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

4.4 Optimal decisions

The optimal strategy, using the parameters described in subsection 4.3, in each of the NR and IR treatments, are illustrated in the panels on the left and right sides of Figure 5, respectively. In each of these panels, we plot the solution for one level of f . 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 and the recall realization imply that recall is not available (so that the value of recalling enters through the rejection payment).

The figure shows that the optimal thresholds in NR are monotonic and concave. The acceptance and freezing regions are increasing over time. More freezing occurs when the freezing fee is low. For example, in the LF condition ($f = 10$), the freezing region is already larger than 10% of the range of possible offers in stage $T - 9$, whereas under HF 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 IR case, optimal thresholds are also monotonic, and are concave once the DRR sub-policy is in effect.

Figure 5: 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 expected continuation 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.

4.5 Hypotheses

The hypotheses guiding the design of our experiment are derived from the theoretical results presented in the last section for the NR treatments, and from the computation of the optimal solution for the IR treatments. 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 length, the percentage of rounds in which it is optimal to freeze an offer, the average earnings of the searcher (which equal $1000 - \text{the price paid} - \text{total search costs} - \text{freezing fee if an offer is frozen}$) 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. A higher freeze cost results in on average a 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.

Hypotheses 1 and 2 are a collection of mostly intuitive relationships that might be expected under a variety of plausible decision rules, not only the optimal decision making that is assumed in our model. However, 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 freezing 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 under T4 when the horizon is relatively short than when it is long in T10. We interpret support for this predictions as strong evidence in favor of our model.

Our third hypothesis is that individuals employ the optimal RR-DRR policy. That is, that individuals’ behavior is consistent with the model’s prediction, as depicted in Figure 5. All three hypotheses are evaluated in the NR and IR data separately.

Beyond these hypotheses about direct treatment effects, we focus on how under-searching depends on the availability of a PFO. We do not have a conjectured direction for this effect. We do predict that subjects tend to under-search and that this has a negligible impact on their earnings.

Table 2: Model predictions, by treatment

q	T	f	Search length	Freezing usage	Earnings	Other party's surplus
0	4	10	2.98	71.67	791.20	188.97
		40	2.58	40.00	772.82	211.35
		∞	2.23	0.00	764.58	223.08
	10	10	4.62	38.33	852.82	111.02
		40	4.05	11.67	853.85	115.65
		∞	3.92	0.00	850.07	120.77
0.5	4	10	2.98	68.33	791.53	188.63
		40	2.63	15.00	782.32	201.35
		∞	2.57	0.00	781.32	203.02
	10	10	4.63	18.33	857.83	105.83
		40	4.37	3.33	859.47	106.87
		∞	4.35	0.00	860.97	105.53

This table includes the models' 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, and Other Party's Surplus is the earnings of a hypothetical individual on the other side of the market. This equals 1000 minus any offer accepted plus any freeze fee paid.

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.

Hypothesis 1: Search length is shorter when f is smaller, and T is greater.

Hypothesis 2: Freezing is less frequent when f is greater and T is greater.

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

5 Results

This section is organized in the following manner. Subsection 5.1 summary statistics, describing overall patterns in the data, aggregated over subjects and rounds. In subsection 5.2 we evaluate our main hypotheses, controlling for the sequences subject's observed. This cleans the effects from sampling error, resulting in interpretable treatment effects. Subsection 5.3 considers hypothesis 3, and studies the extent to which individuals employ the optimal strategy specified in our model. We summarize and discuss our results in subsection 5.4.

5.1 Summary Statistics

Figure 6 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 xs , respectively). The data displayed are the average search length (panel 6a), the percentage of rounds in which an offer has been frozen (panel 6b), the earnings per round (panel 6c) and the percentage of instances in which recall was used (panel 6d). In panel 6b 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 (recall that the model predicts that the only time that a frozen offer would be accepted is in the last stage, T , of a round, conditional on previous optimal behavior).²⁰

The model also predicts that any offer that is frozen would only be accepted, if at all, in the last stage. Indeed, in the short-horizon T4 data, the vast majority of acceptances of frozen offers are in the terminal stage, T , as predicted, whereas this is not the case for the T10 treatment. This raises doubts about the model’s prediction that frozen offers are only accepted in the terminal round.

The numerical solution of optimal behavior under imperfect recall predict that offers would never be recalled when $T = 4$ and freezing is possible. Offers would be recalled in 11% of rounds under T4 when no freezing is possible. We do find that recall is rarely used. The acceptance of a recalled offer occurs in between 5 - 8% of rounds in all treatments. However, it is somewhat more common with a higher freeze fee, and is highest when recall is not possible.

²⁰Figure 13 in Appendix E contains the same information as Figure 6, 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.

Figure 6: Within-group Means and Frequencies

This figure contains the observed within-treatment means and frequencies for key variables (as circles) and the theoretical predictions (as *x*s). Panel (b) includes also a breakdown of whether offers that have been frozen were later accepted and if so, whether the acceptances of frozen offers occur in the last stage.

5.2 Treatment effects

Estimated treatment effects on search length, freezing usage and earnings are reported in Tables 3, 4 and 5, respectively. We obtain these by regressing each of these dependent variables on treatment dummies and round-level fixed effects, thereby controlling for heterogeneity in offer sequences across rounds. Standard errors, clustered by subject, are given in parentheses. For the purpose of evaluating the magnitude of the coefficients from these subject-round-level linear regressions, we also include the model's predictions for the coefficients, taken from Table 2, in italics. The specification allows us to compare the magnitudes of treatment differences to those predicted by the model. The estimated intercepts yield estimates of the mean in the base categories NF, T4, and NR. The other coefficients are estimates of the differences from the baseline category. Note that the stars on the treatment coefficients are from a *t*-test comparing the estimate to the model's prediction (that is, the stars do not carry the typical interpretation of being significantly different from zero).

Result 1. *There is under-searching, compared to the model’s predictions, in almost all treatments. 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 1: Consider the upper panel of Table 3. For readability, we present the number of rounds by which individuals under-search in Table 6. In all treatments, except for T4NRNF, there is under-searching, albeit by negligible amounts for a minority of treatments. In about half of the treatments, the under-searching is by an economically substantial amount, for example, in the T10 treatments, the under-searching ranges between 0.39 and 0.91 rounds. In each (T, q) treatment, the margin by which individuals under-search is largest under the LF treatment. Moreover, in the NR, the extent of under-searching monotonically decreases in f : for example, in the T10NR treatments, individuals under-search on average by 0.74 under LF, by 0.44 under HF and 0.39 under NF.

Result 2. *The length of search is (1) decreasing in the freeze fee and (2) increasing in the horizon length, as predicted by the model. These patterns appear in both NR and IR.*

Support for Result 2: All of the coefficients relating to different freeze fees and horizon lengths are the same sign as predicted by our model, supporting Hypothesis 1. Lowering freeze fees substantially increases search, with the effect being strongest for short horizons, as seen in the upper part of Table 3. For example, estimated search length is higher by 0.35, out of four rounds, in the T4 treatments. The bottom part of the table shows that NR and IR lead to similar search lengths.

Result 3. *There is less usage of the freezing option than is predicted by the model. Usage decreases in the freeze fee and the horizon length, as predicted by the model. Under T4, the large majority of frozen offers are accepted in the last period, as predicted, whereas this is not the case for T10.*

Support for Result 3: Table 4 describes how freezing usage is affected by the treatment. Similarly to search length, all coefficients in the upper two thirds of the table have the same sign as the model predicts, but are smaller in magnitude than the prediction. This possibility of freezing offers is underutilized, compared to the model’s predictions, in each of the treatments.

Result 4. *The magnitude of the treatment effects on earnings are very similar to those predicted by the model. Earnings are mostly unaffected by introducing PFOs and are greater when the horizon is longer.*

Support for Result 4: Table 5 shows that except for the T4NR treatment, where earnings increase modestly when introducing a cheap PFO (by about 4%), earnings remain unchanged when freezing is available. Earnings are substantially higher in the long horizon treatments, with the difference ranging between 7% and 10%.

Table 3: Search length as a function of treatment

	$T = 4$ $q = 0$	$T = 10$ $q = 0$	$T = 4$ $q = 0.5$	$T = 10$ $q = 0.5$		
LF - NF	<i>0.75</i> 0.35*** (0.044)	<i>0.70</i> 0.35*** (0.097)	<i>0.42</i> 0.35* (0.039)	<i>0.28</i> 0.07 (0.129)		
HF - NF	<i>0.35</i> 0.30 (0.045)	<i>0.13</i> 0.08 (0.060)	<i>0.07</i> 0.14** (0.031)	<i>0.02</i> 0.03 (0.069)		
Constant (NF)	<i>2.23</i> 2.25 (0.031)	<i>3.91</i> 3.52*** (0.116)	<i>2.57</i> 2.49*** (0.052)	<i>4.35</i> 3.65*** (0.142)		
Observations	7560	7560	9180	7200		
r2	0.69	0.60	0.55	0.54		

	$f = 10$ $q = 0$	$f = 40$ $q = 0$	$f = \infty$ $q = 0$	$f = 10$ $q = 0.5$	$f = 40$ $q = 0.5$	$f = \infty$ $q = 0.5$
T10 - T4	<i>1.63</i> 1.27** (0.144)	<i>1.47</i> 1.06*** (0.139)	<i>1.68</i> 1.27*** (0.120)	<i>1.65</i> 0.99*** (0.172)	<i>1.73</i> 1.15*** (0.154)	<i>1.78</i> 1.26*** (0.150)
Constant (T4)	<i>2.98</i> 2.60*** (0.048)	<i>2.58</i> 2.54 (0.053)	<i>2.23</i> 2.25 (0.031)	<i>2.98</i> 2.74*** (0.066)	<i>2.63</i> 2.53* (0.059)	<i>2.57</i> 2.39*** (0.052)
Observations	5040	5040	5040	5460	5460	5460
r2	0.52	0.54	0.55	0.45	0.49	0.50

	$f = 10$ $T = 4$	$f = 40$ $T = 4$	$f = \infty$ $T = 4$	$f = 10$ $T = 10$	$f = 40$ $T = 10$	$f = \infty$ $T = 10$
IR - NR	<i>0.00</i> 0.13 (0.082)	<i>0.05</i> -0.01 (0.079)	<i>0.33</i> 0.14*** (0.060)	<i>0.02</i> -0.15 (0.209)	<i>0.32</i> 0.08 (0.192)	<i>0.43</i> 0.13 (0.183)
Constant (NR)	<i>2.98</i> 2.60*** (0.048)	<i>2.58</i> 2.54 (0.053)	<i>2.23</i> 2.25 (0.031)	<i>4.62</i> 3.87*** (0.136)	<i>4.05</i> 3.60*** (0.128)	<i>3.92</i> 3.52*** (0.116)
Observations	5580	5580	5580	4920	4920	4920
r2	0.57	0.60	0.69	0.55	0.58	0.58

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

Results from linear regressions where the dependent variable is search length. Each column fixes two treatment variables from $\{f, T, q\}$ (these are given in the column headers) and includes dummy variables for the third treatment variable. Each regression includes round-level fixed effects to control for heterogeneity across offer sequences. Standard errors clustered by subjects are in parentheses. The italic font stands for the coefficients predicted by the model.

Table 4: Freezing usage as a function of treatment

	<i>T = 4</i> <i>q = 0</i>	<i>T = 10</i> <i>q = 0</i>	<i>T = 4</i> <i>q = 0.5</i>	<i>T = 10</i> <i>q = 0.5</i>
LF - HF	<i>0.32</i> 0.06*** (0.035)	<i>0.27</i> 0.17*** (0.022)	<i>0.53</i> 0.17*** (0.023)	<i>0.15</i> 0.11** (0.020)
Constant (HF)	<i>0.40</i> 0.32** (0.031)	<i>0.12</i> 0.09 (0.019)	<i>0.15</i> 0.14 (0.021)	<i>0.03</i> 0.05 (0.017)
Observations	5040	5040	6120	4800
r2	0.25	0.14	0.16	0.08
	<i>f = 10</i> <i>q = 0</i>	<i>f = 40</i> <i>q = 0</i>	<i>f = 10</i> <i>q = 0.5</i>	<i>f = 40</i> <i>q = 0.5</i>
T10 - T4	<i>-0.33</i> -0.13*** (0.041)	<i>-0.28</i> -0.23 (0.036)	<i>-0.50</i> -0.15*** (0.039)	<i>-0.11</i> -0.09 (0.027)
Constant (T4)	<i>0.72</i> 0.38*** (0.031)	<i>0.40</i> 0.32** (0.031)	<i>0.68</i> 0.31*** (0.031)	<i>0.15</i> 0.14 (0.021)
Observations	5040	5040	5460	5460
R ²	0.18	0.19	0.13	0.07
	<i>f = 10</i> <i>T = 4</i>	<i>f = 40</i> <i>T = 4</i>	<i>f = 10</i> <i>T = 10</i>	<i>f = 40</i> <i>T = 10</i>
IR - NR	<i>-0.03</i> -0.08 (0.044)	<i>-0.25</i> -0.18* (0.038)	<i>-0.20</i> -0.10*** (0.039)	<i>-0.08</i> -0.04 (0.025)
Constant (NR)	<i>0.72</i> 0.38*** (0.031)	<i>0.40</i> 0.32** (0.031)	<i>0.38</i> 0.26*** (0.026)	<i>0.12</i> 0.09 (0.019)
Observations	5580	5580	4920	4920
R ²	0.17	0.17	0.13	0.04

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

Results from the estimation of a linear probability model where the dependent variable is a binary variable indicating freezing usage within a round. Each column fixes two treatment variables from $\{f, T, q\}$ (these are given in the columns' headers) and includes dummy variables for the third treatment variable. Each regression includes round-level fixed effects to control for heterogeneity across offer sequences. Standard errors clustered by subjects are in parentheses. The italic font stands for the coefficients predicted by the model.

Table 5: Earnings as a function of treatment

	$T = 4$ $q = 0$	$T = 10$ $q = 0$	$T = 4$ $q = 0.5$	$T = 10$ $q = 0.5$		
LF - NF	<i>26.62</i> 30.44 (4.130)	<i>2.75</i> -3.25 (6.454)	<i>10.22</i> 6.10 (4.370)	<i>-3.13</i> 3.43 (9.428)		
HF - NF	<i>8.23</i> 6.99 (2.560)	<i>3.78</i> 0.09 (2.953)	<i>1.00</i> -0.67 (2.392)	<i>-1.50</i> 11.35* (6.628)		
Constant (NF)	<i>764.58</i> 741.56*** (4.053)	<i>850.07</i> 819.51*** (7.351)	<i>781.32</i> 749.87*** (6.920)	<i>860.97</i> 806.44*** (14.080)		
Observations	7559	7560	9180	7200		
r2	0.68	0.50	0.67	0.39		
	$f = 10$ $q = 0$	$f = 40$ $q = 0$	$f = \infty$ $q = 0$	$f = 10$ $q = 0.5$	$f = 40$ $q = 0.5$	$f = \infty$ $q = 0.5$
T10 - T4	<i>61.61</i> 44.28* (9.339)	<i>81.03</i> 71.06 (8.834)	<i>85.48</i> 77.91 (8.363)	<i>66.30</i> 53.90 (14.110)	<i>77.15</i> 68.59 (12.210)	<i>79.65</i> 56.57 (15.610)
Constant (T4)	<i>791.20</i> 771.97*** (1.712)	<i>772.82</i> 748.53*** (3.742)	<i>764.58</i> 741.60*** (4.045)	<i>791.53</i> 755.97*** (8.969)	<i>782.32</i> 749.20*** (7.243)	<i>781.32</i> 749.87*** (6.905)
Observations	5040	5040	5039	5460	5460	5460
R^2	0.57	0.53	0.51	0.49	0.50	0.46
	$f = 10$ $T = 4$	$f = 40$ $T = 4$	$f = \infty$ $T = 4$	$f = 10$ $T = 10$	$f = 40$ $T = 10$	$f = \infty$ $T = 10$
IR - NR	<i>0.33</i> -16.00* (9.128)	<i>9.50</i> 0.67 (8.149)	<i>16.73</i> 8.32 (7.997)	<i>5.02</i> -6.38 (14.260)	<i>5.62</i> -1.79 (12.680)	<i>10.90</i> -13.06 (15.810)
Constant (NR)	<i>791.20</i> 771.97*** (1.710)	<i>772.82</i> 748.53*** (3.738)	<i>764.58</i> 741.58*** (4.035)	<i>852.82</i> 816.25*** (9.184)	<i>853.85</i> 819.59*** (8.004)	<i>850.07</i> 819.50*** (7.323)
Observations	5580	5580	5579	4920	4920	4920
R^2	0.75	0.64	0.61	0.42	0.46	0.46

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

Results from linear regressions where the dependent variable is profit. Each column fixes two treatment variables from $\{f, T, q\}$ (these are given in the column headers) and includes dummy variables for the third treatment variable. Each regression includes round-level fixed effects to control for heterogeneity across offer sequences. Standard errors clustered by subjects are in parentheses. The italic font stands for the coefficients predicted by the model.

Table 6: Under-searching

	$T = 4$ $q = 0$	$T = 10$ $q = 0$	$T = 4$ $q = 0.5$	$T = 10$ $q = 0.5$
LF	0.38	0.74	0.15	0.91
HF	0.03	0.44	0.01	0.69
NF	-0.02	0.39	0.08	0.70

This table indicates the estimated number of rounds by which subjects under-search, as compared to the model’s predictions.

5.3 Do individuals use optimal reservation price strategies?

In this section, we evaluate Hypothesis 3, which concerns the use of the optimal RR-DRR policy. Figures 7 and 8 illustrate the fraction of offers of different magnitudes that have been accepted and frozen in the first stage of each round, under NR and IR, respectively.²¹ 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

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

²¹This 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 E (Tables 9, 10, 11 and 12). Another complication for stages after the first is that the threshold in the IR case is history-dependent. Figures 10 and 12 in Appendix E assume that no prices were previously observed.

where m and r are parameters representing the midpoint and the curvature, respectively. Under each panel in figures 7 and 8, 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 can be reasonably 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. 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 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, with the exception of T10HF, in which freezing is predicted never 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 T4LF, and just outside

of it in the other two treatments in which freezing is predicted.

Figure 7: Empirical frequencies of decisions in the first stage - no recall

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 decision thresholds implied by the model.

Figure 8: Empirical frequencies of decisions in the first stage - imperfect recall

The plots contain empirical frequencies of accepting and freezing in the first stage of each round for the six Imperfect 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 decision thresholds implied by the model.

5.4 Discussion

Overall, hypotheses 1 and 2 are strongly supported in the data. The treatment differences are as the model predicts. This is the case even for the seemingly coun-

terintuitive prediction that freezing would be more common when the time horizon is shorter. Hypothesis 3 receives modest support in the sense of a good fit of aggregate decisions to a steep logistic function, although the fitted curve’s center is higher than the predicted threshold, which is consistent with under-searching. The main predictions regarding earnings are supported, namely that they are not significantly affected by the presence of a PFO, but are affected by the horizon length..

Even though the responses of search length to having a PFO are weaker than predicted, the PFO is a useful tool for stimulating search, which may be of interest for policy considerations. This can lead to more favorable price distributions in the long run. Furthermore, one can think of the introduction of a wage freezing mechanism in labor markets, potentially improving allocative efficiency as it induces exploration, increasing expected match values.

Under-searching has been a prominent finding in the experimental search literature. Beyond documenting it, we contribute by making three points that relate to under-searching. Firstly, affordable PFOs increase the extent of under-searching. Secondly, substantial under-searching leads to negligible loss in earnings. This is particularly interesting under equilibrium considerations, such as in a [Burdett and Judd \(1983\)](#) model, because firms’ response 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 behalf of the searchers in the long run.

Thirdly, we evaluate three mechanisms among the explanations of under-searching proposed in previous literature. The first mechanism is a cognitively cheap heuristic in the spirit of [Gabaix *et al.* \(2006\)](#), whereby individuals treat the search problem as having fewer remaining future stages than there actually are. In other words, acting as if they are in stage $s > t$ when in fact they are in stage t . In Appendix F we show that this conjecture does not help in explaining our data. Specifically, a model assuming that individuals are optimizing for a stage $s > t$ has less ability to predict decisions than the model introduced in section 3. The second mechanism we consider is anticipated regret, as proposed by [Loomes and Sugden \(1986\)](#) and operationalized in a search framework by [Weng \(2009\)](#).²² We also briefly consider whether risk aversion explains under-searching, even though we are doubtful about the risk elicitation protocols, as they do not agree with each other. We find no clear

²²Evidence regarding the anticipated regret mechanism will be included in the next version of this paper.

relationship between risk aversion as measured with our two protocols and search length.²³

Individuals behave as if they employ threshold strategies with regard to the range of offers that they accept, freeze and reject. The probability 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 optimal search in structural modeling of demand in markets featuring consumer search.

Our last insight is that most decisions that are not consistent with the model's predictions occur for offers near the acceptance threshold, indicating that errors are more likely the less costly that they are. This is consistent with the central behavioral assumption that generates a Quantal Response Equilibrium (McKelvey and Palfrey (1995)).

6 Conclusions

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 the length of a search. Lower freeze costs also increase search length. These patterns are observed in the data. Furthermore, as predicted, the usage of freezing decreases as it becomes more expensive, as the time horizon increases, and when recall of prior offers is possible. Our results regarding freezing are robust to having imperfect recall, though, as predicted, the impact of the presence of a freeze option is more pronounced when recall is impossible.

Beyond our hypotheses about treatment effects, we took the analysis of under-searching in sequential search models a step further. We did so by firstly showing that while affordable PFOs stimulate search, they strengthen under-searching because the additional search exerted by subjects falls short of the extra search predicted by the model. Furthermore, we have shown that under-searching is individually

²³See Appendix G for a summary of subjects' elicited risk attitudes and their correlation with subjects' choices.

harmless, but in equilibrium searchers may suffer welfare losses because firms respond to the under-searching tendencies of consumers. One avenue for future research is to formalize this notion theoretically within an equilibrium search framework.

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/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 given 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. \(2018\)](#).

Our theoretical model could be extended to allow the freezing of multiple offers, and of 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 for evaluating long run welfare effects. Finally, policymakers and firms may want to experiment with introducing PFOs for wage offers in labor markets.²⁴

References

- BASU, K. (2006). Consumer cognition and pricing in the 9's in oligopolistic markets. *Journal of Economics and Management Strategy*, (15).
- BERRY, S. and JIA, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, **2** (3), 1–43.
- BORENSTEIN, S. (1992). The evolution of us airline competition. *Journal of Economic perspectives*, **6** (2), 45–73.
- BRAUNSTEIN, Y. M. and SCHOTTER, A. (1981). Economic search: an experimental study. *Economic Inquiry*, **19** (1), 1–25.
- BRONNENBERG, B. J., KIM, J. B. and MELA, C. F. (2016). Zooming in on choice: How do consumers search for cameras online? *Marketing Science*, **35** (5), 693–712.

²⁴We thank Ron Oaxaca for discussing this with us.

- BURDETT, K. and JUDD, K. L. (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, pp. 955–969.
- COX, J. C. and OAXACA, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, **2** (3), 301–329.
- DE LOS SANTOS, B., HORTAÇSU, A. and WILDENBEEST, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, **102** (6), 2955–80.
- DEGROOT, M. H. (1970). *Optimal statistical decisions*, vol. 82. New York: McGraw-Hill.
- EINAV, L. (2005). Informational asymmetries and observational learning in search. *Journal of Risk and Uncertainty*, **30** (3), 241–259.
- GABAIX, X., LAIBSON, D., MOLOCHE, G. and WEINBERG, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, **96** (4), 1043–1068.
- HONG, H. and SHUM, M. (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics*, **37** (2), 257–275.
- KARNI, E. and SCHWARTZ, A. (1977). Search theory: The case of search with uncertain recall. *Journal of Economic Theory*, **16** (1), 38–52.
- KASHYAP, A. K. (1995). Sticky prices: New evidence from retail catalogs. *The Quarterly Journal of Economics*, **110** (1), 245–274.
- KIM, J. B., ALBUQUERQUE, P. and BRONNENBERG, B. J. (2010). Online demand under limited consumer search. *Marketing science*, **29** (6), 1001–1023.
- KOGUT, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, **14** (3), 381–392.
- LANDSBERGER, M. and PELED, D. (1977). Duration of offers, price structure, and the gain from search. *Journal of Economic Theory*, **16** (1), 17–37.
- LIPPMAN, S. A. and MCCALL, J. J. (1976). Job search in a dynamic economy. *Journal of Economic Theory*, **12** (3), 365 – 390.

- LOOMES, G. and 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), pp. 113–126.
- MCKELVEY, R. D. and PALFREY, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, **10** (1), 6–38.
- MORAGA-GONZALEZ, J., SÁNDOR, Z. and WILDENBEEST, M. (2018). Consumer search and prices in the automobile market.
- MORAGA-GONZÁLEZ, J. L. and WILDENBEEST, M. R. (2008). Maximum likelihood estimation of search costs. *European Economic Review*, **52** (5), 820–848.
- RUFFLE, B. J. and 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.
- SENGUPTA, A. and WIGGINS, S. N. (2014). Airline pricing, price dispersion, and ticket characteristics on and off the internet. *American Economic Journal: Economic Policy*, **6** (1), 272–307.
- SONNEMANS, J. (1998). Strategies of search. *Journal of Economic Behavior and Organization*, **35** (3), 309 – 332.
- (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, **21** (1), 91 – 102.
- STIGLER, G. J. (1961). The economics of information. *Journal of Political Economy*, **69** (3), pp. 213–225.
- THALER, R. (1985). Mental accounting and consumer choice. *Marketing Science*, **4** (3), 199–214.
- WEITZMAN, M. L. (1979). Optimal search for the best alternative. *Econometrica*, **47** (3), pp. 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. and MUTHUKRISHNAN, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, **22** (4), 503–519.

Six appendices are included with this paper. Appendix A consists of a table listing the cost of options to freeze ticket prices offered by a number of airlines in late 2017. Appendix B describes the procedure used to compute the optimal decision rule in our treatments with imperfect recall. Appendix C contains all proofs for section 3. Appendix D contains the instructions available to participants in the experiment. Appendix E consists of additional tables and figures. Appendix F considers whether a level- z decision structure, to be defined below, can describe the acceptance and freezing thresholds that individuals employ in the experiment. Finally, Appendix G reports the results from the two risk aversion measurement tasks that we implemented and shows the lack of correlation between behavior in these tasks and search behavior.

Appendix A: Price Freezing Options Offered in Late 2017

Table 7: Price freezing options offered by various airlines in the fall of 2017

Airline	Country	Cost	Length (in days)
KLM	The Netherlands	Variable	≤ 14
Air France	France	Variable	≤ 3
Aer Lingus	Ireland	5 EUR	1
Wizz Air	Hungary	Variable	2
Air Baltic	Latvia	5-11 EUR	2
Emirates	UAE	Variable	3
Jet Airways	India	350-700 INR	3
United	US	Variable	Either 3 or 7
British Airways	UK	10 USD	3
Flybe	UK	2.5 GBP	1
Kulula	South Africa	50 ZAR	1
Tiger Air	Australia	Variable	2
Vueling	Spain	Variable	1
Pegasus	Turkey	3 EUR	2
Turkish Airlines	Turkey	Free	1

Appendix B: Optimal behavior under imperfect recall

In this appendix, we outline and discuss the model that we use to compute the optimal policy for the imperfect recall treatments. We formalize the model and then discuss how our modeling of imperfect recall differs from previous literature. Importantly, in our implementation, it is possible that an offer is not available for recall in some period, but is available for recall in a later period.

We denote the lowest offer seen up to and including stage t as $\underline{p}_t = \min\{p_1, \dots, p_t\}$. The recall availability is modeled as a Bernoulli process with probability q , whose realization in stage t is denoted by q_t . The payment and rejection payments in stages after an offer k has been frozen are:

$$\tilde{V}_t^0(p_t, \underline{p}_{t-1}, k) = \min\{p_t, \tilde{R}_t(\underline{p}_t, k), k\} \quad (14)$$

$$\tilde{V}_t^1(\underline{p}_t, k) = \min\{\underline{p}_t, \tilde{R}_t(\underline{p}_t, k), k\} \quad (15)$$

$$\tilde{R}_t(\underline{p}_t, k) = \int_0^{\bar{p}} \left[q\tilde{V}_{t+1}^1(\underline{p}_{t+1}, k) + (1-q)\tilde{V}_{t+1}^0(p_{t+1}, \underline{p}_t, k) \right] dF(p_{t+1}) + c \quad (16)$$

where $\tilde{V}_t^0(p_t, \underline{p}_{t-1}, k)$ is the payment in a stage in which recall is not available, and $\tilde{V}_t^1(\underline{p}_t, k)$ is the payment when recall is possible.

If no offer has been frozen at an earlier stage, the payments $V_t^0(p_t, \underline{p}_{t-1})$ and $V_t^1(p_t, \underline{p}_{t-1})$, rejection payment $R_t(\underline{p}_t)$ and freezing payment $K_t(p_t, \underline{p}_{t-1})$ functions are given by:

$$K_t(p_t, \underline{p}_{t-1}) = \tilde{R}(\underline{p}_t, p_t) + f \quad (17)$$

$$V_t^0(p_t, \underline{p}_{t-1}) = \min\{p_t, R_t(\underline{p}_t), K_t(p_t, \underline{p}_{t-1})\} \quad (18)$$

$$V_t^1(p_t, \underline{p}_{t-1}) = \min\{\underline{p}_t, R_t(\underline{p}_t), K_t(p_t, \underline{p}_{t-1})\} \quad (19)$$

$$R_t(\underline{p}_t) = \int_0^{\bar{p}} \left[qV_{t+1}^0(p_{t+1}, \underline{p}_t) + (1-q)V_{t+1}^1(p_{t+1}, \underline{p}_t) \right] dF(p_{t+1}) + c \quad (20)$$

Note that in the terminal period we have:

$$\tilde{V}_T^1(\underline{p}_T, k) = \min\{\underline{p}_T, k\} \quad V_T^1(p_T, \underline{p}_{T-1}) = \underline{p}_T \quad (21)$$

$$\tilde{V}_T^0(p_T, \underline{p}_{T-1}, k) = \min\{p_T, k\} \quad V_T^0(p_T, \underline{p}_{T-1}) = p_T \quad (22)$$

and therefore the rejection payments are initialized by

$$\tilde{R}_{T-1}(\underline{p}_{T-1}, k) = qh(\min\{\underline{p}_{T-1}, k\}) + (1 - q)h(k) + c \quad (23)$$

$$R_{T-1}(\underline{p}_{T-1}) = qh(\underline{p}_{T-1}) + (1 - q)h(\bar{p}) + c \quad (24)$$

We differ from [Landsberger and Peled \(1977\)](#) in how we introduce imperfect recall. In [Landsberger and Peled \(1977\)](#), imperfect recall is introduced such that in every stage, the highest offer observed so far resets once a realization of $q_t = 0$ is encountered. In contrast, our experiment does not feature such a resetting property. Furthermore, in our experiment, in any given stage, the decision maker may freeze only the current offer, regardless of the recall realization. That is, even if there is a realization of $q_t = 1$, the decision maker cannot freeze the highest offer observed so far, but rather only the current offer.

Because the problem quickly becomes intractable, we use numerical methods to solve for the optimal policy. That is, while we do not solve for the general solution to the model analytically, we implement backward induction computationally, and compute the optimal policy for the parameters of our experimental environment.

Appendix C: Proofs

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.

Proof. First consider the case of $k \leq p^*$. This implies that $h(k) + c \geq k$. Therefore, at $T - 1$, we have $\tilde{R}_{T-1}(k) \geq k$. Suppose $\tilde{V}_t(p_t, k) = \min\{p_t, k\}$ for some $t \in \{2, \dots, T - 1\}$ and use (3) to observe that $\tilde{R}_{t-1}(k) = h(k) + c \geq k$. Therefore, we have that $\tilde{V}_{t-1}(p_{t-1}, k) = \min\{p_{t-1}, k\}$. Thus, at any stage, it is optimal to accept the lowest among the current offer and the outside option, since the rejection payment is greater than at least one of them. Thus, $\tilde{V}_t(p_1, k) = \min\{p_1, k\}$.

Now assume that $k > p^*$, which implies $\tilde{R}_{T-1}(k) = h(k) + c < k$. The outside option is not used at $T - 1$. Now, $\tilde{R}_{T-1}(k) \in [c, h(\bar{p}) + c]$ (which holds by assumption 1) and the independence of $\tilde{R}_{T-1}(k)$ from p , imply reservation play at $T - 1$ with threshold $\tilde{R}_{T-1}(k)$. That is, offers above $\tilde{R}_{T-1}(k)$ are rejected and lower offers are accepted. Suppose $\tilde{V}_t(p_t, k) = \min\{p_t, R_t(k)\}$ for some $t \in \{2, \dots, T-1\}$ and consider stage $t - 1$, in which $\tilde{R}_{t-1}(k) = h(\tilde{R}_t(k)) + c < h(k) + c < k$ where the first inequality holds because $h(\cdot)$ is increasing. Notice that $\tilde{R}_{t-1}(k) \in [c, k] \subset [0, \bar{p}]$ and therefore there is reservation play at $t - 1$ with threshold $\tilde{R}_{t-1}(k)$. Thus, we have shown that when $k > p^*$,

$$\tilde{R}_t(k) = \begin{cases} h(\tilde{R}_{t+1}(k)) + c & \text{for } t < T - 1 \\ h(k) + c & \text{for } t = T - 1. \end{cases} \quad (25)$$

Finally, we show by induction that the rejection payment $\tilde{R}_t(k)$, is increasing over time. Firstly, $\tilde{R}_{T-2}(k) = h(\tilde{R}_{T-1}(k)) + c < h(k) + c = \tilde{R}_{T-1}(k)$. Now, it is immediate from (25) that when $t < T - 1$, assuming $\tilde{R}_t(k) < \tilde{R}_{t+1}(k)$ implies $\tilde{R}_{t-1}(k) < \tilde{R}_t(k)$, as $h(\cdot)$ is an increasing function. Thus, the rejection payment remains lower than k until $T - 1$, so that k is not accepted at any stage from 1 to $T - 1$. Therefore, k may only be accepted at stage T . □

Lemma 1. The sequence $(\tilde{R}_t^T(k))_{T=2}^\infty$ converges uniformly to $\tilde{R}_t^\infty(k) = p^*$ for $k > p^*$.

Proof. The sequence $(a_T)_{T=2}^\infty$ defined by $a_T = \sup_{k \in (p^*, \bar{p}]} \tilde{R}_1^T(k)$ is strictly decreasing because $\tilde{R}_t^T$ is strictly increasing in t when $(p^*, \bar{p}]$, as shown in Proposition 1. We show by induction on $T \geq 2$ that it is also bounded below by p^* . Start by noting that $\tilde{R}_1^2(\bar{p}) = h(\bar{p}) + c > h(p^*) + c = p^*$. Suppose $\tilde{R}_t^T(\bar{p}) > p^*$ and observe that $\tilde{R}_t^{T+1}(\bar{p}) = h(\tilde{R}_t^T(\bar{p})) + c > h(p^*) + c = p^*$, which completes the inductive argument. □

Lemma 2. $R_t \leq \tilde{R}_t(\bar{p})$ for all $t < T$

Proof. By induction, $R_{T-1} = h(\bar{p}) + c = \tilde{R}_{T-1}(\bar{p})$. Suppose $R_t \leq \tilde{R}_t(\bar{p})$ for some $t \in \{2, \dots, T - 1\}$. Then,

$$R_{t-1} = \int_0^{\bar{p}} \min\{p_t, K_t(p_t), R_t\} dF(p_t) + c \quad (26)$$

$$\leq \int_0^{\bar{p}} \min\{p_t, K_t(p_t), \tilde{R}_t(\bar{p})\} dF(p_t) + c \quad (27)$$

$$\leq \int_0^{\bar{p}} \min\{p_t, \tilde{R}_t(\bar{p})\} dF(p_t) + c \quad (28)$$

$$= \tilde{R}_{t-1}(\bar{p}). \quad (29)$$

□

Lemma 4. Under the optimal solution, $t \in T^{RR} \cup T^{DRR}$ for all $t < T$.

Proof. It is sufficient to show that:

(i) $K_t(p_t)$ has a unique fixed point on $[0, \bar{p}]$ located $(p^*, \bar{p}]$, and that

(ii) $R_t \in (0, K_t(\bar{p}))$.

To establish (i), we show that (a) $K_t(p_t) > p_t$ for $p_t \in [0, p^*]$, (b) $K'_t(p_t) > 0$ and (c) $K_t(\bar{p}) < \bar{p}$. (a) holds by Proposition 1, as it implies that $K_T(p_t) > \tilde{R}_t(p_t) \geq p_t$. (b) holds as 25 implies that for $(p^*, \bar{p}]$, we have

$$K'_t(p_t) = h'(p_t) \prod_{j=t+1}^{T-1} h'(K'_j(p_t) - f) \quad (30)$$

where we use the convention that $\prod_a^b(\cdot) = 1$ when $a > b$. Now, as $h'(p_t) = 1 - F(p_t) \in [0, 1]$, we have $K'_t(p_t) \in [0, 1]$. To show (c) holds, it suffices to show that $K_{T-1}(\bar{p}) = h(\bar{p}) + c + f < \bar{p}$ because $\tilde{R}_t(\bar{p})$ is increasing in t . This holds because assumption 2 and the fact that $h(\cdot)$ is a contraction imply $f < h(\bar{p}) - h(h(\bar{p}) + c) < \bar{p} - h(\bar{p}) - c$. To see that (ii) holds, note that lemma 2 implies $R_t \leq \tilde{R}_t(\bar{p}) < \tilde{R}_t(\bar{p}) + f = K_t(\bar{p})$. □

Lemma 4. Any offer that is rejected in some period t is also rejected in period $t - 1$. That is,

$$V_t(p) = R_t \implies V_{t-1}(p) = R_{t-1} \quad (31)$$

Proof. The result follows because $h(\cdot)$ is a contraction. Define $B_t(p) = R_t - \tilde{R}_t(p)$ to be the marginal benefit from freezing over rejecting. When $t \in T^{RR}$,

$$B_{t-1}(p) = h(R_t) - h(\tilde{R}_t(p)) < B_t(p) \quad (32)$$

and when $t \in T^{DRR}$,

$$B_{t-1}(p) = \int_0^{a_t} p_t dF(p_t) + \int_{a_t}^{b_t} K_t(p_t) dF(p_t) + (1 - F(b_t))R_t - h(\tilde{R}_t(p)) \quad (33)$$

$$< h(R_t) - h(\tilde{R}_t(p)) \quad (34)$$

$$< B_t(p) \quad (35)$$

where the first inequality holds because $R_t \in (a_t, b_t)$. Therefore $B_t(p)$ is increasing in t , from which the result follows because f is time invariant. $\square$

Proposition 2. The optimal policy consists of a strictly increasing RR for $t \leq t^*$ and a strictly increasing DRR for $t > t^*$, where $0 \leq t^* < T - 1$.

Proof. Assumption 2 implies $T - 1 \in T^{DRR}$. By lemma 4, $t \in T^{RR} \implies t - 1 \in T^{RR}$. Combined with lemma 3, this implies that the optimal policy is characterized by some $t^* \geq 0$. Now we show that the DRR on $t > t^*$ is increasing over time. Note that $a_t > p^*$ for all $t \in T^{DRR}$ because p^* is the fixed point of $\tilde{R}_t(\cdot) = K_t(\cdot) - f$ and a_t is the fixed point of $K_t(\cdot)$. Therefore, the fact that $\tilde{R}_t(p)$ is increasing over time for $p > p^*$ implies that a_t is increasing in t . Lemma 4 implies that b_t is increasing in t .

To see that the RR is increasing on $t \leq t^*$, we first show that $R_t > p^*$ for all $t < T$. By induction, start with $R_{T-1} = h(\bar{p}) + c > h(p^*) + c = p^*$. Suppose $R_t > p^*$. If $t > t^*$, then $R_{t-1} > h(a_t) + c > h(p^* + f) + c > h(p^*) + c = p^*$. If $t \leq t^*$, then $R_{t-1} = h(R_t) + c > h(p^*) + c = p^*$ by the inductive assumption. Now, if $t \in T^{RR}$, then $R_{t-1} = h(R_t) + c < R_t$ because $R_t > p^*$. $\square$

Proposition 3. There exists a $T^* < \infty$ such that for $T > T^* \implies t^* > 0$.

Proof. It suffices to show that there exists a $T^* < \infty$ such that $\min\{p_1, R_1^T\} < K_1^T(p_1)$ for all $p_1 \in [0, \bar{p}]$. Firstly, R_1^T is strictly decreasing and bounded below by p^* , so we have that $(R_1^T)_{T=2}^\infty$ converges from above to p^* . Therefore, there exists T^* such that $R_1^T - p^* < f$ for all $T > T^*$. Secondly, Lemma 1 implies that $K_1^T(p_1) > p^* + f$

for all $p_1 \in (p^*, \bar{p}]$. Together, these imply that there exists a T^* such that $R_1^T < f + p^* < K_1^T(p)$ for all $T > T^*$ and $p \in (p^*, \bar{p}]$. Offers $p_1 \leq p^*$ are never frozen, as freezing yields $K_1^T(p_1) = h(p_1) + c + f > h(p_1) + c > p_1$ by definition of p^* . $\square$

Appendix D - Experimental Interface and Instructions

First Instructions

Instructions

Welcome. In this experiment, the instructions are simple and you can earn a considerable amount of money by following them. If anything is unclear, please do not hesitate to ask the experimenter. You can do so by raising your hand and the experimenter will come to you. At the end of the experiment, the earnings you have made will be paid to you in cash.

Let us start by describing the general structure of the experiment. The experiment consists of two stages. There are 2 periods in the first stage and 180 periods in the second stage. You will be paid for one of these 182 periods. The period for which you will be paid will be randomly drawn by the computer at the end of the experiment. Therefore, by obtaining high earnings in every period you increase your chances of getting a high final payment.

Stage 1

There are two periods in this stage. In each period, you will see a table. You need to choose between option A and option B in each row of the table. These options can be either a sure amount of money or an uncertain amount, which we call a gamble. A gamble consists of several possible payment amounts and the chances of obtaining each of these amounts. If you choose a gamble, the computer will draw the payment for this row randomly, according to the gamble's chances. If you choose a certain amount, then this amount will be the payment for this row. If the period is chosen at the end of the experiment to count, you will be paid for one of these rows, and the row for which you will be paid will be randomly determined by the computer. All amounts are in dollars.

Instructions for stage 2 will be given after all participants finish stage 1.

Second Instructions No Recall (1)

Stage 2

In this stage, there are 180 periods. Before these 180 periods, there are also five practice periods that cannot count toward your final earnings. The practice periods are included only to make sure you understand the interface. Properly understanding the interface will increase your chances of obtaining a higher payment.

The main element in each period is a sequence of offers you will receive. An offer is a randomly drawn number between 0 and 1000. Any number in this range is equally likely to be chosen by the computer as an offer. Moreover, offers are independent, which means that previously seen offers do not have any relation to the current offer or future offers. The offers are completely random - the experimenter has no effect whatsoever on the offers.

If you accept an offer, you receive a number of points equal to the amount in the offer. Points are converted into dollars at the exchange rate of 60 points to one dollar.

Please look at the figure in the handout, which is a screenshot of one period. To see an offer costs you some money. This is the offer cost, which you can find at the upper-left part of your screen. After every offer you see, you must click on one of the buttons "Accept current offer", "Reject current offer", or "Freeze current offer". We will first explain the "Accept current offer" and "Reject current offer" buttons.

- By clicking "Reject current offer", you refuse the current offer and will see another offer. You will pay the offer cost for seeing the new offer.
- By clicking "Accept current offer" you accept the current offer and the period ends.

Now we explain the "freeze current offer". By clicking on it, you ensure that the current offer will be available throughout the entire period, and you will be allowed to go back and accept this offer whenever you want, after seeing new offers, by clicking on the "accept frozen offer" button. You may either freeze one offer per period or not freeze any offer. That means that after you have frozen one offer, you cannot freeze any more offers in the current period. The cost you pay for freezing an offer can be found in the upper-center part of your screen. This is called the freeze cost. You will be notified if there is any change in this cost.

Second Instructions No Recall (2)

If a period in stage 2 is chosen as your final payment, your earnings are equal to:

$$\text{EARNINGS} = (\text{ACCEPTED OFFER}) - (\text{FREEZE COST IF AN OFFER WAS FROZEN}) \\ - (\text{OFFER COST}) * (\text{NUMBER OF OFFERS REJECTED})$$

Note that in every period, it is not possible to see more than 10 offers. Thus, it is not possible to reject the 10th offer. Remember, we will start with five practice periods that do not count toward your earnings. **During the practice periods, try clicking on all buttons to make sure you understand what they do and raise your hand if you have any questions.**

You must remain in the lab for at least one hour after the experiment started. Therefore, nothing can be gained by excessively rushing your decision. Taking your time to think about your actions will increase your performance, granting you higher earnings.

You may now start.

Second Instructions Imperfect Recall (1)

Stage 2

In this stage, there are 180 periods. Before these 180 periods, there are also five practice periods that cannot count toward your final earnings. The practice periods are included only to make sure you understand the interface. Properly understanding the interface will increase your chances of obtaining a higher payment.

The main element in each period is a sequence of offers you will receive. An offer is a randomly drawn number between 0 and 1000. Any number in this range is equally likely to be chosen by the computer as an offer. Moreover, offers are independent, which means that previously seen offers do not have any relation to the current offer or future offers. The offers are completely random - the experimenter has no effect whatsoever on the offers.

If you accept an offer, you receive a number of points equal to the amount in the offer. Points are converted into dollars at the exchange rate of 60 points to one dollar.

Please look at the figure in the handout, which is a screenshot of one period. To see an offer costs you some money. This is the offer cost, which you can find at the upper-left part of your screen. After every offer you see, you must click on one of the following buttons:

- By clicking "**Reject current offer**", you refuse the current offer and will see another offer. You will pay the offer cost for seeing the new offer.
- By clicking "**Accept current offer**", you accept the current offer and the period ends.
- By clicking "**Accept highest offer**", you accept the highest offer seen so far in the period. This button is not always available. For every offer seen, there is a 50% chance that it will be available.

Second Instructions Imperfect Recall (2)

- By clicking "**Freeze** current offer". , you ensure that the current offer will be available throughout the entire period, and you will be allowed to go back and accept this offer whenever you want, after seeing new offers, by clicking on the "accept frozen offer" button. You may either freeze one offer per period or not freeze any offer. That means that after you have frozen one offer, you cannot freeze any more offers in the current period. The cost you pay for freezing an offer can be found in the upper-center part of your screen. This is called the freeze cost. You will be notified if there is any change in this cost.

If a period in stage 2 is chosen as your final payment, your earnings are equal to:

$$\text{EARNINGS} = (\text{ACCEPTED OFFER}) - (\text{FREEZE COST IF AN OFFER WAS FROZEN}) \\ - (\text{OFFER COST}) * (\text{NUMBER OF OFFERS REJECTED})$$

Note that in every period, it is not possible to see more than 10 offers. Thus, it is not possible to reject the 10th offer.

You must remain in the lab for at least one hour after the experiment started. Therefore, nothing can be gained by excessively rushing your decision. Taking your time to think about your actions will increase your performance, granting you higher earnings.

Remember, we will start with five practice periods that do not count toward your earnings.

During the practice periods, try clicking on all buttons to make sure you understand what they do and raise your hand if you have any questions. You may now start.

Screen shot - No Recall

Offer cost

10

Freeze cost

10

Maximum number of offers per period

10

Offer number	Offer
3	437
2	201
1	947

Current offer	Total cost	Frozen offer
437	20	0

Accept current offer

Reject current offer

Freeze offer

Screen shot - Imperfect Recall

Simple Choice List

Option A	Option B	Choice
15 with a 50% chance ; 0 with a 50% chance	0 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	1 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	2 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	3 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	4 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	5 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	6 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	7 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	8 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	9 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	10 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	11 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	12 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	13 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	14 for sure	A ☐ ☐ B
15 with a 50% chance ; 0 with a 50% chance	15 for sure	A ☐ ☐ B
		<div style="border: 1px solid black; background-color: #ff0000; color: white; padding: 5px; display: inline-block;">OK</div>

Holt-Laury Choice List

Option A	Option B	Choice
12 with 10% chance ; 9.6 with 90% chance	23 with 10% chance ; 0.6 with 90% chance	A ☐ B ☐
12 with 20% chance ; 9.6 with 80% chance	23 with 20% chance ; 0.6 with 80% chance	A ☐ B ☐
12 with 30% chance ; 9.6 with 70% chance	23 with 30% chance ; 0.6 with 70% chance	A ☐ B ☐
12 with 40% chance ; 9.6 with 60% chance	23 with 40% chance ; 0.6 with 60% chance	A ☐ B ☐
12 with 50% chance ; 9.6 with 50% chance	23 with 50% chance ; 0.6 with 50% chance	A ☐ B ☐
12 with 60% chance ; 9.6 with 40% chance	23 with 60% chance ; 0.6 with 40% chance	A ☐ B ☐
12 with 70% chance ; 9.6 with 30% chance	23 with 70% chance ; 0.6 with 30% chance	A ☐ B ☐
12 with 80% chance ; 9.6 with 20% chance	23 with 80% chance ; 0.6 with 20% chance	A ☐ B ☐
12 with 90% chance ; 9.6 with 10% chance	23 with 90% chance ; 0.6 with 10% chance	A ☐ B ☐
12 for sure	23 for sure	A ☐ B ☐ <div>OK</div>

Appendix E: Additional Tables and Figures

Figure 9: Empirical frequencies of decisions in the second stage - No Recall

The plots contain empirical frequencies of accepting and freezing in the second 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 decision thresholds implied by the model.

Figure 10: Empirical frequencies of decisions in the second stage - Imperfect Recall

The plots contain empirical frequencies of accepting and freezing in the second stage of each round, for the six Imperfect 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 decision thresholds implied by the model.

Figure 11: Empirical frequencies of decisions in the third stage - No Recall

The plots contain empirical frequencies of accepting and freezing in the third 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 decision thresholds implied by the model.

Figure 12: Empirical frequencies of decisions in the third stage - Imperfect Recall

The plots contain empirical frequencies of accepting and freezing in the third stage of each round, for the six Imperfect 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 decision thresholds implied by the model.

Figure 13: Within-group Means - Second Half of the Data in Each Block

This figure contains the observed within-group means and frequencies for key variables for the second half of the data, that is, for the last 30 rounds within each block. Panel (b) includes also a breakdown of whether offers that have been frozen were later accepted and if so, it details whether these frozen acceptances occur in the last stage.

Appendix F: Using steps of reasoning to describe observed thresholds

As solving backward induction problems is mentally challenging, it is conceivable that individuals do use a rule similar to the threshold strategy implied by our optimal stopping model, but apply later stages' decision rules. That is, individuals are able to backward induct for a certain number of stages, but for fewer than the actual number of stages remaining. That would imply that in a given stage, individuals would use a lower acceptance threshold than that which is predicted by the model because they are using an optimal threshold from a later period. This argument is consistent with the pattern of early acceptances relative to the optimum for a risk-neutral agent that we observe. Another implication would be more freezing in earlier stages than under the optimal policy, which also appears in the data.

We define a *level- z decision rule* to be the decision rule under which the choice made in stage $t < T$ is optimal decision for stage $\min \{t + z, T - 1\}$. When $z = 0$ decisions are optimal, and when $z > 0$, decisions taken at stage t are those that would be optimal at stage $t + z$. Table 8 provides the fit of such decision rules for $z \in \{0, 1, 2\}$ to the data in terms of the percentage of pre-freezing choices that are consistent with this decision rule, conditional on past within-round behavior and for each treatment. The data are classified by whether the observed decision was to accept, to freeze, or to reject.

Table 8: Optimality of pre-freezing decisions under level- z decision rules

q	T	f	All_0	All_1	All_2	R_0	R_1	R_2	F_0	F_1	F_2	A_0	A_1	A_2
0	4	10	74.4	63.9	55.4	70.5	44.1	21.7	86.6	95.6	94.4	72.2	72.2	79.0
		40	79.7	74.0	66.2	82.4	63.0	46.4	53.6	68.7	63.5	88.4	91.4	94.6
		∞	94.3	91.4	87.6	94.3	87.1	78.8				94.2	96.8	98.7
	10	10	83.4	81.1	77.3	90.3	85.7	78.3	47.9	58.3	61.6	76.1	76.1	79.9
		40	89.8	87.4	83.9	95.8	92.0	85.7	7.8	14.2	32.3	82.6	83.1	84.4
		∞	93.0	92.9	92.2	97.0	96.2	94.5				82.9	84.6	86.3
	0.5	10	70.1	59.6	50.1	64.9	43.7	24.3	73.7	79.9	75.8	76.9	77.0	81.4
		40	80.3	72.5	66.9	87.8	65.4	54.6	9.5	37.2	29.9	80.7	88.6	90.8
		∞	87.2	86.5	83.5	92.6	83.5	76.1				79.7	90.6	93.8
0.5	10	10	84.5	81.3	78.2	96.4	91.1	83.2	10.6	19.4	39.0	67.0	67.1	72.5
		40	87.5	87.8	86.3	95.6	94.9	91.1	1.6	1.6	1.6	70.5	73.3	78.1
		∞	88.4	88.3	86.9	94.9	93.9	90.4				71.1	73.4	77.6

Each column contains the percentage of pre-freezing choices made that are consistent with a level- z decision rule, where z is indicated in the subscript of each choice type.

We see that, overall, level- z decision rules are not better than the model's predictions in rationalizing observed behavior. The overall hit rate does not improve if it assumed that $z > 0$. $z > 0$ implies fewer rejection decisions, and rejection decisions are far more frequent than acceptances and freezing. However, in almost all treatments, the hit rate for of freezing and acceptance choices increases with z . That is, early decisions to freeze and accept are better rationalized by a level z decision rule.

Table 9: Optimality of post-freezing decisions under level- z decision rules

q	T	f	All_0	All_1	All_2	R_0	R_1	R_2	A_0	A_1	A_2	Af_0	Af_1	Af_2
0	4	10	93.58	94.67	94.67	96.66	93.89	93.89	88.52	95.53	95.53	95.06	95.06	95.06
		40	93.48	93.61	93.61	96.09	92.85	92.85	90.44	95.41	95.41	92.31	92.31	92.31
	10	10	78.21	78.69	78.69	83.59	83.43	83.26	82.01	84.58	85.05	41.78	41.78	41.78
		40	82.68	83.31	83.62	91.07	91.07	91.07	79.19	81.50	82.66	35.59	35.59	35.59
	0.5	10	86.03	82.73	82.73	89.33	89.22	89.22	78.51	88.45	88.45	89.56	59.53	59.53
		40	84.20	83.23	83.23	91.02	90.77	90.77	73.62	87.01	87.01	83.93	59.52	59.52
0.5	10	10	79.92	79.75	81.01	85.98	85.98	85.73	69.93	74.83	80.77	57.14	39.56	39.56
		40	78.06	78.61	80.00	84.45	84.45	84.03	73.00	76.00	82.00	31.82	27.27	27.27

Each column contains the percentage of post-freezing choices made that are consistent with a level- z decision rule, where z is indicated in the subscript of each choice type.

Table 9 shows the accuracy of predictions, when it is assumed that $z > 0$, after an offer has been frozen. It shows that after freezing, assuming that $z \in \{1, 2\}$ does not

improve predictive accuracy. While it does predict more accurately when a current offer is accepted, it is not as good at predicting when a frozen offer is accepted.

Appendix G: Risk Aversion

In this appendix we present the data from the two risk aversion elicitation protocols and correlate the results to behavior in the experiment. Figures 14 and 15 show the distributions of measured risk aversion under the two measures. The two panels on the left are the data from the simple choice list and those on the right are for the Holt-Laury protocol. Figure 14 displays a histogram of the number of safe choices made by all participants, while Figure 15 contains only the data from those who exhibited a single switching point. The data show that a majority of individuals are risk averse, in that they make strictly greater than eight safe choices in the simple protocol and strictly greater than four safe choices under the Holt-Laury protocol. Substantial minorities are risk neutral (choosing 8 and 4 safe options in the two protocols, respectively) and another considerable group is risk seeking (making fewer safe choices than the risk neutral level).

Figure 14: Empirical Distribution of Elicited Risk Aversion - All Subjects

Histograms of elicited risk aversion measures using all subjects. The measure used is the number of safe choices made in a risk-aversion elicitation protocol. The left panel is using the simple choice list, whereas the right panel is using the Holt-Laury protocol. See Appendix A for the instructions to participants and the interface.

Figure 15: Empirical Distribution of Elicited Risk Aversion - Single Switchers

Histograms of elicited risk aversion measures using subjects who have made one switch from the risky to the safe option. The measure used is the number of safe choices made in a risk-aversion elicitation protocol. The left panel is using the simple choice list, whereas the right panel is using the Holt-Laury protocol. See Appendix A for the instructions to participants and the interface.

Tables 10 and 11 show the average search length in each treatment for risk averse (RA), risk neutral (RN), and risk seeking (RS) subjects, as measured in the two tasks. In the $f = \infty$ case, it can be shown theoretically that risk averse individuals terminate their searches earlier than risk neutral searchers, and in turn risk seeking individuals search for longer than those who are risk neutral. However, as can be seen in the two tables, this relationship is not consistently supported in the data. There is also no consistent relationship between risk attitude and search length when freezing is possible.

Table 10: Risk Aversion Summary - Simple Choice List

			RA	RN	RS
q	T	f			
0.0	4	10	2.49	2.72	2.70
		40	2.54	2.57	2.58
		∞	2.21	2.23	2.30
	10	10	4.26	3.16	4.02
		40	3.98	2.98	3.33
		∞	3.85	3.12	3.56
0.5	4	10	2.94	2.72	2.63
		40	2.70	2.45	2.54
		∞	2.57	2.35	2.39
	10	10	3.97	3.44	4.31
		40	3.79	3.18	4.08
		∞	3.85	3.20	3.97
Frac. subjects			47.72	29.55	22.73

Table 11: Risk Aversion Summary - Holt-Laury Choice List

			RA	RN	RS
q	T	f			
0.0	4	10	2.55	2.67	2.77
		40	2.57	2.54	2.52
		∞	2.26	2.15	2.25
	10	10	4.17	3.72	3.64
		40	3.87	3.20	3.52
		∞	3.78	3.16	3.70
0.5	4	10	2.83	2.49	2.82
		40	2.59	2.35	2.64
		∞	2.48	2.32	2.45
	10	10	3.78	4.23	4.30
		40	3.70	3.77	4.04
		∞	3.58	3.75	4.14
Frac. subjects			62.88	20.45	16.67