Learning while shopping: an experimental investigation into the effect of learning on consumer search

Ben Casner¹

Received: 6 March 2019 / Revised: 6 May 2020 / Accepted: 9 May 2020 / Published online: 9 June 2020
© Economic Science Association 2020

Abstract

In many search environments, searchers are learning about the distribution of offers in the market. I conduct an experiment exploring a broad class of search problems with learning about the distribution of payoffs. My results support the prediction that learning results in declining reservation values, providing evidence that learning may be an explanation for recall. Theory predicts a “one step” reservation value strategy, but many subjects instead choose to set a high reservation value in order to learn about the distribution before adjusting based on their observations. Under-searching in search experiments may stem from a reinforcement heuristic and lack of negative feedback after using sub-optimal strategies.

Keywords Search · Learning · Information heuristics · Search experiments · Unknown distribution · Reinforcement heuristic

JEL Classification D83 · D9 · L15

1 Introduction

Much of the literature on consumer search behavior assumes that searchers know the distribution of offers available in the market. In many markets the consumers who are most familiar with this distribution are the consumers who search the least (Moorthy et al. 1997). Therefore, it may be unreasonable to assume that consumers know the distribution in circumstances where they purchase infrequently or where a large percentage of consumers are new to the market. Uncertainty about the distribution can have many implications for search behavior depending on the form

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10683-020-09659-7>) contains supplementary material, which is available to authorized users.

✉ Ben Casner
casner.15@osu.edu

¹ The Ohio State University, 410 Arps Hall, 1945 N. High St., Columbus, OH 43210, USA

uncertainty takes. In this article I focus on circumstances where theory predicts that learning will result in declining reservation values similar to those frequently seen in field and experimental search behavior (e.g. Brown et al. 2011).

Most papers in the theoretical literature agree that with well behaved priors, searchers who are uncertain about the distribution of values will use a “one step” reservation (or cutoff) strategy similar to those common in standard search models.¹ “One step” here means that the searcher is deciding between accepting the best offer currently available and one additional search (Burdett and Vishwanath 1988; Dana 1994; Bikhchandani and Sharma 1996; Dubra 2004; Mauring 2017; Kaya and Kim 2018).² Furthermore, the initial reservation value of a purely Bayesian learner should be above that which would be used for the expected distribution from the priors, and reservation values decline as a search spell continues (Burdett and Vishwanath 1988; Bikhchandani and Sharma 1996; Dubra 2004). The high initial reservation value comes not from a desire to gather information about the distribution, but from the fact that a high observation will lead to more optimistic posteriors and therefore a more persistent searcher. It takes a very high draw to cause the Bayesian learner to stop after just one search. Declining reservation values stem from the fact that a low observation is viewed not just as an unlucky draw, but as bad news about the distribution. Consequently this low draw will lead to posteriors which are more pessimistic than the priors. A high draw may be sufficient to make the searcher more optimistic, but any draw which would be high enough to raise the reservation utility above its previous value will also be above that reservation value, meaning that the search spell ends before the reservation value can increase.³ I refer to this elimination of searchers who would raise their reservation values from the search process as the *selection effect*, whereby searchers who would increase their reservation values are filtered out of the data because they cease searching.

This selection effect is important because it could serve as a potential explanation for real world behavior which is inconsistent with standard sequential search theory without learning. Most notably, the decline in reservation values means that a rational searcher will sometimes exercise recall or exit the market in an environment with learning where they would not if they had full information. In a policy context this has important implications for labor market search. In particular, the increasing pessimism of job searchers could exacerbate the standard supply/demand dynamics that lead to low wages in a loose labor market. For managers of firms, particularly firms which are acting as marketplaces, learning means that consumers need to be shown appealing products early in their search process. Otherwise they are likely to believe that the products on offer are low-quality or overpriced, which can lead to consumers exiting the market entirely (e.g. Nosko and Tadelis 2015).

The selection effect has been thoroughly explored in the theoretical literature, but it has not previously been tested in an experimental environment. In this article I

¹ “Well behaved” here meaning that all distributions allowed by the priors should have a shared support, and one observation should not move the posteriors too much.

² The term from the search literature is “myopic”, but I use “one step” here to avoid accidentally conveying pejorative connotations. Especially given that the strategy is theoretically optimal.

³ Note that this refers to reservation values after the initial draw. It takes a high *initial* draw to get the Bayesian learner to stop after just one draw, but subsequent reservation values must decline.

attempt to fill this gap in the experimental literature by presenting an experiment which compares subjects' behavior in a simulated consumer search environment with a known distribution (the "full information treatment") to their behavior when they are updating over a set of priors (the "learning treatment"). I elicit reservation values using a method similar to Jhunjhunwala (2018) and Brown et al. (2011). Subjects exhibit declining reservation utility in both treatments, consistent with the previous experimental literature, but the rate of decline is much stronger in the learning treatment, suggesting that learning is a plausible explanation for the recall seen in real world search data.⁴

While the selection effect appears to work well in this experimental environment, the prediction that subjects will use a one-step reservation value does not. The change in the *initial* reservation values across treatments is inconsistent across subjects. Higher ability subjects—as measured by numeracy metrics and academic data—set higher reservation values than in the full information treatment and often maintained these values for several searches. Lower ability subjects tended to decrease their initial reservation values in the face of uncertainty. In a post-session questionnaire, those subjects who raised their initial reservation values expressed a desire to gain information about the distribution. This heuristic (gathering information and then deciding how to react) requires much less computation than the theoretically optimal strategy and performs quite well in terms of search payoffs. Additionally, it results in behavior which is quite similar to the "search funnel" described in the marketing literature, whereby consumers first engage in broad searches within a market to learn about the products on offer and their prices before narrowing their search once they have this information. Previous discussions of this search funnel have suggested that it may result from searchers' imperfect knowledge of their own tastes (Blake et al. 2016). However, taste variation is not a factor in my experiment. The similarity of behavior between the subjects in this article and the description of the search funnel suggests that learning about the market may be sufficient explanation for this phenomenon.

This information heuristic is likely to be of interest to researchers studying learning in complex environments. The general form of separating an information gathering "explore" stage from payoff relevant "exploit" behavior seems like it would apply to many economic contexts. For example, subjects in a market experiment who are setting prices against uncertain demand might prefer to test out several prices to get a sense of the market before they focus on maximizing profits. The heuristic performs well in this search environment, but the penalty for deviating from optimal behavior might be much larger in a different setting.

Despite some subjects setting higher initial reservation values in the learning treatment, most subjects in this experiment drastically under-search relative to the theoretical optimum. I provide evidence that this is a result of the structure of the game failing to give them negative feedback when they set a reservation value which is too low. Subjects who draw a value close to their reservation value are much more likely to increase their reservation value in the subsequent search. This draw

⁴ Subjects were 50% more likely to exercise recall in the learning treatment than the full information treatment. The *p* value on a Mann–Whitney test for similarity of the treatments is less than 0.1%.

makes the possibility of settling for a low value more salient and encourages them to become more persistent. Stopping with a value high above the reservation value does not intuitively feel like negative feedback, but stopping with a value close to or below the elicited reservation value (the latter due to exercising recall) does. Therefore the only significant change between search spells is subjects who reduce their initial reservation values in the following spell after receiving this “negative” feedback. This pattern is suggestive of a situation similar to Charness and Levin (2005), where the obvious reinforcement heuristic leads to actions opposite of those which would be indicated by Bayesian rationality.

2 Related literature

My experimental design emulates models where the selection effect causes reservation values to decline (Burdett and Vishwanath 1988; Bikhchandani and Sharma 1996), but it also contributes to a broader literature on search and learning. Dubra (2004) emphasizes the importance of priors in search behavior, and shows that with some priors a specific observation can cause the reservation value to increase drastically. Dana (1994) explores the implications of strategic firms in a consumer search environment with learning, but consumers’ rational expectations about firm pricing strategies mean the selection effect is not present. Mauring (2017) and Kaya and Kim (2018) describe models with alternative forms of learning which allow reservation values to increase or decrease. Yang and Ye (2008) create a theoretical model explaining “rocket and feather” pricing in markets where consumers learn about the value of search which has subsequently been used to explain gasoline pricing (Chandra and Tappata 2011; Lewis 2011). However Yang and Ye’s model is similar to that of Varian (1980) in that searchers know all the prices in the market, meaning that the learning dynamics do not allow for a reservation value. These examples illustrate that, while quite robust within the standard sequential search environment, the selection effect is dependent on the primary source of learning being consumers’ own observations.

This article also contributes to a small but growing experimental search literature. The seminal experimental economics work in search was conducted by Schotter and Braunstein (1981), whose main finding was that consumers tend to under-search relative to the theoretical optimum for a risk neutral agent. Most experiments have been similar to the design of Schotter and Braunstein in that they model the search process by allowing subjects to either continue searching or stop after observing an offer rather than eliciting reservation values (Rapoport and Tversky 1970; Kogut 1990; Zwick et al. 2003; Caplin et al. 2011). To my knowledge Jhunjhunwala (2018) and Brown et al. (2011) are the only other experiments which elicit reservation values from subjects. Both papers find that reservation values decline over the course of search spells despite subjects having full information about the distribution from which they are drawing prospects. Some of this decline may be due to the fact that subjects do not fully understand the distribution even after having it explained to them. Rapoport and Tversky (1970) show that subjects fail to converge to the theoretical optimum level of search even after weeks of training in

a distribution. However, Brown et al. show that time spent searching has a stronger effect on the decline in values than the number of searches, so their fatigue hypothesis seems more plausible.

A few other experiments have considered uncertainty. Falk et al. (2006) consider search in a labor market frame with uncertainty as to the probability of getting an offer upon searching.⁵ Cox and Oaxaca (2000) investigate search behavior in an environment with uncertainty and unknown distributions. Their environment is one of the most similar from the previous literature to that in this article. However, they use discrete uniform distributions with a finite search horizon. Importantly, their distributions often allow for observations which fully reveal the distribution. Because of the nature of their priors, they find that behavior is inconsistent with a one-step reservation value strategy. Indeed, one of their main research questions is whether theoretical models can successfully predict behavior in environments where theory predicts that a reservation value strategy will not be used, while the priors in my experiment are selected specifically to ensure that the theory predicts a reservation value strategy. Because they do not elicit reservation values, the majority of their analysis consists of tests evaluating whether observed stopping behavior is significantly different from theoretical predictions.

I am not aware of any other experiments which elicit reservation values from participants in a setting with uncertainty about the distribution of match values. This is significant since most of the theoretical predictions I test are stated in terms of reservation values, so I can approach this problem more directly than previous work. Eliciting reservation values does come at the sacrifice of making the assumption that subjects would use a reservation value strategy given freedom to choose. While I show that there is some evidence that this may not be true with an uncertain distribution, I believe that it is generally a safe assumption. Sonnemans (1998) shows that reservation value procedures are among the most frequently chosen when subjects are given the freedom to design their own strategies. Additionally, even if consumers would not use a one-step reservation value strategy of their own accord, the elicitation method provides a more nuanced picture of their perceived value for an additional search than implicit estimation methods using only stopping behavior would. For example, the information demand heuristic described in Sect. 6.2 would be much harder to study and might even have gone unnoticed without the elicited reservation values in the learning treatment.

There is a small empirical literature on search with learning. De los Santos et al. (2017) is the only empirical paper I know of which addresses the selection effect directly. They use data from online MP3 player purchases to demonstrate that ignoring the possibility that consumers learn while they search can lead to severe statistical bias in estimations of search cost. Weisbuch et al. (2000) describe consumer search behavior in a Marseille fish market, and show that consumers expend

⁵ This situation is strategically more similar to uncertainty about the cost of search than the uncertainty about the distribution described in the articles which form the theoretical basis for my experiment. See Yang (2013) or Casner (2020) for a more in depth discussion of search environments with a probability of failure and the similarity to an increase in search costs.

significant effort making sure that they have a general idea of the prices being charged on the market even though they generally only purchase from one store. Nosko and Tadelis (2015) show that consumers learning about the reputation of sellers on eBay exhibit increasing pessimism similar to the selection effect.

3 Theory

The class of search models explored in this paper is a one-sided, undirected, and sequential consumer search environment with recall, unit demand, and non-strategic firms.⁶ Consumers' utility from consuming a product is a stochastic consumer-firm match value u which is drawn randomly from some distribution and which consumers only observe upon visiting a firm. Consumers pay a search cost c to visit a firm, and they can only purchase from a firm they have visited. Since the firms are non-strategic, we can think of them as all charging the same price, or equivalently we can think of a match value as being the consumption value for a firm's product net of its price, in which case with the distribution of values stems from different prices being charged in the market. Regardless of interpretation, a consumer who searches k times and stops with match value u receives final utility

$$u - c * k \quad (1)$$

From Weitzman (1979), the optimal strategy in this environment is to use a one-step reservation value strategy. This means that the searcher will choose a reservation value, and if the highest value observed so far is below that reservation value then they will search once more and make a new decision based on the new draw. If their draw is above the reservation value then they will stop searching and consume the highest option they have observed. Given a distribution of utility values $F(\cdot)$ with support $[u, \bar{u}]$ and density $f(\cdot)$, the reservation utility x will be the minimum utility at which the cost of searching (c) equals or exceeds the expected value of an additional search. For a continuous distribution with free recall this is the x which solves⁷

$$\int_x^{\bar{u}} (u - x) dF(u) = c \quad (2)$$

In a stationary environment, neither the cost of searching nor the distribution of values changes as the consumer continues to search, so the reservation value should be constant, rather than declining as we see in actual behavior. If consumers are learning about the distribution of prospects as they search, then the stopping rule changes slightly, but a one-step rule is still optimal assuming the prior distributions have the

⁶ Non-strategic firms are not strictly necessary for the selection effect, but for the purposes of this experiment I am interested in consumer rather than firm behavior, so considering firm strategy would introduce counterproductive complexity. This model is also equivalent to one sided labor market search models with non-strategic wages.

⁷ Recall in my experiment is costly, but the stopping condition with costly recall is more complex and the intuition is identical, so I use free recall for the explanation of the one-step stopping rule.

same support and a single observation does not change beliefs too much (Burdett and Vishwanath 1988). In the experiment all uncertainty is about the mean of the distribution, so I focus on that environment here. Let μ be the mean of the distribution, $\mu^* \in [\underline{\mu}, \bar{\mu}]$ the searchers' prior mean and $g(\mu|\mu^*, x_1, x_2, x_3, \dots, x_n)$ the posterior density assigned to μ given the priors and previous observations $x_1, x_2, x_3, \dots, x_n$. The stopping rule for a one-step searcher is then stopping after observing any $u > x(\mu^*, x_1, x_2, x_3, \dots, x_n)$ where $x(\mu^*, x_1, x_2, x_3, \dots, x_n)$ solves

$$\int_{\underline{\mu}}^{\bar{\mu}} \int_{x(\mu^*, x_1, x_2, x_3, \dots, x_n)}^{\bar{\mu}} (u - x(\mu^*, x_1, x_2, x_3, \dots, x_n)) dF(u|\mu) g(\mu|\mu^*, x_1, x_2, x_3, \dots, x_n) d\mu = c \quad (3)$$

The selection effect says that there is no observation which could place a high posterior weight on a high μ without also being above the reservation value. Consumers who continue to search must become increasingly pessimistic as long as the priors are well behaved. Therefore $x(\mu^*, x_1, x_2, x_3, \dots)$ must decline as the number of searches increases. My experiment is designed to test if the selection effect appears in actual behavior. Previous experimental work (Brown et al. 2011) has shown that reservation values decline even in environments with a known distribution, so it is clear that learning about the distribution cannot be the only factor leading to declining reservation values.

Two other common explanations for declining reservation values include failure to account for the sunk cost fallacy (Kogut 1990), and fatigue (Brown et al. 2011). Failure to account for the sunk cost fallacy means that the searcher fails to treat expended search costs as sunk, instead adding the them to the right hand side of Eq. (2). Fatigue is the idea that continuing to search is tiring and/or boring, so that the cost of searching (c) actually increases over time. All three processes should result in a monotonically declining reservation value and the effects are difficult to separate in an empirical setting, but they are distinguishable in a laboratory environment. I control for fatigue in my experiment by tracking how long subjects have been searching when they set reservation values. I do not explicitly control for sunk cost, but Kogut's description of sunk cost mostly focuses on the induced search cost rather than effort. Since the induced search cost is constant whether or not subjects know the distribution, any difference between the treatments should be entirely the result of the difference in information structures.

4 Design of the experiment

The experiments for this article were conducted in the Experimental Economics Laboratory at the Ohio State University using zTree software (Fischbacher 2007). The design of the experiment was within-subject, with subjects participating in seven search spells with a known distribution and seven search spells with an uncertain distribution. I refer to the search spells where they knew the distribution as the *full information* treatment and the search spells with an uncertain distribution as the

learning treatment. To test for order effects, some subjects experienced the learning treatment first and some the full information treatment first. The 81 participants were recruited through The Ohio State University Economics Department's "Online Recruitment System for Economic Experiments" (ORSEE) (Greiner 2015). 38 of these subjects experienced the full information treatment first, and 43 saw the learning treatment first. The average payoff for the experiment was \$13.68, with a minimum of \$7.55 and maximum of \$18.75. Each session took approximately 1 h.

4.1 Anatomy of an experimental search spell

To simulate the search process, subjects drew values denoted in experimental currency units from a truncated normal distribution with support $[0, 700]$, a standard deviation of 100, and a mean which depended on the treatment.⁸ To simulate the cost of searching each draw required subjects to pay a cost of 5 experimental currency units. Subjects' final payoffs for a search spell were the maximum match value observed minus the sum of search costs accrued over the course of the spell. For example, a subject who searched 5 times and stopped with a match value of 400 would have a final payoff of 375 experimental currency units for that search spell. This payment structure creates a finite horizon for search. I impose a non-negativity constraint on payout for the search task, but a subject who searched sufficiently many times could theoretically have accrued a final payoff of 0. I included a notification at 25 searches warning subjects of this possibility, but the longest search spell in all of the data was 19 searches, so this warning was not seen by any subject, nor did any subject come close to a 0 payoff in any search spell. It is worth noting that even at 20 searches the total accumulated search costs would be 100, meaning that a subject would have to be incredibly unlucky to not receive a positive payoff after that many searches. Given the ubiquity of under-search in previous experiments and the small search cost, it does not seem likely that this finite horizon had any impact on search behavior.

The theory makes specific predictions about how the path of reservation values should differ when the distribution of values is and is not known. The main outcome of interest in this experiment is whether subjects' reservation values decline within search spells and if that decline is steeper when they are learning about the distribution (see Sect. 5 for a more detailed description of the hypotheses in the context of this experimental design). In order to elicit reservation values at every point in a search spell, subjects were asked to input a "minimum acceptable value" (i.e. a reservation value) before each draw. If the draw was above the subject's minimum acceptable value, then that would be interpreted as an "acceptable offer" and the search spell ended with the payoff as described above. If the draw was below this reservation value, then subjects had the option to either pay the search cost to accept their highest draw so far, equivalent to exercising costly recall in a traditional search

⁸ For clarity, the standard deviation of the original distribution was 100. It is therefore slightly smaller for truncated distribution, but the amount of mass removed by the truncation is quite small, so the difference is negligible.

environment, or they could state a new reservation value and pay the search cost to receive another draw. Each search spell continued until the subject either drew a match value above their current minimum acceptable value, or accepted their highest draw. Costly recall ensures strict incentive compatibility of stating the true reservation value for every search. A value below the true reservation value only changes the payoff if it causes the subject to stop after observing a value that is, by definition, lower than the payoff at which they would wish to stop. Stating a value above their true reservation value and recalling leads to a strictly lower payoff than stating the true reservation value in the first place because of the cost of recall. Each treatment consisted of two practice search spells with no payoff, and five payoff relevant spells. Payment for the search task was calculated by randomly selecting the results of one of the five payoff relevant search spells in each treatment.

4.2 Treatments

In the “full information” treatment, the mean of the truncated normal distribution was 350 for every search spell and subjects were aware of this fact. In addition, subjects were shown a histogram with the results of several thousand simulated draws from the distribution. Once subjects began to search, the values they drew appeared as red dots on the axis of the histogram. Figure 1 shows the user interface in the full information treatment.

In the learning treatment, the support of the distribution was still 0 to 700 units, but the mean of the distribution was determined randomly at the beginning of each search spell. The mean was drawn from a $U[150,550]$ distribution, and subjects were aware of both the range and shape of the prior distribution, but did not know the realized value. Figure 2 shows the user interface in the learning treatment.

Figure 3 shows the density plots for two possible distributions that subjects might have faced in the learning treatment. Importantly, the endpoints of the parameter distribution were chosen to be sufficiently far away from the endpoints of the truncation that a high or low mean had minimal effect on the shape of the distribution. Subjects guessed the actual value of the mean after each search spell in the learning treatment. If their guess was within 100 of the actual mean, then they had 100 minus the absolute value of the difference between the correct number and their guess added to their payoff for that search spell. This data is used in Sect. 9 to evaluate the determinants of the distance between subjects’ guess and the actual value. Experimental currency units were converted to dollars at a rate of $1\text{ ECU} = \$0.014\text{ USD}$.

4.3 Supplemental data

The experiment included several additional tasks to gather supplemental data. Subjects underwent a risk aversion assessment similar to Holt and Laury (2002) to track risk preference. To track numeric competency they completed a symbolic number mapping line task (Peters and Bjalkebring 2015). For the number line task, subjects were shown an otherwise unmarked number line with endpoints between 0

and 1,000 and they were asked to mark the position of a sequence of numbers on that line. The closer their marks to the actual positions of the numbers on the line, the more numerate they were considered to be. Numeracy is correlated with better ability to process information, and so the expectation was that more numerate subjects would exhibit behavior closer to the optimal strategy. Additionally, Schley and Peters (2014) show that higher symbolic number mapping numeracy is associated with more risk neutral behavior. This comes from the association between symbolic number mapping numeracy and ability to form accurate intuitive estimates of numeric information. Subjects with low numeracy scores tend to estimate the value of experimental lotteries more conservatively and so exhibit more risk averse behavior. There was also a subjective numeracy questionnaire which assessed subjects' self-perception of their ability with numbers. Falco (2019) hypothesizes that subjective numeracy may have effects on behavior separate from its correlation with actual numeric ability.⁹ Those subjects who have higher subjective numeracy have been shown to have a desire to use more numeric information than those who are less confident with numbers. The hypothesis behind measuring subjective numeracy was that this behavior could manifest as greater persistence in search spells with an unknown distribution due to a desire for more information about the distribution.

5 Hypotheses from the theory

Applying Eqs. (2) and (3) to the experimental environment, the optimal reservation value in the full information treatment is 475—which implies that the average search spell should be 9.5 searches long—and the optimal initial reservation value in the learning treatment is 593, although the expected length would not be trivial (or indeed particularly informative) to calculate because of path dependence between the reservation values and observations. According to the theory, reservation values should be constant in the full information treatment, while the learning treatment will lead to declining reservation values. The rate of decline is also path dependent as a searcher who sees middling values will be less pessimistic than one who has seen a number of low draws. These predictions are summarized in Hypothesis 1:

Hypothesis 1 Reservation values will decline as search spells continue in the learning treatment but will be constant in the full information treatment.

From the previous literature we tend not to see constant reservation values with a known distribution. However, the selection effect from the theory should only be present in the learning treatment. Therefore the main comparison of interest is

⁹ The data from this project was shared with David Falco for a separate research project. I would like to thank him both for suggesting the subjective numeracy assessment and providing the questionnaire. See his paper (Falco 2019) for a psychological investigation of the relationship between numeracy and consumer search.

Fig. 1 An example of the user interface subjects saw in the full information treatment. The horizontal axis of the graph is the range of possible value draws and the histogram shows 10,000 simulated draws from the distribution (an approximation of the density function). The dots on the axis are values that the subject drew themselves

Fig. 2 An example of the user interface subjects saw in the learning treatment. It is mostly identical to the full information treatment, but the subjects did not receive full information about the distribution of payoffs and only saw their own draws as dots on the axis

Fig. 3 The density plots for two possible distributions in the learning treatment. The solid line shows the density for a mean of 300 and the dashed line shows a mean of 450

whether the rate of decline is significantly steeper in the learning treatment than in the full information treatment.

Based on the predictions of the theory and behavior in previous experiments with full information, the design of this experiment assumes that subjects will use a one-step reservation value strategy in both treatments. While the previous literature has provided ample evidence that this is true when subjects know the distribution (e.g. Sonnemans 1998; Brown et al. 2011), it is important to check that this is indeed the case when interpreting the results of the learning treatment.

Hypothesis 2 Subjects will use a one-step reservation value strategy in the learning treatment.

Testing Hypothesis 2 is not quite as straightforward as Hypothesis 1, and is partly a qualitative exercise. The key feature of a one-step strategy is that the Bayesian searcher will be using the information conveyed by searches, but will not search specifically to gain information. If a rational Bayesian agent observes an initial draw of 550 in the learning treatment, then they will continue searching not because they want to gather more information but because their posterior beliefs put high weight on favorable distributions. If the first observed value is above the theoretically optimum reservation value of 563 then the posteriors are optimistic, but the expected marginal value of one additional search is smaller than the search cost. If the searcher observes 550 and then 520, then the posterior beliefs would still be quite high, but the reservation value would be approximately 540 and the searcher would exercise recall. An example of behavior in this experiment that is inconsistent with a one-step strategy is entering a reservation value of 700 because there is no possibility of drawing a value which would exit the search spell. This ensures that the searcher must either exercise costly recall or search at least once more, meaning that there is no possibility of completing the search process in *one step*. Another strategy that some subjects use is that they will set a high but not extreme reservation value, e.g. 600. This reflects the reasoning that if they get a draw above 600 then it

does not matter so much what the actual distribution is, but they will then maintain that reservation value for several draws before adjusting based on their observations. The fact that they are waiting to make use of the informational content of the draws is again inconsistent with deciding between accepting the current highest value or drawing once more. Results 1 and 2 in Sect. 6 give more detailed descriptions of how Hypotheses 1 and 2 fare when evaluated against the data.

Hypotheses 3 and 4 are behavioral hypotheses and so I have deferred them to Sect. 7 later in the paper.

In Sect. 9, I analyze subjects' behavior in comparison to the theoretically optimal behavior. I did not enter this project with a strong prior about how behavior would differ from the optimum, so Result 5 in that section does not have a corresponding hypothesis here.

Finally, Sect. 10 describes a robustness check evaluating whether subjects behavior is similar given a higher search cost. Result 6 does not have a corresponding hypothesis apart from the implicit null prediction that behavior will not be significantly different from behavior in the low search cost sessions apart from shorter search spells.

6 Main results

Because the practice spells were not payoff relevant, I drop those spells from all analysis. The results in the rest of the paper use the data from the ten payoff relevant spells.

Table 1 shows summary statistics for the full information and learning treatments. The average search length is longer with the learning treatment, but the average payout for the searchers is lower. This is because the average length of search spells in the learning treatment is inversely correlated with the mean of the distribution (as can be seen in Fig. 8). I discuss the implications of this correlation in Sect. 6.1.

Figure 4 shows a kernel density representation of the number of searches in each search spell. Most search spells lasted less than 5 searches, and 97% of all spells lasted less than 10 searches, I therefore drop search spells longer than 10 from the graphical analysis in this section. If risk neutral subjects were searching according to the theoretical optimum in the full information treatment then the expected number of searches would be 9.5. Subjects are therefore under-searching.

Figure 5 shows the reservation values of subjects at each point in their search spells aggregated across all search spells. Thus the left most box and whisker describes the initial reservation values in all of the full information search spells, whereas the 2nd from the left collects 2nd reservation values from the full information spells for all subjects who searched twice and so on. The boxes represent the middle 50% of the data, with the upper end of each box being the 75th percentile of the data and the lower end the 25th percentile. The whiskers contain all data that is less than 1.5 times the inter-quartile range away from the upper or lower quartile of the data.¹⁰ Dots represent outliers. The optimum reservation value in the full

¹⁰ The inter-quartile range is the distance between the 75th and 25th percentile.

information treatment is 475 and the optimum initial reservation value for a risk neutral searcher using Bayes' rule in the learning treatment is 593. Consistent with under-searching in the previous experimental search literature, most subjects stated reservation values well below the optimum in both treatments, although there is more dispersion in the learning treatment than in the full information treatment.¹¹ It is possible that some of this dispersion can be attributed to differences in values observed by different subjects in the learning treatment causing different inferences about the mean, but this cannot explain the difference in initial reservation values.

Figure 6 shows that subjects had disparate responses to the change in environment. There is a much wider spread in the distribution of initial reservation values in the learning treatment. Some subjects increased their reservation values, while others *reduced* their initial reservation values. Of particular interest is the cluster of initial reservation values at or near the top of the possible outcomes in the learning treatment, including three observations *above* the support of possible draws. This is the most extreme version of the information demand heuristic described in the introduction, with subjects setting reservation values that will almost certainly not be drawn in order to get a better sense of the distribution. This behavior is completely inconsistent with a one-step strategy as the one-step strategy is only concerned with the expected marginal benefit—conditional on beliefs—of one additional draw. Drawing a value with the intent of making even more draws afterward should not occur according to the theory.

Subjects also had differing reports about their own reaction to the learning environment. In a post session questionnaire asking subjects to describe their strategies in the learning treatment, the two most common responses were variations on either “Chose high until enough draws were done to see a distribution” or “The unknown distribution was harder so I chose to be even more conservative and opt for a lower minimum because I didn’t want to draw multiple times” (both quotes are actual responses). Thus there appear to be two behavioral factors at play. The first is a desire to learn about the distribution before making a decision, which increases reservation values. The second is aversion to the increased complexity in the uncertain environment which leads to more conservative search behavior.

The demand for information from the first factor contrasts with the high reservation values in the one-step strategy from the literature in that many of the subjects will pick a high reservation value and stick with it until they have enough information to begin adjusting, rather than adjusting with every draw. While this demand for information is not strictly optimal, it is a well performing heuristic and much easier to calculate than the true optimal reservation value. For context, calculating the optimal reservation values given data on subjects' observations took a computer nearly an hour. Thus, this heuristic is likely boundedly optimal given that the theoretically optimal strategy is too computationally intense to actually implement. It separates the complex problem of simultaneously exploring the environment and exploiting

¹¹ The average distance from the optimum was approximately 85 in the full information treatment and 185 in the learning treatment. I discuss distance from the optimum more thoroughly in Sect. 9.

Table 1 Summary statistics for each treatment

	Full information	Learning
Average search length	2.785 (0.127)	3.217 (0.158)
Average search profits	418.17 (3.876)	409.77 (6.376)

Search profits are stated in experimental currency units. Standard errors are shown in parentheses

Fig. 4 Kernel density representation of the length of the search spells **Fig. 5** The reservation values by depth in search spell and treatment

Fig. 6 Initial reservation values by treatment

previously gathered information into an “explore” stage and an “exploit” stage. This explore stage typically lasted for 2–3 searches, but the length of the exploit stage was much more variable since it depended on the information gathered in the initial searches. A subject who had unusually high initial draws from an unfavorable distribution would be much more likely to continue searching for a long time than one who had low draws from a favorable distribution. I discuss this heuristic more thoroughly in Sect. 6.2.

6.1 The reservation value path

The main interest in this project is in the change in reservation value paths when subjects are learning as opposed to when they have full information about the distribution. Using the unmodified reservation values masks patterns in the reservation value paths. For example, the apparent flat or upward trend of reservation values in the full information treatment in Fig. 5 is an illusion caused by the fact that subjects with low reservation values end their searches earlier. This creates a survivorship bias that appears to give a positive correlation between depth in a search spell and reservation values. To correct for this bias, much of my analysis in this paper uses subjects’ *normalized* reservation values. That is, I subtract the initial reservation value from the later reservation values in each search spell so that only the change in reservation values remains. Formally, for the k th search by subject i in the j th search spell, the normalized reservation value is given by

$$\text{Normalized Reservation Value}_{ijk} = \text{Reservation Value}_{ijk} - \text{Reservation Value}_{ij1} \quad (4)$$

I drop the first search from graphs involving the normalized reservation values since this value is always 0. As Fig. 7b shows, the selection effect is much more apparent when looking at the normalized reservation value. There is a declining trend in both treatments, but it is much stronger when subjects are learning about the distribution. The reservation values are so strongly concentrated around 0 in the full information treatment that the means and the upper end of the inter-quartile ranges very nearly coincide—making them difficult to distinguish in the graph. Similarly, Fig. 7a shows the kernel density of the normalized reservation values in each treatment after the initial search. The peak of the density for the full information treatment is firmly over 0, with little deviation, again reflecting the fact that many subjects picked a reservation value and did not vary far from it. In the learning treatment, not only is there a long tail reflecting declining reservation values, but the peak of the kernel density is noticeably leftward of 0, reflecting the much stronger decline in reservation values along search spells in the learning treatment. 90% (76 out of 81) subjects lowered their reservation values more often than they raised them in both treatments. The difference is more striking when we look at behavior spell by spell, with approximately half of the spells in the full information treatment having a decreasing trend in the reservation values and most of the remaining being flat. However, in the learning treatment approximately 70% of the reservation value paths showed a decreasing trend.

The graphical evidence strongly indicates that this decline is primarily due to the selection effect. We can see in Fig. 8 that subjects with high draws for the mean tend to have much shorter search spells, and even then the tendency is for the reservation value to decline since continuing with a high mean is correlated with a high initial reservation value. When the distribution is less favorable, search spells tend to be longer and reservation values exhibit a much stronger decline. This pattern demonstrates that welfare implications for searching from an unknown distribution are doubly negative. Not only is the optimal strategy more difficult to approximate, but a poor distribution leads to increased expenditure of search costs in addition to a low expected payoff. This pattern is strikingly similar to Blake et al.'s (2016) finding that successful search spells on eBay tend to be significantly shorter than unsuccessful ones.

The results of a fixed effects ordinary least squares regression of the normalized reservation values on predictive variables are shown in Table 2.^{12,13} “# of searches” is the number of times a subject has searched within a given search spell. The number of searches is not predictive of changes in the reservation value, but both time spent searching and the interaction of time with the learning treatment are significant and negative. Time is measured in seconds since the beginning of the search spell.¹⁴ This result is consistent with Brown et al.'s (2011) result that time spent searching has a greater impact on reservation values than the number of searches.

¹² I do not include subject level controls here because the fixed effects regression makes them redundant.

¹³ All regressions in this article use robust standard errors (clustered on subject where appropriate). Insignificant covariates are dropped from regression tables to aid readability.

¹⁴ A regression using time since the beginning of the experiment rather than since the beginning of the search spell yielded broadly similar results.

(a) The kernel density of the normalized reservation values in each treatment. The initial searches have been dropped from this graph

(b) The normalized reservation values by depth in search spell and treatment. Note that the data starts at the 2nd search since the first value is always 0 due to the normalization

Fig. 7 The kernel density and approximate paths of the normalized reservation values

“Learning” is a dummy for the learning treatment, and it is not significant on its own, but the interaction between time spent searching and learning is significant, consistent with the selection effect on reservation values. However, this interaction could also be caused by greater fatigue in the more complex learning environment. Period refers the number of search spells completed, so data associated with a subject’s third search spell would have a period value of 3. This coefficient is marginally significant and positive, which in the context of this regression is evidence suggesting that subjects became more persistent as they gained more experience with the environment.

To separate the fatigue explanation from the selection effect and to explore further the factors driving the changes in reservation values, I create an “update” variable, which takes on a value of 1 if a subject increases their reservation value above the previous one, 0 if the subject does not change their reservation value, and -1 if the subject lowers their reservation value.¹⁵ Table 3 shows the results of a random effects ordered probit on this update variable. “Previous Observation” is the last value the subject drew before setting a new reservation value, if subjects are learning using an approximation of Bayes’ rule, then this value should be the only information they are using to update their beliefs since all previous observations are already incorporated into the priors. The previous observation only has a significant effect in the learning treatment, and the coefficient is positive, supporting the idea that the decline in reservation values in the learning treatment is driven significantly by changes in beliefs as a result of the observed values. In the context of an ordered probit with three outcomes, a positive coefficient implies a positive correlation with the event of raising the reservation value and a negative correlation with lowering the reservation value.¹⁶ Therefore the significant positive coefficient on the previous

¹⁵ Although the theory predicts strictly declining reservation values, I avoid making this restriction since subjects did occasionally raise their reservation value, albeit rarely.

¹⁶ The implications for the middle event (maintaining the same reservation value) are ambiguous. I thank an anonymous referee for pointing this out.

Fig. 8 The normalized reservation values in the learning treatment by depth in search spell, broken down by the true mean of the distribution. A low mean is less than 280, a high mean is greater than 415

Table 2 Fixed effects OLS predicting the reservation value path

Dependent variable: normalized reservation value	
Covariates	Coefficients (SE)
Learning	9.831 (21.0405)
# of searches	1.1402 (1.8340)
Learning * # of searches	− 3.8343 (3.3525)
Time	− 0.5268** (0.2230)
Learning * time	− 0.9405** (0.3956)
Period	3.3311* (1.7622)
Number of subjects	81
Number of observations	2428
Fixed effects grouped on subject and period	
* ** ***Indicates significance at the 10, 5, and 1 percent levels respectively	
Standard errors are shown in parentheses	

observation means that subjects who observe a high (low) value are more (less) likely to raise their reservation value in the following search and less (more) likely to lower their reservation value, while the impact on the probability of maintaining the

Table 3 Random effects ordered probit regression predicting the direction of change in reservation values

Dependent variable: change in reservation value		
Covariates	Coefficients (SE)	
	(1)	(2)
Previous observation	– 0.0004 (0.0004)	– 0.000 (0.0005)
Learning * previous observation	0.0019*** (0.0003)	0.0018*** (0.0005)
SMAP numeracy score	0.0090*** (0.0044)	0.0047 (0.0083)
Time	– 0.0082* (0.0046)	– 0.0109** (0.0050)
Learning * time	– 0.0085** (0.0040)	– 0.0081* (0.0045)
Academic controls	No	Yes
Upper cutoff	0.9268	0.5974
Lower cutoff	– 0.7165	– 1.2226
Number of subjects	69	53
Number of observations	1670	1220

None of the academic controls (ACT, GPA etc.) in the 2nd regression were significant, therefore I omit their coefficients to save space

* ** ***Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

same reservation value is ambiguous. Although these results do support the selection effect, time spent searching has a significant effect both on its own and when interacted with the learning treatment. This means that subjects are more likely to lower and less likely to raise their reservation value as they spend more time searching. Therefore it is likely that fatigue is playing a role as well.

The SMAP numeracy score is the mean absolute error of the subjects' performance on the number line task described above, so a higher SMAP score indicates a less numerate subject.¹⁷ Numeracy is significantly predictive of behavior when there are no other proxies for ability, but insignificant when GPA and ACT scores are introduced to the regression. This is likely an increase in standard error caused by loss in efficiency of the regression because the academic data and the SMAP score are both serving as proxies for numeric ability.¹⁸ Less numerate subjects are less likely to adjust their reservation values downward and more likely to raise their

¹⁷ The number line task involved completing the task 10 times, so

$$SMAP = \frac{1}{10} \sum_{i=1}^{10} |i\text{th input value} - i\text{th actual value}|.$$

¹⁸ The academic controls were not significant when regressed without numeracy. Also note that academic records were not available for all subjects, so the second regression also has fewer subjects.

reservation values with any given search, which is partly due to the fact that they tended to set lower reservation values in the first place.

Result 1 Reservation values decline in both treatments, only partially supporting Hypothesis 1. However, the rate of decline is much steeper in the learning treatment. The informational content of the draws in the learning treatment is one of the main drivers of this additional decline, which is consistent with the prediction of increasing pessimism within search spells. Therefore the selection affect from the theoretical literature does appear in subject behavior.

6.2 The information demand heuristic

Hypothesis 2—that searchers will use a one-step reservation value strategy with well behaved priors—does not fare as well as the selection effect. I refer to the strategy of setting a high initial reservation value and adjusting only after a few observations in the learning treatment as the “information demand heuristic”. I define any search spell with an initial reservation value over 400 which is maintained at least through the second search as using this strategy. I define any subject who used this strategy in at least two of their payoff relevant learning treatment search spells as a high information demand subject or “information demander”; 25 out of the 81 subjects fit this description. While these subjects are setting high initial reservation values in response to uncertainty much as Bayesian searcher would, this heuristic is not a one-step strategy. The reason *why* the initial reservation values are higher differs drastically from the factors in the theory. The initial reservation value is high in the theory because a high initial draw means that the one-step Bayesian searcher is more optimistic about continued search, with subsequent draws leading to declining values. With this heuristic, the initial value is high because the searcher wants to search more than once before deciding whether or not to stop. More restrictive definitions for the heuristic or type (e.g. maintaining the same reservation value through the 3rd search or using the strategy for three of the learning search spells) yield similar qualitative results to those described here.

Table 4 shows the results of a probit predicting whether a subject is classified as a high information demand subject. “SNS Score” represents how subjects rated their comfort and ability with numeric information in the subjective numeracy survey, with a higher score indicating more comfort with numbers and a higher self-assessment of numeric ability. Interestingly, numeric competency is not correlated with a subject being an information demander, and subjective numeracy is *negatively* correlated with use of the heuristic, which is the opposite of Falco’s (2019) hypothesis that increased subjective numeracy will lead to more demand for numeric information. One potential explanation for this discrepancy is that subjects with greater subjective numeracy may have a greater tendency to use numeric information, and so

they feel a desire to adjust their reservation values based on their observations. This hypothesis is merely speculation without further research to test it. While numeric competency does not predict information demand, academic ability metrics are positively correlated with its use, which is not terribly surprising given that it seems to be a well performing strategy (see below).

Use of the heuristic within a given search spell does not predict a higher search payoff, but frequent use of it across search spells does. Table 5 shows the results of a regression of payoffs from the search task against numeracy and the categories described above.¹⁹ The learning treatment is positively correlated with payoffs in the first regression, but this correlation disappears with the introduction of ability proxies, suggesting that it is mainly driven by high ability subjects exploiting search spells with favorable distributions. Subjects with more numeric ability are better able to use information in the learning treatment, and also tend to set higher reservation values overall.²⁰ The heuristic is not significantly correlated with payoffs from the search task, but the information demander type is significant and positive. This difference likely comes from sub-optimality of rigidly sticking to the heuristic. The information demander types set higher reservation values even when they do not stick to the heuristic (an average of 446 as opposed to 390 for the non-demanders) and they were less likely to stick to the heuristic if their first draw was quite low. Being classified as an information demander type is only borderline significant when proxies for ability are introduced. Ability proxies, especially academic data, are more predictive of payoffs than frequent use of the heuristic.²¹ The heuristic is a reasonable strategy used by high ability subjects to deal with the uncertainty in the learning treatment. However, if the information in the first draw is sufficiently compelling then the information demanders recognize that it is optimal to deviate from the heuristic.

Result 2 Hypothesis 2 is not supported by the data. Many subjects do not use a one-step reservation value strategy in the environment with learning. They instead use a heuristic where they separate the search process into an “explore” stage and an “exploit” stage. While the use of the strategy itself is not significantly predictive of payoffs, those subjects who regularly use the heuristic (i.e. the information demanders) perform significantly better on the search task than subjects who did not.

¹⁹ I use ordinary OLS here rather than a panel estimator as a Hausman test rejects use of random effects, and a fixed effects estimator would cause the control variables (which are the variables of interest) to drop out. I therefore believe this specification to be the most informative.

²⁰ There are some endogeneity concerns with the regression in Table 5 as subjective numeric ability is determined at least partially by actual ability, so SMAP numeracy will partially determine the SNS Score. Three stage least squares analysis does suggest that this is valid, but the qualitative results of the 3SLS regressions are essentially identical to those presented here except that significant variables become more significant. I use the simple OLS here because it conveys the same information in an easier to read format.

²¹ The numeracy variables lose significance after the addition of academic controls, but, as I think a reviewer for pointing out, this is likely due to collinearity with the numeracy scores. A similar regression with academic controls and no numeracy scores is qualitatively similar.

Table 4 Probit predicting whether a subject would be classified as a high information demand type

Dependent variable: information demand type		
Covariates	Coefficients (SE)	
	(1)	(2)
SNS score	− 0.4679** (0.2361)	− 0.5453** (0.2652)
ACT score	—	0.1936*** (0.0728)
Other academic controls	No	Yes
Number of subjects	81	55
Number of observations	81	55

*, **, ***Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

Table 5 OLS regression predicting payoffs for the search task in the learning treatment

Dependent variable: search task payoffs			
Covariates	Coefficients (SE)		
	(1)	(2)	(3)
Information demand type	36.2192** (17.3214)	34.9088* (18.1015)	29.6689* (17.4039)
Information demand heuristic	17.5541 (23.4486)	17.1551 (25.8476)	12.0834 (21.0770)
Learning	23.6106** (11.5433)	116.8047* (68.8906)	49.6895 (71.4848)
SNS score	—	20.1837** (10.1061)	11.4456 (11.9451)
SMAP numeracy score	—	− 0.9528 (0.0704)	0.3558 (.3271)
Learning * SMAP numeracy score	—	− 0.4788** (0.1503)	0.7491 (0.6669)
ACT score	—	—	7.8939*** (2.2524)
Other academic controls	No	No	Yes
Number of subjects	81	81	61
Number of observations	477	477	367

*, **, ***Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

This heuristic is quite similar to descriptions of real world search behavior. Blake et al. (2016) find that search patterns on eBay follow a pattern quite similar to the heuristic described in this paper. Blake et al. describe a “search funnel”, which is a concept from the marketing literature that involves initial exploratory searches to get a sense of the products and prices available in the market, followed by more directed precise search strings once this information has been gathered. Previous research has assumed that the search funnel is a result of consumers’ imperfect knowledge of their own taste. According to this framework consumers leave the information gathering stage once they have decided what it is they want to buy. The similarity of the search pattern in my results absent any uncertainty about taste suggests that the search funnel may also be a way for consumers to deal with the complexity of learning prices in an unfamiliar market.

7 Behavioral hypotheses

Section 8 was a reaction to finding under-searching at levels similar to those in other experiments. One possible explanation for this under-searching is that, similar to overbidding in second price auctions, subjects are usually not punished in an obvious way for setting a reservation value below the optimal level. Even when they are forced to settle for a low value as a result of entering a low reservation value they still receive a positive payout. If on the other hand, a subject experiences a “near miss”, where they receive a draw that is close to, but still below their sub-optimal reservation value, then this may make them aware of the possibility that they could exit the search spell with a low draw and cause them to increase their reservation value.

Hypothesis 3 Subjects will be more likely to increase their reservation value if they see a draw which is below, but close to their reservation value.

Charness and Levin (2005) show that subjects tend to respond to a high payout by repeating the strategy that led to that payout even when doing so is not optimal. If subjects use a similar reinforcement heuristic in this environment, then we would expect a subject who sets a low reservation value and receives a draw far above that reservation value to imitate the strategy that received this result in the following search spell.

Hypothesis 4 Subjects who draw a value high above their reservation value in one search spell will set an initial reservation value in the following search spell which is similar to the initial reservation value in the search spell which received the high draw.

Results 3 and 4 summarize how Hypotheses 3 and 4 fared in the data respectively.

8 Feedback and under-search

In theory, risk averse preferences might be an explanation for the ubiquity of under-search. The reservation value is determined by a comparison of a certain cost to a stochastic benefit, so a more risk averse searcher should search less. In practice, the degree of risk aversion as measured by the Holt and Laury (2002) risk preference elicitation was not significant for any regression, including regressions predicting the length of search spell by subject. While this does not rule out risk aversion playing a role in subject behavior, it does seem unlikely to be the primary explanation for under-search.

One potential alternative explanation is the lack of negative feedback for suboptimal strategies. A subject setting a constant reservation value of 350 in the full information treatment will still stop with a positive payoff, and therefore may not realize that their reservation value is too low. An exception to this rule is a situation where a subject draws a value near, but not above their reservation value. This event will make the possibility of settling for a low reservation value more salient, and can cause subjects to revise their reservation values upward.

Table 6 describes a random effects probit analyzing the event that a subject's reservation value increases. Only 11% of the 1724 observed changes in reservation values were increases, nevertheless the mechanism hypothesized above appears to work. I create the "Distance" variable by subtracting the last observed value from the last reservation value for every search. It is significant and negative, indicating that the further the last observation is below its corresponding reservation value the less likely a subject is to raise their subsequent reservation value. This effect is stronger in the learning treatment because of the informational content in the observations. Receiving a draw close to the reservation value in the learning treatment not only increases the salience of settling for a low draw, but also indicates that the subject may be facing a favorable distribution. Subjects with higher ACT scores are somewhat less likely to raise their reservation values, but this is in large part a mechanical effect because higher ability subjects generally set higher reservation values in the first place.

Result 3 Hypothesis 3 is supported by the data. While other factors may explain the observed behavior, lack of negative feedback for sub-optimal strategies appears to be a plausible explanation for the prevalence of under-search and lack of inter-spell convergence toward the optimum value. When the possibility of settling for a low draw is made salient subjects are likely to raise their reservation values in response.

If we look at feedback *between* search spells, the picture is slightly different. If a subject consistently stops with values far above their reservation value, then this is feedback that their reservation values are too low. However, subjects may also view this as positive feedback because they are receiving a high payoff relative to their reservation value, which may act as reinforcement (in the psychological sense) and encourage them to maintain or even lower their reservation values. As Charness and Levin (2005) establish, people have difficulty updating when Bayesian rationality

Table 6 Random effects probit predicting the probability of reservation value increasing

Dependent variable: event that reservation value increases

Covariates	Coefficients (SE)		
	(1)	(2)	(3)
Distance	− 0.0036*** (0.0005)	− 0.0036*** (0.0005)	− 0.0030*** (0.0007)
Learning * distance	− 0.0043*** (0.0006)	− 0.0045*** (0.0006)	− 0.0034*** (0.0007)
SMAP numeracy score	−	0.0135** (0.0063)	0.0159 (0.0125)
Learning * SMAP numeracy score	−	− 0.0144* (0.0076)	0.0016 (0.0113)
ACT score	−	−	− 0.0969* (0.0577)
Numeracy measures	No	Yes	Yes
Other academic controls	No	No	Yes
Number of subjects	69	69	53
Number of observations	1724	1724	1257

"Distance" = Last reservation value – Last observed value

*, **, *** Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

and reinforcement heuristics are in opposition. They show that people tend to stick with strategies that yield positive payoffs even when the information conveyed by that payoff suggests that they should change their action.

This experiment is a particularly interesting environment in which to explore this possibility because we would expect any response in the full information treatment to be significantly dampened in the learning treatment. A subject who receives a draw significantly above their reservation value in the full information treatment could interpret it as positive reinforcement, but would likely just attribute it to a favorable distribution in the learning treatment.

The data support this reinforcement hypothesis, but in a very specific way. As Fig. 9 shows, the only significant correlation between the final observation/final reservation value gap and the initial reservation value in the following search spell is that subjects tend to lower their initial reservation value if their observations are too close to their reservation values or if they find themselves exercising recall and settling for a value below their stated reservation value. This is particularly true of high ability subjects.²²

²² High ability is defined as having an effective ACT Composite over 28, or an average error on the numeracy task less than 25 if this data was not available. These values are the integers closest to the median.

Fig. 9 The difference between the final observation and the corresponding reservation value by whether subjects lower, maintain, or raise the initial reservation value in the next search spell

Similar to the update variable, I create an “initial update” variable which takes a value of 1 if the initial reservation value is higher than the one in the previous period, 0 if it is the same, and -1 if it decreases. The “Difference” variable is the final observation minus the final reservation value in the previous search spell. As Table 7 shows, this difference between the final reservation value and the final observation is consistently the most important factor determining whether the initial reservation value changes relative to the previous search spell. Ability controls and the number of searches in the previous search spell (which controls for potential frustration with an unluckily long search spell) also play a role. Naive interpretation of the coefficient for the difference variable would suggest that subjects who saw a value high above their reservation are more likely to raise their initial reservation value in the following search spell while those who stop with a low final observation relative to their reservation value are likely to lower their initial reservation value in the following spell. However Fig. 9 suggests that this correlation is driven entirely by the latter effect, and even then only in the full information treatment. If I instead run a regression against a dummy that takes a value of 1 whenever the reservation value increases then difference is never significant.²³ Restricting the regression to the learning treatment similarly suppresses any significance. This behavior is similar to Charness and Levin (2005), but instead of maintaining a strategy after positive feedback, subjects are adjusting their strategy after negative feedback.

Result 4 Hypothesis 4 is partially supported by the data. Subjects appear to use a reinforcement heuristic similar to that proposed by Charness and Levin (2005) when choosing their initial reservation values. Subjects rarely become more persistent, but often adjust their reservation values to under-search further in response to a search spell with an unfavorable result. This adjustment does not occur in the learning treatment, suggesting that subjects do not transfer information across search spells as readily when the distribution is uncertain.

²³ I omit these relatively uninteresting regressions for the sake of space.

Table 7 Random effects ordered probit regression predicting the direction of change in initial reservation values

Dependent variable: change in initial reservation value			
Covariates	Coefficients (SE)		
	(1)	(2)	(3)
Difference	0.0013*** (0.0005)	0.0025*** (0.0007)	0.0026*** (0.0009)
Learning * difference	0.0000 (0.0004)	− 0.0008 (0.0008)	− 0.0008 (0.0008)
High ability * difference	− −	0.0010* (0.0005)	0.0015* (0.0008)
Learning * high ability * difference	− −	0.0003 (0.0004)	− 0.0001 (0.0005)
Previous initial reservation value	− 0.0014 (0.0009)	− 0.0015 (0.0009)	− 0.0017** (0.0007)
# Searches in previous spell	− 0.0242 (0.0148)	0.0246 (0.0154)	− 0.0335* (0.0173)
SMAP numeracy score	− −	− 0.0036*** (0.0009)	− 0.0057 (0.0038)
SNS score	− −	0.1145 (0.0714)	0.2542*** (0.0760)
ACT score	− −	− −	0.0149 (0.0160)
Learning * ACT score	− −	− −	0.0448** (0.0192)
Academic controls	No	No	Yes
Upper cutoff	− 0.1301	0.3126	1.4040
Lower cutoff	− 1.3396	− 0.901	0.1258
Number of subjects	81	81	61
Number of observations	805	805	610

*, **, *** Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

Based on the responses to the post-session questionnaires, it appears that subjects often equated low reservation values with caution. It is therefore likely that the main mechanism behind this result is a desire to increase caution in response to negative feedback. Given how important feedback and information design has been shown to be for optimal auction behavior (e.g. Kagel et al. 1987), the design of feedback in search environments seems like it may be fertile ground for further research.²⁴

²⁴ See Jhunjunwala (2018) for a more detailed exploration of the role feedback plays in search behavior.

Fig. 10 The percentage error in reservation values (compared to the theoretical optimum for a Bayesian learner) by depth in search spell and treatment

9 Distance from the theoretical optimum

In this section I compare subject's behavior to that of a risk neutral rational Bayesian agent using a one-step stopping rule. In the full information treatment this simply involves solving Eq. (2) using the given distribution. The process for the learning treatment is roughly equivalent, except that $F(\cdot)$ is a function of the priors and posterior beliefs as a result of observing x . The rational agent sets a reservation value by anticipating what the effects of observing x would be on their beliefs, then calculating the marginal benefit of an additional search given those posteriors. Anticipating the effect of an observation on future beliefs is not an easily intuited concept, so it is unlikely that most subjects are engaged in this level of forward thinking. This complexity may explain why subjects' reservation values are farther from the optimum in the learning treatment than in the full information treatment (as shown in Fig. 10).

The optimal reservation value in the full information treatment is 475, and the optimal initial reservation value in the learning treatment is 593. Later optimal reservation values in the learning treatment depend on the observations and therefore cannot be summarized simply. I calculate the percentage error using the formula

$$\text{Percentage Error} = \frac{\text{Reservation Value} - \text{Optimal Value}}{\text{Optimal Value}} \quad (5)$$

so a negative value indicates that subjects are setting a reservation value lower than the theoretical optimum. It is apparent from Fig. 10 that most subjects are setting reservation values below the optimum, and they are farther below the optimum in the learning treatment. The amount of error decreases as subjects continue to search, and this rate of decrease is faster with the learning treatment, such that the

percentage errors are roughly equivalent between treatments by the time subjects reach their 3rd search. Note that some of the convergence to the optimal reservation value occurs because subjects who input low reservation values drop out early. However, even with this fact we do see evidence for convergence toward the optimal value later in search spells for the learning treatment. This convergence has a number of potential causes: First, as a search spell continues, the optimal reservation value decreases. Subjects systematically set their reservation values too low, so as the optimal value decreases it will mechanically approach stated values. Second, the optimal reservation values decrease more quickly than actual reservation values, which could be caused by subjects using a non-Bayesian updating process, or just updating more slowly. Finally, subjects are closer to the optimum when they know the distribution, and as they continue to search within a search spell in the learning treatment they will eventually have a reasonably good estimate of the true distribution.²⁵ For later draws within a spell they are approximating the stationary search problem, which is much easier to solve.

Result 5 The absolute value of error decreases as subjects continue to draw values within a given search spell. This within-spell convergence between actual and optimum behavior is slightly stronger in the learning treatment. This suggests that subjects do use information about the distribution to inform their strategies, and that better information leads to a strategy closer to the theoretical optimum.

While there is strong intra-spell convergence, the convergence between spells is not nearly as strong, but we can still see some interesting patterns. Figure 11 shows progression of the mean absolute error in the first 3 searches of each spell.²⁶ Formally, I define the variable *M3Error* by

$$M3Error_k = \frac{1}{N_k * 3} \sum_{n=1}^N \sum_{i=1}^3 |\text{Reservation Value}_{nki} - \text{Optimal Value}_{nki}|$$

Where *k* is the *k*th search spell, *N_k* the number of subjects who searched at least 3 times within the *k*th search spell, and *i* the depth in the search spell. I use the mean of the absolute value because it provides a clearer interpretation of the graph (whether or not there is convergence to the optimum across search spells) than simply using the mean error would. Just using the initial reservation value for each search spell yields similar results.

We can see two trends in the error:

²⁵ The average absolute value of error for the initial search is 85 in the full information treatment and 185 in the learning treatment. A Mann–Whitney test for difference of behavior in the two treatments is significant at the 0.001% level. The single most important determinant (as determined by OLS regression) of how close subjects' elicited beliefs about the mean of the distribution were to the true mean was the number of searches.

²⁶ I use mean absolute error rather than mean absolute percentage error here to give the reader a sense of scale for the deviations from the theoretical optimum. The graphs with percentage error are nearly identical.

- 1 There is a downward trend in the error in the first few search spells of each session, regardless of the order.
 - While this may simply indicate that the experimental design should have allowed for more practice spells at the beginning, it is noteworthy that we see some evidence for gradual mastery of the environment in the full-information treatment. Subjects are told the distribution in the full information treatment, but they need practice with the problem in order to understand how to use this information.
- 2 Once this initial environmental learning curve is accounted for, there does not appear to be a strong trend in the error across periods.

Regressions looking at the driving factors behind error largely replicate Result 5. Additional results, such as a negative correlation between ability measures and error, are largely unsurprising, so I do not include these regressions here.

10 High cost sessions

I ran additional sessions with a search cost of 25 rather than 5 to test the robustness of these findings. The results are largely, though not entirely similar. The 72 subjects were again recruited from the Ohio State ORSEE system with 42 seeing the learning treatment last and 30 seeing it first.²⁷ The average payout was \$12.13 with a minimum of \$6.45 and a maximum of \$18.75. All tasks in the high cost sessions were identical to the main part of the experiment aside from the difference in search cost.

10.1 Results

It is immediately obvious from Table 8 that there is much less differentiation in search behavior between treatments in the high cost search environment. The average search length is actually *shorter* in the learning treatment as opposed to longer, though this average is not significantly different from the average search spell length in the high cost full information treatment.

Despite the similarity in search lengths between treatments, we can still see some evidence for the selection effect. The kernel density of the normalized reservation values in Fig. 12a still shows the same peak over 0 for the full information treatment and left skew for the learning treatment. This pattern is indicative of learning causing reservation values to decline, though the difference is much less striking than in the low cost data. In Fig. 12b we can see a weak decline as subjects get further into their search spells in the learning treatment. Repetition of

²⁷ The imbalance in learning-last versus learning-first here came from declining show-up rates by the time the learning-first sessions were run.

Fig. 11 Mean absolute error for the first three reservation values of each search spell

the regression in Table 2 yields no significant coefficients. Based on the graphical evidence above I believe that this does not demonstrate absence of a decline but instead comes from the increased search cost diminishing the effect size.

This belief is reinforced by examining the results of an ordered probit similar to Table 3 looking at the factors influencing the direction in which subjects update their reservation values in the high cost sessions. We can see in Table 9 that the previous observation is still one of the most influential factors in the learning treatment. The informational content of the observations in the learning treatment is still playing a significant role. There is a *negative* correlation between the previous observation and the direction of update in the full information treatment, which would suggest that subjects are lowering (raising) their reservation values after observing high (low) values. However the correlation is only borderline significant when I restrict the regression to search spells where subjects did not exercise recall, suggesting that the correlation is partly driven by frequent use of recall when the observed value is close to the stated reservation value. In other words, subjects appear to be engaging in satisficing behavior in the face of higher search costs (see Caplin et al. 2011 for a discussion of search and satisficing behavior).

While many subjects described strategies similar to the information demand heuristic, only 4 out of the 72 high cost subjects satisfied the definition of information demander compared to 25 out of 81 in the low cost sessions. While this is too few subjects for any statistical analysis to be meaningful, the difference in use of the heuristic is itself an interesting result. The information gathering stage of the heuristic involves exchanging search costs for additional information, and it appears that the additional expense in the high cost session made this information gathering stage too expensive for most subjects.

Result 6 The main results from the low-cost treatments are somewhat robust to a higher cost environment. However, subjects are more prone to satisficing behavior when the search cost is larger. Additionally, the high cost of gathering information reduced use of the information demand heuristic.

Table 8 Summary statistics for each treatment in the high cost sessions

	Full information	Learning
Average search length	2.200 (0.098)	2.079 (0.079)
Average search profits	415.461 (3.785)	388.737 (6.511)

Search profits are stated in experimental currency units. Standard errors are shown in parentheses

(a) The kernel density of the normalized reservation values in each treatment for the high cost sessions. The initial searches have been dropped from this graph

(b) The normalized reservation values by depth in search spell and treatment for the high cost sessions. Note that the data starts at the 2nd search since the first search is always 0 due to the normalization

Fig. 12 The kernel density and approximate paths of the normalized reservation values in the high cost sessions

11 Conclusion

Search with learning about the distribution has been under-explored in the experimental literature. For this experiment I restrict the setting to learning through observations from the distribution where the priors have overlapping support. This environment is quite common in the theoretical literature, and has some of the most consistent predictions. Namely, that searchers will use a one-step stopping rule a la Weitzman (1979), and that the reservation value for this stopping rule will monotonically decline because any observation high enough to cause an increase in the reservation value will also be above that reservation value. I show that this selection effect does indeed hold for experimental subjects in a simulated search environment. But I also have evidence that they may not be using a one-step stopping rule in the learning treatment.²⁸ Instead, many subjects seem to have a desire to learn about the

²⁸ Given that the elicitation method assumes a one-step strategy I had intended to run an additional treatment as suggested by a referee to see if behavior is broadly similar in a setting without it. Unfortunately, because of the 2020 Covid-19 outbreak, the Ohio State University laboratory suspended experiments during the spring of 2020 and so I was not able to run these additional sessions before leaving that institution.

Table 9 Random effects ordered probit regression predicting the direction of change in reservation values in the high cost sessions

Covariates	Dependent variable: change in reservation value	
	Coefficients (SE)	
	(1)	(2)
Previous observation	− 0.0015** (0.0006)	− 0.0016** (0.0006)
Learning * previous observation	0.0029*** (0.0007)	0.0033*** (0.0008)
SMAP numeracy score	0.0070 (0.0053)	0.0055 (0.0063)
SNS score	0.3274*** (0.1105)	0.4249*** (0.1271)
Learning * SNS score	− 0.5170*** (0.1882)	− 0.5710*** (0.2107)
# Of searches	0.0432 (0.0314)	0.0332 (0.0362)
Learning * # of searches	0.1210** (0.0473)	0.1376*** (0.0519)
Time	− 0.0007 (0.0027)	0.0000 (0.0026)
Learning * time	− 0.0033 (0.0045)	− 0.0058 (0.0055)
GPA	− −	− 0.2496* (0.1355)
Other academic controls	No	Yes
Upper cutoff	2.2706	2.5359
Lower cutoff	1.1296	1.5507
Number of subjects	63	47
Number of observations	841	675

*, **, ***Indicates significance at the 10, 5, and 1 percent levels respectively

Standard errors are shown in parentheses

distribution before stopping. The complexity and unintuitive nature of the optimal strategy may make this heuristic optimal in real world situations with limited computational power. Additionally, this heuristic may be closer to being unboundedly optimal in situations where it is possible that a searcher may encounter the same distribution again because additional information about the distribution would carry over to the next search spell. Subjects consistently under-search, and the persistence of this behavior appears to be at least in part due to the unintuitive nature of the feedback which the search game provides. Given Jhunjhunwala's (2018) result that under-searching can be counteracted by the mere prospect of additional feedback, this area seems like a fruitful avenue for future research.

Acknowledgements I would like to thank Jim Peck, Huanxing Yang, PJ Healy, John Kagel, Marc Bellemare, Mark Dean, Ignacio Esponda, David Falco, Tanushree Jhunjhunwala, Jason Kerwin, Kirby Neilsen, attendees at the Midwest Economic Theory Conference (Vanderbilt 2018) and seminar participants at the University of Minnesota Applied Economics Department for valuable feedback on this project. All mistakes are my own. This research was reviewed and approved by the Ohio State University Institutional Review Board. Funding for this project came from the Journal of Money Credit and Banking small research Grants fund and the Ohio State University Decision Sciences Collaborative.

References

- Bikhchandani, S., & Sharma, S. (1996). Optimal search with learning. *Journal of Economic Dynamics and Control*, 20(1–3), 333–359.
- Blake, T., Nosko, C., & Tadelis, S. (2016). Returns to consumer search: Evidence from ebay. In *Proceedings of the 2016 ACM conference on economics and computation*, pp. 531–545. ACM.
- Brown, M., Flinn, C. J., & Schotter, A. (2011). Real-time search in the laboratory and the market. *American Economic Review*, 101(2), 948–74.
- Burdett, K., & Vishwanath, T. (1988). Declining reservation wages and learning. *The Review of Economic Studies*, 55(4), 655–665.
- Caplin, A., Dean, M., & Martin, D. (2011). Search and satisficing. *American Economic Review*, 101(7), 2899–2922.
- Casner, B. (2020). Seller curation in platforms. Working Paper.
- Chandra, A., & Tappata, M. (2011). Consumer search and dynamic price dispersion: An application to gasoline markets. *The RAND Journal of Economics*, 42(4), 681–704.
- Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect. *American Economic Review*, 95(4), 1300–1309.
- Cox, J. C., & Oaxaca, R. L. (2000). Good news and bad news: Search from unknown wage offer distributions. *Experimental Economics*, 2(3), 197–225.
- Dana, J. (1994). Learning in an equilibrium search model. *International Economic Review*, 35(3), 745–71.
- De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2017). Search with learning for differentiated products: Evidence from e-commerce. *Journal of Business & Economic Statistics*, 35, 626–641.
- Dubra, J. (2004). Optimism and overconfidence in search. *Review of Economic Dynamics*, 7(1), 198–218.
- Falco, D. (2019). Economic search and numeracy. Working Paper.
- Falk, A., & Huffman, D. B., & Sunde, U. (2006). “Self-Confidence and Search,” IZA Discussion Papers 2525, Institute of Labor Economics (IZA).
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1), 114–125.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644–1655.
- Jhunjhunwala, T. (2018). Searching to avoid regret: Experimental evidence and an application to charitable giving. Working Paper.
- Kagel, J. H., Harstad, R. M., & Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica: Journal of the Econometric Society*, 55, 1275–1304.
- Kaya, A., & Kim, K. (2018). Trading dynamics with private buyer signals in the market for lemons. *The Review of Economic Studies*, 85, 2318–2352.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization*, 14(3), 381–392.
- Lewis, M. S. (2011). Asymmetric price adjustment and consumer search: An examination of the retail gasoline market. *Journal of Economics & Management Strategy*, 20(2), 409–449.
- Mauring, E. (2017). Learning from trades. *The Economic Journal*, 127(601), 827–872.
- Moorthy, S., Ratchford, B. T., & Talukdar, D. (1997). Consumer information search revisited: Theory and empirical analysis. *Journal of Consumer Research*, 23(4), 263–277.

- Nosko, C., & Tadelis, S. (2015). The limits of reputation in platform markets: An empirical analysis and field experiment. Working Paper.
- Peters, E., & Bjälkebring, P. (2015). Multiple numeric competencies: When a number is not just a number. *Journal of Personality and Social Psychology*, 108(5), 802.
- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5(2), 105–120.
- Schley, D. R., & Peters, E. (2014). Assessing “economic value”: Symbolic-number mappings predict risky and riskless valuations. *Psychological Science*, 25(3), 753–761. PMID: 24452604.
- Schotter, A., & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, 19(1), 1–25.
- Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, 35(3), 309–332.
- Varian, H. R. (1980). A model of sales. *American Economic Review*, 70(4), 651–59.
- Weisbuch, G., Kirman, A., & Herreiner, D. (2000). Market organisation and trading relationships. *Economic Journal*, 110(463), 411–36.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 47, 641–654.
- Yang, H. (2013). Targeted search and the long tail effect. *RAND Journal of Economics*, 44(4), 733–756.
- Yang, H., & Ye, L. (2008). Search with learning: Understanding asymmetric price adjustments. *The RAND Journal of Economics*, 39(2), 547–564.
- Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, 22(4), 503–519.

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