Sequential Search and Satisficing

Mark Dean

Lecture Notes for Spring 2015 Behavioral Economics - Brown University

1 Satisficing

One of the earliest attempts to formalize a model of bounded rationality came in 1955, when Herbert Simon introduced the concept of 'satisficing'. The name is designed to distinguish it from the process of 'maximizing': rather than trying to find the maximal available alternative, a satisficer tries only to find a satisfactory alternative - one that is 'good enough'.

The broad idea of satsificing is that the decision maker, when faced with a set of options, will search through them one by one until they come across one that matches a set of predetermined criteria. They then stop searching and choose that alternative. In fact, you have already come across a variant of the satisficing model in homework 1.

Simon introduced the concept of satisficing as a psychologically plausible method of choice. His aim was not to justify it as an optimal procedure. However, it turns out that satisficing behavior *is* optimal in a simple model of sequential search with costs. We will begin by proving this result. We will then talk about how to test the satisficing model by identifying its observable implications. In fact, it will turn out that if the only data that you have is the choices that people make, testing the satisficing model is difficult (you should get a hint of this from homework 1). It is therefore worth considering richer data sets. We will look at two. The first is an experimental data set which records how subject's choices change as they think about what to choose. Broadly speaking, this data is consistent with the satisficing model. The second data comes from the search behavior of customers on the web who are looking to buy books. This data allows us to understand what options consumers examined but rejected, as well as those that were eventually chosen. This data is better described by a model in which people choose a priori how many options to search through, rather than satisficing.

2 Satisficing as Optimal Choice

Think about a decision maker who is faced with a set of options from which they can choose one: for example, they want to buy a bar of soap from the supermarket, and they need to choose which one to buy. The decision maker knows how many bars of soap are for sale, but before looking at a particular bar of soap, the decision maker knows nothing about it - just that it is a bar of soap that they could purchase. Once they have looked at the bar of soap, they know everything they need to know about it in order to assess its utility. However, finding out this information is costly (perhaps in terms of time, or mental effort).[1]

We can think of this as a simple optimal stopping problem. At any given point, the decision maker is going to know the value of the best option they have already seen, the number of alternatives left to search, and the cost of searching one more alternative. We want to solve the problem of when it is optimal for the decision maker to keep on searching, and when it is optimal for them to stop.

Let's set up the problem formally. The basic elements are the following

- A set A containing M items
- A utility function u: $X \rightarrow \mathbb{R}$ that represents the value of each option
- A probability distribution f that represents the decision maker's beliefs about the value of each option

[1] This set up is simplistic in a number of ways. On the one hand, it might be more realistic to think that the DM gets some information about the available soap 'for free' - for example, they have bought one particular brand of soap before, and so know all about it. On the other, it could be that the DM could find out some partial information about one of the options before moving on to another (for example, they could look at the prices of all the available soaps, then look at whether or not they are anti bacterial, and so on). However, one could think of examples where this does not seem like a good description of the way people search for information (for example, if there is a large cost of switching attention from one item to another.) It also turns out that this set-up is going to make things relatively simple for us to analyze.

- A cost k that the decision maker has to pay in order to understand the value of the next available alternative.

At any given point, the decision maker can take one of two options

1. Stop searching, and choose the best available alternative that they have looked at (we allow recall, so the DM can choose any of the objects that they have already seen)
2. Search another item and pay the cost k

If the DM does choose to search another item, then they will face the same choice next period (stop, or search one more item) until they have searched all the available alternatives.

We can solve this problem by backwards induction. Let's imagine that the decision maker has searched all but one item, and the best item they have found has utility $\bar{u}$. They have two choices:

1. Stop searching, and receive utility $\bar{u} - (M-1)k$
2. Search the final item. In this case, the final item will either have utility u less than $\bar{u}$, in which case they will just receive $\bar{u} - Mk$, or it will have utility u greater than $\bar{u}$, in which case they will receive utility $u - Mk$. Thus, the expected utility of searching one more item is

$$\int_{-\infty}^{\bar{u}} \bar{u} f(u) du + \int_{\bar{u}}^{\infty} u f(u) du - Mk$$

So, when will the decision maker choose to search the next item? They will do so when the expected utility of doing so is higher than the expected utility of stopping. In other words, when

$$\begin{aligned} \bar{u} - (M-1)k \leq & \\ & \int_{-\infty}^{\bar{u}} \bar{u} f(u) du + \int_{\bar{u}}^{\infty} u f(u) du - Mk \end{aligned}$$

Which, if we rearrange, gives us the following:

$$k \leq \int_{\bar{u}}^{\infty} (u - \bar{u}) f(u) du$$

On the left hand side is the cost of searching one more item. On the right is the expected benefit of doing so.

The precise solution to this problem will depend on the nature of the function f. However, we know what form the solution will take. Notice that the left hand side of this equation is fixed, while the right hand side is decreasing in $\bar{u}$ (this makes sense - the better the value of the alternative that you have already seen, the less value there is to carrying on searching. Thus, we can deduce that there is some value u^* such that it is optimal to search the final object if $\bar{u} \leq u^*$, but not if $\bar{u} > u^*$. Moreover, we know that u^* must be the solution to

$$k = \int_{u^*}^{\infty} (u - u^*) f(u) du \tag{1}$$

This is what we call a *reservation stopping rule.* There is a reservation level of utility u^* such that search stops if and only if the DM has found an object which has value higher than u^*.

Okay, so that deals with the case where there is only one more item to search. What about in other cases? Somewhat surprisingly, the optimal strategy is exactly the same: stop if you have seen an object with value above u^*, and carry on otherwise. To see this, think about the case where there are two more objects that could be searched, let $\bar{u}$ be the best valued object that the decision maker has seen, and consider two cases

1. $\bar{u} < u^*$ (where u^* is the reservation level for the last period). In this case, it is clearly optimal for the DM to carry on searching. Remember that for values $\bar{u} < u^*$, it is worthwhile for the consumer to search one more alternative. Here they can search at least one more alternative. Thus, it cannot be optimal for them to stop searching

2. $\bar{u} \geq u^*$ Here it is not optimal for the DM to carry on searching. Remember, we said that, for $\bar{u} > u^*$, it is not worthwhile for the DM to search one more alternative. However, we know that, despite the fact there are potentially two more alternatives to search, a DM who has seen an object that has value above u^* will never search the final object. Why? Because we have solved for the optimal strategy of a decision maker that has only one more item to search, and it is to stop if the best item they have seen is above u^*. Thus, at most, the DM will search one more item, which we know is not worthwhile if $\bar{u}$ is bigger than u^*

Iterating on this argument, we can show that the DM has a fixed reservation strategy, which is, for any period, to stop searching if $\bar{u} > u^*$ and carry on otherwise.

We have therefore obtained satisficing behavior as the optimal solution to a particular problem: we can think of objects with utility above u^* as 'good enough', while those below as not good enough. It should be noted that this fixed reservation strategy is a result of some of the specific assumptions that we have made - in particular that the decision maker is not learning about the distribution f.

One advantage of thinking of satisficing as the result of optimizing behavior is that we can make some predictions about how the reservation level u^* will change with the environment. Looking at equation 1, it should be obvious that the reservation level

1. falls as the cost k rises
2. rises with the variance of f, for a fixed mean
3. changes one for one with the mean of f for a fixed variance
4. does not change with the size of the choice set.

3 Testing the Satisficing Model

3.1 Standard Choice Data

Unfortunately, one downside of the satisficing model is that it is hard to test using standard choice data. The model of choice that we have in mind is as follows:

1. The decision maker searches through each choice set according to some order
2. If they come across an object x that has utility level $u(x) > u^*$ they stop searching and choose that object
3. Otherwise they continue to search the whole set and choose the best available option.

This procedure should sound familiar - you thought about it in homework 1. At that stage you showed that, if the decision maker always searches through things in the same order, then, in fact, it is observationally equivalent to utility maximization: a decision maker can be modeled as using the decision process above if and only if they satisfy properties α and β. On the other hand, if we assume that the the order in which people search changes from choice to choice, then *any* choice behavior can be explained by this model. How? Just assume that all objects have above-reservation utility, and the object that is chosen in every case is just the first object that has been searched. Thus, it is very difficult for us to spot a satisficer.

3.2 Choice Process Data

So how can we test whether or not people are satisficing? One way would be to make further refinements to the satisficing model to the point that it does make predictions that we can test. This is certainly one way to go, but it does mean that any test we do is not a 'pure' test of the satisficing model - it would be a joint test of the satisficing model and these additional assumptions.

Another way would be to expand the data set that we collect so that we can get some idea of how people are actually searching for information. On simple way to do this would be to record not only the choices that people make, but also how their choices change over time. Let's imagine that we were watching someone who was trying to choose a stereo. Rather than observing the final choice of stereo they make, we could ask them to tell us, during the course of their search, which stereo they would choose if they had to choose *right now*. So, rather than observing $C(A)$, the choice they make from a set of alternatives A, we observe $C(A, t)$, the choice that they would make from the set of alternatives having thought about the problem for length of time t. Could we use this data (which we call choice process data) to test the satisficing model?

The answer is yes. On this extended data set, our model of behavior is now as follows: the decision maker searches through objects one at a time. At any given time they will choose the best object of the ones that they have seen. They stop searching if they find an object that is above the reservation level of utility, otherwise the search the full set. To make things easier, we are going to assume that we can observe the utilities of the different alternatives. (In fact, we can do the same analysis without observing this, we just need more data). In this case, our data is going to come in the form of sequences of numbers, which is the utility of the object that they choose in each period, so for example

Observation	Available options	Sequence of Choices	Final Choice
1	$\{1, 2, 3, 4\}$	$\{3, 1, 4\}$	4
2	$\{2, 4, 6, 10\}$	$\{2, 4, 6\}$	6
3	$\{2, 4, 6, 8, 10\}$	$\{2, 4, 8\}$	8
4	$\{2, 4, 6, 8, 10\}$	$\{2, 6, 8, 10\}$	10

where $\{3, 1, 4\}$ means that the subject chose an object with utility 3, then an object with a utility 1, then one with utility 4, then stopped (and so finished with the object with utility 4 - remember the subject is reporting what they would choose if forced to stop at each point in time.

So which of these sequences is consistent with satisficing? Remember, the satisficing model has two elements (i) subjects search through options one by one, and at any given time report the best one that they have seen (ii) they stop searching if and only if they come across an object that is above the reservation level.

Condition (i) rules out data 1: if the subject searches through objects one by one and always chooses the best of things that they have seen, then they cannot possibly switch from choosing an item with utility 3 to one with utility 1. If they chose the utility three object, then they must know that it is available, so there is no way that they would switch to the utility one object. This gives us our first condition

Condition A Subjects can only switch to higher utility objects

What about the remaining sets? 2 and 3 together are fine: they could both be explained by a satisficer that has a reservation level of (say) 5. Though note that these choice could not be explained by standard utility maximization. However, it could not be the case that 2, 3 and 4 could jointly be explained as the result of satisficing. The fact that, in choice sets 2 and 3, the option with utility 10 is not chosen, but was available, indicates that both 6 and 8 must be above the reservation level of utility, as search must have stopped when these options were found. However, in 4, we see search continue after an object of utility 6 was found, meaning that the reservation level must be *above* 6. This gives us our second condition

Condition B We must be able to find a u^* such that search stops for objects such that $u(x)$ is above u^* and carries on when $u(x)$ is below u^*

Note that, as per our discussion above, it might be the case that u^* may vary depending on the environment.

It turns out that this model can explain choice quite well, even in circumstances in which standard utility maximization does badly, as you will see in the class presentation.

4 Suggested Reading

Andrew Caplin & Mark Dean & Daniel Martin, 2011. "Search and Satisficing," American Economic Review, American Economic Association, vol. 101(7), pages 2899-2922, December.

Caplin, Andrew & Dean, Mark, 2011. "Search, choice, and revealed preference," Theoretical Economics, Econometric Society, vol. 6(1), January.

Simon, H (1955) "A behavioral model of rational choice" Quarterly Journal of Economics, Vol. 69, No. 1 (Feb., 1955) , pp. 99-118