A simple satisficing model

Erlend Dancke Sandorf[a], Danny Campbell[a] and Caspar Chorus[b]

[a]Economics Division, Stirling Management School, University of Stirling
[b]Transport and Logistics Groups, Faculty of Technology, Policy and Management, Technical University of Delft

International Choice Modelling Conference, 19–21 August 2019, Kobe

Background, motivation and contribution

Satisficing model

Threshold utility

Secondary decision rule

Properties of the model

Synthetic data generating process and results

Data generating process

Results

Conclusion

Limitations

Take home

Background, motivation and contribution

In many (most) choice situations alternatives are evaluated ***sequentially***.

In many (most) choice situations alternatives are evaluated ***sequentially***.

Background

In many (most) choice situations alternatives are evaluated ***sequentially***.

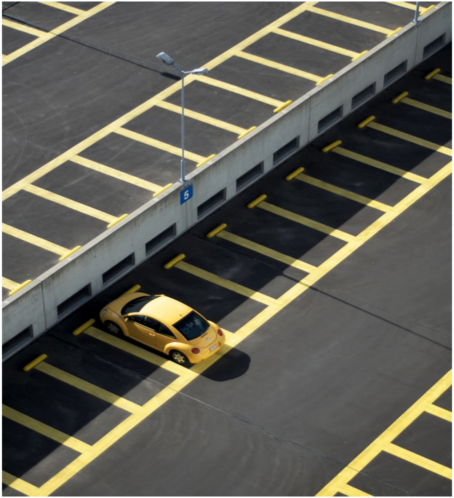

Background

In many (most) choice situations alternatives are evaluated ***sequentially***.

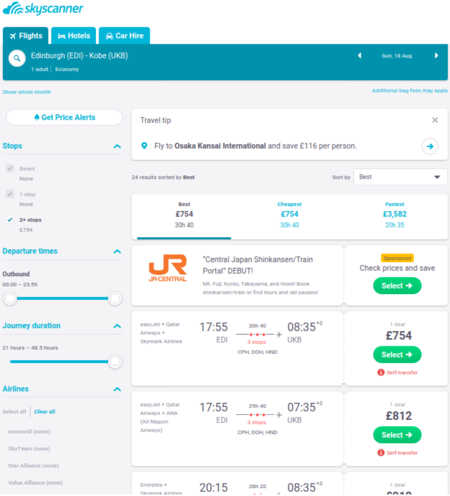

In many (most) choice situations alternatives are evaluated ***sequentially***.

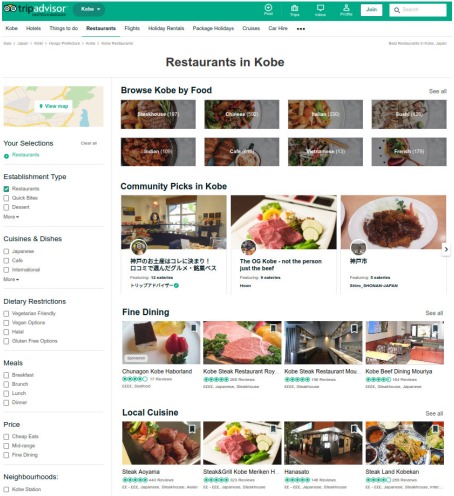

Satisficing

The sequential evaluation of alternatives means that the decision process may be one in which the first alternative exceeding some threshold utility is chosen.

- In this case, the decision maker does not continue to evaluate all available alternatives.
- Consequently, the choice may not be utility maximizing.
  - If the first satisfactory alternative encountered happens to be the one that gives the highest global utility, then that choice is also utility maximizing.
  - Any other choice is, by definition, satisfactory, but not maximizing.

Satisficing

The sequential evaluation of alternatives means that the decision process may be one in which the first alternative exceeding some threshold utility is chosen.

- In this case, the decision maker does not continue to evaluate all available alternatives.
- Consequently, the choice may not be utility maximizing.
  - If the first satisfactory alternative encountered happens to be the one that gives the highest global utility, then that choice is also utility maximizing.
  - Any other choice is, by definition, satisfactory, but not maximizing.

This type of behaviour is referred to as ***satisficing*** (i.e., choosing the first alternative that is satisfactory).

Only a few papers have developed models that can identify satisficing behaviour in more traditional discrete choice data.

- González-Valdés and Ortúzar (2018)[1].
- Sandorf and Campbell (2019)[2].

[1] Journal of Choice Modelling 27: 74–87.

[2] European Review of Agricultural Economics, 46(1): 133–162.

Satisficing based on attribute acceptability

Both papers explored satisficing based on the acceptability of attribute levels.

- But accommodating this behaviour based on attribute acceptability may have shortcomings.

Suppose that the first product evaluated by an individual has one attribute that is just below their acceptable level and that all the other attributes far exceed their acceptable levels to extent that the alternative itself exceeds their acceptable level.

- Based on acceptability of attribute levels, the individual would be predicted to not choose this product.
- Whereas, in reality, since the overall utility surpasses their threshold, the individual would be predicted to choose this product.

Motivation

A | C | R | G
Applied Choice Research Group

We wanted to find a better way to accommodate this type of behaviour.

- That focuses on utility rather than acceptability of attribute levels.

Contribution

We develop a satisficing model that involves choosing the first alternative with utility exceeding some threshold level of utility.

- An important feature of the model is that the reservation utility is estimated alongside the marginal utility parameters.
- Crucially, the model explicitly accounts for situations where none of the available alternatives exceed the threshold and another decision rule is then employed.
- We show that the model retrieves the true parameters under various assumptions about the level of the threshold utility and under a range of behavioural rules.

Satisficing model

We assume that a decision maker faces a choice between J different alternatives provided in the complete and exhaustive choice set C.

Decision makers are indexed by $n \in \{1, \ldots, N\}$ and alternatives by $j \in \{1, \ldots, J\}$.

The utility, u, decision maker n receives from choosing the j^{th} alternative is given by:

$$u_{nj} = v_{nj} + \varepsilon_{nj} = \boldsymbol{\beta}\mathbf{x}_{nj} + \varepsilon_{nj},$$

where $\boldsymbol{\beta}$ is a row vector of parameters, $\mathbf{x}_{nj}$ is a column vector of attribute levels and ε_{nj} is an iid error term from a type I extreme distribution with variance $\pi^2/6$.

Threshold utility

When people make choices, they do not always choose the utility maximizing alternative.

One possibility is that they choose the first one exceeding some minimum level of ***acceptable*** utility.

Let us define the minimum level of utility, or threshold utility, as t.

Defining the threshold utility

Just as we cannot observe an individual's utility function, we cannot fully observe their threshold utility.

We are reduced to making probabilistic statements about whether or not utility of the alternative exceeds the threshold.

Let us define the threshold as being comprised of an observable component τ to be estimated and an unobservable component ϵ, such that:

$$t = \tau + \epsilon,$$

where ϵ_{nj} is an iid error term from a type I extreme distribution with variance $\pi^2/6$.

Under the assumption that the differences in the unobserved parts are logistically distributed, the probability that alternative j yields utility greater than this threshold is of the logit form:

$$
\begin{aligned}
\Pr\left(u_{nj} > t \mid \mathbf{x}_{nj}, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) &= \Pr\left(v_{nj} + \varepsilon_{nj} > \hat{\tau} + \epsilon\right) \\
&= \Pr\left(\varepsilon_{nj} - \epsilon > \hat{\tau} - v_{nj}\right) \\
&= \frac{1}{1 + \exp\left(\hat{\tau} - \hat{\boldsymbol{\beta}}\mathbf{x}_{nj}\right)}.
\end{aligned}
$$

Given the sequential manner in which individuals consider alternatives, the choice probability of an alternative being chosen in a satisficing model (S) must account for the probability that all subsequent alternatives were not chosen:

$$\Pr\left(j_n \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}, \mathrm{S}\right) = \begin{cases} \Pr\left(u_{nj} > t \mid \mathbf{x}_{nj}, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) & \text{if } j = 1\text{; or,} \\ \Pr\left(u_{nj} > t \mid \mathbf{x}_{nj}, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) \prod_{j \in \{1,\ldots,j-1\}} \left(1 - \Pr\left(u_{nj} > t \mid \mathbf{x}_{nj}, \hat{\boldsymbol{\beta}}, \hat{\tau}\right)\right) & \text{if } j > 1. \end{cases}$$

Probability that no alternative exceeds the threshold

The probability that none of the alternatives in the choice set yield utility that exceeds the threshold utility is simply one minus the sum of the choice probability of an alternative being chosen in a satisficing model over all alternatives:

$$\Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) = 1 - \sum_{j \in \{1,\ldots,J\}} \Pr\left(j_n \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}, \mathsf{S}\right),$$

where $0 < \Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) < 1$.

Given the strict inequality $\Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) > 0$, there remains a probability that the choice task contains no satisfactory alternative.

- $\Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right)$ can be interpreted as the probability of individual n switching to a secondary decision rule after they have evaluated all J alternatives in choice set C and established that none of them meet their acceptable threshold utility.

After evaluating all possible alternatives, individuals must switch to another, ***secondary***, decision making strategy.

The overall choice probability then becomes the satisficing probability plus the choice probabilities derived conditional on the secondary decision rule weighted by the probability that this rule is enacted:

$$\begin{aligned}\Pr\left(j_n \mid \mathbf{X}_n, \hat{\beta}, \hat{\tau}, 1^{\text{st}}\text{:S},\ 2^{\text{nd}}\text{:}\cdot\right) = &\Pr\left(j_n \mid \mathbf{X}_n, \hat{\beta}, \hat{\tau}, \text{S}\right)\\ &+ \Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\beta}, \hat{\tau}\right)\Pr\left(j_n \mid \cdot\right),\end{aligned}$$

where 1^{st}:S and 2^{nd}:$\cdot$ signify the primary and secondary decision making rules, respectively, and $\Pr\left(j_n \mid \cdot\right)$ is the probability of choice conditional on the secondary decision making strategy.

Considered secondary decision rules

The secondary decision rule may entail a combination of decision making strategies and possible heuristics.

Here we consider four strategies.

The secondary decision rule may entail a combination of decision making strategies and possible heuristics.

Here we consider four strategies.

1. Choose the last alternative:

$$\Pr\left(j_n \mid \text{Last}\right) = \begin{cases} 1 & \text{if } j = J\text{; and,} \\ 0 & \text{otherwise.} \end{cases}$$

Considered secondary decision rules

The secondary decision rule may entail a combination of decision making strategies and possible heuristics.

Here we consider four strategies.

2. Choose a random alternative.

$$\Pr\left(j_n \mid \text{Random}\right) = \frac{1}{J}.$$

The secondary decision rule may entail a combination of decision making strategies and possible heuristics.

Here we consider four strategies.

3 Choose the utility maximizing alternative:

$$\Pr\left(j_n \mid \mathbf{X}_n, \hat{\beta}, \mathrm{RUM}\right) = \frac{\exp\left(\hat{\boldsymbol{\beta}}\mathbf{x}_{nj}\right)}{\sum_{j \in \{1,\ldots,J\}} \exp\left(\hat{\boldsymbol{\beta}}\mathbf{x}_{nj}\right)}.$$

The secondary decision rule may entail a combination of decision making strategies and possible heuristics.

Here we consider four strategies.

4. Choose to opt-out or the explicitly offered status-quo alternative.

$$\Pr\left(j_n \mid \text{Opt-out}\right) = \begin{cases} 1 & \text{if } j = \text{opt-out or status-quo; and,} \\ 0 & \text{otherwise.} \end{cases}$$

Note: For the opt-out/status-quo alternative, we assume that it is the first considered alternative, since they must first consider their current offering and decide if they are in the market.

As the threshold, τ, goes to $-\infty$ every single alternative will have a utility higher than the threshold:

$$\lim_{\tau \to -\infty} \Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) = 0.$$

Choosing the first alternative that exceeds the threshold involves choosing the first encountered alternative.

- If search costs are considered, this is analogous to the ***no deliberation strategy*** outlined by Manski (2017)[3].
- Every choice is identified as a satisficing choice.
- The choice probability approaches one meaning the log-likelihood will tend to zero.

[3] Theory and Decision, 83(2): 155–173.

Properties of the model: $\tau \to \infty$

As τ goes to $+\infty$, none of the alternatives will give a utility that is higher than the threshold:

$$\lim_{\tau \to +\infty} \Pr\left(\mathbf{u}_n < t \mid \mathbf{X}_n, \hat{\boldsymbol{\beta}}, \hat{\tau}\right) = 1.$$

The model will, therefore, collapse to the model associated with the secondary decision rule.

- In this case, the model has the same fit and retrieves the same parameters, but is less parsimonious.

The satisficing choice probability and, thus, the joint choice probability is affected by the order in which alternatives are evaluated.

- Therefore, the evaluation order must be known (or assumed).

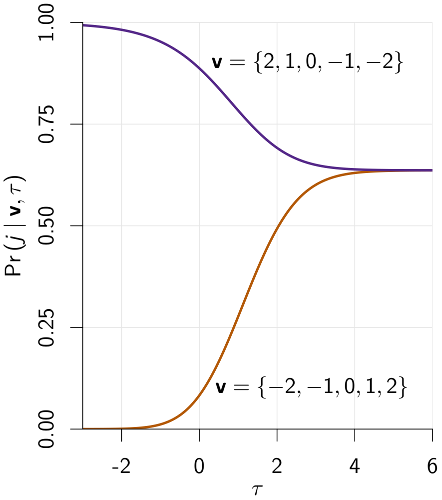

{Note: Shows the probability of choosing alternative $v = 2$ assuming the secondary decision rule is utility maximization.}

Properties of the model: number of alternatives

{Note: $\mathbf{v} = \{-2, \ldots, 2\}$ with J equidistant intervals.}

The probability of switching to a secondary decision rule depends on the number of alternatives in the choice set.

- As one would expect, as the number of alternatives increases the probability that the secondary decision rule is needed reduces.
  - In other words, with more alternatives the likelihood of satisficing increases.

Synthetic data generating process and results

Data generating process: Monte Carlo settings

To test the performance of our model and how well it retrieves the true parameters under varying experimental conditions we run a series of Monte-Carlo simulations.

Our Monte-Carlo strategy involves a variety of generation processes.

Data generating process: Monte Carlo settings

To test the performance of our model and how well it retrieves the true parameters under varying experimental conditions we run a series of Monte-Carlo simulations.

Our Monte-Carlo strategy involves a variety of generation processes.

We generate data for $J \in \{2, 3, 4, 5, 6, 8, 10, 25, 50\}$.

To test the performance of our model and how well it retrieves the true parameters under varying experimental conditions we run a series of Monte-Carlo simulations.

Our Monte-Carlo strategy involves a variety of generation processes.

Threshold utilities are derived by generating the full factorial design and for each profile we simulate a distribution of possible utilities.

- The minimum, maximum and the intermediate ventile utility values are used to represent τ.
- This ensures we have a wide spread of threshold utilities.

Data generating process: Monte Carlo settings

To test the performance of our model and how well it retrieves the true parameters under varying experimental conditions we run a series of Monte-Carlo simulations.

Our Monte-Carlo strategy involves a variety of generation processes.

This leads to 189 different simulation treatments.

- 9 settings relating to the number of alternatives times 21 settings relating to τ.

To test the performance of our model and how well it retrieves the true parameters under varying experimental conditions we run a series of Monte-Carlo simulations.

Our Monte-Carlo strategy involves a variety of generation processes.

A random experimental design is generated randomly for every replication.

- We use 1,000 replications for each of the 189 treatments.

Each treatment consists of 1,000 individuals answering a single choice task.

Each alternative is described by four generic attributes:

- AttA and AttB, which have binary (0, 1) levels; AttC, which takes levels between 0 and 1 in 0.01 increments; and, Cost, which has levels between €5 and €30 in €0.50 increments.
- We assume that the true parameters were: 0.5 for AttA, 0.8 for AttB, -1.6 for AttC, and -0.1 for Cost, and that the alternative specific constants are all zero.

Based on these experimental settings, the data was generate based on $-7.74 \leq \tau \leq 23.83$ with intermediate ventile utility values for τ.

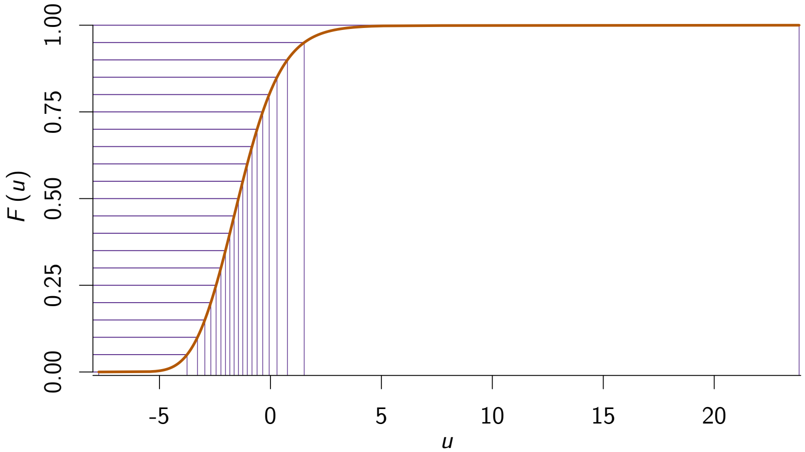

Candidate models

For every dataset generated, we estimate two candidate models:

1. The naïve specification based solely on the respective secondary decision rule; and,
2. The specification where satisficing is used as the primary decision rule and the respective strategy as the secondary decision rule.

Estimating both candidate models allows us to compare the effects under correctly specified and misspecified cases and to make inferences regarding the consequences of the naïve assumption.

- We estimate alternative-specific constants for the opt-out and J^{th} alternatives.

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

Consistent with ***satisficing decision rule***.

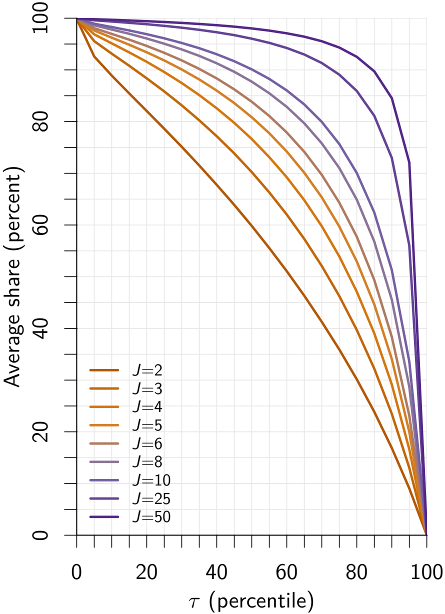

Results: Observed choice shares by decision rule

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

Consistent with ***choosing the last alternative decision rule***.

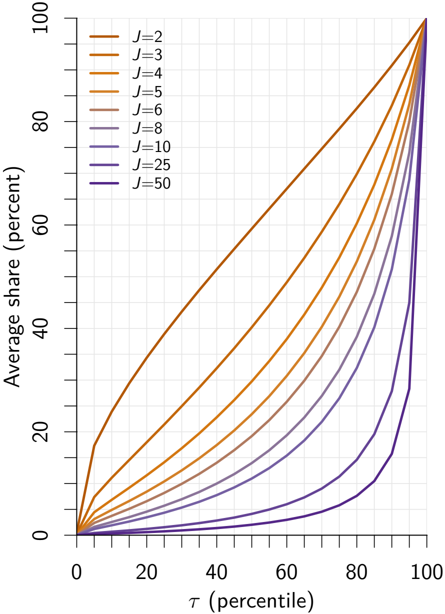

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

Consistent with ***choosing a random alternative decision rule***.

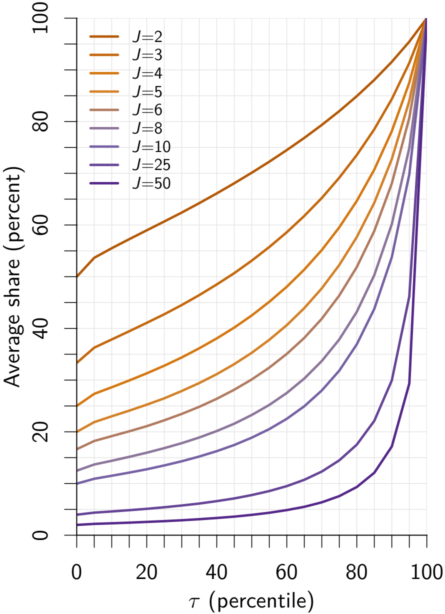

Results: Observed choice shares by decision rule

ACRG Applied Choice Research Group

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

Consistent with ***choosing the utility maximisation alternative decision rule***.

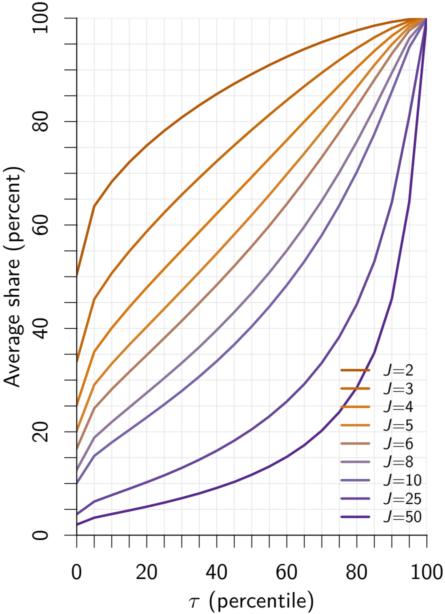

We begin with a comparison of the average share (averaged over the 1,000 sample simulations) of simulated choices that are observed to be consistent with each decision rule.

Consistent with ***choosing the opt-out or the explicitly offered status-quo alternative decison rule***.

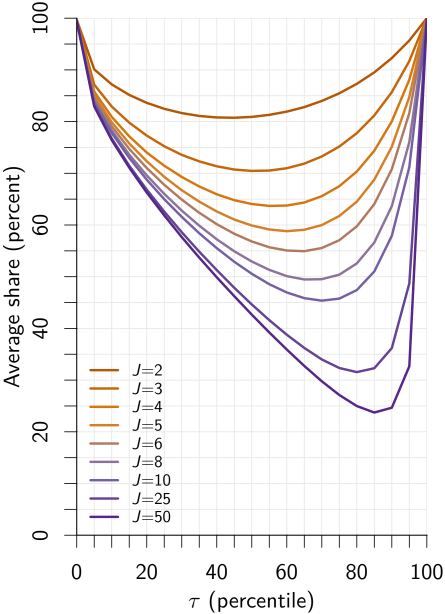

Results: Observed choice shares by decision rule

A | C | R | G
Applied Choice Research Group

Key findings

- When the threshold utility is low, a high share of choices are consistent with satisficing.
    - This declines as the threshold increases, but to a lesser extent as the number of alternatives increase.
- The share of choices that are consistent with all but one of the secondary decision rules increases with the threshold.
    - This stems from the fact that increases in the threshold increases the need to switch to the secondary decision rule.
    - A different pattern is observed for choosing the opt-ou/status-quo alternative decison rule, since this is effectively the first considered alternative.

We next compare the difference in log-likelihoods (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

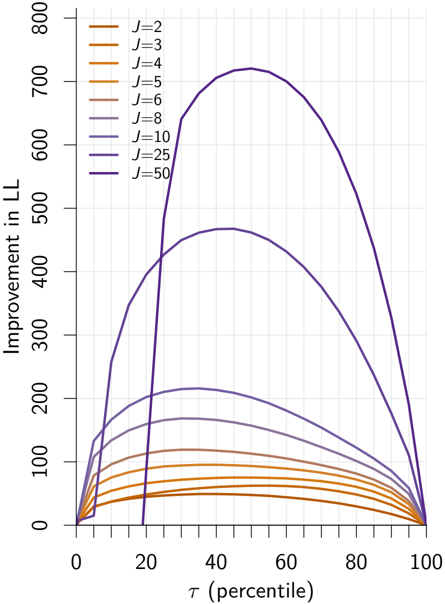

We next compare the difference in log-likelihoods (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Last alternative secondary decison rule.

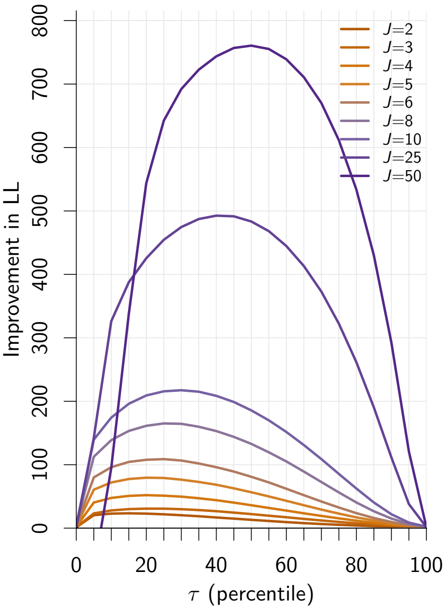

We next compare the difference in log-likelihoods (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Random alternative secondary decison rule.

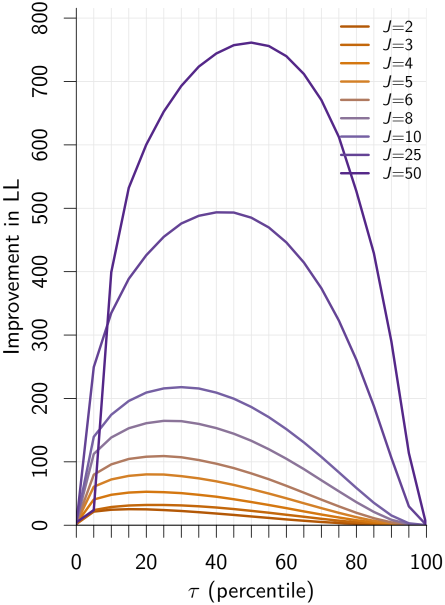

We next compare the difference in log-likelihoods (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Utility maximizing alternative secondary decison rule.

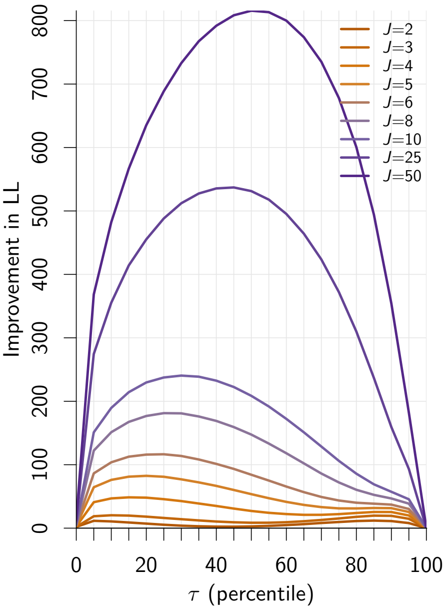

We next compare the difference in log-likelihoods (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Opt-out or choose the explicitly offered status-quo alternative secondary decison rule.

Key findings

- With extreme (lower and upper) threshold utilities the naïve specification and the satisficing model both produce equivalent model fits.
  - Recall that the satisficing model collapses to the model associated with the secondary decision rule as the threshold goes to the upper extreme.
  - The inclusion of alternative-specific constants is what ensures the model fits are equivalent at the lower extreme.
    - Note though, that the alternative-specific constants will be biased.
- It is interesting to note that the largest gain in fit is observed with increasing thresholds as the number of alternatives grow.

Results: Differences in correctly predicted

We next compare the difference in correctly predicted (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

We next compare the difference in correctly predicted (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Last alternative secondary decision rule.

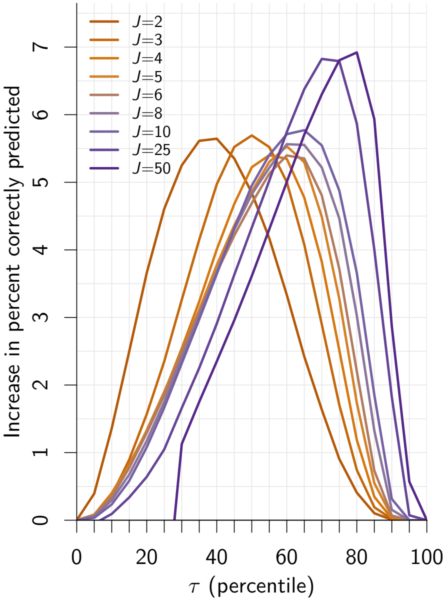

Results: Differences in correctly predicted

We next compare the difference in correctly predicted (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Random alternative secondary decison rule.

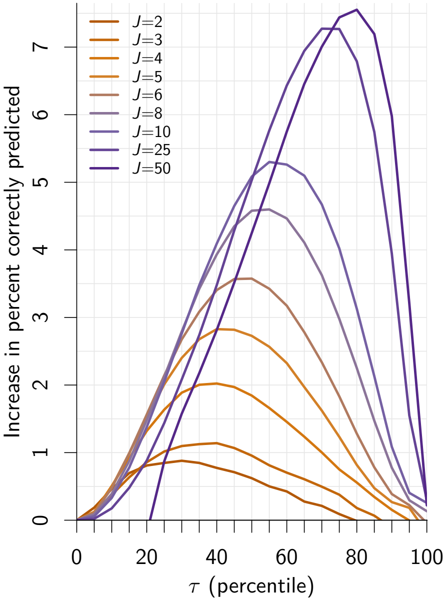

We next compare the difference in correctly predicted (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

***Utility maximizing alternative secondary decison rule*.**

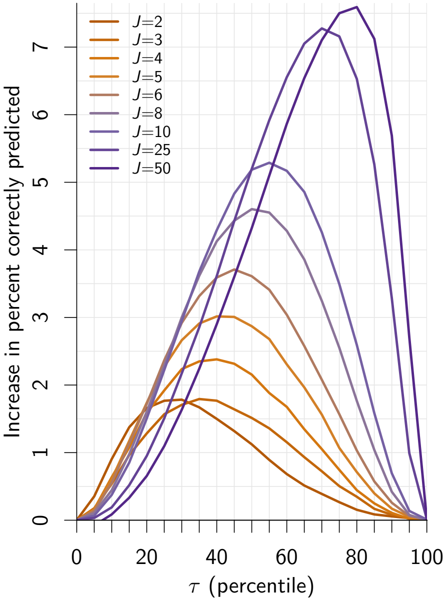

Results: Differences in correctly predicted

We next compare the difference in correctly predicted (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

Opt-out or choose the explicitly offered status-quo alternative secondary decison rule.

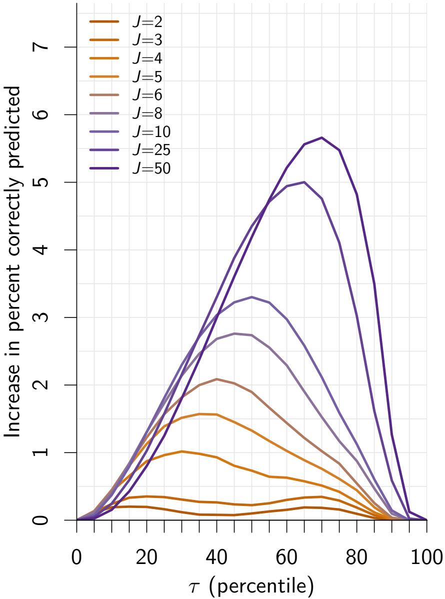

Results: Differences in correctly predicted

Key findings

- For non-extreme thresholds the satisficing model predicts a higher share of choices correctly.
  - Increases in improved prediction are linked with the number of alternatives.
  - The threshold associated with the maximum improvement in prediction inreases with the number of alternatives.

As an indicator of estimation performance of τ we calculate the root-mean-squared-error.

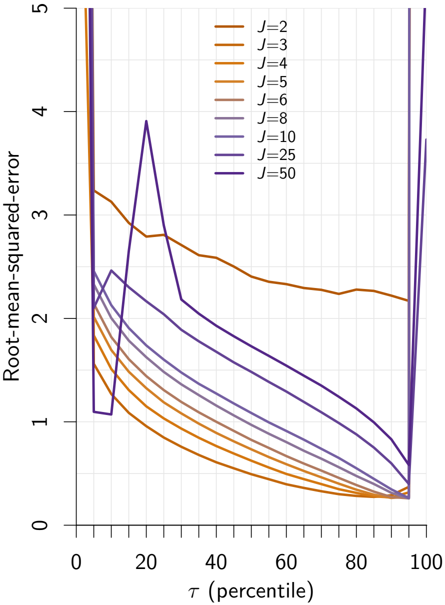

As an indicator of estimation performance of τ we calculate the root-mean-squared-error.

Last alternative secondary decison rule.

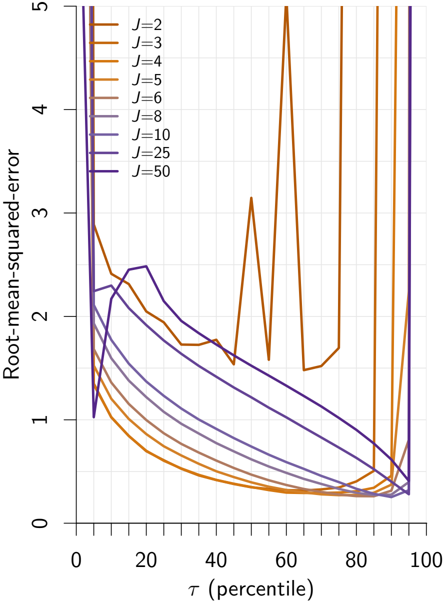

As an indicator of estimation performance of τ we calculate the root-mean-squared-error.

***Random alternative secondary decison rule*.**

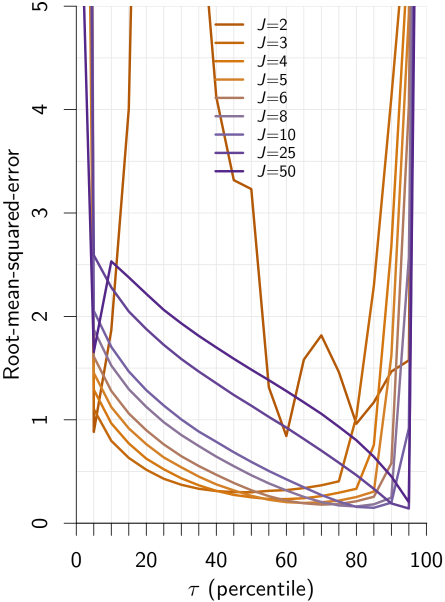

As an indicator of estimation performance of τ we calculate the root-mean-squared-error.

Utility maximizing alternative secondary decison rule.

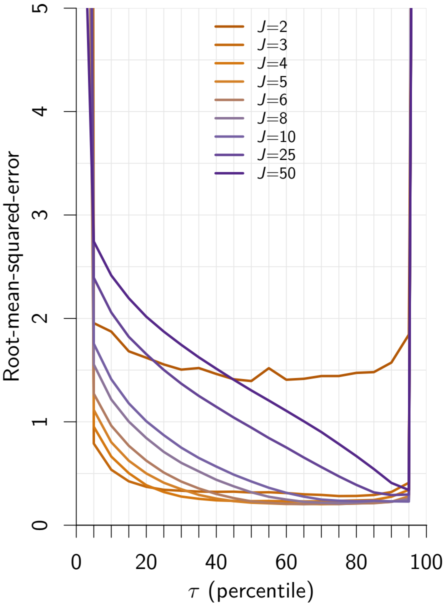

As an indicator of estimation performance of τ we calculate the root-mean-squared-error.

Opt-out or choose the explicitly offered status-quo alternative secondary decison rule.

Results: Bias in $\hat{\tau}$

Key findings

- The ability to retreive accurate estimates of the threshold depends on the threshold.
  - It is not well estimated at the extremes (since any extreme value will produce the same result.)
  - Inbetween the extremes, the estimated threshold becomes less biased as the threshold increases.
- The ability to retreive accurate estimates of the threshold also depends on the number of alternatives.
  - Estimated values of the threshold become increasingly biased as the number of alternatives grow.

Results: Differences in bias in marginal utilities

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

We present just for the setting where the secondary decision rule is utility maximisation. The other settings exhibit similar results.

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***AttA***:

$0.475 < \hat{\beta} < 0.525$.

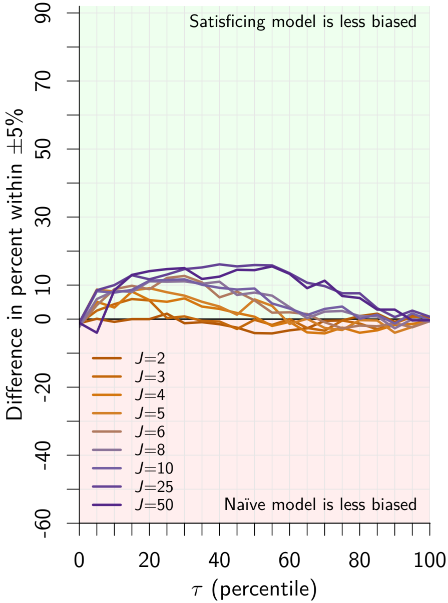

A | C | R | G
Applied Choice Research Group

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***AttB***:
$0.760 < \hat{\beta} < 0.840.$

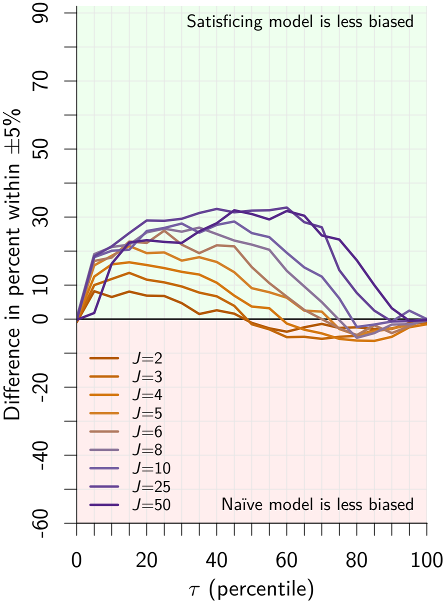

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***AttC***:
$-1.680 < \hat{\beta} < -1.520.$

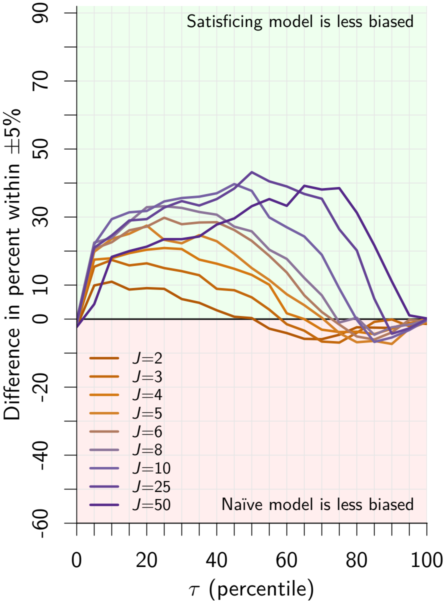

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***Cost***:
$-0.105 < \hat{\beta} < -0.095$.

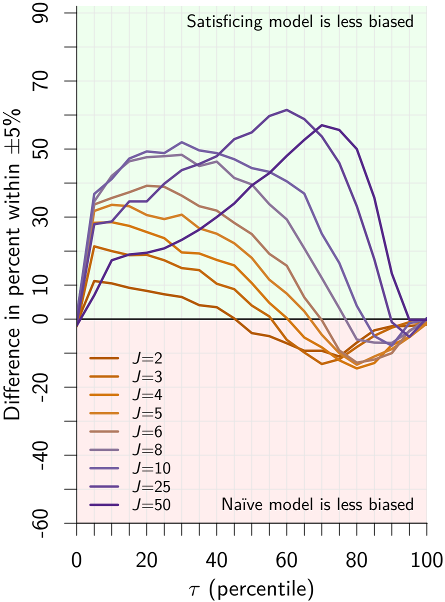

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***opt-out ASC***:
$-0.050 < \hat{\beta} < 0.050.$

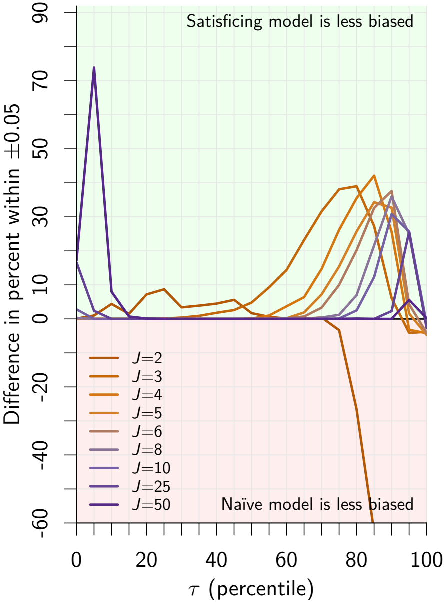

As an indicator of estimation performance we compare the percentage of parameter estimates within a given range (averaged over the 1,000 sample simulations) for the satisficing versus the naïve specification.

For ***alternative J ASC***:
$-0.050 < \hat{\beta} < 0.050.$

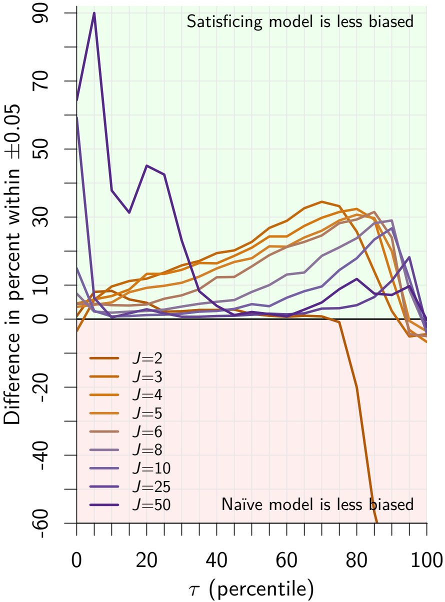

Key findings

- Generally the satisficing model produces less biased marginal utilities, especially so as the number of alternatives increases.
- This said, as the threshold goes beyond a certain point the naïve specification produces more accurate marginal utilities, especially for datasets with a small number of attributes.

Conclusion

Overall, our ***simple*** satisficing model appears to correctly identify and accommodate threshold levels of utility.

We find that it is robust under a variety of settings.

- Number of alternatives.
- Threshold levels of utility.
- Secondary behavioural rules.

We assume that all individuals use satisficing as their primary decision making rule and that they use one of four decision rules as their secondary rule.

- Admittedly, in reality every individual will use a strategy (or combination of strategies) that may be unique to them and that is likely to be highly dependent on the choice context.
- This limitation could, of course, be potentially relaxed through the use of probabilistic decision rule process models that accommodate heterogeneity in decision making strategies across individuals.

Related, we assume a constant τ, which implies that everyone has the same observable threshold utility.

- A pure satisficing strategy lies where τ uniquely identifies all choices in the data, which may require τ to be individual-specific.
- An easy extension is to reparameterize τ to accommodate individual ability, motivation and a range of other, perhaps unobserved, factors.

Of course, there is also scope for further specifications to accommodate preference heterogeneity.

- This may, in fact, be a necessary step to, at least partially, alleviate potential confounding concerns between $\hat{\boldsymbol{\beta}}$ an $\hat{\tau}$.

In some (but not all) cases it may make sense to impose $\tau \geq U_{\mathsf{SQ}}$.

Take home

When the threshold is not very low nor very high, the importance of capturing satisficing is greatest.

When the choice set is relatively big, the importance of capturing satisficing increases.

When choice sets get bigger, the importance of modelling satisificing increases with increases in the threshold.

A C R G

Applied Choice Research Group

www.acrg.site

Economics Division
Stirling Management School
University of Stirling
Scotland