



Strategic bias in discrete choice experiments

DOI:

[10.1016/j.jeem.2018.08.010](https://doi.org/10.1016/j.jeem.2018.08.010)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Meginnis, K., Burton, M., Chan, R., & Rigby, D. (2018). Strategic bias in discrete choice experiments. *Journal of Environmental Economics and Management*, 0, [102163]. <https://doi.org/10.1016/j.jeem.2018.08.010>

Published in:

Journal of Environmental Economics and Management

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

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

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Strategic Bias in Discrete Choice Experiments

June 5, 2018

Abstract

An induced value laboratory experiment is conducted to explore the vulnerability of discrete choice experiments to strategic misrepresentation of preferences. We consider strategic behaviour to arise when an agent: (i) believes the choice experiment will be used to determine a provision decision over a discrete set of alternatives; and (ii) has expectations about the relative likelihood of those alternatives being selected and delivered. In the experiment, agents receive induced values for the discrete set of provisioning alternatives. In treatments where agents receive information that their first best outcome is unlikely to win, we investigate the extent to which their choices change, in a manner consistent with them seeking to deliver their second best outcome in the provisioning decision. We find that 27% of respondents misrepresent their preferences and reveal evidence of strategic bias. We find that this behaviour is sufficient to change inferences about preferred provision at the aggregate level.

Keywords: Discrete choice experiments, strategic bias, provision rule, response strategies

1 Introduction

Strategic bias has long been a concern in stated preference studies. The phenomenon occurs when an individual “deliberately misrepresents their preferences in order to influence the decision making process” (Bennett and Blamey, 2001). Most of the research regarding strategic bias has concerned the contingent valuation method, with very few studies exploring whether, and to what extent, the bias might manifest itself in discrete choice experiments (DCE).¹

The main objective of this study is to examine strategic bias in discrete choice experiments within a realistic representation of the end use of DCEs in a decision making process. We conduct an induced value laboratory choice experiment that exhibits essential characteristics of when strategic bias might occur in the field. First, respondents believe an agency is deciding on a provision decision (e.g. yes/no development; high/medium/low conservation); second, respondents’ choices in the DCE are used in the provisioning decision; and finally, respondents anticipate that true preference revelation will lead to a unfavourable provision outcome. To address the first element, we present a set of possible provision outcomes to respondents. Additionally, we induce values for each of the outcomes, such that respondents know their preference ordering over the provision outcomes. To address the second, respondents are told that the DCE will be used to determine which provision outcome, and hence payoff, will occur. To address the final element, some respondents are given information on the likelihood of each of the provision outcomes being implemented. This final aspect reflects situations when individuals have expectations of the likely provision outcomes through, for example, media outlets, opinion polls, social networks, etc.

We compare choice behaviour between those respondents who do, and do not, receive information on the probabilities of the provision outcomes being implemented. We investigate the extent to which those who, in the expectation of an unfavourable outcome, target a non-demand revealing option and misrepresent their unconditional preferences in the DCE in order to deliver their second best outcome.

¹See Hanley et al. (2001) for an overview and example of DCE.

The final provision outcomes are each described in terms of attribute levels. Respondents are told that the survey will first be used to evaluate their likes and dislikes over the different attributes and then that information will be used to make a decision on which provision outcome they prefer. More formally, the DCE choice data are used to estimate marginal utilities for the attributes and calculate the choice probabilities for the predefined provision outcomes. The provision outcomes are not seen as a choice set within the DCE, in which case the task would revert to a simple referendum. As such, respondents need to establish and use attribute-based decision rules consistent with whichever provision outcome the respondent desires. The full set of DCE choice data are used to determine which provision outcome (and hence payoff) will be provided.

Our experimental set up creates a challenging task for respondents, yet one we believe has characteristics that allow for strategic behaviour found in the field. Despite the complex task, we find a significant portion of respondents misrepresent their preferences when faced with the likely implementation of an undesirable provision outcome. We find evidence of the use of lexicographic heuristics by those opting to misrepresent their preferences. Finally, we find that by misrepresenting their preferences agents were able to influence the provision outcome, such that an outcome more favourable to them is delivered.

In contingent valuation studies (e.g. Bohm (1972), Whittington et al. (1990), Wheeler and Damania (2001), Taylor et al. (2001)) there is mixed evidence regarding the impact of strategic behaviour. Motivation for strategic responses in this case, is driven by two types of reasoning. The first occurs if the respondent thinks that the good will be provided given substantial positive responses; however, in actuality they will be able to pay less than their stated amount or someone else will pay for the service (Whittington et al., 1990). In this case, respondents have incentive to overbid. The second reason is driven by a respondent's belief that the decision to provide the good has already been made and the survey is aimed to elicit at what price the good should be provided (Carson and Groves, 2007). In this consideration, a respondent has incentive to underbid. The contingent valuation framework is such that it is possible for respondents to bias their

responses because it may be easy to discern the research agenda and implement a strategy in the open ended or dichotomous framework. Even when contingent valuation studies are set up as an advisory referenda, responses may not always be incentive compatible if a respondent believes that the survey will be inconsequential (Vossler and Watson, 2013; Carson et al., 2014).

One of the (many) factors that prompted the widespread adoption of DCEs was that the dominant stated preference method, contingent valuation, had some weaknesses which DCEs were argued to be less vulnerable to. These included the scope for protest bidding and strategic behaviour (Hanley et al., 2001). Bennett and Blamey (2001) argue there are two main reasons as to why the DCE framework mitigates the possibility of strategic responses. First, a choice experiment environment can hide the true purpose of the survey and by asking respondents to make trade-offs between multiple attributes suggests certain ambiguity regarding the optimal response strategy. Second, the repetitive nature of DCEs provide an environment that is difficult to be consistently strategic. Despite the claims put forth that DCEs are less susceptible to strategic behaviour than other valuation techniques, there has been little testing of these assertions.

Carson and Groves (2007) argue that the DCE structure violates the Gibbard-Sattherwaite necessary condition for incentive compatibility, by asking respondents multiple questions. A response to this shortcoming is to stress consequentiality, which Carson and Groves (2007) define as the agent caring about the outcome and believing that the results from the survey will influence the decision making process. Using homegrown and induced value experiments, previous literature has explored how to achieve incentive compatibility in DCEs (e.g. Collins and Vossler (2009), Taylor et al. (2010), Carson et al. (2015), Interis et al. (2016)).² One approach in the literature has been to randomly select one choice set from the experimental design to be binding to achieve incentive compatibility. In the binding set, a respondent is asked to pay for the good/bundle once the outcome is determined by some provision rule (e.g. plurality or majority rule). While a useful

²Homegrown experiments use actual goods where values are known to individuals, while induced value experiments use arbitrary goods where the good itself is worth nothing to a respondent except through the value induced by the experimenter for the good.

mechanism to achieve incentive compatibility, this approach does not reflect how DCE data is used in practice. In practice, a single set from the experimental design would not be chosen randomly and the relative popularity of the constituent options used to determine the provision outcome. Our approach seeks to move the consideration of strategic bias closer to a more plausible use of DCE data, in which all choice data are aggregated, is used to inform the provision outcome.

This paper is structured as follows. Section 2 summarises the literature on strategic bias in discrete choice experiments. Section 3 outlines the experimental design and testable hypotheses. Section 4 discusses the results from the experiment and Section 5 concludes the paper.

2 Literature Review

The DCE framework, which involves eliciting preferences for multi-attribute goods by asking repeated discrete choice questions, makes DCEs susceptible to untruthful responses. The need for surveys to be seen as consequential allows for respondents to have an incentive to misrepresent their preferences. This is driven by expectations about (i) the provision rule and (ii) other subjects' responses. Those in combination create an environment where it is not always in one's best interest to tell the truth. For example, Carson and Groves (2007) highlight that in a simple plurality provision rule, an agent's best strategy is to pick between the two alternatives they believe to have the highest probability of winning; this strategy may exclude an agent from ever choosing their most preferred alternative.

Two main experimental approaches have been used to examine preference misrepresentation. The first is to look at differences in homegrown values between hypothetical and binding treatments. In this case, the binding treatment is theoretically incentive compatible and thus any deviations from the hypothetical treatments measure some degree of bias (Vossler and Evans, 2009). A second approach is to look at differences between choices when values are induced and therefore the true preferences of individuals are

known. In both approaches, a common theme driving misrepresentation of preferences is the stated provision rule. Carson and Groves (2007) state that in analysing consumer responses, it is important to determine how a respondent believes the provision outcome will be affected by their survey responses.³ Vossler and Evans (2009), Taylor et al. (2010) and Vossler et al. (2012) use homegrown experiments to explore how changes to provision rules drive strategic misrepresentation of preferences when binding treatments are compared to hypothetical ones. In all three studies, bias was defined as the difference in willingness to pay (WTP) between the binding treatment with a simple majority provision rule, i.e. respondents would actually pay at the end of the experiment, and treatments where the provision rule was more opaque or hypothetical.

In Vossler and Evans (2009) participants respond to a referendum regarding a classroom recycling programme. Using different treatments, respondents' votes will either (1) be binding and the outcome will be implemented through a majority vote provision rule; (2) only advise an agency on the referendum outcome; (3) be completely hypothetical; or, (4) be added to moderator votes. Bias was largest in the latter two treatments. When participants' votes correspond to only 25% of the deciding votes, and 75% of the votes were by moderators, a respondent's probability of influencing the referendum outcome falls. Thus, when a respondent believes that all moderators will vote yes (no) the provision rule is no longer incentive compatible. This highlights the need for agents to believe their answers will actually affect the outcome. If agents anticipate having little probability of influencing the agency's decision, the incentive to tell the truth diminishes.

However, differing provision rules do not always lead to bias. Vossler et al. (2012) run an experiment for a public tree planting programme. Respondents answer a DCE and are told that one choice set will be randomly selected as binding. The experiment included three real payment treatments and one hypothetical treatment. The real treatments gave varying degrees of information on how responses in the binding choice set would influence how the project would be implemented. In the hypothetical treatment respondents were

³Scheufele and Bennett (2012) argue that learning may also be an important driver of strategic responses in DCEs. They conduct an online survey to test for strategic learning for a nature reserve. They find that respondents are more cost sensitive as they progress through the multiple DCE questions.

not given a provision rule, such that their choices were completely hypothetical and they were simply told that results from the survey will be provided to the government. Significant differences in WTP were observed only between the real and hypothetical treatments. No significant difference was found in the real treatments which varied the stated provision rule. This suggests that in the absence of a provision rule, bias tends to be most apparent. These findings are supported by Taylor et al. (2010) who also find that when provision rules have been explicitly mentioned, WTP is closest to the binding treatment while bias is largest in the absence of a provision rule.

Deviations from true preference revelation are also strongly driven by a respondent's expectations about the choices and preferences of other participants (Carson and Groves, 2007). Vossler et al. (2012) directly investigate this phenomenon by asking respondents *ex post* whether they considered the decisions of other participants. By doing so they create a platform for identifying strategic respondents. However, they find little evidence of strategic behaviour; few respondents indicated considering how others might vote. This is likely due to an individual having little reason to expect that the preferences of others will lead to an unfavourable tree planting project. As the survey gave no indication regarding the preferences of others, it is not surprising that few respondents exhibited strategic behaviour in this study.

Diverging from homegrown studies, we turn to induced value experiments where researchers can clearly test the vulnerability of DCEs to strategic bias, because the true preferences are known. In induced value experiments, the researchers can discern true preferences from non-demand revealing preferences, i.e. those which deviate from the induced values. Collins and Vossler (2009) run an induced value DCE where participants receive value for arbitrary attributes and are presented with choice sets that vary the levels of these attributes. Like the homegrown studies, one choice set is randomly selected to be binding. Depending on the treatment, the provision rule determines which alternative from the binding set is implemented. Deviations from induced preferences in a choice set stem from (i) the provision rule, (ii) other respondents' support for an option and (iii) expectations about the decision makers' preferences. Respondents knew that

the values of other players might differ from their own values; however, the distribution was unknown. Overall, there was little evidence of deviation from induced preferences across all treatments. These findings differ from Luchini and Watson (2014) who, in a simple posted-price provision rule (where respondents receive the option they selected in the binding set), find only 2/3 of respondents make choices that are consistent with their induced values. This may be due to the cognitive effort needed in Luchini and Watson (2014) being more difficult; their study had four attributes including a cost attribute, while Collins and Vossler (2009) had only one attribute and one cost attribute.

A similar study conducted by Carson et al. (2015) found that individuals chose options that are not aligned with the induced values, over 10% of the time. Respondents were given values for arbitrary attributes and presented with choice sets that vary the levels of the attributes. Like Collins and Vossler (2009), one choice set is randomly selected to be binding and using a plurality provision rule, whichever option has the most votes is paid out to all individuals. They find that respondents usually cast these non-demand revealing choices for the status-quo option, which was usually the second best alternative in each choice set. They offer several explanations for this phenomenon. First, selecting the status-quo option is due to loss aversion, i.e. respondents prefer the option that is offered at zero cost. Second, computation errors in calculating marginal differences in attributes may lead to non-demand revealing choices. Finally, selecting the second best alternative, which happens to be the status-quo a significant portion of the time, is due to an individual's priors about the distribution of other participants' preferences and not a behavioural bias for the status-quo. Furthermore, when respondents continuously face choice sets where many of the alternatives are bad outcomes (i.e. low monetary pay offs), respondents gravitate toward the safety of the status-quo option, as the preferences of others are unknown and could result in a low pay off outcome.

So far, the literature has focused on the notion that incentives to strategically bias is from the expected outcome of a provision tournament. Burton (2010) examines strategic bias as instead trying to manipulate the value of one attribute within the choice experiment. Burton runs a two-stage DCE concerning university housing. In the first stage,

subjects are asked to make choices according to their preferences, providing unbiased parameter estimates. In the second stage, subjects are informed that they will be rewarded for their ability to distort the values of certain attributes. Using a conditional logit model, Burton estimates the marginal utilities in the first phase as well as in the second phase. Bias is measured as the difference in choice probabilities for simulated scenarios. Despite respondents not knowing the mechanism by which the DCE choices will be used to elicit values for the different attributes, individuals are able to manipulate the choice probabilities through their biased choices in the second stage.

A common element of the literature examining strategic bias in a DCE, using home-grown or induced value experiments, is the random selection of one choice set to be binding (except from Vossler and Evans (2009) who present a binding one-shot dichotomous choice and Burton (2010) who ask respondents to distort parameter values through their choices.). This differs from how choice experiments are used in the real world. Single sets are not treated as an individual tournament; rather agencies are provided with results and values based on all choices. Additionally, subjects are likely to interpret the repetitive nature of the choice experiment to imply some aggregation of preferences from all questions (Vossler et al., 2012).

The literature provides a foundation for when strategic bias might occur. Respondents who are likely to strategise are those who (i) believe the survey is influential in deciding the provision outcome; (ii) have an expectation on the likelihood of the provision outcomes being implemented; and (iii) are better off not behaving truthfully. We consider all these elements in our experimental design. Furthermore, we investigate the strategic agent under plausible conditions of how DCEs may be used to inform the decision making process. We contribute to the literature by inspecting the vulnerability of DCEs to strategic bias under the assumption a provisioning agency will use all of a subjects' responses and an agent will expect some aggregation of preferences to inform the provision outcome. We first elicit choices and outcomes under truthful preference revelation conditions and compare them to choices made when information on the likelihood of the provision outcomes being implemented creates a scenario where respondents must

choose: continue revealing truthfully and likely end up with a “bad” provision outcome or deviate in the hopes that by doing so, the “bad” provision outcome is out favoured by a second best outcome.

Section 3 explains how the experiment builds upon the existing literature with an induced value choice experiment to determine the provision outcome.

3 Experimental Design and Testable Hypotheses

Our experimental design deviates from previous studies in three main aspects. First, we explore strategic bias in choice experiments using a predetermined set of three possible provision outcomes. Establishing the provision outcomes prior to the DCE differs from Collins and Vossler (2009) and Carson et al. (2015) who randomly choose one choice set to be binding. We assume the provision outcomes are a predetermined set such that we do not rely on the statistical design to determine which bundles of attributes make up the possible outcomes. Additionally, we emulate a field DCE where it is reasonable that respondents would assume only a finite set of provision outcomes are ever being considered by decision makers. For example, decision makers could be deciding between high, medium or low development options (e.g. wind farm development, housing options, etc.).

Second, we use all the preference information collected by the choice experiment. Consider a development project where the DCE will be administered to elicit preferences for attributes of the development, which decision makers will then use to decide the level of development that should be implemented. In our experiment, a respondent will be asked a series of trade off questions involving a number of attributes. Those trade offs will be analysed to identify the values held for the attributes. Based on those values, we will calculate the probability of preferring each of the possible provision outcomes for the individual. Of the three possible provision outcomes, the one with the highest probability will be chosen, hence a plurality provision rule. The payment to the respondent will depend on which option is selected.

Third, respondents are given information on the the likelihood of each provision option being delivered. As previously mentioned, this emulates the case when individuals have a general idea about their community’s preferences through opinion polls, market research, media outlets, etc.⁴ In the case of three provision outcomes, each respondent has a first, second and third best outcome. When given information that suggests one’s first best outcome has the lowest probability of being implemented, a respondent will have incentive to switch to target the second best outcome. Thus, the incentive to misrepresent one’s preferences is a direct result of receiving information on the probabilistic outcome of the different provision outcomes.

Previous studies by Collins and Vossler (2009) and Carson et al. (2015) offer no information to respondents on others’ preferences, they simply suggest that others may have differing values. Their approach allows for strategic behaviour to manifest, but it relies on two assumptions. First, that an individual will expect the values of others to be different than their own; and second, that these expectations will be profiled such that it diverts them from representing their true preferences. Additionally, without the respondent having explicit information on others’ preferences, the researcher cannot hypothesize as to the direction of the bias, should there be one. In our study, we provide respondents with the likelihood of the provision outcomes being implemented which will also be taken into consideration when determining the provision outcome. The likelihoods are presented as probabilities that each provision outcome will be the most preferred; for example, Option A has a 50% chance of being preferred by others. In order for a respondent to misrepresent their preferences, it must be the case that their most preferred outcome has little chance of winning (Scheufele and Bennett, 2012). This allows us to test whether, in the expectation of an unfavourable outcome, respondents switch to target a non-demand revealing option.

⁴This has been shown in the voting literature where public opinion can be inferred through the media, but also through conversations with friends, speeches, or the actions of public figures (Gunther, 1998).

3.1 Experimental Design

We now derive the experimental setup more formally. Consider three provision outcomes $B = b_1, b_2, b_3$ that make up the different provision outcomes. Each outcome b_k is composed of attributes $X = x_1, x_2, x_3$ indexed i . An outcome b_k , will be composed of attribute levels x_{ki} . Each agent n receives an induced value for each provision outcome b_k denoted as $V_{k,n}$. Table I outlines the profiles of the three possible outcomes.

Insert Table I Here

Each provision is composed of specific attribute levels. Given a simple plurality provision rule, strategic bias manifests when an agent anticipates their most preferred outcome to be unlikely to win given the likelihood of the provision outcomes being implemented. To illustrate this, suppose we have a sample of N agents. Within N agents, there exists different types of people with heterogeneous preferences for outcomes B . Consider the simplest case with two different types of agents, labelled Type I and Type II. Type I individuals have preference ordering $b_3 > b_2 > b_1$, while Type II individuals have preference ordering $b_1 > b_2 > b_3$.

The expected probability of any one outcome b_k being most preferred is $Pr(b_k)$. Blais and Nadeau (1996) argue that there are two factors that will increase the probability that an agent will behave strategically. The first concerns the relative intensity of one's preference ranking. The closer one's preference is between the first and second best option and the further one's preference is between the second and third best option, the greater a respondent's probability of behaving strategically. The second factor is regarding the expected probability of each option winning. The larger the difference is between the probability of one's first and second best option winning and the smaller the difference is between the probability of one's second and third best option winning, the greater the probability of strategic voting.

For example, if the expected probability of the final provision outcome is $Pr(b_3) > Pr(b_2) > Pr(b_1)$ then Type II respondents have incentive to behave strategically: their worst option is the most likely outcome, while their first best option is least likely to "win". Assuming the probability between $Pr(b_2)$ and $Pr(b_1)$ is large and the probability

between $Pr(b_3)$ and $Pr(b_2)$ is small, a Type II respondent has the incentive to behave strategically.⁵ In this case, a respondent is faced with a high probability of receiving a bad outcome and therefore has incentive to target their second best outcome, b_2 .⁶

In a simple voting exercise, a strategic Type II respondent would simply cast a vote for their second best option, b_2 . However, our experiment emulates the realistic environment where a DCE is administered to advise on a predetermined set of provision outcomes. First, respondents are given a DCE of C choice sets and asked to select their preferred option in each choice set. The agency will establish the marginal utility estimates β for all attributes in X through a conditional logit model. Second, the β utilities will be used to calculate the logit probabilities for B. Whichever provision is probabilistically most preferred, will be implemented.

In this experimental design an agent can only indirectly influence the provision outcome by revealing preference for the attributes seen in the DCE. In Collins and Vossler (2009) and Carson et al. (2015) respondents can switch from a first best to a second best alternative in an individual choice set to reveal preference for a different provision outcome. In our set up, choice sets do not contain actual provision possibilities, but solely attributes, attributes which happen to make up the provision outcomes. In order to shift the provision outcome in their favour, respondents must devise attribute based decision rules based on the composition of the provision outcome they are targetting. The DCE will advise on which provision outcome, from the predetermined set, is most preferred and hence provided. Exhibiting strategic bias in this experiment is much more difficult than in the experiments of Collins and Vossler (2009) and Carson et al. (2015) as (i) it is not a simple voting exercise; (ii) our provision rule incorporates the repeated nature of DCEs; and, (iii) the unlabelled DCE implies decision rules must be attribute based in order to shift the provision outcome.

⁵According to Blais and Nadeau (1996), we also need to assume that the utility difference between b_2 and b_1 is small, while the utility difference between b_2 and b_3 is large.

⁶It can be argued that in anticipation of Type II respondents behaving strategically, Type I respondents will retaliate. However, for simplicity we do not explore this possibility and focus on whether Type II respondents will act strategically under these conditions.

3.2 Testable Hypotheses

This experimental design is to test four main hypotheses.⁷ The first hypothesis concerns the robustness of the experiment. We want to first confirm that when agents expect to receive a bad outcome from the final provision outcomes, we see a switch to the second best outcome. To do so, we consider control and treatment groups where the control group have no priors about others' preferences and therefore have no incentive to misrepresent. In the treatment groups, respondents are given the probabilistic outcome of the final provisions being implemented. Therefore, we can compare outcomes and strategies in the presence and absence of information that creates incentives to misrepresent one's preferences.

We consider that strategic bias is likely to occur when an agent's most preferred option is least likely to be the provision "winner"; therefore, the information given to the treatment group regarding the likelihood of the provision outcomes being implemented reflects this. We induce agents to have preference ordering $b_1 > b_2 > b_3$. Respondents in the treatment group are given information on the preferred outcome of N-1 agents through provision outcome probabilities, such that $Pr(b_3) > Pr(b_2) > Pr(b_1)$. The probabilities of the second and third best outcomes are close, such that the second best outcome is a realistic contender. Similarly, the probability of the first best outcome is small, such that the probability of it being implemented is negligible. In this case, respondents in the treatment groups can anticipate an unfavourable outcome b_3 as the likely winner, and will switch to target the second best option. Agents in the control group, who do not have information on the outcome probabilities, will more likely target the first best option.

H1: Significantly more respondents in the treatment group will target the second best option than in the control group.

Our second hypothesis concerns whether the relative difference in $V_{k,n}$ between outcomes in B, affects an individual's willingness to switch to their second best outcome.

⁷The following hypotheses will compare respondents targeting the first and second best outcomes. This is because it is never in one's interest to vote for one's least preferred outcome (Collins and Vossler, 2009) and only a small percentage of our sample targeted the third and in our case, worst option.

Recall Blais and Nadeau (1996) who state that two differences in utility matter for a strategic agent. First, the utility difference between the first and second best option and second, the utility difference between the second and third best option. As this is an induced value experiment, the utility for each option is expected to be the given values $V_{k,n}$. Our first treatment group has values for the first, second and third best outcomes of £5, £4 and £0, respectively. We refer to this treatment as Treatment(0,4,5). In order to test the mental anchoring effects of the given values for the outcomes in B, we include an additional treatment, referred to as Treatment(0,3,6) whose values for the three provision outcomes are £6, £3 and £0. Within Treatment(0,3,6), respondents face a higher pay off for the first best option and a slightly smaller pay off for the second best option.⁸ We hypothesise that even when the difference in pay offs between one's first and second best option, d_{12} , increases, respondents will still target the second best option. We selected a range of values for V_k such that strategic behaviour is expected to manifest in both treatment groups.

H2: As d_{12} increases, respondents in Treatment (0,4,5) and Treatment (0,3,6) are equally as likely to behave strategically.

Our third hypothesis considers the decision rules taken by respondents. Respondents receive pay off for each of the final provision contenders $V_{k,n}$, but not for the individual attributes within b_k . Respondents have to conjecture how the agency will use their responses in the DCE to implement one of the outcomes, a common behavioural problem of choice experiments (Carson and Groves, 2007). Individuals must therefore define a decision rule to use within the DCE to reveal preference for b_k by making tradeoffs between x_1 , x_2 and x_3 . One way for respondents to interpret this task is to choose a simple decision rule based on the attribute composition of the desired outcome. We consider it likely that given the complexity of the task they face, that individuals will use “fast-and-frugal heuristics” (Mousavi et al., 2016). At its most simple, this reduces to a lexicographic decision rule. Lexicographic decision rules are one way for respondents

⁸Within each group, all individuals see the same values for each outcome. Such that V_k^c , $V_k^{t(0,4,5)}$, and $V_k^{t(0,3,6)}$ is the same across individuals in the control group, Treatment(0,4,5) and Treatment(0,3,6).

to simplify the task, while making reasonable choices and solving the problem presented in the experiment. To implement this, respondents must choose one attribute that most clearly distinguishes their target outcome from the other two outcomes. For example, if we consider an agent who wants to reveal preference for $b_{k=2}$, then we must examine all the outcome-attribute combinations and choose the attribute where $x_{2i} > x_{ki}$ for $k = 1, 3$, or $x_{2i} < x_{ki}$ for $k = 1, 3$, holds. These can be referred to as a maximum/minimum attribute x_i strategy. Respondents targeting their first best outcome and respondents targeting their second best outcome will exhibit different decision strategies on attribute x_i .

H3: An agent targeting the second best outcome b_2 will use a maximum(minimum) strategy on x_i more frequently than an agent targeting the first best outcome b_1 is using a maximum(minimum) strategy on x_i .

Our final hypothesis compares the stated objective of individuals with the choice probability predictions found through the conditional logit model. We ask respondents after the DCE to state the outcome they were trying to reveal as their preferred outcome.⁹ In this regard, we will observe a respondent n wanting the provision outcome b_k making choices $c_{n,k}$ in C choices sets. We estimate a respondent's preference parameters via a random utility model on their DCE choices and find the predicted probabilistic ranking of contenders in B . Based on Lancaster (1966) and McFadden (1974) we use a conditional logit model to estimate the marginal utilities for x_i . We estimate the conditional logit model for the aggregate sample, as well as for sub-sample of individuals based on their stated objective, such that all respondents targeting b_k consist of one sub-sample. Using the parameter estimates, we calculate the logit probabilities for each outcome. We hypothesize that conditional on which outcome b_k an agent n is targeting and the choices

⁹We have no reason to believe that when answering this question, a respondent would respond strategically by stating an objective other than the objective they were targeting in the DCE. If a respondent believes this question will have no effect on the provision outcome, then there is no incentive to lie. However, even if a respondent believes this question will influence the provision outcome, there is still no incentive to lie. Consider a respondent in the control. If a respondent has preference ordering $b_1 > b_2 > b_3$, there is no logical reason to state they were trying to reveal preference for b_2 or b_3 in case b_2 or b_3 is then provided. In the treatment, the same logic holds for a respondent revealing preference for b_1 . In the case where the respondents has made the decision to be strategic and target the second best outcome, b_2 , then stating b_1 or b_3 in this question who negate that effort.

$c_{k,n}$ they made, the estimated choice probability for outcome b_k will be largest.¹⁰

H4: Given that a respondent is targeting the provision outcome b_k , the choice probability calculated *ex post* for outcome b_k will exceed the choice probability for $b_h \forall h \neq k$. Such that, $Pr(b_k|c_{k,n}) > Pr(b_h|c_{k,n}) \forall h \neq k$.

3.3 Experiment participants and procedures

The sample consisted of one hundred and thirty individuals. They were all undergraduates, postgraduates and staff at the University of Manchester. Each respondent answered twelve choice sets, with three possible alternatives in each choice set. The DCE survey was designed and implemented in Sawtooth Software using a balanced overlap design with five different blocks. Individuals participated one at a time on a laptop via a computer assisted personal interview (CAPI). Respondents were told that their responses to the choice sets would be used to determine the experimental outcome between three possible outcomes. Below we describe the treatments and experimental procedure, the exact instructions can be found in Appendix C.

3.3.1 Treatments

The experiment was conducted in two waves. In the first wave, respondents were randomly allocated to a control or treatment group (referred to as Treatment(0,4,5)). Thirty individuals were in the control and thirty-nine individuals were in Treatment(0,4,5). In both groups, the values for the provision outcomes were £5, £4 and £0 for respondents first, second and third best outcomes. Respondents allocated to the treatment received information about the provision outcome probabilities, such that $Pr(b_3) > Pr(b_2) > Pr(b_1)$ was 50%, 45% and 5% respectively. These probabilities are fixed and are a feature of the experiment controlled by the research. However, it was explained to participants that these probabilities represented the preferences of others, to give the illusion that respondents were part of a pool. We varied the sample size of ‘others’ from which these

¹⁰The effect of strategic switchers on the predicted provision outcome was extensively investigated in simulation work prior to conducting the lab experiment. One example of this simulation exercise can be shown in Appendix A.

probabilities were found. The stated sample was a random number between 10 and 20. These values were selected such that each respondent represents between 5% and 10% of the overall sample.¹¹ We propose this to be a small enough sample that an individual will believe they can increase the probability of their second most preferred outcome to be the largest, but a large enough sample such that they believe they cannot increase the probability of their first best option to be the provision winner.¹² In actuality respondent's choices in the treatment groups were pooled with a fixed simulated sample of 10 agents whose predicted outcome over the three contenders was 50%, 45% and 5%.

Additionally, as mentioned in Section 3.1, the chosen probabilities reflect the case when the difference between the second and third best option is small, while the difference between one's first and second best is substantial. In this case, the worst outcome (i.e. the outcome that paid £0) is predicted to be the provision winner, the best outcome is unlikely to win, while the second best outcome is a serious contender. Given the conditions of the information on the provision outcome probabilities, a respondent will have incentive to behave strategically and switch to target their second best outcome. In the instructions, respondents in the treatment groups were given the size of the sample and the provision outcome probabilities. The experimental set-up in the control group were given no information on the provision outcome probabilities and the instructions used language to imply that only their individual responses were being considered and no other preferences would be taken into account in deciding the provision outcome.

In the second wave of the experiment, we changed the induced values for the first, second and worst provision outcomes to be £6, £3 and £0, respectively. In this second treatment group (Treatment(0,3,6)), all respondents were given information about the provision outcome probabilities. This differs from the first wave as we do not

¹¹There is anecdotal evidence that in collecting data for DCE studies, communities become aware of the study and discuss (via Internet forums, word of mouth, etc.) the merit and overall objective of the study. Thus it may not be that an individual needs to believe themselves to be able to influence the provision, but a constituency who collectively agrees to bias their answers to attain a different provision result.

¹²Blamey (1998) argue that a respondent's willingness to provide positive WTP for a contingent valuation study is also contingent on a respondent's belief regarding the size of the population. Therefore, we chose to vary the stated size of the population (N-1). We use a multinomial regression to test whether there are any effects of varying the stated N-1 on respondents choice of provision outcome. We found no such effects. Results available upon request.

need an additional control group who does not see the probabilities. The results from Treatment(0,3,6) can be compared to the control from the first wave as well as Treatment(0,4,5).¹³ Treatment(0,3,6) saw information that the provision outcome probabilities were $Pr(b_3) = 50\% > Pr(b_2) = 45\% > Pr(b_1) = 5\%$, such that the experimental set up was exactly that of Treatment(0,4,5).

3.3.2 Experimental procedures

In order to simplify the cognitive burden of the task, we chose to label the provision outcomes and attributes. We selected three possible provision outcomes: Leisure Centre, Community Park or Market. These outcomes were composed of three attributes: roads, jobs and trees.¹⁴ We deliberately have all ‘positive’ attributes, such that there is no cost attribute. This is first because we are not giving respondents any initial balance with which they could use to ‘pay’ for the provision outcome. Second, any inclusion of a price attribute would conflate the task at hand with making trade-offs between an alternative’s attributes and the associate cost.

The DCE would lead to one of the development options being implemented; however, within the choice sets themselves, the alternatives are unlabelled.¹⁵ The different provision outcome profiles are shown in Table II. Only respondents in Treatment(0,4,5) and Treatment(0,3,6) saw information on the probabilities of each outcome.

Insert Table II Here

Respondents saw the final provision outcomes and the associated monetary values prior to answering the DCE. A reminder was given under each choice set as to the

¹³We note that different monetary pay offs in a control treatment with £0, £3, and £6 may have differed from the control with £0, £4 and £5 payoffs. However, one’s decision to deviate from the first best outcome can only be *reduced* when the first best outcome payoff is increased (from £5 to £6) and the second best outcome payoff is decreased (from £4 to £3). Therefore, if any differences are observed between Treatment(0,3,6) and control(0,4,5) then they would likely be present, if not even more apparent, between Treatment(0,3,6) and a control(0,3,6).

¹⁴A variety of different attributes were considered, including generic attributes such as different colour buttons, alphabetic characters, etc. Initial trials indicated that using such arbitrary generic attributes increased the cognitive challenge for respondents who had to translate their target provision outcome into an attribute based decision rule. Hence the use of the more intuitive jobs/roads/trees. We acknowledge that this creates the risk of inadvertently introducing ‘real’ preferences into the choices, something we tested for and failed to detect in pilot debrief questions.

¹⁵The experiment instructions as well as an example choice set can be seen in Appendix C.

attribute levels in each possible provision outcome, the monetary payoff, and the provision outcome probabilities (only for the treatment groups). All respondents are expected to have a preference ordering of $M > CP > LC$ based on the induced values for each of the possible provision outcomes. Respondents received information that the agency will evaluate how much they like/dislike the attributes based on their responses to the twelve choice sets; they will then translate this into which provision option is most preferred.

There are three components to this experiment. First, respondents must understand the relative value of the attribute levels for roads, jobs and trees in the three labelled provision contenders. Second, they must decide which option to target based on V_k as well as $Pr(b_k)$ for the treatment groups. This involves targeting one's first best option or understanding that the second best option has a higher chance of being the provision outcome. Third, respondents need to develop a decision rule based on jobs, roads and trees to implement in the unlabelled DCE, that will reveal preference for their target provision outcome.

After the choice sets, respondents answered subsequent follow up questions. These included asking participants which provision outcome they were targeting through their choices, their strategy during the DCE, whether or not they considered the provision outcome probabilities (only for the treatment groups) and how difficult they found the task. As respondents played one at a time, payment needed to be immediate. Therefore, after the responses had been recorded, a conditional logit model was estimated using R and the subsequent choice probabilities were calculated for the final provision (R Core Team, 2013). Respondents saw the probabilistic outcome of the provision based on their individual responses. For those in the treatment groups, R collated the individual's responses with those of the simulated sample to emulate the environment explained to respondents. All participants received £3 for their participation and an additional amount according to which outcome was estimated to have the largest choice probability.¹⁶

¹⁶This experiment received ethical approval in 2016 from the University of Manchester Research Ethics Committee 6, reference number 16275. It was funded by the University of Manchester School of Social Sciences.

4 Results

The follow up questions regarding strategies and targets within the choice experiment included explicitly asking respondents which provision outcome they were targeting. In the case of the treatment groups, they were asked whether they considered the information on the provision outcome probabilities. Table III shows the stated objectives for all individuals in each treatment. Respondents in the treatment groups are subdivided into whether or not they reported having considered the information on the provision outcome probabilities.

Insert Table III Here

Across the control and treatment groups, respondents indicated targeting both the first, second and even the third best outcomes. We define “First best” respondents (FBR) as the participants who targeted the Market, i.e. the provision contender which provided the largest monetary payoff. “Second best” respondents (SBR) are defined as all participants who targeted the Community Park, i.e. the provision contender which provided the second largest monetary payoff. In Treatment(0,4,5), 49% of respondents indicated having considered the information on the provision outcome probabilities, while 44% did not take into account this information. In Treatment(0,3,6), 34% of respondents reported taking into consideration the provision outcome probabilities while 49% did not.

4.1 Hypothesis 1

Hypothesis 1 states that respondents will switch to a second best outcome when faced with an expected unfavourable provision outcome. In other words, considering the voting behaviour of other respondents may be signs of strategic bias (Vossler et al., 2012). Looking at Table III, 57% of individuals in the control group, which did not receive information on the provision outcome probabilities, targeted the first best provision option and 27% targeted the second best option despite receiving less payment. This is consistent with previous studies which find that even under the simplest incentive structure, participants deviate from payoff maximising choices due to measurement error and/or cognitive bur-

den (e.g Luchini and Watson (2014), Carson et al. (2015)). In Treatment(0,4,5), 44% target the first best option and 44% of respondents target the second best provision option. In Treatment(0,3,6), 39% and 49% of respondents reported targeting the first and second best outcomes respectively. Hypothesis 1 states that respondents in the treatment groups will target the second best option more than respondents in the control. Comparing first SBR respondents in the control group (8/30) to SBR in Treatment(0,4,5) (17/39), a non-parametric Fishers Exact Test does not indicate that significantly more respondents are targeting the second best option in the treatment (FE p-value 0.207). Comparing the SBR in the control group (8/30) to SBR in the Treatment(0,3,6) (30/61), we find evidence at the 5% level that more respondents are targeting the second best outcome (FE 0.045).

However, only respondents who take into account the information on the provision outcome probabilities have incentive to target the second best provision contender. In that case, we examine the subset of respondents who considered the probabilities (49% in Treatment(0,4,5) and 34% in Treatment(0,3,6)). Comparing SBR in the control (8/30) to SBR in Treatment(0,4,5) conditional on having taken into account the provision outcome probabilities (12/17), significantly more respondents will target the second best option in the treatment than in the control group (FE p-value 0.017). Similarly, conditional on considering the probabilities, significantly more respondents target the second best in Treatment(0,3,6) (15/19) than the control (FE p-value 0.003).

4.2 Hypothesis 2

Hypothesis 2 states that individuals in Treatment(0,3,6) will be equally as likely to switch to the second best option as respondents in Treatment(0,4,5). Table III shows that in Treatment(0,3,6), 34% of participants said they took into consideration the provision outcome probabilities, while 49% did not consider this information. Recall that a strategic individual will target their second best outcome because they take into consideration the provision outcome probabilities. Therefore, fifteen out of the nineteen respondents who considered the information in Treatment(0,3,6) can be categorized as strategic, while

twelve out of seventeen from Treatment(0,4,5) fit the criteria. Using a Fisher's Exact Test we fail to reject equality among these proportions (FE p-value 0.706). This suggests that increasing the pay off difference between the first and second best option d_{12} , does not change the probability that an individual will switch to target their second best outcome when faced with a likely unfavourable provision outcome.¹⁷

4.3 Hypothesis 3

Hypothesis 3 concerns choice strategies. As described in Section 3.1, respondents are faced with a difficult task and need to derive a choice strategy that is based on the attributes of their provision target. Hypothesis 3 is therefore two fold. First, we must establish individual behavioural decision rules. Second, we look at whether the decision rules across respondents with different objectives, differ. In other words, as there is no way to determine whether what a respondent says they are doing (i.e. through the self reported statement at the end) with what they are actually doing, we can determine if decision strategies are consistent with their reported target.

We expect FBR to be making lexicographic choices on the attribute that distinguishes the Market from the other two provision outcomes, in this case attributes jobs or roads. Recall that individuals targeting an outcome are expected to look for the attribute that most distinguishes the target outcome from the other two provision contenders. For FBR, this may either be the attribute jobs or the attribute roads as the Market has the highest provision of jobs and the lowest provision of roads compared to the Leisure Centre or Community Park. The attribute jobs satisfies the maximum attribute criteria: $x_{M,jobs} > x_{h,jobs}$ for $h = LC$ and $h = CP$. In addition, roads satisfies the minimum attribute criteria: $x_{M,roads} < x_{h,roads}$ for $h = LC$ and $h = CP$.

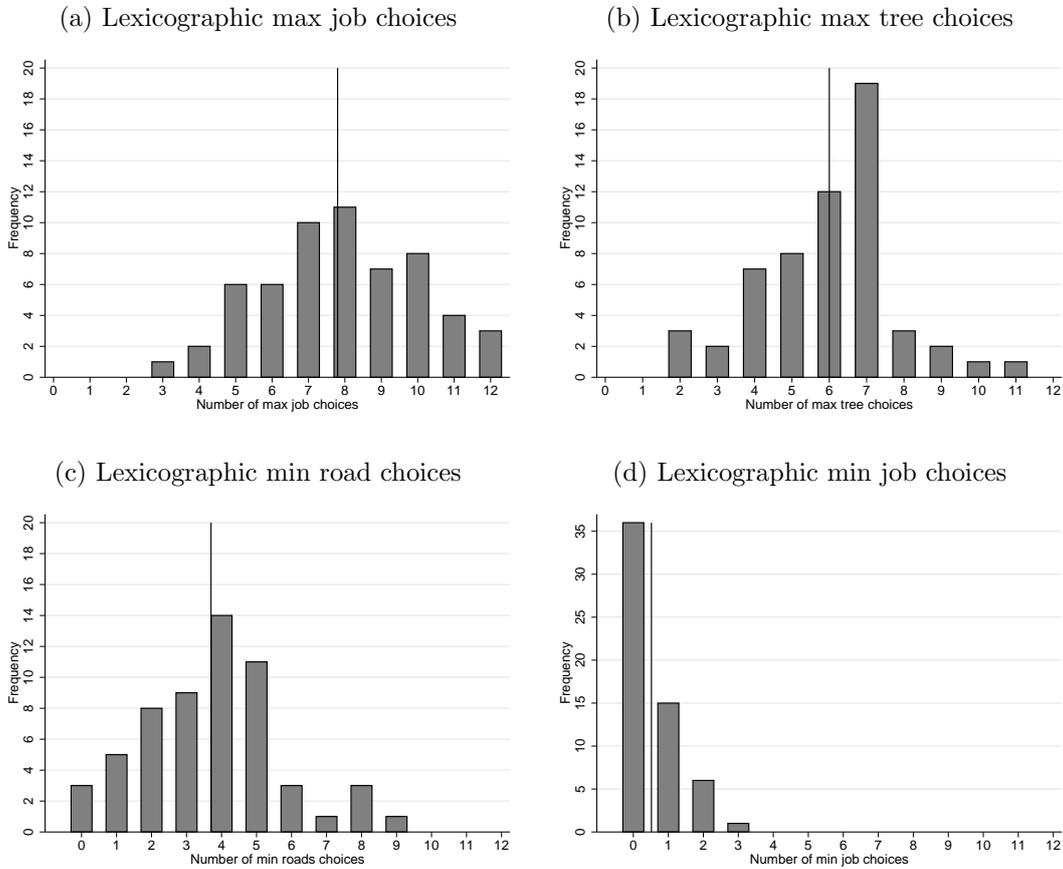
For SBR, the attribute trees satisfies the maximum attribute criteria while jobs satisfies the minimum attribute criteria. The Community Park has the highest provision of trees relative to the Leisure Centre or Market, such that $x_{CP,trees} > x_{h,trees}$ for $h = LC$

¹⁷We note that given that these samples sizes are small, with 30 in the control, 39 in Treatment(0,4,5) and 61 in Treatment(0,3,6), only very large effects are likely to be statistically significant. A larger sample would potentially invalidate this finding.

and $h = M$. The provision of trees is 7, 9 and 20 for the Leisure Centre, Market and Community Park, respectively. In addition, jobs satisfies the minimum attribute criteria: $x_{CP,jobs} < x_{h,roads}$ for $h = LC$ and $h = M$.

Across all FBR, we plot the number of times they used a maximum jobs, maximum trees, minimum roads and minimum jobs decision rule out of the twelve choice sets.¹⁸ This is presented in Figure 1.

Figure 1: Lexicographic choices by attribute – First best respondents (average out of 12 is indicated)

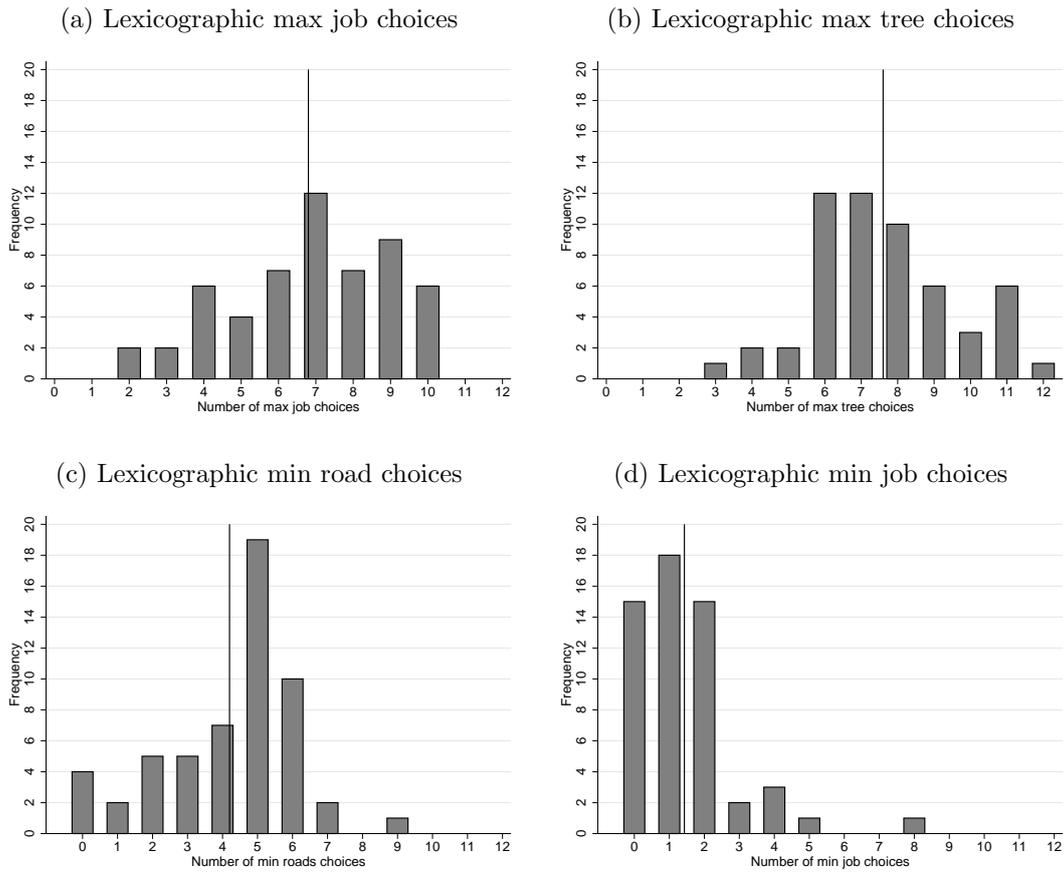


The mean number of choices characterised by the maximum jobs, maximum trees, minimum roads and minimum job decision rule is 7.8, 6, 3.7 and 0.5 respectively. Using a Mann Whitney U test for equality of the distributions, we reject that the distributions are the same between maximum jobs (7.8) and maximum trees (6) choices (MW p-

¹⁸Appendix B outlines the number of times out of twelve that a choice set includes a choice set where the lexicographic strategies overlap.

value 0.000) and between maximum jobs (7.8) and minimum road (3.7) or minimum job (0.5) choices (MW p-value 0.000, 0.000) This suggests that FBR are using a maximum lexicographic strategy on the attribute jobs more than one on roads or trees. Figure 2 illustrates the lexicographic choices for each attribute across SBR.

Figure 2: Lexicographic choices by attribute – Second best respondents (average out of 12 is indicated)



The mean number of lexicographic choices characterised by maximum jobs, maximum trees, minimum jobs and minimum roads is 6.8, 7.6, 4.2 and 1.4 out of 12, respectively. We reject the hypothesis that the distributions are the same for maximum trees (7.6) and maximum jobs (6.8) as well as for maximum trees (7.6) and minimum jobs (4.2) (MW p-value 0.000, 0.000). We also reject that the distributions are the same for maximum trees (7.6) and minimum roads (1.4) (MW p-value 0.000). Targeting the second best provision outcome suggests that respondents choose lexicographically on the attribute trees, more than on the attribute jobs or roads. This further suggests that respondents

are using decision rules that are a direct result of the composition of the final alternatives.

Having concluded that SBR are implementing a decision rule on trees more often than any other strategy, we question whether SBR are more likely to use the maximum tree decision rule than FBR. Using the Fishers Exact Test for unmatched data, we reject equality and find that SBR are using a decision rule on trees (7.6) more than FBR (6) (FE p-value 0.000). Similarly, FBR are using a maximum job decision rule (7.8) more than SBR (4.2) (FE p-value 0.026). This further confirms that choice strategies are consistent with the provision contenders and are significantly different across respondents with different provision objectives.

4.4 Hypothesis 4

Hypothesis 4 concerns whether the predicted choice probabilities for the provision contenders are in line with the stated objectives of respondents from whose choices the probabilities are estimated. We first estimate a conditional logit model to obtain the parameter estimates β for the three attributes using all the DCE responses. We run a conditional logit model for the aggregate sample using all 130 responses and collect the marginal utilities for the attributes β_A , as well as for each sub-sample of respondents based on their stated provision objective. Respondents are grouped solely by their stated objective b_k and therefore we estimate marginal utilities for each group of respondents targeting the Market (β_M), Community Park (β_{CP}), Leisure Centre (β_{LC}) and those unsure (β_U). These marginal utilities can be seen in Appendix B. After estimating the marginal utilities, we calculate the logit choice probabilities for the three provision contenders. By doing this we can examine the indirect effect choice strategies have on the predicted provision outcome. Table IV outlines the choice probabilities for the aggregate sample as well as for each sub-sample of individuals based on their stated objective.¹⁹

Insert Table IV Here

The results in Table IV show the predicted outcome given the responses $c_{k,n}$ in the DCE. Under a plurality provision rule, both FBR and SBR have delivered on their stated

¹⁹Appendix B shows a simulated outcome of the choice probabilities if respondents were to choose randomly within the DCE.

task. Respondents targeting the Market (β_M), have a probability of preferring the Market by 47%. Respondents targeting the Community Park (β_{CP}) have a probability of preferring the Community Park with 56% probability. Based on their stated provision objective and their choice decision rules, both FBR and SBR would receive the desired outcome under a plurality provision rule. Overall, 45% of respondents targeted the first best outcome, while 42% indicated targeting the second best outcome. Despite more respondents preferring the first best outcome, using the aggregate sample (β_A) we find a predicted preference of the Market to be 39% while the Community Park is most preferred with 43%. Under a plurality decision rule, the second best outcome would be provided. This suggests that strategic bias could lead to a provision outcome that differs from the unconditionally preferred outcome.

5 Conclusion

The experimental results reported here indicate that significant numbers of respondents are behaving strategically. Despite the complexity required to translate targeting one's second best outcome into an attribute based decision rule, we find that many respondents are able to do so. They comprehend the payoff scheme, understand the information about the relative likelihoods associated with the provision outcomes and translate that information into an attribute based decision rules which differs from those they would employ to target their first best outcome. This would be much simpler for them if the choice sets in the DCE were labelled. In our unlabelled experiment the work involved to strategically bias results was considerably more demanding.

Of the one hundred individuals in the treatment groups, we categorize 27% as acting strategically, such that they report considering the final provision outcome probabilities and resolve to switch to target their second best outcome. We find that even though respondents are using lexicographic decision rules in less than 100% of the choice sets, the choice probabilities estimated on their data reflect their stated provision target.

The provision outcome probabilities were constant in both treatment groups. They

were designed to convey that respondents' first best outcome was unlikely to occur and that the second best might be delivered— it was a serious contender. This combination incentivised respondents to switch choice behaviour to target their second best outcome. The belief that their choices can have an effect on the provision outcome is a probable requirement before a person will exert the mental effort to adjust their choice behaviour. Additional requirements, identified at the outset of the paper, are that the respondent believes the DCE will affect which outcome will be provided (over a discrete set of alternatives) and that the respondent has expectations about the relative likelihood of those provision options being selected. We note that in field choice experiments, the substantive case study may not be one where the likelihood of different provision outcomes is well known by communities or is clearly defined. A better understanding of the conditions under which people will (not) make the effort to adjust their choice behaviour is warranted; for example by varying the priors the respondents have about the probabilities associated with each of the provision outcomes or by introducing uncertainty on the actual provision outcomes. Respondents may use the DCE in order to affect provision outcomes they believe are under consideration, but the true set may differ. In this case, the DCE would still be susceptible to strategic bias and yield unreliable estimations of preferences. As strategic misrepresentation is a direct result from unfavourable priors regarding the preferences of others, field studies should be aware of how opinion polls on current issues may bias respondent choices and whether information can be given prior to the task to combat those priors to elicit truthful responses.

This study shows that strategic bias is a potential issue in choice experiments. However, we note that the task required a significant degree of cognitive effort which not all respondents chose to exert, either because they were not bothered or because they did not understand. As the sample was university staff and students, we anticipate the former as an explanation. The cognitive effort required to deviate from one's true preference under the complexity of choice experiments may be an issue in the real world. It would depend on how much a respondent cares about the outcome, the work involved, and their perception regarding the effect their effort would have on the overall outcome. Further research

may extrapolate the experimental findings to the field. Additional research on complexity of the choice experiment would advise whether design complexity (i.e. increased attributes, choice sets, alternatives per choice set, etc.) can address this behavioural inefficiency without compromising statistical efficiency. Additionally, our study eliminated the burden of cost; research which reintroduces the cost attribute would further shed light on respondents' willingness to strategically bias and its effect on value estimates.

The implications of this research are that individuals may not respond truthfully in choice experiments if they (i) envision the choice experiment as potentially affecting the provision outcome; (ii) have expectations of the final provision options; and, (iii) have an expectation of the likelihood of the provision outcomes being implemented. Carson and Groves (2007) indicate that a DCE must be considered a consequential mechanism by respondents if the results are to be used in economic analysis. The results reported here suggest that under certain, plausible conditions, the strategic misrepresentation of preferences is a direct consequence of that consequentiality requirement.

References

- Bennett, J. and R. Blamey (2001). *The choice modelling approach to environmental valuation*. Edward Elgar Publishing.
- Blais, A. and R. Nadeau (1996). Measuring strategic voting: A two-step procedure. *Electoral Studies* 15(1), 39–52.
- Blamey, R. (1998). Contingent valuation and the activation of environmental norms. *Ecological Economics* 24(1), 47–72.
- Bohm, P. (1972). Estimating demand for public goods: An experiment. *European Economic Review* 3(2), 111–130.
- Burton, M. (2010). Inducing strategic bias: and its implications for choice modelling design. *Working Paper*.
- Carson, K., S. M. Chilton, W. G. Hutchinson, and R. Scarpa (2015). Do choice experiments generate reliable preference estimates? *Working Paper*.
- Carson, R. and T. Groves (2007). Incentive and informational properties of preference questions. *Environmental and Resource Economics* 37(1), 181–210.
- Carson, R. T., T. Groves, and J. A. List (2014). Consequentiality: A theoretical and experimental exploration of a single binary choice. *Journal of the Association of Environmental and Resource Economists* 1(1/2), 171–207.
- Collins, J. P. and C. A. Vossler (2009). Incentive compatibility tests of choice experiment value elicitation questions. *Journal of Environmental Economics and Management* 58(2), 226–235.
- Gunther, A. (1998). The Persuasive Press Inference. *Communication Research* 25(5), 486–504.

- Hanley, N., S. Mourato, and R. E. Wright (2001). Choice modelling approaches: A superior alternative for environmental valuation? *Journal of Economic Surveys* 15(3), 435–462.
- Interis, M., C. Xu, D. Petrolia, and K. Coatney (2016). Examining unconditional preference revelation in choice experiments: a voting game approach. *Journal of Environmental Economics and Policy* 5(1), 125–142.
- Lancaster, K. J. (1966). A new approach to consumer theory. *The Journal of Political Economy*, 132–157.
- Luchini, S. and V. Watson (2014). Are choice experiments reliable? evidence from the lab. *Economics Letters* 124(1), 9–13.
- McFadden, D. (1974). Conditional logit analysis of qualitative choices. *Zarembka, P.(eds): Frontiers*.
- Mousavi, S., G. Gigerenzer, and R. Kheirandish (2016). Rethinking behavioral economics through fast-and-frugal heuristics. *Routledge Handbook of Behavioral Economics*, 280–296.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Scheufele, G. and J. Bennett (2012). Response strategies and learning in discrete choice experiments. *Environmental and Resource Economics* 52(3), 435–453.
- Taylor, L. O., M. McKee, S. K. Laury, and R. G. Cummings (2001). Induced-value tests of the referendum voting mechanism. *Economics Letters* 71(1), 61–65.
- Taylor, L. O., M. D. Morrison, and K. J. Boyle (2010). Exchange rules and the incentive compatibility of choice experiments. *Environmental and Resource Economics* 47(2), 197–220.

- Vossler, C. A., M. Doyon, and D. Rondeau (2012). Truth in consequentiality: theory and field evidence on discrete choice experiments. *American Economic Journal: Microeconomics* 4(4), 145–171.
- Vossler, C. A. and M. F. Evans (2009). Bridging the gap between the field and the lab: Environmental goods, policy maker input, and consequentiality. *Journal of Environmental Economics and Management* 58(3), 338–345.
- Vossler, C. A. and S. B. Watson (2013). Understanding the consequences of consequentiality: Testing the validity of stated preferences in the field. *Journal of Economic Behavior and Organization* 86, 137–147.
- Wheeler, S. and R. Damania (2001). Valuing New Zealand recreational fishing and an assessment of the validity of the contingent valuation estimates. *Australian Journal of Agricultural and Resource Economics* 45(4), 599–621.
- Whittington, D., J. Briscoe, X. Mu, and W. Barron (1990). Estimating the willingness to pay for water services in developing countries: A case study of the use of contingent valuation surveys in Southern Haiti. *Economic Development and Cultural Change*, 293–311.

Table I: Final provision outcome profiles

Attribute	Provision Outcomes		
	b_1	b_2	b_3
x_1	x_{11}	x_{21}	x_{31}
x_2	x_{12}	x_{22}	x_{32}
x_3	x_{13}	x_{23}	x_{33}
Value	$V_{k=1,n}$	$V_{k=2,n}$	$V_{k=3,n}$

Table II: Final provision outcome profiles and induced values

Attribute	Contender		
	Leisure Centre (LC)	Community Park (CP)	Market (M)
Roads	4	3	2
Jobs	5	4	8
Trees	7	20	9
Value			
Control	£0	£4	£5
Treatment(0,4,5)	£0	£4	£5
Treatment(0,3,6)	£0	£3	£6
Probability	50%	45%	5%

Note: Treatment(0,4,5) and Treatment (0,3,6) refer to the two treatments with pay off levels for the three provision options of £0, £4, £5 and £0, £3, £6 respectively.

Table III: Self reported provision target by treatment

Group	Considered Information	provision Target				Total
		Market	Community Park	Leisure Centre	I Don't Know	
Control	-	17 (57%)	8 (27%)	2 (6%)	3 (10%)	30 (100%)
Treatment(0,4,5)	Yes	5 (26%)	12 (63%)	1 (5 %)	1 (5%)	19 (49%)
	No	11 (64%)	4 (24%)	1 (6%)	1 (6%)	17(44%)
	I don't know	1 (33.3%)	1 (33.3%)	0 (0%)	1 (33.3%)	3 (8%)
	Total	17 (44%)	17 (44%)	2 (5%)	3 (7%)	39 (100%)
Treatment(0,3,6)	Yes	4 (19%)	15(71%)	1(5%)	1(5%)	21 (34%)
	No	17 (57%)	10 (33%)	2 (7%)	1 (3%)	30 (49%)
	I don't know	3 (30%)	5 (50%)	0 (0%)	2 (20%)	10 (16.3%)
	Total	24 (39%)	30 (49%)	3 (5%)	4 (7%)	61(100%)
	Total	58(45%)	55 (42%)	7 (5%)	10 (8%)	130 (100%)

Note: Percentages may not add to 100 due to rounding. Treatment(0,4,5) and Treatment (0,3,6) refer to the two treatments with pay off levels for the three provision options of £0, £4, £5 and £0, £3, £6 respectively.

Table IV: Predicted choice probabilities for the final provision outcomes

	Provision Outcomes			N^o indiv.
	Leisure Centre	Community Park	Market	
$\Pr(b_k \beta_M)$	20.3%	32.1%	47.5%	58
$\Pr(b_k \beta_{CP})$	12.4%	56.0%	31.5%	55
$\Pr(b_k \beta_{LC})$	29.9%	37.4%	32.6%	7
$\Pr(b_k \beta_U)$	23.4%	32.0%	44.5%	10
$\Pr(b_k \beta_A)$	18.0%	42.2%	39.7%	130

Appendices

A Strategic Simulation Exercise

Consider two types of simulated agents, Type I and Type II, who have heterogeneous preferences.

$$U_I = -20 * x_1 - 13.32 * x_2 + 33.33 * x_3 + 46.66 * x_4 - \epsilon_n$$

$$U_{II} = -20 * x_1 + 8 * x_2 - 20 * x_3 + 20 * x_4 - \epsilon_n$$

Where ϵ is a Gumbel distributed error term. Given these utility functions, respondents have different preference ordering for three provision outcomes b_j , b_k and b_m shown in Table A.1.

Insert Table A.1 Here

For the outcomes in Table A.1, Type I agents have preference ordering $b_j > b_k > b_m$ while Type II agents have preference ordering $b_m > b_k > b_j$. Agents are presented with a choice experiment, regarding attributes three discrete variables x_1 , x_2 , x_3 and one dummy variable x_4 . The design was created in Sawtooth Software using balanced overlap. The DCE will ask agents to make trade offs based on attributes x_i and will be used to elicit the preference ordering of the final provision. We consider a sample of 500 agents made up of 75% Type I agents and 25% Type II agents.

Using Monte Carlo simulations, we simulate choices to the DCE under true preference revelation based on the utility functions, a sample of 500 agents and 500 replications. Using the average marginal utilities found from these simulations, we predict the choice probabilities for each provision contender. When all agents respond to the DCE truthfully, the predicted preference ordering is $Pr(b_j) > Pr(b_k) > Pr(b_m)$. Expecting an unfavourable outcome, we turn to simulate when Type II respondents choose to target their second best provision outcome, b_k . We assume a lexicographic decision rule on the attribute x_2 is implemented by strategic respondents, as it satisfies the minimum/maximum attribute rule of Hypothesis 3. The predicted choice probabilities for the three provision contenders is shown in Table A.2.

Insert Table A.2 Here

Again, under 0% strategic respondents, we have a predicted probability that b_j is most preferred. In this case, Type II respondents switch to target their second best outcome b_k and choose lexicographically on the attribute x_2 . As the sample of strategic respondents increases, the predicted probability of the preferred outcomes change. At 10% strategic respondents, the largest predicted probability is now for b_k . Under a plurality decision rule, Type II respondents have switched the provision outcome.

Table A.1 provision Contenders

Attribute	Provision outcomes		
	b_j	b_k	b_m
x_1	13	5	8
x_2	5	11	8
x_3	15	10	4
x_4	1	0	1

Table A.2 Simulated outcomes

Percent Strategic	Predicted Probability		
	b_j	b_k	b_m
0%	0.54	0.39	0.07
5%	0.48	0.44	0.08
10%	0.43	0.47	0.09
25%	0.37	0.54	0.09

B

Table B.1 Number of choice sets with lexicographic overlaps (out of 12)

	Block				
	1	2	3	4	5
max jobs= max trees	3	6	4	7	7
max jobs = min roads	5	5	5	4	2
max jobs = min jobs	0	0	0	0	0
max trees = min roads	6	5	5	4	4
max trees = min jobs	7	4	2	6	5

Table B.2 Parameter estimates

	Provision Outcome Target				
	Aggregate (1)	Market (2)	Community Park (3)	Leisure Centre (4)	Unsure (5)
Roads	-0.0033 (0.0199)	0.0337 (0.0307)	-0.0557 (0.0325)	0.160* (0.0779)	-0.0397 (0.0663)
Jobs	0.207*** (0.0103)	0.268*** (0.0172)	0.189*** (0.0160)	0.111** (0.0346)	0.165*** (0.0322)
Trees	0.0813*** (0.0054)	0.0586*** (0.0082)	0.126*** (0.0093)	0.0382 (0.0205)	0.0339* (0.0168)
Observations	4680	2088	1980	252	360
<i>N</i> ^o indiv.	130	58	55	7	10

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.3 Provision outcomes under randomised selection in the DCE

	Leisure Centre	Community Park	Market
Roads	3	2	2
Jobs	4	5	6
Trees	7	20	10
Choice Probabilities	0.336	0.328	0.335

C Experiment Instructions

The following instructions show the pay offs for the control and treatment groups such that first and second best outcomes correspond to a pay off of £5 and £4, respectively. Respondents in Treatment(0,3,6) would have seen pay offs where the first best value is £6 and the second best is £3.

The Choice Game

Thank you for playing The Choice Game.

You will receive £3 for your participation.

You can receive an additional amount between £0 and £5, depending on the outcome of the game, which will be described shortly.

Imagine that you are a citizen of Toonsville.

The local authority is going to develop a plot of land.

There are three options for the development: a leisure centre, a community park, or a market.

Each of the three possible options differ in terms of the amount of roads, jobs and trees they will provide, as shown here:

Leisure Centre	Community Park	Market
4 Roads	3 Roads	2 Roads
5 Jobs	4 Jobs	8 Jobs
7 Trees	20 Trees	9 Trees

The local authorities decide to run a survey to inform their planning decision.

In that survey people are presented with 12 questions that look like this:

	Option 1	Option 2	Option 3
Roads 	3 Roads	0 Roads	1 Roads
Jobs 	11 Jobs	13 Jobs	5 Jobs
Trees 	5 Trees	0 Trees	10 Trees
Make your choice from the 3 options:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

↓ ↓ ↓

Respondents are asked to select their most preferred option in each choice set.

In each question, the number of roads, jobs and trees will change.

The local authority will:

- Analyse people’s choices to determine how much they like or dislike roads, jobs and trees.
- Use the responses to determine which of the three policies is most preferred in the community.

Your Task

Your task in this game is to complete the survey for the local authorities.

You will make a series of choices out of three options, as shown in the example question.

The local authorities will use your responses to determine whether the leisure centre, community park. or market will be built.

We (i.e. the local authority) will analyse your choices and determine which project will go ahead.

Strategic Bias in Discrete Choice Experiments

You will be paid £3 simply for playing.

You will be paid an additional amount according to which project is found to be most preferred based on your survey responses in this way:

- If the market is most preferred you will receive £5
- If the community park is most preferred you will receive £4.
- If the leisure centre is most preferred you will receive £0.

	Leisure Centre	Community Park	Market
	4 Roads	3 Roads	2 Roads
	5 Jobs	4 Jobs	8 Jobs
	7 Trees	20 Trees	9 Trees
Your £ pay-out	£0	£4	£5

We are now ready for you to make your choices from the set comprising 3 options.

We will remind you when you make each choice of:

1. The three options the local authority is considering
2. The payment to you depending on which one is chosen by the local authority, using this picture:

	Leisure Centre	Community Park	Market
	4 Roads	3 Roads	2 Roads
	5 Jobs	4 Jobs	8 Jobs
	7 Trees	20 Trees	9 Trees
Your £ pay-out	£0	£4	£5

(Begin choice sets)

If these were your only options, which would you choose?
Choose by clicking one of the buttons below:

(1 of 12)

	Option 1	Option 2	Option 3
 Roads	4 Roads	5 Roads	0 Roads
 Jobs	13 Jobs	9 Jobs	0 Jobs
 Trees	15 Trees	20 Trees	5 Trees
Make your choice from the 3 options:	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>

Reminder:

The local authority will use your answers to determine which of these three options is implemented.

	Leisure Centre	Community Park	Market
4 Roads	3 Roads	2 Roads	
5 Jobs	4 Jobs	8 Jobs	
7 Trees	20 Trees	9 Trees	
Your £ pay-out	£0	£4	£5

In the treatment groups, prior to answering the twelve choice sets, respondents additionally saw:

News just in...

An opinion poll appears in the local paper.

It shows that, among the other community leaders:

- 50% prefer the leisure centre
- 45% prefer the community park
- 5% prefer the market

Strategic Bias in Discrete Choice Experiments

	Leisure Centre	Community Park	Market
	4 Roads	3 Roads	2 Roads
	5 Jobs	4 Jobs	8 Jobs
	7 Trees	20 Trees	9 Trees
Probability of being chosen	50%	45%	5%
Your £ pay-out	£0	£4	£5

We will remind you when you make each choice of:

1. The three options the local authority is considering
2. The popularity of each option among the other community leaders
3. The payment to you depending on which one is chosen by the local authority, using this picture:

	Leisure Centre	Community Park	Market
	4 Roads	3 Roads	2 Roads
	5 Jobs	4 Jobs	8 Jobs
	7 Trees	20 Trees	9 Trees
Probability of being chosen	50%	45%	5%
Your £ pay-out	£0	£4	£5

(Begin choice sets)

Strategic Bias in Discrete Choice Experiments

If these were your only options, which would you choose?
Choose by clicking one of the buttons below:

(1 of 12)

	Option 1	Option 2	Option 3
 Roads	4 Roads	5 Roads	0 Roads
 Jobs	13 Jobs	9 Jobs	0 Jobs
 Trees	15 Trees	20 Trees	5 Trees
Make your choice from the 3 options:	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>

Reminder:

The local authority will use your answers to determine which of these three options is implemented.

	Leisure Centre	Community Park	Market
	4 Roads	3 Roads	2 Roads
	5 Jobs	4 Jobs	8 Jobs
	7 Trees	20 Trees	9 Trees
Probability of being chosen	50%	45%	5%
Your £ pay-out	£0	£3	£6

Thank you for completing this questionnaire.

We will inform you which project was most preferred shortly.

First, can you please answer the following questions to the best of your ability.

1. Can you please indicate which project were you trying to increase influence over by your choices in the survey?

- Market
- Community Park
- Leisure Centre
- None of the above. I was confused

2. Briefly describe the strategy you took in order to influence the local authority to increase the popularity of ‘answer from question 1’.

3. How difficult did you find the task?

- Very difficult
 - Difficult
 - Neutral
 - Easy
 - Very easy
-