



UNIVERSITY OF WARSAW

Faculty of Economic Sciences

WORKING PAPERS

No. 20/2016(211)

MIKOŁAJ CZAJKOWSKI

NICK HANLEY

JACOB LARIVIERE

WILLIAM S. NEILSON

KATHERINE SIMPSON

INFORMATION AND LEARNING IN STATED-PREFERENCE STUDIES

WARSAW 2016



Information and Learning in Stated-Preference Studies

MIKOŁAJ CZAJKOWSKI

Faculty of Economics
University of Warsaw
e-mail: miq@wne.uw.edu.pl

NICK HANLEY

Department of Geography and Sustainable
Development, University of St. Andrews
e-mail: ndh3@st-andrews.ac.uk

JACOB LARIVIERE

Department of Economics
University of Tennessee
e-mail: jlarivi1@utk.edu

WILLIAM S. NEILSON

Department of Economics
University of Tennessee
e-mail: wneilson@utk.edu

KATHERINE SIMPSON

Economics Division
University of Stirling
e-mail: k.h.simpson@stir.ac.uk

Abstract

We use experimental variation to influence how people learn a given amount of objective, scientific information about an unfamiliar public good. We then estimate the impact of treatment on valuations for that good in a stated preference survey. Our main treatment, a pre-survey multiple choice quiz about objective public good attributes, increased learning rates by over 60%. We find that despite increasing learning and retention rates, treatment had no statistically significant impact on mean nor variance of the distribution of valuations. We show with a very simple theoretical model this result is consistent with a model of confirmatory bias used by agents in stated preference surveys and inconsistent with other models of preference formation.

Keywords:

Information, Updating, Preferences, Public Goods

JEL:

D01; D83; Q41

Acknowledgements:

We thank Scottish Natural Heritage and the Scottish Environmental Protection Agency for funding part of this work, along with the Marine Alliance Science and Technology Scotland. The second author gratefully acknowledges the support of the Polish Ministry of Science and Higher Education and the Foundation for Polish Science.

Working Papers contain preliminary research results.

Please consider this when citing the paper.

Please contact the authors to give comments or to obtain revised version.

Any mistakes and the views expressed herein are solely those of the authors.

1 Introduction

Stated-preference studies have become an established tool for assessing values for non-market goods, both for weighing the cost and benefits for new public projects and for assessing damages when public goods are harmed. In a well-designed stated-preference study, the researcher provides information about the public good, project, or resource under consideration, establishes consequentiality for the decision including identifying the likely source of revenue to pay for the project, and uses an incentive-compatible method for eliciting the subject's willingness to pay (WTP) for the good (Hanley, Shogren, and White (2013)). Because of their importance, stated-preference techniques have been the subject of numerous academic studies, and most of this literature has devoted attention to the elicitation of subjects' WTP values.

The information-presentation stage of state preference studies has also received attention. The basic rationale for providing subjects with information about the good or project under consideration is straightforward: providing subjects with information reduces noise in the WTP data by pinning down more precisely what subjects have been asked to value, thereby ensuring that they all assign values to the same collection of characteristics. In addition, providing information is also thought to change preferences and WTP. Presenting information mimics the situation people would be in if there was a referendum on a public good issue - even if they do not know much about it, they could be expected to learn something before the referendum, e.g. from public debate on the issue (Mitchell and Carson (2013)). Two classes of studies in this area examine the framing of information (Smith and Johnson (1988), Smith, Desvousges, Fisher, and Johnson (1988), Smith and Desvousges (1990), and Czajkowski, Hanley, and LaRiviere (2015)) and how familiarity with a good affects WTP (Cameron and Englin (1997) and Czajkowski, Hanley, and LaR-

iviere (2014)).¹ The goal of most of these studies is to identify normative criteria about best practices during the information stage of a stated preference study.

This paper tests the above rationale for including objective, easily understood information explicitly given that it already satisfies best practices.² We begin by introducing a simple theoretical setting in which consumers have heterogeneous but unbiased priors about the levels of the attributes embodied in the good, and providing information to consumers reduces the variance of their WTP values but leaves the mean unchanged. This model, then, provides the underpinnings of the noise-reduction rationale for sharing information at the beginning of a stated-preference study. Because the mean WTP value remains unchanged, though, the model cannot explain informative advertising. We modify the model using the attribute-unawareness framework of Schwartzstein (2014) and show that in the modified setting increased information raises the mean WTP.³ Another modification to the unbiased-priors model allows it to capture confirmatory bias (Rabin and Schrag (1999)) in which agents confronted with new information misinterpret it as supporting their previously-held beliefs. In our model, confirmatory bias takes the form of interpreting new information as confirming prior beliefs about the levels of different attributes embodied in the good.

The three models lead to testable predictions: in the unbiased-priors model information leaves mean WTP unchanged but reduces the variance; the attribute-unawareness model increases the mean and either reduces the variance or leaves it unchanged; and the

¹Huffman, Rousu, Shogren, and Tegene (2007) finds that uninformed subjects are more likely to be affected by product information, a result we discuss below.

²Recent evidence shows, though, that economic agents can learn information about product characteristics, or meaningful economic decisions in general, but they do so imperfectly perhaps due to cognitive load or lack of attention (Gabaix, Laibson, Moloche, and Weinberg (2006), Schwartzstein (2014), Caplin and Dean (2015) and LaRiviere, Czakowski, Hanley, and Simpson (2015)).

³Hanna, Mullainathan, and Schwartzstein (2014) introduces a similar model in the context of firms.

confirmatory-bias model leaves both the mean and variance unchanged. Importantly, the ensuing patterns in the data are mutually exclusive, allowing one to test between the three models of information use.

We then present results from novel experimental variation embedded in a stated preference survey in which treatment varies the rate at which consumers learn a given amount of neutral, scientifically verifiable information about the attributes of the public good. Next we investigate the effect of increased learning on the distribution of WTP values in a stated preference survey. To do so, we leverage recent evidence that stated-preference surveys common in environmental economics lead consumers to truthfully reveal their valuation for a good if the results of the survey will be used to influence policy and are therefore consequential (Vossler, Doyon, and Rondeau (2012) and Carson, Groves, and List (2014)). Our survey is consequential: we value a coastal community's willingness to pay to replace the familiar seawall approach to preventing floods by reclaiming wetlands.⁴ Because seawalls are already prevalent, their benefits are fairly well-known. Wetland restoration is new to the area and can generate additional benefits besides flood control, such as increased bird habitats and wildlife. Consumers might not understand these features of restored wetlands, and so might not know how to properly value the replacement of seawalls with wetlands. As with all stated-preference surveys we give all subjects a set of information about the product being valued and the policy change considered, in our case flood protection and restored wetlands.⁵

The key experimental variation we provide is giving treated subjects a short multiple

⁴Because our survey is consequential, there is an argument for calling this a field experiment. The natural activity is a stated preference survey which has direct policy implications which affect both flood mitigation strategies and taxation levels.

⁵Because we are only considering a single well-defined policy, we cannot use a choice experiment to elicit valuations. We discuss this in detail below.

choice quiz about historical flood protection strategies and objective information about wetland restoration at the beginning of the survey. We then provide all subjects identical information about the historical flood protection and attributes of restored wetlands. Importantly, each piece of information we provided subjects during the information stage of the survey corresponds to a single question on the quiz. After eliciting WTP, we give all subjects the exact same multiple choice quiz given to treated subjects before the survey. If treated subjects score higher on the post survey quiz, then we can plausibly argue that treatment provided exogenous variation subjects' information retention/learning rates.

We are interested in two key questions: whether treatment affects learning and whether treatment affects the mean and variance of WTP. If the answer to the first question is affirmative then our design creates exogenous variation in learning rates. The second question relates to whether the noise-reduction rationale governs the provision of information in stated-preference studies, or if providing information does something else.

We find that treatment caused a significant 62% increase in the learning rate.⁶ Differences in quiz scores, and therefore differences in learning rates, are attributable to the causal impact of the quiz since treatment was randomized. The learning effect is significant and very robust.

Having established that treating with a quiz affects learning, we can then attribute effects of the treatment on WTP to learning. The results are very robust to functional forms and inclusion of sociodemographic controls, and they show that learning has no causal effect on the variance of WTP values. Thus, our data are inconsistent with the

⁶There are three plausible explanations: the initial quiz attunes treated subjects to paying attention to information in a more helpful way, treated subjects see similar information more than once and therefore have more opportunities to learn, or that an initial quiz causes subjects to take the entire informational portion of the survey more seriously.

unbiased-prior model and the noise-reduction theory. The data also show no causal effect of learning on the mean of WTP values, making them inconsistent with the attribute-unawareness model underlying the informative-advertising explanation. The model most consistent with the data is confirmatory bias, with consumers absorbing the new information but not using it to update their values of the good.

It is somewhat surprising that increased information retention rates have a zero effect on both the mean and variance of the WTP distribution. The variance effect in particular is striking because treated subjects had more knowledge about product attributes than control subjects. In that sense treated subjects were better-informed, but despite this valuations were just as dispersed as for the less-informed control group. This result is different than Huffman, Rousu, Shogren, and Tegene (2007) which finds that information provided to some types of subjects in a stated preference survey can influence valuations. Unlike that study, we have experimental variation on learning testing exclusively for the effect of objective- rather than subjective third party information- on valuations.

Our data suggest that rather than reducing heterogeneity across subjects or leading them to have new assessments, the new information served instead simply to justify their prior assessments. Our results, then, open questions regarding whether our findings extend to all stated-preference surveys, or if they depend on the types of information provided, on the nature of the good being studied, or on something else entirely. As with issues of consequentiality and incentive-compatibility, the information-provision stage of stated-preference elicitation deserves more exploration.

Our results indicate that increased retention of neutral, scientific information related to attributes of a public good do not impact willingness to pay for that good. There is a related literature which evaluates how information is conveyed (e.g., qualitatively or

quantitatively) affects risk attitudes, risk mitigation behavior and WTP for expert advice about risk levels (Smith and Johnson (1988), Smith, Desvousges, Fisher, and Johnson (1988), and Smith and Desvousges (1990)). Other studies find that information provided by competing and opposed stakeholders during a survey can affect WTP (Hanley, Czajkowski, Hanley-Nickolls, and Redpath (2010) and Czajkowski, Hanley, and LaRiviere (2015)). This literature is similar in that we find different ways of conveying information affects learning. However, this paper is also fundamentally different: we condition on a set of information constructed to make the survey as objective as possible, including verifying the language and content of our information with third party experts. Our treatment then provides exogenous variation in how much learning of that information occurs, rather than varying the framing, content and type of information like the previous literature. As a result, we condition on a set of information and vary learning over it rather than varying the information itself. We are not aware of any study which asks this basic question in the literature.

The remainder of the paper is organized as follows: Section 2 presents the theoretical framework that leads to the hypotheses on WTP distributions. Section 3 describes the experiment and addresses issues raised by an unbalanced sample. Section 4 presents the results on learning, and Section 5 presents the main results for the mean and variance of WTP. The final section offers discussion and concluding remarks.

2 Theoretical Framework

Our model uses a setting with a multi-attribute good in the spirit of Lancaster (1966). To allow for learning, the model assumes that consumers may not know the amounts of attributes embodied in the good. This leaves two possible approaches, and rather than model

the uncertainty at the individual level using probability distributions over the possible levels of the attributes, we model uncertainty at the aggregate level. More specifically, we assume identical consumers who have heterogeneous, point-valued assessments of the level of each attribute, thereby capturing uncertainty through heterogeneity of point estimates across consumers rather than uncertainty of point estimates for each individual.

A benefit of this approach is that it provides a clear role for new information. On the individual level it can change the point estimates of attribute levels away from the initial assessments, and on the aggregate level it can change the heterogeneity of the point estimates. The former effect can change the mean of the ensuing WTP measure, and the latter can change the variance. Importantly, both the impact of learning and all of the heterogeneity in the model arise from individual differences in beliefs and not differences in preferences across consumers.

2.1 Baseline Framework

To make all of this precise, begin with a representative consumer and assume the good has n attributes, $1, \dots, n$. Conditional on a set of priors, the baseline value of the good is V_0 and it comes from the individual assigning priors to the levels of the attributes. Let the initial attribute level beliefs be $\hat{x}_1, \dots, \hat{x}_n$. Let $v_i(\cdot)$ be a function which maps levels of attribute i to valuations such that the baseline valuation for the good by agent k is:

$$V_0 = \sum_{i=1}^n v_i(\hat{x}_i).$$

This representation of WTP is consistent with quasilinear utility in the public good

and that there is a zero WTP for a public good containing no attributes.⁷ Quasilinearity is a non-trivial assumption over the marginal utility of income with respect to the good. The point of this model, though, is to focus on the dynamics of how learning affects WTP making the assumption reasonable in our context.

There are K individuals in the sample, each with their own initial assessments. Let $\hat{x}_i^k$ denote the baseline level of attribute i according to agent k 's prior. Conditional on agent specific underlying functions v_i for attributes $i = 1, \dots, n$, average initial WTP is:

$$\mu_0 = \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n v_i(\hat{x}_i^k). \quad (1)$$

Learning takes a strong form. All subjects who learn about attribute i observe the same true value of that attribute. Specifically, a subject can learn the true embedded levels of m attributes where $m < n$. Learning serves to provide new values of the attributes, $\bar{x}_1, \dots, \bar{x}_m$, which replace the baseline values for m of the n attributes. After learning, individual k 's new valuation is

$$V_1^k = \sum_{i=1}^m v_i(\bar{x}_i) + \sum_{i=m+1}^n v_i(\hat{x}_i^k).$$

If all subjects learn the same things, averaging across the population leads to an average WTP of

$$\mu_1 = \frac{1}{K} \sum_{k=1}^K \left(\sum_{i=1}^m v_i(\bar{x}_i) + \sum_{i=m+1}^n v_i(\hat{x}_i^k) \right).$$

⁷For example, assume that m is income and consider the expression: $m - V_0 + \sum_i^n v_i(\hat{x}_i) = m + \sum_i^n v_i(0)$ where $v_i(0) = 0$. Also, in this discussion we'll generally assume that attributes are goods but they could also be bads.

The difference between average valuations before and after learning is

$$\mu_1 - \mu_0 = \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^m \left[v_i(\bar{x}_i) - v_i(\hat{x}_i^k) \right] \quad (2)$$

The sign of this difference depends on whether the baselines over- or underestimated the actual levels of the attributes. Consequently, without further structure this model makes no prediction about the direction of the WTP change.

The model can also be used to explore the effect of learning on the variance of WTP values in the sample. Assume that all priors are independently distributed across attributes, and let $\hat{\sigma}_i$ denote the variance of the baseline valuations for attribute i , that is,

$$\hat{\sigma}_i = \frac{1}{K} \sum_{k=1}^K \left(v_i(\hat{x}_i^k) - \frac{1}{K} \sum_{k=1}^K v_i(\hat{x}_i^k) \right)^2. \quad (3)$$

If subjects do not learn about attribute i its variance remains equal to $\hat{\sigma}_i$, but if they do learn about attribute i its variance is reduced to zero because they all share the same valuation $v_i(\bar{x}_i)$. Overall pre-learning variance in WTP is $\hat{\sigma} = \hat{\sigma}_1 + \dots + \hat{\sigma}_n$ and post-learning variance in WTP is $\bar{\sigma} = \hat{\sigma}_{m+1} + \dots + \hat{\sigma}_n$. Consequently, this model allows for reductions in variance but not increases, with $\hat{\sigma} - \bar{\sigma} = \hat{\sigma}_1 + \dots + \hat{\sigma}_m$.

2.2 Using the Model to Interpret Results

We first use the model to provide a basis for the noise-reduction rationale for providing information. Assume that subjects draw their prior assessments $\hat{x}_i^k$ from an unbiased distribution with mean $\bar{x}_i$, so that the expected average valuation is equal to what they would learn if they did so. The change in variance is as above, leading to a straightforward prediction based on learning with unbiased priors: learning leads to no change in average

WTP but a reduction in its variance. Below we refer to this as the unbiased-priors model.

The rationale behind informative advertising is that it can increase WTP, so it requires a different model. One obvious assumption is that prior assessments are biased and low, but for consistency we continue, to the extent possible, with the assumption that all priors are unbiased. An alternative model comes from considering attribute unawareness. In these settings subjects do not know that the good embodies a subset of the attributes, and therefore assigns the unknown attribute i a prior value of $v(\hat{x}_i^k) = 0$. We interpret zero valuations of an attribute as unawareness. If $\bar{x}_i^k > 0$, then new information will cause mean WTP to rise as subjects learn that the good does, in fact, possess that attribute.

Suppose that $\hat{x}_i^k = 0$ for all $k = 1, \dots, K$, so that everyone thinks the attribute is missing from the good. Then the new information increases mean WTP. However, there would be no effect on the variance of WTP because when everyone has the prior $v(\hat{x}_i^k) = 0$, it follows that $\hat{\sigma}_i = 0$. Thus, the assumption of uniform (across subjects) attribute unawareness leads to a prediction that learning increases mean WTP but does not change the variance. We refer to this as the uniform attribute-unawareness model.

More generally, there could be two types of agents in the population, some characterized by attribute unawareness and others aware of the attributes and having unbiased priors. In such a case some agents would have $v_i(\hat{x}_i^k) = 0$ but others have $v_i(\hat{x}_i^k) > 0$, and learning leads to an increase in average WTP driven entirely by the agents characterized by attribute unawareness. Learning also leads to a reduction in variance, driven entirely by the subjects with unbiased priors. We refer to this as the partial attribute-unawareness model.

A final variant of the baseline framework captures confirmatory bias (Rabin and Schrag (1999)). Under this bias, when individuals are confronted with new information about a

good they abide by the information when it confirms their prior beliefs about the good but discount it heavily when it does not. While there are a variety of frameworks which could operationalize this, the simplest one is that individuals update from $\hat{x}_i^k$ to $\hat{x}_i^k + \theta(\tilde{x}_i^k - \hat{x}_i^k)$, where $0 \leq \theta \leq 1$.⁸ When $\theta = 0$ the individual ignores new information or fully discounts it, while when $\theta = 1$ she fully incorporates it. Small values of θ , then, are consistent with confirmatory bias.

Let θ vary by the individual, and denote the individual-specific parameter by θ_k . Assume that individuals $1, \dots, \bar{k}$ are subject to complete confirmatory bias with $\theta_k = 0$ and the rest of the population, $\bar{k} + 1, \dots, K$ update fully with $\theta_k = 1$. Moreover, continue assuming that all priors are unbiased. These assumptions make it straightforward to derive predictions. Learning has no impact on average WTP because of the unbiased priors. If $1 \leq \bar{k} \leq K - 1$ variance falls due to the learning by those not prone to confirmatory bias, but if $\bar{k} = K$ so that everyone exhibits confirmatory bias, learning has no impact on variance because it does not lead to any change in individual WTP values. We refer to this last case, with $\bar{k} = K$, as the full confirmatory bias model.

Table 1: Predictions of the Different Models

	Variance falls	Variance unchanged
Mean WTP unchanged	Unbiased priors	Full confirmatory bias
Mean WTP rises	Partial attribute unawareness	Uniform attribute unawareness

NOTE: The table shows the predicted data pattern of an increase in learning for the mean and variance of WTP under the different models in this subsection. Full confirmatory bias and full attribute unawareness correspond to the versions of the models where all subjects have the same biases. Partial attribute unawareness corresponds to the case where some, but not all, subjects are unaware of a subset of the attributes.

⁸This setting does not allow for the use of Bayes rule because, as modeled, both the prior assessment and the learned value are treated as degenerate distributions. Because of this we use an adaptive rule to reflect a possibly-partial update from the prior assessment to the learned, true value.

Table 1 summarizes the predictions of the different models. It can also be used to identify the appropriate model from the patterns in the data. If, for example, the data show no change in average WTP but a fall in the variance, the data are consistent with the model based on the assumption of unbiased priors, but if the data show an increase in the mean accompanied by a fall in the variance, it would be consistent with a model in which some, but not all, of the subjects exhibit attribute-unawareness. The primary purpose of the experiment is to determine which cell of Table 1 organizes the data.

Lastly, there are some similarities to this model and one of attribute non-attendance (Scarpa, Gilbride, Campbell, and Hensher (2009) and Scarpa, Thiene, and Hensher (2010)). The difference between those models and this one relates to beliefs about attributes versus attending to those attributes. In this model, beliefs impact $\hat{x}_i$. A model of attribute non-attendance could weight multiplying the functions $v_i(\hat{x}_i)$ describing the likelihood they are considered by the subject or not.

3 Flood Protection Experiment

3.1 Setting

We conducted an experiment embedded in a stated preference survey in Scotland during 2013. The survey was related to current efforts by a local government and the national regulator (the Scottish Environmental Protection Agency, SEPA) in Scotland to improve flood defenses along the Tay estuary in Eastern Scotland. Local councils and SEPA were concerned that current defenses are not sufficient to prevent major flooding episodes, given changes in the incidence and magnitude of extreme weather events. Residents also are concerned: we find that many people in the area purchase flood insurance.

In considering their options for decreased risk of flood, one option for regulators is to

encourage the conversion of land currently used for farming to re-build the estuarine and coastal wetlands which once characterized many of Scotland's east coast firths and estuaries. Such wetlands serve two major roles. For flood protection, wetlands offer a repository for temporary episodes of high tides, and mitigate flow rates from the upper catchment which otherwise may cause flooding. The amount of flood protection is commensurate with the size of the wetlands created. Second, wetlands are a rich habitat for wildlife. As a result, wetlands offer a non-market benefit in the form of increased recreation (wildlife viewing) to the local community, as well as providing a range of other ecosystem services such as nutrient pollution removal.

Historically, Scotland constructed large seawalls and other hard structures to provide flood protection rather than reclaiming wetlands. Restored wetlands, then, are effectively a new good offering one attribute for which there is a known alternative (seawalls). They also offer an additional and possibly less understood attribute in wildlife viewing and ecosystem services. We note that both attributes- flood protection and increased wildlife habitat- are public (e.g., non-excludable and non-rivalrous).

In order to gauge the public's willingness to pay for restoring wetlands as a method of flood defense, the government of Scotland commissioned a stated preference survey. Subjects were invited to participate in the survey via repeated mailings and radio and newspaper advertisements. We elicited valuations by subjects clicking on the maximum increase in yearly taxes they would be willing to pay to engage in a specific managed realignment plan described in the survey. These taxes would be levied at the city level in nominal amounts, rather than a percentage of income, for all households. Amounts started at zero and increased in £10 increments to £150. Options included a greater than £150 option. Subjects who completed the survey were given a £10 (\$16) Amazon gift card.

The survey was conducted online through a website we designed and operated. Average completion times were on the order of 15 minutes.

We use a payment card to elicit preferences because there was only one particular policy being considered rather than a set of possible policies. While choice experiments are currently more common in the literature, in our setting the local government was only considering a single policy. As a result we use a payment card to evaluate this single policy. We note, though, that a payment card is a special case of a multiple bounded discrete choice elicitation technique (Vossler, Poe, Welsh, and Ethier (2004)).

Every stated preference survey discusses the good or policy change being studied before eliciting willingness to pay estimates. Part of this discussion includes conveying relevant information about the good or policy change. According to the stated preference literature, subjects should be informed about the good or policy in a stated preference survey in order to simulate the informed state of a representative agent should the policy change go to a referendum (Hanley, Shogren, and White (2013) and Mitchell and Carson (2013)). Similar methods are standard in marketing surveys when introducing a new product (Hair, Wolfinbarger, Ortinau, and Bush (2010)). A stated preference survey, then, is an ideal place to embed an experiment related to information and learning.

3.2 Experimental Design

We embedded experiment variation in the stated preference survey describe above. Figure 1 shows our experimental design visually. The experimental design includes one treatment arm and one control arm. The treatment arm is defined by subjects taking a nine question multiple choice quiz over historical information about flood defense in Scotland and the science behind restoring wetlands both before and after WTP elicitation. In the control

arm, subjects were not given the quiz before elicitation; they were only given the quiz after. In both arms, all subjects received nine pieces of information relevant to the decision task and context before eliciting willingness to pay measures. This type of contextual, background and scientific information is commonly used and considered best practice in stated preference work.




Figure 1: Schematic of the timing of the experiment. Randomized assignment occurs as indicated in Figure 1. Numbers indicate the order of each step. All subjects receive quiz 2.

Each question in the nine question quiz corresponded to a single piece of information, or bullet point, that subjects could subsequently be provided with in the informative portion of the survey. We convey each piece of information on a single screen with a figure and one or two simple sentences. The appendix gives several examples. After the informative portion of the survey, we elicit subjects willingness to pay for reclaiming a particular wetland area. We end the survey by giving all subjects the identical quiz we gave to the treatment arm before the survey before asking a series of demographic and debriefing questions. The ex post quiz given to all subjects provides a measure of ex post knowledge for each subject. Comparing ex post quiz scores by treatment status identifies the causal effect of treatment on learning. Similarly, comparing stated valuations by treatment status identifies the causal effect of learning on willingness to pay.

Each question on the quiz concerns specific attributes of extant flood defense, possible flood defense benefits of restoring wetlands and possible wildlife benefits of restoring wetlands. Restoring wetlands increases flood defense functions above pre-existing flood

defense policy. In that sense, pre-existing flood defenses like sea walls are more likely to be known flood defense attributes. The additional benefits of restoring wetlands- for example, providing wildlife habitat and enhanced recreation opportunities- are plausibly less likely to be familiar to respondents.

We picked each piece of information in order to inform subjects of historical flood protection techniques, flood protection characteristics of restored wetlands, and well understood co-benefits of wetlands like their ability to provide habitat for wildlife. Each bullet point was selected to provide objective and relevant information about the policy being considered. These are exactly the types of background and contextual information which are commonly given in stated preference surveys. Each bullet point was then vetted by University of Stirling scientists (engineers and ecologists) to ensure accuracy. Finally the set of information was vetted by us in order to ensure that language was neutral, objective and easy to understand for voting adults. In this way, we implemented best practices in selecting information presented to subjects. We then created a single multiple choice question to be associated with each informative bullet point.⁹

We measure treatment effects on actual learning and willingness to pay (WTP) by comparing ex post quiz scores and WTP across the treatment and control arms. This is of course only a single form of learning. For example, it is not reflected upon knowledge over a long time period. That said, it is a useful learning metric for two reasons. First, we aren't aware of a better learning metric and it is a commonly used metric to evaluate learning in schools. Second, it is the right metric for our study in the sense that information provided in all stated preference studies is indeed provided and shortly thereafter the economist

⁹We chose not to add any information about the value of the farmland being considered for reclamation. Part of the reason relates to uncertain market prices. Indeed, this motivates our subsequent statement about uncertain costs of the policy when eliciting WTP later on.

elicits willingness to pay. As a result, this is the right type of knowledge metric to test.

We also record time spent on answering questions. By recording time spent on each screen during the informative portion of the survey, this provides two measures of how treatment could affect learning. First, a scale effect: more time implies more learning. Second, a technique effect: time spent on each screen is the same across treatment and control but results are different.¹⁰ We briefly address these two possible explanations in the text and present more detailed results in the appendix.

In the analysis we don't distinguish between different question topics (e.g., co-benefits versus traditional flood defenses) because we'd like to measure the impact of the type of learning which occurs in stated preference surveys generally. Parsing questions by topic would reduce the signal to noise ratio of increased learning for any given information type. From a statistical power perspective, we don't have a large enough variation in learning across questions due to treatment to estimate separate effects by question. Therefore we weight each question equally in the analysis.

Lastly, there is certainly selection into taking the survey. Despite repeated mailings, households had to choose to go online to take the survey. However, treatment was orthogonal to selection. Conditional on taking the survey treatment was random. As a result, the internal validity of our design is maintained. We comment on the external validity of the study in detail below.

¹⁰This second possibility is akin to dedicating either more working memory or more processing power to the task.

3.3 Balancing Concerns

All participants for the survey were selected from the Scottish Phone Directory. Only people living within the local authorities affected by the flood defense scheme were selected to take part. In total 4000 households were contacted by mail and invited to take part in an online survey, with a reminder card sent two weeks after the first contact attempt. Of the 4000 people invited, 749 completed or partially completed the online survey with 562 of the responses completed in sufficient detail to be used in the analysis. Such response rates are typical of mail-out stated-preference surveys in the UK. Although this participation rate is somewhat low, our results are still externally valid so long as treatment effects are orthogonal to self-selected response determinants. Furthermore, because treatment is random attrition from subjects due to partial completion is a second order concern.¹¹

In this paper we use a subset of subjects who received exactly the same information set. In the experiment not all subjects were shown all nine information slides. In a different paper we look at the effect of additional information, as opposed to the effect of additional learning (or retained information) on WTP values (LaRiviere, Czajkowski, Hanley, and Simpson (2015)). To correct for this extra step of sampling in this paper, we therefore weight each observation in order to reconstruct the composition of ex ante information levels observed in our entire subject pool across both this paper and LaRiviere, Czajkowski, Hanley, and Simpson (2015).¹² We note that weighting in this way relative

¹¹The possible violation of this would be if treatment itself caused attrition asymmetrically in groups which both correlates with heterogeneous treatment effects of treatment and WTP, which is of course an unknowable. Note this is very different than correlation between attrition and WTP. While we don't view this as plausible, we cannot rule it out nor can any experimental design in which subjects are able to opt out of completing an experiment.

¹²This experiment was one part of a larger experiment in which we also exogenously vary the amount of new information provided to them within the treatment arm. That allows us to test for the causal effect of additional information, learning and the amount of learning on the distribution of WTP estimates in LaRiviere, Czajkowski, Hanley, and Simpson (2015). After treated agents completed the quiz and their answers were recorded, we grouped them into low (L), medium (M) or high (H) ex ante information

to running OLS without weights does not alter the qualitative results of the analysis.

Table 2: Observations by Information & Treatment Status

Ex Ante Info	Completed	
	All	Debriefing Consequential
L	72	60 48
M	94	82 65
H	12	10 8
Control	89	75 49
n	267	227 170

NOTE: For all treated subjects who completed the initial quiz, *L* info implies quiz score 1 of 0-3, *M* implies 4-6, *H* implies 7-9. “All” column includes all subjects completing survey. “Complete Debriefing” column includes only subjects who answered all debriefing questions upon survey’s conclusion. “Consequential” column includes only subjects who stated they believed survey was somewhat, likely, or very likely to be used to inform policy.

The counts of treatment and control groups by ex ante information for both treatment and control groups are shown in Table 2. We present three columns. The first column shows sample sizes for subjects that completed the survey but may not have answered debriefing questions. The second column shows sample sizes for subjects that both completed the survey and answered the debriefing questions. One of the debriefing questions asked about consequentiality, and the third column shows how many subjects in each category stated they thought the results of the survey were reasonably likely to be used by policy makers to inform policy. Previous research shows both theoretically and empirically that

groups. Each treatment corresponds to answering up to a particular number (3, 6 or 9 for L, M or H respectively) of questions correctly. In the complete survey, we also exogenously vary the amount of *new* information provided to treated subjects since we observe precisely which questions the subject answered correctly. To do so, we vary the number of slides shown to subject since each slide contains exactly one piece of objective information about flood defense and/or reclaimed wetlands corresponding exactly to one question asked on the multiple choice quiz. The quiz and complete set of bullet points are in the Appendix. As a result, treated subjects can be summarized as an ex ante type and information pair. For example, a type-treatment pair could be MH: a subject who answers between four and six questions correctly and who is then given all nine bullet points of information (e.g., the high information treatment). Since in this experiment we are concerned with learning and valuation conditional on an information set we isolate attention to only subjects given all nine information bullet points, and must weight the treated subjects to reflect population ex ante information levels in all analyses. As a result, we overweight ex ante L and M information types while underweighting type H information types in the treated set.

subjects who believe their responses have a chance to be used by policy are more likely to truthfully state their willingness to pay for a policy change (Vossler, Doyon, and Rondeau (2012)).

There are two important pieces of information in Table 2. First, it is important to note that only twelve subjects scored between 7-9 on the first quiz. As a result, there are only 12 a priori type H subjects. Of those 12 subjects, there were only eight who believed the survey was consequential. Because of this, we trim the sample of H subjects in some specifications to mitigate noise.

Second, in order to verify that treatment is balanced on observables and that the subject pool we observe is representative of the larger population, we asked subjects to answer socio-demographic questions upon completion of the second quiz. We lose roughly 16% (36%) of our sample when restricting to subjects who completed debriefing socio-demographic questions (stated perceiving survey as consequential).

Table 3: Differences Between Treatment and Control Groups

	Treatment	Control	Difference between treatment and control	p-value for difference in means
	(1)	(2)	(3)	(4)
Age	51.4	54.4	3	.126
Male	57.3%	61.7%	4.4%	.537
HH Income	50,398	50,609	-211	.954
Flood Insurance	67.3%	61.8%	5.5%	.432
Complicated (1 - 5)	2.4	2.33	.07	.603
Enviro Group Member	38.7%	25%	13.7%	.038
Confidence (1 - 5)	2.92	3.21	-.29	.05

NOTE: HH income is measured in £. Flood insurance is an indicator if subject states that they own some type of flood insurance. “Complicated” indicates how complicated subjects thought the information was (1 = “strongly disagree” to 5 = “strongly agree” the information was too complicated). “Confidence” indicates how likely subject thought survey was going to be used by policy makers (1 = “very unlikely” to 5 = “very likely”). “Enviro Group Member” is an indicator variable equal to one if subject is member of an environmental support group.

To ensure we randomized appropriately, Table 3 shows the balancing table for all 227

subjects who answered debriefing questions. For all but one sociodemographic characteristic was treatment randomized. Table 3 shows that environmental group membership was significantly higher in the treated group relative to the control group. Also, the treated group were significantly more likely to view the results of the survey as consequential for policy.

There are two possible reasons why there would be differences between the treatment and control groups for different characteristics. First, it could be that due to our small sample size the statistical differences are real and the result of chance. Second, it could be that treatment actually caused subjects to change their self-reported characteristics. The significant difference in environmental group membership seems likely due to chance and reflect actual differences in the treatment and control groups. It seems unlikely that treatment caused subjects to claim environmental membership. On the other hand, the significant difference in consequentiality between treatment and control groups seems likely to be due to a direct effect of treatment. For example, subjects might take a survey more seriously if it is preceded by a quiz on the subject.

Table 3 therefore raises two concerns. First, environmental group membership is likely to be correlated with WTP for restored wetlands since wetlands harbor wildlife. Second, treatment seems likely to have caused an increase in consequentiality. As a result, treatment was likely to affect the probability that subjects truthfully answered questions directly. This is problematic: true WTP for the good could be different from strategically stated WTP. As a result, our treatment effect for WTP captures a joint effect of the “pure treatment” effect and the consequentiality effect. Lastly, these two concerns could be related: it could be that environmental group members purposefully stated they were not environmental group members in order to affect the average WTP of non-environmental

group members.

We employ two techniques to address these concerns. First, we allow for the treatment effect to vary by environmental group membership for both the learning and valuation hypotheses. Even with an unbalanced sample, allowing for a direct effect of environmental group membership and the treatment effect to vary by environmental group membership mitigates this issue. Second, we implement a trimming procedure in order to restrict the estimating sample to only subjects who viewed the survey as being possibly consequential. Since treatment was random, the likelihood of a subject who has a low or high value for the project is identical. Keeping only subjects who believe the survey was consequential therefore serves to increase the percentage of control subjects relative to treated subjects in the estimating sample. While this trimming affects the power of the estimated treatment effect, it will not affect the estimated level of the treatment effect unless there is an interaction of treatment, confidence and stated valuation. We discuss this possibility below. We’ve also estimated the model by adding consequentiality directly as a control and find qualitative results.

4 Learning Results

At the start of the survey each treated subject answered identical nine question multiple choice quizzes concerning objective information about the good. This quiz was given to all respondents (both treated and control) after stating WTP. Figure 2 shows the histogram of subjects’ scores in quiz one and quiz two for all subjects who took both quizzes. Figure 2 shows that there was a significant increase in the scores from quiz one (mean= 3.08, SD=1.76) to quiz two (mean=5.19, SD=2.23) for the treated subjects.

In order for our experimental design to be valid, treatment must increase subjects’




Figure 2: Quiz score histograms by test for treated group. Ex post quiz scores include only the treated group. Treated subjects take quiz before information provided in survey and again after WTP elicitation. Control subjects take quiz only after WTP elicitation.

retention rates for information (i.e., the amount of information treated subjects learn relative to control subjects). Since treatment in our experiment is the act of taking the initial quiz, we can take a difference in means in the average score of subjects on the second quiz by treatment status to estimate the average treatment effect. As a result, we run the following regression:

$$Score_i = \alpha + 1\{treated\}\beta + 1\{Env\ Group\}\gamma_e + 1\{treated\}1\{Env\ Group\}\gamma_{et} + \epsilon_i. \quad (4)$$

In equation (4), $Score_i$ represents the score on the second quiz of subject i and ϵ_i is noise.¹³ Weighted appropriately, the estimated coefficient on β is the causal effect of the pre-survey quiz on learned information. We allow the treatment effect to vary by environmental group membership since it could be correlated with willingness to learn and the sample was unbalanced over environmental group membership. The interpretation of the coefficient γ_{et} is the difference in the causal impact of treatment on quiz scores for environmental

¹³We do not report specifications with sociodemographic controls in for learning. The treatment effect for learning is very robust to inclusion of controls. We are happy to provide them upon request.

group members relative to the control group. In addition to estimating equation (4) using ordinary least squares (OLS) we estimate it using a negative binomial model since the left hand side variables are count variables. In every case we use White robust standard errors.

Table 4 shows the results from estimating equation (4) excluding and including controls for both OLS and negative binomial models. The causal effect of treatment on retained knowledge is significant and positive. To put the parameter estimate into context, the average ex ante knowledge level was roughly 3.06 and the average ex post knowledge score for the control group was 4.95 for subjects who stated they believed that the survey was consequential. As a result, the treatment effect point estimate of 1.27 translates to a 62% increase in the rate of learning.¹⁴ The treatment effect is positive and significant for both the OLS and negative binomial models both with and without controls regardless whether the estimating sample is restricted to subjects stating they view the survey as consequential.

A 62% increase in the rate of learning is large. It amounts to the average subject knowing more than one additional piece of information about the attributes of the public good (an average of knowing under five to an average of over six). An ideal treatment effect would lead to every treated subject knowing every piece of information provided in the survey. This type of treatment is not realistic, though, because participants in a stated preference survey don't learn each piece of information perfectly. As a result, we take this effect to be the precise kind of exogenous variation which identifies the key comparative static of interest in our study: the causal impact of learning more information about good

¹⁴This is calculated as follows: $.62 = 1.27 / (4.95 - 3.06)$. Recall that 3.06 is the weighted quiz score average of the treated group score on their first quiz.

attributes on stated willingness to pay. Put another way, treatment provides a significant, large and internally valid source of variation.

Table 4: The Effects of Treatment on Knowledge Score

	(1)	(2)	(3)	(4)	(5)	(6)
Treated	.74** (.35)	1.12*** (.42)	1.27*** (.49)	.143** (.07)	.198*** (.077)	.228** (.091)
Env Group		-.175 (.804)	-.13 (1.05)		-.035 (.16)	-.027 (2.15)
Treated * Env		-1.20 (.94)	-1.33 (1.17)		-.215 (.185)	-.241 (.237)
Constant	4.80*** (.28)	5.12*** (.33)	4.95*** (.39)	1.57*** (.059)	1.63*** (.064)	1.60*** (.078)
Model	OLS	OLS	OLS	Neg Bi	Neg Bi	Neg Bi
Only Consequential	N	N	Y	N	N	Y
n	267	232	175	267	232	175
LL				-647.11	-559.9	-412.5

NOTES: LHS variable is second quiz score. ***= significant at 1%, ** = significant at 5%, * = significant at 10%. All regressions weighted and robust standard errors are reported.

There are several explanations for a significant effect of treatment on learning rates. It could be that the initial quiz attunes subjects to be more efficient with their attention. Alternatively, the quiz may reduce the cost of learning, for example, if the quiz uses words that are similar to the words in the informative slides. It could also be that the quiz causes subjects to take the information portion of the survey more seriously and spend more time learning the information. Regardless of the channel that leads to the improved learning rates, the key finding of this section is that the treatment caused subjects to increase their knowledge.¹⁵ Because of the causal relationship, we can interpret any change in the

¹⁵Still, our data proved insight into the open question of why the quiz led to increased learning. We recorded the amount of time spent subjects spent on the second quiz, and the appendix reports a regression similar to equation (4) but with time as the dependent variable. We find that in no case is there a significant treatment effect on time spent on quiz 2. The point estimates are negative but the standard errors around the estimates are large in every case. We take this as evidence of treatment leading to a “technique” learning effect (a pre-quiz increases retention of information by subjects) rather than a “scale” learning effect (subjects re-allocate time following the pre-quiz). This result is consistent with models of focusing (Kőszegi and Szeidl (2013)).

WTP distribution between the treatment and control groups as driven by differences in knowledge.

5 Valuation Results

The motivating question behind our paper is how treatment- and therefore exogenous increases in the rate of learning- affects a consumer’s valuation for the good. The theoretical framework of Section 2 sketches three models of learning and valuation, all starting with consumers have unbiased prior assessments of the attributes embodied in the good. In one model learned information removes heterogeneity in assessments across consumers, in the second the learned information makes consumers aware of attributes they did not know the good contained, and in the third consumers learn but are subject to confirmatory bias. These models predict different combinations of increased or unchanged mean WTP and reduced or unchanged variance of WTP, and they are summarized in Table 1. This section presents our results in the context of that theoretical framework.

Figure 3 shows a histogram of all subjects who completed the survey. The important feature of Figure 3 is the presence of significant heterogeneity in valuation for restoring wetlands. The histogram also shows some anchoring around £50, £100, and £150 for both the treatment and the control subjects. There are three reasons why this is not a major concern for our study. First, the control group is smaller than the treatment group. As a result, it is difficult to know if there are actually anchoring differences across the treatment and control group. Second, the raw data in the histogram doesn’t control for unbalanced sociodemographic characteristics like environmental group membership. Third, treatment is randomized and we are interested in estimating an average treatment effect. In order for anchoring to affect our findings, treatment must interact with anchoring effects to affect




Figure 3: Weighted histogram of WTP for all subjects. $n = 267$.

the mean and variance of stated WTP.

One intuitive way to visualize the effect of treatment on valuation is to plot the CDF of valuations for both treated and control subjects. Using a payment card, we asked subjects to state what is the most they would definitely be willing to pay per year in increased taxes for managed realignment flood protection. The response interface offered the subject £5-10 increments from £0 to £160 or greater. This allows us to plot a CDF of WTP by treatment status. We then fitted stated WTP levels to Kaplan-Meier and Weibull distributions in Figures 4 and 5, respectively.¹⁶ A one-sided, two-sample Kolmogorov-Smirnov test rejects equality of the valuation CDFs with a p-value of .026 for the Kaplan-Meier distribution. Similarly, the scale coefficient of the Weibull distribution varies significantly with treatment: treatment increases the scale parameter by a highly significant .702 (standard error .0716). Using a one-sided Wald test, treatment leads to a significant increase

¹⁶We tested for other possible distributions as well but the Weibull was the best fit for the classes of continuous distributions we tested. See Appendix for details.




Figure 4: Fitted Kaplan-Meier CDF by treatment and control status.

in the mean WTP (p-value .022). As a result, both CDFs show that treatment leads to a significant increase in stated willingness to pay. This result is robust to different weighting procedures as well.

There are two problems with relying on the above results in testing if treatment affected WTP. First, our balancing table shows that we must control for environmental group membership since our treatment and control groups were not balanced along that dimension. Second, in a subset of specifications, we restrict the sample to include only those subjects who report viewing the survey as consequential in an attempt to control for consequentiality similar in spirit to Vossler and Watson (2013). Without the restricted sample, our WTP estimates are subject to hypothetical bias which could interact with treatment in unexpected ways. We therefore test for an effect of treatment on WTP by




Figure 5: Fitted Weibull CDF by treatment and control status.

estimating the following model using the restricted sample:

$$Valuation_i = \alpha + 1\{treated\}\beta + 1\{Env\ Group\}\gamma_e + 1\{treated\}1\{Env\ Group\}\gamma_{et} + X_i'\delta + \epsilon_i. \quad (5)$$

The main feature of equation (5) is that we let the treatment effect vary across environmental group membership. Since Table 3 shows that the treated group had a larger share of environmental group members therefore we allow the treatment effect to vary by environmental membership. This is important since environmental preferences are likely to be correlated with valuations for managed realignment. The interpretation of the coefficient γ_{et} is the difference in the causal impact of treatment for environmental group members relative to the control group. However, our primary coefficient of interest is β : since treatment was random, the estimated coefficient on β is the causal effect of the pre-survey quiz on valuation for “managed realignment” flood protection. From the theoretical model in

Section 2, a positive coefficient is compatible with a model of attribute inclusion, while the models of movement around a baseline and confirmatory bias imply a zero coefficient.

We also control for sociodemographic characteristics and the date of subject response by including controls X_i . In every specification we trim the sample to exclude subjects who stated they thought the results of the survey were either very unlikely or unlikely to be used in policy. This serves an important purposes by mitigating any possible impacts of a causal effect of treatment on consequentiality leading to different strategic incentives for varying levels of WTP.¹⁷

Table 5: Treatment on Stated Valuation

	(1)	(2)	(3)	(4)	(5)	(6)
Treated	11.015*	11.75*	6.25	6.01	1.19	3.64
	(5.79)	(6.99)	(6.54)	(7.98)	(7.33)	(9.29)
Env Group			18.6	24.58	27.34*	51.16**
			(12.68)	(17.9)	(14.85)	(20.83)
Treated * Env			-2.19	-7.45	-10.78	-35.45
			(15.24)	(20.44)	(16.81)	(22.79)
Constant	33.15***	34.84***	30.35***	32.69***	52.57***	68.27**
	(4.36)	(4.77)	(4.77)	(6.20)	(26.64)	(29.9)
Model	OLS	OLS	OLS	OLS	OLS	OLS
Date FEs	N	N	N	N	Y	Y
Demographic Controls	N	N	N	N	Y	Y
Only Consequential	N	Y	N	Y	N	Y
n	267	210	232	175	216	164

NOTES: LHS variable is Stated Willingness to Pay in £. ***= significant at 1%, ** = significant at 5%, * = significant at 10%. All regressions weighted and robust standard errors are reported. Date FEs include indicator variables to the mailing wave preceeding the observed survey response. Demographic controls were gender, age group indicators, household income group indicators, and an indicator if the subject owned flood insurance.

Table 5 shows the results from estimating equation (5). In the specification without

¹⁷Put another way, it gives an appropriate apples to apples comparison of the treatment group to the control by restricting the sample. If consequential subjects are more likely to state higher WTP levels because those values better represent actual preferences, for example, trimming the sample for both the treated and control groups gives the correct counterfactual. More generally, so long as treatment is orthogonal to any strategic incentives governed by consequentiality, our design remains internally valid.

controls, we find the same result as in the CDFs. However, when we add an indicator for environmental group membership, let the treatment effect vary by environmental group membership, and add in control variables the point estimate of the treatment effect decreases in magnitude and becomes insignificant.¹⁸ As the sample is restricted to include only subjects who answered all demographic questions and only subjects who view the survey as consequential, the sample size decreases and the estimated standard errors increase only modestly so our zero result does not appear to be driven by a lack of power. As the sample is restricted and controls are added, environmental group membership also becomes significant and positively correlated with stated valuation for the project. This is expected: a major component of the project is that it increases wildlife habitat. However, treatment causes no significant effect on either non-environmental group members nor environmental group members.¹⁹

The results in Table 5 show that treatment had no effect on the mean valuations even though it significantly increased the rate of learning by over 60%. Even though there is a weak positive impact of treatment on WTP in OLS specifications with no controls, that impact goes to zero when we control for (unbalanced) demographic variables likely to be correlated with preferences for the good. In the context of the theoretical model this finding is consistent with either unbiased priors or confirmatory bias. In order to parse between those two models, we must test for the effect of treatment on the variance of the WTP distribution.

Estimating the causal impact of treatment on the variance of WTP is a bit more challenging than using OLS. In order to do so we estimate the impact of treatment on the

¹⁸This finding is replicated when we break out learning by treated high versus low scorers.

¹⁹While suppressed in Table 5, other control variables had expected signs. For example, purchasers of flood insurance had a positive willingness to pay for reclaimed wetlands.

Table 6: Treatment on Distribution of Stated Valuations

	(1)	(2)	(3)	(4)	(5)	(6)
Mean	28.37*** (5.999)	34.10*** (6.59)	36.87*** (7.01)	2.77*** (0.18)	2.95*** (.20)	2.99*** (.22)
Treated	15.3778** (7.41)	7.46 (8.67)	3.88 (8.86)	.45** (.21)	0.22 (.27)	.181 (.28)
Env Group		12.30* (6.86)	11.85** (5.96)		0.35** (.17)	.441** (.177)
Treated * Env		-1.61 (7.45)	-2.63 (6.59)		-0.06 (.18)	-.183 (.198)
Variance	53.42*** (5.14)	48.56*** (5.57)	48.63*** (6.42)	1.57*** (.15)	1.43*** (.17)	1.41*** (.195)
Treated	3.06 (6.29)	5.20 (7.38)	.754 (8.35)	-.120 (.18)	-.007 (.23)	-.064 (.249)
Env Group		4.01 (5.55)	-3.40 (4.50)		-.140 (.138)	-.245* (.149)
Treated * Env		-2.25 (6.04)	-.31 (5.12)		.020 (.15)	.082 (.167)
Model	Spike	Spike	Spike	Lognormal	Lognormal	Lognormal
Demographic Controls	N	N	Y	N	Y	Y
n	267	232	210	267	232	210
log likelihood	-724.04	-636.12	-574.10	-711.81	-615.03	-562.23

NOTES: LHS variable is Stated Willingness to Pay in £. ***= significant at 1%, ** = significant at 5%, * = significant at 10%. All regressions weighted. Demographic controls were gender, age level, household income level, and an indicator if the subject owned flood insurance.

variance of WTP when we fit WTP to both a lognormal and a spike distribution. In both cases, we control for environmental group membership and let the treatment effect interact with environmental group membership as before. We also, though, allow for treatment to affect both the mean and the variance of WTP simultaneously.

Table 6 shows estimation results for fitting lognormal and spike distributions to the data for three different sets of controls.²⁰ The top set of results are for the mean and the bottom are for the variance. For both the spike and the lognormal distribution there is no effect of treatment on the mean once environmental group membership is controlled for, thus confirming the findings in the OLS model.

For every specification, there is also no effect of treatment on the variance of stated valuations, even when not controlling for environmental group membership and other sociodemographic characteristics. Further, for both distributions when the full set of sociodemographic controls are included the point estimate for treatment on variance becomes very small in magnitude. These results are robust to trimming the sample to only include subjects who view the survey to be consequential. We note, though, that standard errors remain relatively unchanged and even slightly increase. We take Table 6 as strong evidence that treatment has no effect on either the mean nor the variance of WTP in our survey.

In the context of our model, the treatment effects shown in Tables 5 and 6 are consistent with a model of confirmatory bias. Treatment caused a significant increase in knowledge of good attributes. However that increased knowledge did not affect the mean nor the variance of valuations. We therefore conclude that subjects learned new information but

²⁰Note that for the lognormal distribution all results must be raised to e to be compared directly to the lognormal results.

they ignored that information when forming and/or updating their valuations.

That objective information about public good characteristics do not affect the distribution of stated preference valuations it is a puzzle. The impetus for informing subjects about good characteristics is deeply engrained in stated preference demand estimation for public goods. To that end, there is little work on identifying the effects of learning objective attribute information on valuations. We don't argue that subjects should not be given information in stated preference surveys or that state preference methods are flawed in some way. Rather, we view these results as motivating additional work to fully understand exactly what the causal impact of the information provision portion of a stated preference survey are on valuations. For example, if subjects exhibit confirmatory bias in both stated preference surveys and when making private good purchases, it provides evidence that stated preference methods elicit behavior similar to revealed preference methods. We note, though, that this result is only a single study: while exogenous variation and careful design provide internal validity, external validity is always a concern with any single study. This further motivates additional work in this area.

6 Discussion and Conclusions

In this paper, we embed experimental variation into a stated-preference survey about a public good to test for the causal impact of exogenous increases in learning rates on WTP for the good. Importantly, the good (flood protection through wetland restoration) has one attribute that is likely to be well-understood (flood protection) and another that is likely to be less-understood (wetland restoration and environmental amenities). As a result, our setting bears similarities to firms who advertise information about the attributes of new

products.²¹

Treated subjects received a short multiple choice quiz before receiving nine pieces of information about an actual flood defense plan in Scotland. Both treated and control subjects took the same quiz after we elicited WTP allowing us to identify the effect of the quiz on learning. Our identifying assumption is that randomizing subjects into treatment in which they receive a quiz about traits of the public good before that good is described serves to provide exogenous variation in the learning rates. We find that treatment increased learning rates by 62%.

Our simple theoretical framework makes predictions about how an exogenous increase in the learning rate could impact the estimated mean and variance of the willingness-to-pay distribution. Our main result is that we find no significant impact of treatment on either the mean or the variance of WTP. As highlighted in the theoretical model, that result is consistent with new knowledge serving to confirm whatever “bias” subjects had concerning the good before the experiment.²² Taken as a whole, our findings are consistent with a model of confirmatory bias rather than movement around a baseline or unknown attribute classes.

These findings are somewhat striking. Without them the assumption that more knowledgeable consumers make more efficient decisions would have been uncontroversial. However, we find that when we exogenously vary the rate at which consumers learn objective information about attributes of a public good, there is no change in their willingness to pay or in the heterogeneity of their valuations. An analogy would be that voters do not

²¹For example, Amazon, Google and Microsoft advertise security features for cloud storage or cloud computing.

²²This result is similar to recent work which finds that the causal effect of additional knowledge about product attributes on WTP is zero (LaRiviere, Czajkowski, Hanley, and Simpson (2015)).

deviate from their initial “gut reactions” to political candidates or voter propositions. In that case, discourse targeted toward emotional appeals rather than characteristics of a good or candidates may be more useful for affecting decision making.

There are some other possible explanations for our results than confirmation bias. It could be that our information concerns attributes that subjects do not care about. For example, it could be that subjects only care about what specific species would populate the restored floodplains or what would have happened to a friend’s nearby farm. Alternatively, it could be that consumers must reflect on newly gained knowledge before adjusting their valuations. In that case, any short run experiment or focus group designed to infer preferences would be subject to the same design flaw of our experiment.

More important are the implications of our results on stated-preference valuation methods and public good valuation. A very common feature of stated-preference methods is that subjects are provided with information about the public goods project being valued. To that end a large amount of effort goes into ensuring the information is as easily understood as it can be. For the particular good that we study, though, effort spent on reducing the cost of gaining objective knowledge about good attributes would have absolutely no effect on estimated willingness to pay.

Perhaps even more surprising than there being no effect on mean WTP levels, treatment did not affect the precision of WTP estimates. This creates a puzzle given that the usual rationale for providing information in WTP elicitation experiments is to increase precision. Our results, then, even question the value of informing subjects to increase the signal-to-noise ratio of responses and the power of stated-preference demand estimation methods. We do not claim that our results imply subjects should not be informed about public good attributes, but more work is needed to understand exactly what the role of

the informative part of surveys is.

References

Cameron, T., Elgin, J. 1997. Respondent experience and contingent valuation of environmental goods. *Journal of Environmental Economics and Management* 33 (3), 296–313.

Caplin, A., Dean, M. 2015. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review* 105 (7), 2183–2203.

Carson, R., Groves, T., List, J. 2014. Consequentiality: A theoretical and experimental exploration of a single binary choice. *Journal of the Association of Environmental and Resource Economists* 1 (1), 171–207.

Czajkowski, M., Hanley, N., LaRiviere J. 2014. The effects of experience on preferences: Theory and empirics for environmental public goods. *American Journal of Agricultural Economics* 97 (1), 333–351.

Czajkowski, M., Hanley, N., LaRiviere, J. 2015. Controlling for the effects of information in a public goods discrete choice model. *Environmental and Resource Economics*, doi: 10.1007/s10640-014-9847-z.

Gabaix, X., Laibson, D., Moloche, G., Weinberg, S. 2006. Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review* 96 (4), 1043–1068.

Hair, J., Wolfinbarger, M., Ortinau, D., Bush, R. 2010. *Essentials of Marketing Research*. McGraw-Hill-Irwin, New York.

Hanley, N., Czajkowski, M., Hanley-Nickolls, R., Redpath, S. 2010. Economic values of species management options in human-wildlife conflicts: Hen Harriers in Scotland. *Ecological Economics* 70 (1), 107–113.

- Hanley, N., Shogren, J., White, B. 2013. Introduction to Environmental Economics. Oxford University Press, Oxford-New York.
- Hanna, R., Mullainathan, S., Schwartzstein, J. 2014. Learning through noticing: Theory and experimental evidence in farming. *The Quarterly Journal of Economics* 129 (3), 1311-1353.
- Huffman, W.E., Rousu, M., Shogren, J.F., Tegene, A. 2007. The effects of prior beliefs and learning on consumers' acceptance of genetically modified foods. *Journal of Economic Behavior and Organization* 63 (1), 193-206.
- Kőszegi, B., Szeidl, A. 2013. A model of focusing in economic choice. *The Quarterly Journal of Economics* 128 (1), 53-104.
- Lancaster, K. 1966. A new approach to consumer theory. *Journal of Political Economy* 74 (2), 132-157.
- LaRiviere, J., Czajkowski, M., Hanley, N., Simpson, K. 2015. The Marginal Impact of Learning Information on Public Good Valuation. University of Tennessee Working Paper.
- Mitchel, R., Carson, R. 2013. Using Survey to Value Public Goods: The Contingent Valuation method. Routledge, New York.
- Rabin, M., Schrag, J. 1999. First impression matter: A model of confirmatory bias. *The Quarterly Journal of Economics* 114 (1), 37-82.

Scarpa, R., Gilbride, T.J., Campbell, D., Hensher, D.A. 2009. Modelling attribute non-attendance in choice experiments for rural landscape valuation. *European Review of Agricultural Economics* 36 (2), 151-174.

Scarpa, R., Thiene, M., Hensher, D.A. 2010. Monitoring choice task attribute attendance in non-market valuation of multiple park management services: Does it matter? *Land Economics* 86 (4), 817-839.

Schwartzstein, J. 2014. Selective attention and learning. *Journal of European Economic Association* 12 (6), 1423-1452.

Smith, V.K., Desvousges, W.H. 1990. Risk communication and the value of information: Radon as a case study. *The Review of Economics and Statistics* 72 (1), 137-142.

Smith, V.K., Desvousges, W.H., Fisher, A., Johnson, F.R. 1988. Learning about radon's risk. *Journal of Risk and Uncertainty* 1 (2), 233-258.

Smith, V.K., Johnson, F.R. 1988. How do risk perceptions respond to information? The case of radon. *The Review of Economics and Statistics* 70 (1), 1-8.

Vossler, C., Doyon, M., Rondeau, D. 2012. The truth in consequentiality: Theory and field evidence on discrete choice experiments. *American Economic Journal: Microeconomics* 4 (4), 145-171.

Vossler, C., Poe, G., Welsh, P., Ethier, R.G. 2004. Bid design effects in multiple bounded discrete choice contingent valuation. *Environmental and Resource Economics* 29 (4), 401-418.

Vossler, C., Watson, S. 2013. Understanding the consequences of consequentiality: Testing the validity of stated preferences in the field. *Journal of Economic Behavior and Organization* 86, 137–147.

A Appendix

Distribution Selection

Respondents' were asked to select the maximum bid they would be willing to pay from the payment cards they were provided. This means that we know their maximum WTP is at least as high as the selected bid but lower than the next bid (since it was not selected). We use this information to fit 3 parametric distributions of respondents' WTP-normal, lognormal and Weibull. CDF at the unselected bid less the CDF at selected bid gives us the probability that respondents' WTP is in the interval described by these two bids, conditional on the parameters of the fitted parametric distribution. We estimate these parameters by maximizing the LL function given observations. We used weighted LL function approach (with q1 score weights). The results are for all T0 (control) and T1 (treatment) observations pooled together. Figure 6 shows the results.

The Weibull distribution offers the best fit since it has the lowest log-likelihood, while the number of estimated parameters is the same. To decide if the difference is statistically significant, since the models are not nested we used the Vuong test. Weibull distribution fits best and the difference is statistically significant. Those results are available from the authors upon request.

To provide an illustration of the fit of the fitted parametric distribution we compare it with the simplest possible non-parametric (Kaplan-Meier) estimator, which treats mid-point of selected interval as if it was respondents' true maximum WTP. To compare it with non-parametric distribution, we use the simplest possible Kaplan-Meier estimator, using mid-point of selected interval and taking into account that the highest bid is censored. Figure 7 shows the results.

Distribution	Normal			Lognormal			Weibull		
	coef.	s.e.	p-value	coef.	s.e.	p-value	coef.	s.e.	p-value
Mean / scale	38.616	3.558	0.000	3.078	0.095	0.000	41.212	3.570	0.000
S.d. / shape	56.088	3.005	0.000	1.509	0.082	0.000	0.758	0.044	0.000
LL	736.1697			714.0119			703.4012		

Figure 6: Candidate valuation distributions




Figure 7: Fitted valuation distributions

Analysis of Time Spent on Quiz 2

$$Time_i = \alpha + 1\{treated\}\beta + 1\{Env\ Group\}\gamma_e + 1\{treated\}1\{Env\ Group\}\gamma_{et} + X_i'\delta + \epsilon_i. \quad (6)$$

Table 7: Treatment on Time Spent on Quiz 2

	(1)	(2)	(3)	(4)	(5)
Treated	-12.73 (16.63)	-21.16 (24.8)	-28.54 (35.37)	-29.72 (27.53)	-33.60 (37.72)
Env Group		-8.44 (29.4)	-14.2 (42.2)	.96 (22.15)	10.91 (32.43)
Treated * Env		-3.64 (31.4)	-3.67 (44.1)	-17.99 (24.19)	-46.90 (36.68)
Constant	86.5*** (15.84)	97.75*** (24.84)	111.56*** (34.44)	17.26 (68.22)	38.02 (77.39)
Model	OLS	OLS	OLS	OLS	OLS
Only Consequential	N	N	Y	N	Y
Controls	N	N	N	Y	Y
n	267	232	175	216	164

NOTES: LHS variable is total number of seconds spent on quiz 2. *** = significant at 1%, ** = significant at 5%, * = significant at 10%. All regressions weighted and robust standard errors are reported. Controls include indicator variables to the mailing wave preceding the observed survey response, gender, age group indicators, household income group indicators, and an indicator if the subject owned flood insurance.

Quiz A Short Questionnaire for You

Please answer the following nine questions about flood defence and the Tay Estuary to the best of your knowledge. We would really like to find out how much people know about the Tay Estuary. This will make it easier for the Scottish Government and local authorities to let you know what is taking place in your area now and in the future.

1. In the Tay Estuary what percentage of homes are at risk from flooding?
 - a. Less than 3%
 - b. **Between 3% and 5%**
 - c. Between 6% and 8%
 - d. More than 9%
 - e. I don't know
2. How much money is invested annually in river and coastal defence in Scotland?
 - a. Between £10 million and £30 million
 - b. **Between £30 million and £50 million**
 - c. Between £50 million and £70 million
 - d. Between £70 million and £90 million
 - e. I don't know
3. Historically, the main type of coastal flood protection in Scotland has been:
 - a. Beach replenishment and nourishment
 - b. Planning regulations to limit development on flood plains
 - c. **Concrete sea walls and rock armouring**
 - d. Managed realignment

- e. I don't know
- 4. Managed realignment schemes have the potential to provide:
 - a. A lower level of protection from flooding
 - b. No protection from flooding
 - c. **A greater level of protection from flooding**
 - d. The same level of protection from flooding
 - e. I don't know
- 5. Coastal wetlands are beneficial to fisherman because:
 - a. Wetlands do not benefit fisherman
 - b. Wetlands provide a food source for fish
 - c. **Wetlands provide spawning grounds for fish**
 - d. Wetlands act as a 'no take zone' thereby helping to preserve fish stocks
 - e. I don't know
- 6. Coastal wetlands are beneficial to wildlife because:
 - a. Wetlands do not benefit wildlife
 - b. Wetlands are less polluted than other coastal habitats
 - c. **Wetlands provide a food source for wildlife**
 - d. Wetlands are less likely to be disturbed by humans
 - e. I don't know
- 7. Managed realignment schemes involve the loss of land to the sea. The land most likely to be lost is:
 - a. **Agricultural land**
 - b. Residential land
 - c. Disused brownfield land
 - d. Seafront land
 - e. I don't know
- 8. The Scottish Government has a legal duty to the European Union to protect coastal wetlands because:
 - a. Wetlands are important recreational assets
 - b. Wetlands are important fishing grounds
 - c. **Wetlands are important habitats for waterbirds**
 - d. Wetlands are important natural flood defences
 - e. I don't know
- 9. Which of the following is one of the main causes of decline of shelduck (a waterbird) in the Tay Estuary?
 - a. Commercial fishing
 - b. **Coastal erosion**
 - c. Port operations
 - d. Oil spills
 - e. I don't know

ID 2

In the Tay Estuary what percentage of homes are at risk from flooding? *

- ☐ Less than 3%
- ☐ Between 3% and 5%
- ☐ Between 6% and 8%
- ☐ More than 9%
- ☐ I don't know

ID 3

How much money is invested annually in river and coastal defence in Scotland? *

- ☐ Between £10 million and £30 million
- ☐ Between £30 million and £50 million
- ☐ Between £50 million and £70 million
- ☐ Between £70 million and £90 million
- ☐ I don't know

ID 4

Historically, the main type of coastal flood protection in Scotland has been: *

- ☐ Concrete sea walls and rock armouring
- ☐ Beach replenishment and nourishment
- ☐ Managed realignment
- ☐ Planning regulations to limit development on flood plains
- ☐ I don't know

Figure 8: Three questions as seen by subjects.

Selected screenshots from survey.

In the Tay Estuary and Montrose Basin approximately 4% of homes lie within a floodplain.




Figure 9: Bullet point slide for homes at risk.

Each year the Scottish Government makes £42 million available to support local authorities' flood prevention and coastal protection programmes.




Figure 10: Bullet point slide for flood defense spending.

Historically, the main form of flood defence in Scotland has been hard engineered structures such as sea walls, rock armour and groynes.




Figure 11: Bullet point slide for historical defenses.




Figure 12: Physical location of the proposed site.


Future Flood Defence

Coastal flooding accounts for 29% of the overall flood risk in the Tay Estuary. Coastal flooding is caused by abnormally high spring tides and storm surges. The Scottish Government and local authorities are considering the future river and coastal defence options for the Tay Estuary as part of a commitment to deliver flood risk management plans by 2015. Managed realignment is one of the potential flood defence options for the Tay Estuary.

What is Managed Realignment?

Managed realignment is where part of the existing sea defence is removed and the area behind it is allowed to flood. This flooding creates coastal wetlands known as saltmarshes.

These coastal wetlands can protect the land against storm surges, coastal erosion and sea level rise. By allowing an area of land to flood constructed sea walls needed to protect residential areas and businesses can be shorter in height and thus constructed at a lower cost (see picture below).



On the next page you will learn more about flood defence in Scotland and what managed realignment means for you.

Figure 13: Slide describing wetlands restoration visually.

Future Flood Defences in the Tay Estuary

Newburgh Managed Realignment Scheme

Newburgh (a town on the south bank of the Tay Estuary) is at risk from coastal flooding due to predicted sea level rise:

- The local authority needs to look at future flood defence options.
- A managed realignment scheme is being proposed to protect homes and businesses in the town from flooding.
- It is predicted that the managed realignment scheme would increase flood protection for at least 100 homes in Newburgh (homes shown on the map below).
- Full flood defence benefits will be realised in 15-20 years and then last for at least 100 years.
- Your local authority is responsible for funding 20% of the scheme's cost. The extra income needed for the council to fund this would be raised by increasing your council tax.

Cost of the Managed Realignment Scheme

We would now like you to think about the value to you personally of developing this managed realignment scheme for Newburgh on the Tay Estuary:

- On the next page you will be shown a table of prices that would be added to your council tax annually to cover the costs and maintenance of the scheme.
- You are asked to choose amongst a variety of price options as the precise costs of going ahead with the managed realignment scheme at present are unknown.
- The price you choose will be used to inform the local authorities and the Scottish Government when deciding future flood defence options in the Tay Estuary.
- Before you answer carefully consider the cost to you. Think about your household budget and what you would have to trade off to pay for the increase in council tax e.g. what you like to buy or a reduction in your planned savings. The average household council tax bill in Scotland is £984 per year.

What happens if there is no Managed Realignment Scheme?

- If the managed realignment scheme does not take place the existing flood defences (sea walls) will continue to be maintained by the local authorities at no additional cost on your council tax bill.
- However there will be no additional flood protection and additional benefits of managed realignment will not be realised.

Remember that your preferences will be used in conjunction with costs of the scheme, when they are known, by local authorities and the Scottish Government to inform which flood defence policy is chosen.

Figure 14: General project information given to all subjects.

Please select either yes or no to each of the prices presented in the table that your household would definitely be willing to pay in additional council tax each year to fund the development of a managed realignment scheme at Newburgh on the Tay Estuary. *

	Yes	No
£0 *	☐	☐
£5 *	☐	☐
£10 *	☐	☐
£15 *	☐	☐
£20 *	☐	☐
£30 *	☐	☐
£40 *	☐	☐
£50 *	☐	☐
£60 *	☐	☐
£70 *	☐	☐
£80 *	☐	☐
£90 *	☐	☐
£100 *	☐	☐
£110 *	☐	☐
£120 *	☐	☐
£130 *	☐	☐
£140 *	☐	☐
£150 *	☐	☐
I am prepared to be in excess of £150 *	☐	☐

Figure 15: Payment card used in the survey.



FACULTY OF ECONOMIC SCIENCES
UNIVERSITY OF WARSAW
44/50 DŁUGA ST.
00-241 WARSAW
WWW.WNE.UW.EDU.PL