UNIVERSITY OF STIRLING

Stirling Management School

What is the Causal Effect of Information and Learning about a Public Good on Willingness to Pay?

Mikolaj Czajkowski

Nick Hanley

Jacob LaRiviere

Katherine Simpson

Stirling Economics Discussion Paper 2014-05

April 2014

Online at:

http://www.stir.ac.uk/management/research/economics/working-papers/

What is the Causal Effect of Information and Learning about a Public Good on Willingness to Pay?

February, 2014

Mikolaj Czajkowski[1], Nick Hanley[2], Jacob LaRiviere[3] and Katherine Simpson[2]

1. Department of Economic Sciences, University of Warsaw. miq@wne.uw.edu.pl
2. Division of Economics, University of Stirling. n.d.hanley@stir.ac.uk; Katherine.simpson@stir.ac.uk
3. Department of Economics and Baker Center for Public Policy, University of Tennessee, Knoxville, TN 37996-0550. (865) 974-8114. Jlarivi1@utk.edu

Abstract

In this study we elicit agents' prior information set regarding a public good, exogenously give information treatments to survey respondents and subsequently elicit willingness to pay for the good and posterior information sets. The design of this field experiment allows us to perform theoretically motivated hypothesis testing between different updating rules: non-informative updating, Bayesian updating, and incomplete updating. We find causal evidence that agents imperfectly update their information sets. We also field causal evidence that the amount of additional information provided to subjects relative to their pre-existing information levels can affect stated WTP in ways consistent overload from too much learning. This result raises important (though familiar) issues for the use of stated preference methods in policy analysis.

Keywords: Bayesian, Public Goods, Behavioral Economics, Stated Preference

JEL Codes: Q51, D83, D81

Accepted for presentation at World Congress of Environmental and Resource Economics, Istanbul, June 2014.

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.

I. Introduction

It is important for the economics literature to understand how individuals assimilate new objective information about a good, asset or strategic choice and make subsequent decisions. Such decisions include their maximum willingness to pay for an improvement in environmental quality, or for a reduction in expected damages. While the economics literature often models agents as Bayesian updaters, other models of limited attention, confirmatory bias and asymmetric updating have been developed and supported with experimental evidence (Sims 2003, Rabin and Schrag 1999, Eil and Rao 2011, LaRiviere et. al. 2013). Even though there is some support for each of these models in the literature we are not aware of any study which is explicitly designed to test for which model wins out in an experimental horserace.

Significant design challenges exist in a lab environment to conduct such a horserace. An experiment which tests for different models of updating would need to have agents receiving objective information about a good, asset or strategic choice in a manner which the economist can exogenously vary. One way to introduce such conditions is to utilize stated preference methods in which agents are asked about their willingness to pay for a good or service immediately after receiving detailed information about it. Agents must incorporate this information into their existing ex ante information sets when making stated preference decisions about their willingness to pay for some public good. As a result, stated preference methods are well suited to test updating rules of economic agents in an economically meaningful context. Further, while the stated preference literature acknowledges the importance a previous experience and information, there is no paper which is able to identify the causal effect of information on WTP isolated specifically through learning.

This study uses a large field experiment embedded within a stated preference survey to test different models of information updating and identify the causal effect of information on WTP isolated specifically through the learning channel. The survey concerns a population's willingness to pay for a mixed public good: the regeneration of coastal wetlands acts as a form of flood protection but also provides biodiversity benefits. One aspect of the good is that individuals living in flood plains are less likely to have their property damaged by flooding. Another public aspect of the good is that reclaimed wetlands offer benefits in the form of increased wildlife abundance and biodiversity. Importantly, the results of this survey are being used as part of the policy and management process for coastal flood defenses in Scotland, which makes the survey consequential for respondents, thus incentivizing them to truthfully reveal their demand (Vossler et. al. 2012, Vossler and Evans 2009 and Carson and Groves 2007).

In the experiment, we compare three models of updating: Bayesian, confirmatory bias and bounded rationality. We use a nine question multiple choice test over objective facts about flooding, flood protection and wetlands to elicit prior information levels from subjects. We then randomly assign each subject to an information treatment (low, medium or high) based upon the number of questions answered correctly. We then elicit agents' willingness to pay for a single wetlands restoration project which is uniform across all subjects. We test the subjects' retention of this objective information and identify an objective updating mechanism by giving them the same (identical) quiz at the end of the experiment. Importantly, we are also able to isolate how additional information, and subsequent updating, affects willingness to pay for the good over different levels of ex ante information. Due to this *exogenous* variation in information, we are able to identify the causal effect of information on willingness to pay.

Several important results emerge from the field experiment. The results of the purely informational part of the experiment show that higher information treatments cause significantly more learning in subjects, even though that observed learning is incomplete. This is somewhat surprising because due to our experimental design we are able to isolate the effect not only of providing information to subjects, but also that subjects learn and retain the additional information. However, uninformed subjects who receive and learn little to no new information state similar WTP levels as well-informed individuals who receive and learn all information. Uninformed individuals who receive the most information, and who do in fact learn the most, are willing to pay the least for the good (e.g., significantly less than uninformed individuals who receive and learn little). However, the behavior exhibited by subjects in the WTP portion of the experiment show that providing additional information to subjects does not alter their WTP in a way that is consistent with any of the proposed hypotheses considered in this paper.

This study contributes to the literature in several ways. First, we use a short questionnaire to elicit subjects' previous knowledge levels about a public good and a second identical questionnaire which allows us to estimate the effect of treatment not only on WTP but also on actual updating behavior. Second, we are able to show in a field experiment that agents' updating behavior is consistent with incomplete learning. Third, we show that the amount of information given to agents relative to their pre-existing information sets can have significant effects on willingness to pay estimates and that this occurs via the learning channel. This final result is novel in the economics literature: while previous studies have identified choice overload of economic subjects, we identify "learning overload".

This last point has significant implications for stated preference valuation studies: we find that the amount of information given to subjects matters even though their exhibited behavior is not statistically different from Bayesian updating. This could explain why many behavioral studies are at odds with exhibited updating behavior. We find that the behavioral result doesn't manifest in how the subject is learning information, but rather is a direct function of the amount of information received by the agent. While our design is not capable of parsing the reason behind this finding, we catalogue several possible reasons in the discussion section of the paper.

The remainder of the paper is organized as follows: section two describes the survey and the experimental design in the context of the previous literature. Section three presents results. Section four discusses the results and concludes.

II. Survey, Experimental Design, and Hypotheses

Our experiment has three key components. First, the design allows us to test for how much information respondents possess about the good in question at the outset of sampling: that is, to measure their ex ante knowledge. Second, the design also allows us to test how much of the new information is *retained*. Third, we are able to observe how a priori and new (retained) information affect willingness to pay for a public good. This section describes the design in context of these components.

Survey

We conducted a field experiment in a stated preference survey in Scotland performed during 2013. We set the survey in the context of current efforts by 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 are 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.

In order to gauge the public's willingness to pay for restoring wetlands, we undertook a stated preference survey. Subjects were invited to participate in the survey via repeated mailings and radio and newspaper advertisements. 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. Each subject who participated was given a unique identifier code. In the stated preference survey we embedded the field experiment described below.

Experimental Design

The design of the stated preference survey was as follows: subjects were told that their responses would help inform policy and management of flooding in the area. They were then given a 9 question multiple choice quiz. The quiz was justified so as to inform policy makers as to how well this topic was being communicated to the community. Respondents were then given objective information about flooding, flood protection and wetlands[1]. We then elicited willingness to pay for one specific wetlands restoration project uniform to all participants. Finally, the subjects were given the exact same nine question quiz followed by a series of debriefing questions.

In this survey, we embedded the following field experiment: we gave respondents of the survey an identical nine question multiple choice exam in order to elicit prior information sets regarding their knowledge about flood, flood protection and the public good aspects of increased size of wetlands. The number of correct answers, the specific questions answered correctly and the specific questions answered incorrectly were recorded. We grouped respondents into *a priori types* as a function of the number of correct answers: low (L), medium (M) and high (H). *A priori* type L corresponds to 1-3 correct answers, type M corresponds to 4-6 correct answers and type H corresponds to 7-9 correct answers.

[1] This information is available on request

After agents completed the exam and their answers were recorded, each agent was randomly assigned a *treatment*. A treatment in our case was an amount of information about the attributes of the good. Treatments can be low (L), medium (M) or high (H). Each treatment corresponds to a number (3, 6 or 9 for L, M or H respectively) of bullet points and/or figures conveying precise and objective information about the issue or good. Each bullet point and/or figure corresponds exactly to one question asked on the multiple choice questionnaire. The quiz and complete set of bullet points are in the Appendix A. As a result, after treatment assignment each agent can be summarized as a type/treatment pair in addition to information about their correct and incorrect answers. 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). The experimental design is displayed graphically in Figure 1.

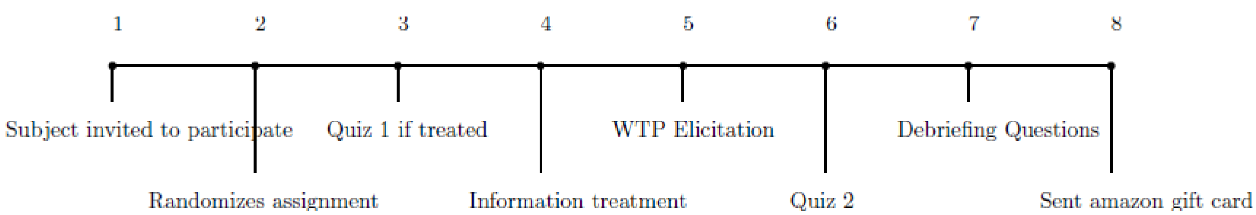

Figure 1: Experimental design

Importantly, respondents were always given information they answered correctly first before any additional information was given as dictated by treatment. For example, assume respondent A gets questions 2 and 7 correct and are in the L treatment. Respondent A is type L since they only got two out of 9 questions correct. The information set they would be provided consisted of two bullet points associated with questions 2 and 7 and, additionally, one information bullet point selected at random from the remaining 7. Alternatively, assume respondent B gets questions 7-9 correct and they are in the M treatment. They are type L since they scored three out of nine. Their bullet points would be the three bullet points associated with questions 7-9 and three randomly chosen bullet points which correspond to questions 1-6.

The reason for not randomly selecting information is that we are concerned with how agents are updating to new information. In order for the experimental design to be valid, we must make sure that, on average, a type-treatment pair of LL is the proper counterfactual for type-treatment pair LM. If the information treatment does not span the agent's a priori information set (e.g., an individual's type), then the proper counterfactual cannot be ensured. Specifically, imagine the situation above in which respondent A gets questions 2 and 7 correct but their L treatment are bullet points associated with questions 3, 4 and 5. In that case, respondent A could test as a type M ex post when their information set is elicited later in the protocol.

Treatment			
H	LH	MH	HH
M	LM	MM	--
L	LL	--	--
	L	M	H

Ex Ante Information

Table 1: Type Treatment Pairs. The x axis of the table represents the groupings (H, M, L, 0) by the first test score and the y axis represents the groups based upon treatment.

One useful way to represent the type-treatment pairs and treatment information sets is shown in Table 1. The x axis of the table represents the types (L, M, H) defined by the a priori test score, and the y axis represents the groups based upon treatment. In general, there are nine potential type-treatment pairs. It is important to note, however, that some of these pairs may be uninformative. For example, if someone has a high information level ex ante (type H) then they will learn no new information when given the low treatment. Alternatively, if someone has a low information level ex ante (type L) then they could learn new information when given the high treatment and subsequently have any ex post information level (L, M, or H). We therefore restrict ex ante H information types to receive only H information treatments and ex ante M information types to receive only M and H treatments to maximize power of the experiment and focus on the effect of additional information. We also use one control in which respondents were not given a quiz beforehand, given the high information treatment and then quizzed afterward in order to have a baseline which corresponds to most "standard" stated preference exercises.

After the quiz, and before the information treatments, subjects were all given identical baseline information as to the potential cost of the policy and other background information given to all survey participants. The information treatments were displayed after this uniform information. At this point all agents were asked to select their maximum willingness to pay for the good – wetlands restoration – from a payment card of 20 different prices starting at zero and increasing to >$150. They were only allowed to choose one of these values. After eliciting willingness to pay estimates the agents were given the exact same test as initially and their answers were recorded. Finally, each agent was given the exact same quiz as at the beginning of the survey in addition to a set of personal characteristic and debriefing questions. Thus, at the end of the survey each respondent was summarized by an initial set of quiz answers (a priori information set), a type-treatment pair, a treatment information set (bullet points), two WTP responses, and a second set of quiz answers (ex post information set).

Hypotheses

Combining the initial quiz, the information treatments and second quiz allows us to test what information updating procedure individuals are using in forming their willingness to pay estimates. This subsection introduces each type of updating procedure and shows how the information treatments allow us to identify the updating rule used in the practice. The baseline to be used as benchmark is that agents do not perform any updating, i.e., new information is irrelevant for both mean and variance of willingness to pay estimates.[2]

Changing the information set of agents could possibly affect various aspects of WTP estimates. Previous literature shows that mean and variance of WTP have been shown to be influenced by changes in agents' information set and that properly controlling for how information can influence such changes with stated preference data is important (MacMillan et al. 2006; Aadland et al. 2007; Hoehn et al. 2010, Czajkowski et. al. 2013 and LaRiviere et. al. 2013). While previous experience with a good and information provided during the survey may be important, the underlying updating mechanism is unclear.

There are an increasing number of alternatives to Bayesian updating which have found support in the literature. For example, models of bounded rationality, rational inattention and cognitive load have agents not completely adsorbing new information into their ex ante information set (Sims 2003, and Gabaix et. al. 2006). Alternatively, there is evidence that individuals may filter additional information through priors in a way that confirms whatever bias they may have previously had (Rabin and Schrag 1999, Eil and Rao 2011, and Grossman and Owens 2012). While previous levels of familiarity have been shown to affect mean and variance WTP in a way that is not necessarily inconsistent with Bayesian updating, those studies are not designed to be able to parse between alternative updating models (Christie and Gibbons 2011, Czajkowski et. al. 2013, and LaRiviere et. al. 2013).

[2] New information could also change preference (taste) of some attribute (and so the mean of random utility function coefficient associated with this attribute) or change variance of the taste for this attribute. Put another way, new information could result in changing standard errors of a random utility model. It could also influence many preference parameters, possibly in different ways. For example, it could simultaneously change variance of ALL parameters in the same way, or influence utility function error term - this is scale – so the choices become more / less random if scale is heterogeneous in the population. While these are important issues, in the current paper we restrict our analysis to updating behavior and simple mean WTP for the mixed good and leave these issue to future work.

Now consider the implication of the survey design on the ability to parse between which updating procedure agents are using. For simplicity we estimate the following two equations:

$$Score_i = X'\gamma + 1\{LL_i\}\Gamma_{LL} + 1\{LM_i\}\Gamma_{LM} + 1\{LH_i\}\Gamma_{LH} + 1\{MM_i\}\Gamma_{MM} + 1\{MH_i\}\Gamma_{MH} + 1\{HH_i\}\Gamma_{HH} + \varepsilon_i \quad (1)$$

$$WTP_i = X'\gamma + 1\{LL_i\}\omega_{LL} + 1\{LM_i\}\omega_{LM} + 1\{LH_i\}\omega_{LH} + 1\{MM_i\}\omega_{MM} + 1\{MH_i\}\omega_{MH} + 1\{HH_i\}\omega_{HH} + \varepsilon_i \quad (2)$$

In equations (1) and (2), *X* is a vector of subject specific control variables.[3]

In equations (1) and (2), as before, the two capitals stand for the ex-ante score and the information treatment. There are two left hand side variables which are considered separately. The first is *score* and that specification measures actual learning that occurs conditional on ex ante information levels and treatment. The second is willingness to pay (*WTP*) conditional on ex ante information levels and treatment.

Score

It is important to understand the implications of treatment on information retention in order to reject the feasibility of various hypotheses dictating updating behavior. One useful way to think about information pairs is summarized in Table 2. Table 2 has ex ante information levels on the x axis and ex post information levels on the y axis. There are three important features about Table 2. First, there are three information pairs that should not be feasible if individuals can recall information: ML, HL and HM. For example an individual with an ex ante high information set should never lose information because they are reminded of a subset of information they already knew. This is equivalent to assuming perfect

[3] In the first specification controls act to verify that assignment is random. Put another way, the average effect of additional information on scores (e.g., the various treatment effects) should not be affected by demographic control variables. Conversely, when estimating the effect of WTP on the controls, it could be the case that the effect of additional information on willingness to pay could vary systematically with demographic characteristics. If those demographic characteristics are also correlated with preferences for the good, then adding in controls could affect the estimated coefficients of treatment on WTP.

recall and can be tested empirically. Second, there are three information pairs in which minimal or no learning occurs: LL, MM and HH. The effect of these information pairings on learning (e.g., the first equation) are the increase in score given by the estimated coefficients Γ_{LL}, Γ_{MM}, and Γ_{HH}. There are three information pairings in which learning must have occurred: LM, LH and MH. The effect of these information pairings on WTP (or variance of WTP) are given by Γ_{LM}, Γ_{LH}, and Γ_{MH}. Third, it is possible for information acquisition to be incomplete. For example, it could be the case that an individual of type L is given treatment H and has ex post information M. Put another way, we cannot rule out respondents of summarized by treatment status LH but information status LM.

Ex Post Information

H_{info}	Γ_{LH}	Γ_{MH}	Γ_{HH}
M_{info}	Γ_{LM}	Γ_{MM}	--
L_{info}	Γ_{LL}	--	--
	L	M	H

Ex Ante Information

Table 2: Ex ante information and ex post information levels. Importantly, the cells in this table do not necessarily correspond to any particular treatment. This table represents all possible scenarios, assuming perfect recall, for how much information a subject can have after treatment assuming that each updating rule is feasible.

Now consider the significance of coefficients which would be consistent with three different updating rules introduced above:

No Updating – H_0: $\Gamma_{LL} = \Gamma_{LM} = \Gamma_{LH} > 0$, $\Gamma_{MM} = \Gamma_{MH} > 0$, $\Gamma_{HH} > 0$

In this case, only a priori information matters.

Complete updating – H_0: $\Gamma_{LM} = \Gamma_{MM} > 0$, $\Gamma_{LH} = \Gamma_{MH} = \Gamma_{HH} > 0$, $\Gamma_{LL} \neq \Gamma_{LM} \neq \Gamma_{LH}$

In this case, the information treatment fully determines ex post information levels.[4]

Bounded Rationality – H_0: $\Gamma_{LL} < \Gamma_{LM} < \Gamma_{LH}$, $\Gamma_{MM} < \Gamma_{MH}$

In this case, type L individuals can learn but they can't fully learn in the high information treatment.

WTP

Conditional on learning, there is still a question of how prior information affects WTP relative to how additional information (e.g., being in one treatment versus another) affects WTP. For example, it is not necessarily the case the two individuals that have the same amount of retained information after treatment have the same WTP for the good. Given the design of this experiment, we can horserace different models of how additional information affects WTP. To do so, we consider the three models below that use the WTP estimating equation (2).

Bayesian updating – $\omega_{\cdot LH} = \omega_{\cdot MH} = \omega_{\cdot HH}$, $\omega_{\cdot LM} = \omega_{\cdot MM}$

Subjects' WTP is determined by the ex post level of information, assuming information is retained. This assumes that information has a uniform effect (e.g., prior information levels don't matter, only information levels at time of WTP elicitation).

[4] Note here that we are only concerned with updating behavior. As a result, the motivation for having different ex ante information levels is irrelevant.

Confirmatory Bias – $\omega_{\cdot LL} = \omega_{\cdot LM} = \omega_{\cdot LH} > 0$, $\Gamma_{MM} = \Gamma_{MH} > 0$, $\Gamma_{HH} > 0$

In this case, the endogenous acquisition of information ex ante fully dictates how additional information affects WTP of agents.

"Behavioral" Cognitive Load – $\omega_{\cdot LH} \neq \omega_{\cdot MH} = \omega_{\cdot HH}$, $\omega_{\cdot LM} = \omega_{\cdot MM}$

Regardless of the updating rule, there is a distinct behavioral reaction to being given significantly more information than the subject already has which is different from the what occurs during updating when the marginal amount of information is not as great.

III. Results

Survey and Questionnaire

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. A reminder card was sent two weeks after the first contact attempt. Of 4000 people invited 749 people completed or partially completed the online survey with 562 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.

Summary Statistics

Self-reported socio-demographic statistics that the sample was representative of the local authority areas sampled in terms of age (χ^2 (6) = 63.04, $p < 0.01$) and gender (χ^2 (1) = 6.71, $p < 0.01$). The mean income band was £20,000 - £30,000 and half the respondents worked full time (Table 1). Not shown in the table, 69% of the respondents reported being insured for flood damages.

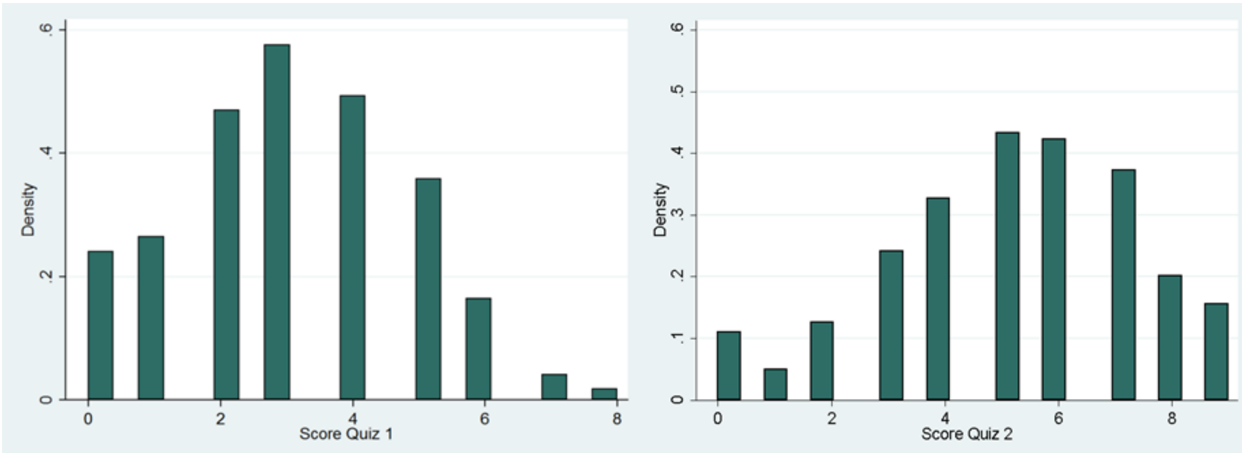

Figure 2: Quiz score histograms by test.

TABLE 3: A Priori Type – Treatment Pairs

LL	153
LM	56
LH	55
MM	91
MH	82
HH	10
Control	79

Note: n = 485 total subjects

At the start of the survey each respondent in a treated group answered identical nine question quizzes concerning objective information about the good. This quiz was then repeated for all respondents after they stated their WTP for all subjects. 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 difference in the scores for quiz one (mean= 3.08, SD=1.76) and quiz two (mean=5.19, SD=2.23). It is important to note that only ten subjects scored between 7-9. As a result, there are only 10 *a priori* type H subjects meaning there are only subjects in the HH treatment. The complete composition of treatment and control groups is shown in Table 3. We over sample from the LL type-treatment group in order to balance the power in estimating this treatment effect relative to the larger information treatments which are more common by the nature of the experimental design.

Information, Learning and Updating

Table 4 shows the results of estimating equation (1) using both OLS and a left hand censored Tobit regression including and excluding various control variables. The estimated coefficients represented the causal effect of being in a particular treatment group on a subject's score on the second quiz. The control group, all of whom received the high information treatment but did not take a quiz before the survey began, is the baseline. In each specification, the control variables do not significantly alter the estimated treatment effects. We take this as evidence that we properly randomized treatment.

TABLE 4: Second Quiz Score on Treatment Group

VARIABLES	(1) OLS	(2) OLS	(3) Tobit
LL	-1.616***	-1.763***	-1.804***
	(0.336)	(0.334)	(0.341)
LM	-0.403	-0.590	-0.570
	(0.425)	(0.430)	(0.434)
LH	0.0689	-0.00115	-0.00786
	(0.476)	(0.445)	(0.453)
MM	0.256	0.107	0.132
	(0.340)	(0.321)	(0.322)
MH	1.053***	0.999***	1.023***
	(0.360)	(0.345)	(0.346)
HH	2.797***	2.340***	2.357***
	(0.384)	(0.373)	(0.371)
Constant	5.403***	6.068***	6.077***
	(0.301)	(0.449)	(0.450)
Controls	no	yes	yes
Observations	485	458	458
R-squared	0.212	0.280	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Dependent Variable is second quiz score. Control group is baseline. Demographic control variables include categorical age, education, gender and flood threat variables. Additional control variable for timing of when subject completed survey.

Table 4 shows that the control group, all of whom received the high information treatment, average between 5.4-6.1 questions correctly. However, the MH and HH treatment groups scored significantly higher than the control group. This is evidence that a key element of the treatment, taking a quiz before elicitation, increased attentiveness and retention, which we address in the appendix.

Turning to the hypothesis tests for updating, we can reject the hypothesis that no informative updating occurs. The null hypothesis that H_o: $\omega_{LL} = \omega_{LM} = \omega_{LH}$ is rejected at the 1% level (F-stat of 11.09). Similarly, we can reject the null hypothesis that subjects exhibit complete retention. The null hypothesis H_o: $\omega_{LH} = \omega_{MH} = \omega_{HH}$ is rejected at the 1% level (F-stat of 18.03). We fail to reject, though, the null hypothesis of bounded rationality by the strict definition given above. It is clear from the coefficients on LL, LM and LH that information monotonically increases scores (similarly for MM and MH). We take this as evidence that our information treatments cause subjects to learn, but that learning is incomplete. This is evidence that the experimental design for the causal effect of not just information, but also learning on WTP for the public good is valid.

Willingness to Pay

Figure 3 shows a histogram of all subjects who completed the survey. Figure 3 shows that demand for this good is downward sloping and that subjects exhibit non-trivial anchoring around 50, 100 and 150 pounds. Importantly, though, there is significant heterogeneity in WTP for this good.

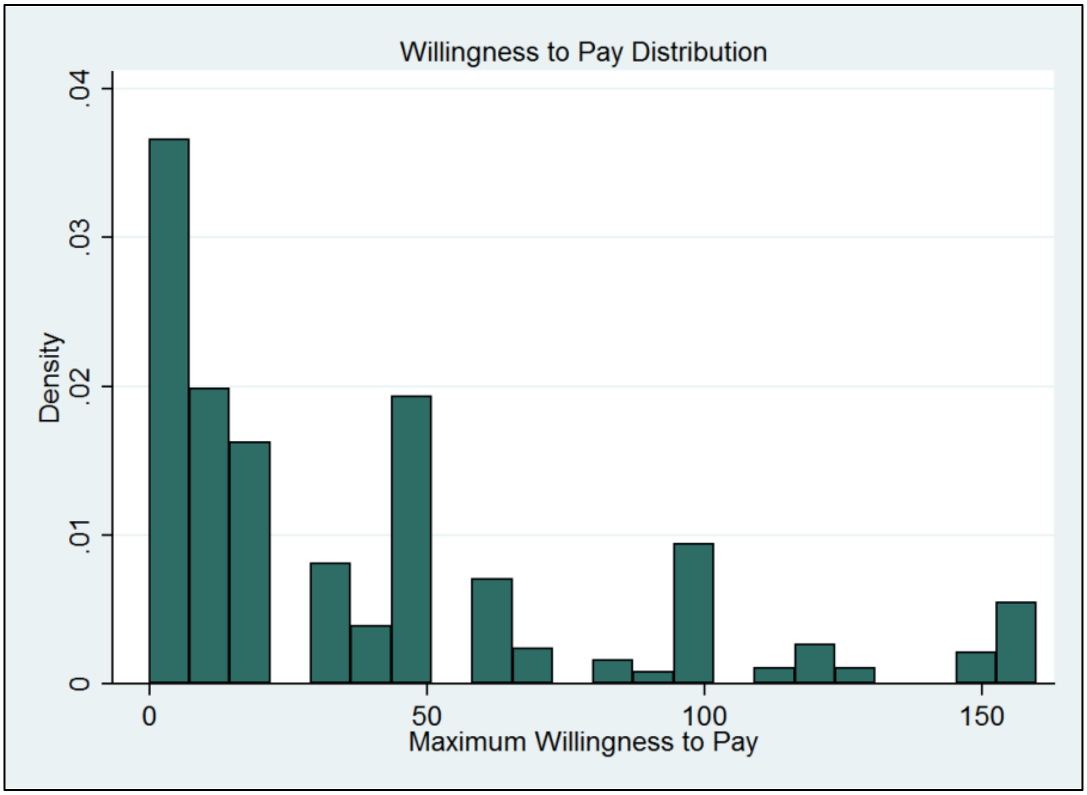

FIGURE 1: Histogram of WTP for all subjects. N = 562.

Table 5 shows the coefficient estimates of regression (2) with WTP as the dependent variable and treatment group (and controls) as independent variables. As before the baseline is the control group, who received the H information treatment.[5] The WTP for three treatment groups never differed from the control group: LH, MM and HH. However subjects in both the LL and MH treatments exhibit significantly higher WTP than the control group across all specifications. Further, adding in control variables alters the coefficient estimates of treatment on WTP. This implies that the causal effect of information and learning on WTP is correlated with demographic characteristics which are also correlated with preference for the public good.

[5] Also as before, some observations are dropped when control variables are included since some subjects chose to not respond to questions about where they lived and education attainment.

Turning to hypothesis testing, the null hypothesis that only information treatment matters for WTP is H_0: $\omega_{LH} = \omega_{MH} = \omega_{HH}$, $\omega_{LM} = \omega_{MM}$. We fail to reject this null hypothesis (p-value = .25, F-stat = 1.39). The null hypothesis of confirmatory bias is H_0: $\omega_{LL} = \omega_{LM} = \omega_{LH} > 0$, $\Gamma_{MM} = \Gamma_{MH} > 0$. We also fail to reject that null hypothesis (p-value = .18, F-stat = 1.61). We also fail to reject the null hypothesis that all treatment groups are the same (p-value .37, F-stat 1.09) and or that the effect of treatment on WTP is jointly zero. Individual t-tests show, though, that the MH and LL treatment groups exhibited systematically higher WTP for the good.

TABLE 5: WTP on Treatment Group

VARIABLES	(1) maxwtp	(2) maxwtp	(3) model
LL	10.79*	14.38**	18.60**
	(5.687)	(6.034)	(7.501)
LM	11.60	15.22*	18.76*
	(7.323)	(8.194)	(9.648)
LH	9.824	2.286	3.989
	(7.665)	(7.678)	(9.310)
MM	5.815	5.524	8.860
	(5.760)	(6.196)	(7.602)
MH	14.43**	13.12*	18.80**
	(6.834)	(7.046)	(8.261)
HH	7.051	11.10	19.75
	(10.16)	(12.42)	(12.86)
Constant	30.45***	42.65***	29.69**
	(4.301)	(11.42)	(13.45)
Controls	No	Yes	Yes
Observations	525	458	458
R-squared	0.011	0.169	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Dependent Variable is WTP. Control group is baseline. Demographic control variables include categorical age, education, gender and flood threat variables. Additional control variable for timing of when subject completed survey.

IV. Discussion

The results of the purely informational part of the experiment show that higher information treatments cause significantly more learning in subjects, even though that observed learning is incomplete. The results of the WTP portion of the experiment show that providing additional information to subjects does not alter their WTP in a way that is consistent with any of the proposed hypotheses considered in this paper. This is somewhat surprising because due to our experimental design we are able to isolate the effect not only of providing information to subjects, and how subjects learn and retain the additional information. In fact, uninformed subjects who receive and learn little to no new information (e.g., LL treatment) exhibit similar preferences to well-informed individuals who receive and learn all information (e.g., MH treatment). Recall that we observe only 10 subjects in the HH treatment making inference challenging). Subjects in all other treatment groups do not state a WTP that is significantly different than the control group, all of whom received the high information treatment. Put another way, the subjects who learned the least from the survey and the subjects who learned the most from the survey were significantly more likely to be willing to pay more for the public good. Subjects willing to pay the least for the good are those in the

This non-linear causal effect of information on learning on WTP is remarkable. Recalling that information causes learning in our experiment, subjects in the LL treatment have the least amount of information and MH have the most. However both groups are willing to pay significantly more for the good. Further, the causal effect of information and learning on WTP is correlated with demographic characteristics which are also correlated with preference for the public good. There are two intuitive implications of these results. First, additional information and learning can alter WTP for a good if agents are not well-informed. We find that in our case additional information and learning caused a decrease in WTP for the good. This could be due to updating of preferences as a function of additional information. However, we do observe that well-informed individuals (MH treatment) are also WTP significantly more for the good. An alternative is "information overload" in incorporating new information on making economic decisions. There is a well-developed literature on choice overload (Iyengar and Lepper 2000). However, it is less clear as to what the mechanism for this result. For example, search costs have been posited as one explanation (Kuksoc and Villas-Boas 2010). We find evidence for the channel associated with learning about the good *in and of itself*. We are not aware of any work in the economics literature which finds evidence for this channel.

Second, we find evidence that information and learning can affect WTP for a good but that those effects are systematically correlated with demographic characteristics. It makes sense that experienced decision makers, or decisions makers more accustomed to quickly incorporating newly learned information into their preference structure, are not randomly distributed throughout the population. As a result, in field experiments, lab experiments and surveys in general, it is very important to take proper subsamples of the population in order to perform inference.

More generally for stated preference studies, we find that agents do indeed learn about the good on stated preference surveys. Our results suggest that learning this objective information about a good during a survey can significantly affect WTP estimates. This is the first evidence we are aware of in the literature which isolates the effect of *learning* information on a stated preference survey on WTP estimates. However, the way in which agents learn does no accord well with any well-formed theoretical model in the literature. More work along these lines is clearly needed.

References

Aadland, D., Caplan, A., and Phillips, O., 2007. A Bayesian examination of information and uncertainty in contingent valuation. *Journal of Risk and Uncertainty*, 35(2):149-178.

Carson, R., and Groves, T., 2007. Incentive and informational properties of preference questions. *Environmental and Resource Economics*, 37(1):181-210.

Christie, M., and Gibbons, J., 2011. The effect of individual 'ability to choose' (scale heterogeneity) on the valuation of environmental goods. *Ecological Economics*, 70(12):2250-2257.

Czajkowski, M., Hanley, N., and LaRiviere, J., 2013. The Effects of Experience on Preference Uncertainty: Theory and Empirics for Public and Quasi-Public Environmental Goods. *University of Tennessee Working Paper*.

Eil, D., and Rao, J. M., 2011. The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. *American Economic Journal: Microeconomics*, 3(2):114-138.

Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S., 2006. Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model. *The American Economic Review*, 96(4):1043-1068.

Grossman, Z., and Owens, D., 2012. An unlucky feeling: Overconfidence and noisy feedback. *Journal of Economic Behavior & Organization*, 84(2):510-524.

Hoehn, J. P., Lupi, F., and Kaplowitz, M. D., 2010. Stated Choice Experiments with Complex Ecosystem Changes: The Effect of Information Formats on Estimated Variances and Choice Parameters. *Journal of Agricultural and Resource Economics*, 35(3):568-590.

Iyengar, S. S. and Lepper, M.R., 2000. When Choice is Demotivating: Can One Desire Too Much of a Good Thing? *Journal of Personality and Social Psychology,* 79(6): 995-1006.

Kuksoc, D. and Villas-Boas, J.M., 2010. When More Alternatives Leas to Less Choice. *Marketing Science*, 29(3): 507-524.

LaRiviere, J., Czajkowski, M., Hanley, N., Aanesen, M., Falk-Petersen, J., and Tinch, D., 2013. The Value of Familiarity: Effects of Experience, Knowledge and Signals on Willingness to Pay for a Public Good. *University of Tennessee Working Paper*.

MacMillan, D., Hanley, N., and Lienhoop, N., 2006. Contingent valuation: Environmental polling or preference engine? *Ecological Economics*, 60(1):299-307.

Rabin, M., and Schrag, J. L., 1999. First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics*, 114(1):37-82.

Sims, C. A., 2003. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665-690.

Vossler, C. A., Doyon, M., and Rondeau, D., 2012. Truth in Consequentiality: Theory and Field Evidence on Discrete Choice Experiments. *American Economic Journal: Microeconomics*, 4(4):145-171.

Vossler, C. A., and Evans, M. F., 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.

Appendix : Modeling Updating Rules

We now introduce a simple theoretical model which shows the difference between these three classes of models. Consider a model in which there are two potential states of the world $x \in \{A, B\}$. An agent trying to infer the true state of the world has a prior that the state of the world is A: $pr(x = A) = \rho$ such that $pr(x = B) = 1 - \rho$. In our case, ρ might be the probability an agent believes they are a high value or low value consumer for a good they have little experience consuming. In that case, A may represent the level of utility from consumption if the agent is a high type and B if the agent is a low type such that $A > B$. There are signals at a time period s_t which inform agents as to their true type such that $s_t \in \{a, b\}$ and $pr(s_t = a \mid A) = pr(s_t = b \mid B) = \theta \in (0.5, 1)$. Note that signals are informative but not perfectly informative given the support of θ.

First, consider the properties of utility conditional on receipt of a signal given the above model given a prior ρ that an agent is a high type consumer and they are a Bayesian updater. The expected utility conditional on the signal for a Bayesian updater is $E(U \mid \rho) = \rho A + (1-\rho)B$ with variance $AB\rho(1-\rho)$. This variance is single peaked with a maximum at $\rho = 0.5$ and equal to zero for $\rho \in \{0,1\}$. As a result, the more informative signals an agent receives through consumption, the lower the variance from consumption as ρ is updated based upon consumption experience.[6]

Second, while it is tempting to think that more information will always reduce the significance of noise relative to the deterministic component of utility, recent work suggests that might not always be the case. Gabaix et al. (2006) shows both theoretically and experimentally that agents making decisions under cognitive load use “directed cognition”. Directed cognition means that agents compute the expected

[6] For example, in the model presented here, the posterior probability an agent is a high type conditional on observing a signal $s_t = a$ is $\rho_{t+1} = \rho\theta \big/ \left(\rho\theta + (1-\rho)(1-\theta)\right)$.

benefits of gathering new information relative to focusing on accurately processing the set of existing information in models of rational inattention (Sims 2003). Because of cognitive load, new information would not always decrease error variance and, more importantly, the importance of new information can vary across agents in the economy if they have heterogeneous a priori information. Taken together, the implication is that agents less familiar with the good will have a larger error variance and the relative decrease in their error variance due to new information may vary as a systematic function of previous levels of information. Further, agents with a relatively complete information set may ignore new information in order to process their existing information.

One reduced form model of rational inattention is to explicitly assume that there is a utility cost to processing new information. More formally, if using a signal s_t in forming a posterior distribution occurs at a utility cost C then the expected utility in future periods conditional of processing a signal represented by a value function $V_{t+1}\left(\rho_{t+1} \mid s_t\right) - C$ versus the value function without processing the additional signal $V_{t+1}\left(\rho_t\right)$. Sims (2003) shows that curvature of the utility function will dictate that it is rational to sometimes disregard additional signals due to the cost C depending on the properties of ρ. It could be that additional information has asymmetric effects across the population as a function of the expected variation in utility (e.g. $AB\rho\left(1-\rho\right)$).

Third, additional information may not always affect agents in the same way conditional on the set of a priori information. Rabin and Schrag (1999) show that over-confidence can lead to agents responding to identical information in different ways. For example, if there is any chance that an agent can interpret information incorrectly and update their beliefs subject to the misperceived information, subsequent information may be perceived incorrectly with increasing frequency. The resultant "mis-informed" equilibrium be driven by chance and as a result, the researcher may wish to be agnostic as to how information can affect changes in error variances for individual agents across the economy.

Formally, in the Rabin and Schrag (1999) model of confirmatory bias, agents may misperceive signals and instead of observing the true signal s_t, they actually observe $\sigma_t \in \{\alpha, \beta\}$. More precisely, if the agent's prior shades toward state A as the true state, $\rho > 0.5$, there is a probability $q > 0$ that the agent misreads a signal $s_t = b$ as $\sigma_t = \alpha$ with probability q. As the probability of falsely interpreting information in favor of state A increases (e.g., q increases) so does the confirmation bias. Given a path of perceived signals $\{n_\alpha, n_\beta\}$, if at any point $pr(x = A \mid n_\alpha, n_\beta) > pr(x = B \mid n_\alpha, n_\beta)$ then further signals are biased. In the context of contingent valuation, then, the same information can lead to varying levels of relative certainty over preferences conditional on heterogeneous priors.

All the behavioral models outlined above have the property that the variance in utility from a good conditional on a signal can be a function of amount of incoming information. In a random utility framework, such as those often used in contingent valuation in which new information is presented to respondents, the implication of accounting for these behavioral issues is letting the idiosyncratic component of utility vary across the population. Formally, the implication is that the variance of utility conditional on an information set provided in a stated preference survey, I, and prior experience with the good, $\{n_\alpha, n_\beta\}$:

$$\begin{aligned} \operatorname{var}\left(U_i\left(\mathbf{X}\boldsymbol{\beta} \mid I, n_\alpha, n_\beta\right)\right) &= \operatorname{var}\left({}'_{i} + \varepsilon_i \mid I, n_\alpha, n_\beta\right) \\ &= 0 + \operatorname{var}\left(\varepsilon_i \mid I, n_\alpha, n_\beta\right) \end{aligned} . \tag{1}$$

Since there is no way to control for previous levels of experience, (n_α, n_β), the above work in behavioral economics suggests that the variance of utility should be allowed to vary across the population.

PART C: Effect of Treatment on Second Quiz Score

We run the follow regression to identify the causal effect of taking a quiz before being surveyed on subject. Since the control group received the H information treatment, we restrict the sample to treated individuals who also received the H information treatment (e.g., LH, MH and HH).

$$Score_i = X'\gamma + 1\{\text{treated}\}\Gamma \quad + \varepsilon_i \qquad \text{(A.c.1)}$$

$$WTP_i = X'\gamma + 1\{\text{treated}\}\omega \quad + \varepsilon_i \qquad \text{(A.c.2)}$$

Output from these regressions is displayed in tables A1 and A2 below. From Table A1, there is clearly an effect of being treated (e.g., taking a quiz before the survey) on second quiz score. The implication is that the quiz causes subjects to learn more. From Table A2, there is some evidence that being treated causes an increase in WTP for the good. However, it is unclear if the increase is due to additional learning or if it is due to being treated. Furthermore, adding in controls attenuates the effect of treatment on WTP in the OLS specification. This implies that either learning or being treated is systematically correlated with demographic characteristics which are correlated with WTP. Lastly, in order for our experimental results to be valid, we only need the effect of being in a treated group to be uniform across the population, which is a plausible assumption.

TABLE A1: Score on Treatment Status Given High Info Treatment

VARIABLES	(1) OLS	(2) OLS	(3) Tobit
treated	0.808**	0.760**	0.770**
	(0.353)	(0.335)	(0.338)
Constant	5.403***	6.592***	6.626***
	(0.300)	(0.659)	(0.654)
Controls	No	Yes	Yes
Observations	214	201	201
R-squared	0.026	0.175	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Dependent Variable is second quiz score. Control variables include categorical age, education, gender and flood threat variables.

TABLE A2: WTP on Treatment Status Given High Info Set

VARIABLES	(1) OLS	(2) OLS	(3) Tobit
treated	12.20**	8.874	13.40*
	(5.762)	(6.458)	(7.454)
Constant	30.45***	31.90*	15.74
	(4.292)	(17.28)	(20.67)
Controls	No	Yes	Yes
Observations	225	201	201
R-squared	0.017	0.161	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Dependent Variable is WTP. Control variables include categorical age, education, gender and flood threat variables.