

Essays on Individual Choice and Behavior

A Dissertation

Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Caleb Andrew Siladke

August 2012

© by Caleb Andrew Siladke, 2012
All Rights Reserved.

To the love of my life,

Tiffany

Acknowledgements

I would like to express my deep gratitude to my adviser, Christian Vossler. His support, guidance, and advice were indispensable to my development. He has been patient, reliable, and challenging and I am extremely thankful for the opportunities I've had to work with him. I would like to thank William Neilson who is responsible for my knowledge of non-expected utility theory, and the other members of my committee, Michael Price and Christopher Clark for their valuable feedback and overall support. For their encouragement, support, and unfailing love, I would like to thank my parents, Richard and Kathleen. Finally, I would not have been able to accomplish what I have without the love and encouragement of my best friend and fiancée Tiffany.

Abstract

This dissertation consists of three chapters that explore individual choice and behavior. Chapter 1 investigates the incentive properties of advisory referenda using a particular form of non-expected utility theory which replaces the independence axiom assumed in expected utility theory with two less-restrictive assumptions: betweenness and fanning-out. Betweenness replaces the independence axiom and allows for context dependent risk attitudes. The fanning-out hypothesis then governs the precise way in which risk preferences change given the unique circumstances in which values are elicited. When the assumption of independence is relaxed, an individual's response to an advisory referendum depends on how consequential she believes her response to be, and her beliefs regarding the existence of alternative proposals.

Chapter 2 utilizes laboratory experiments to investigate the behavioral dynamics pertaining to information acquisition and tax evasion. In recent years, a “service” paradigm, whereby tax authorities provide information about correct tax reporting to taxpayers, has shown the potential to further “encourage” correct tax reporting. The results show that the overall effect of a helpful information service is to decrease tax evasion. Further, an audit has the behavioral effect of lowering information acquisition rates and increasing evasion immediately after experiencing a penalty. This effect persists (although diminishes) in subsequent tax reporting decisions.

Chapter 3 builds upon the existing experimental literature regarding “decoy effects through an innovative design, which preserved the fundamental features of a consumer choice setting in the laboratory. The design involves choices with financial

consequences, real consumer goods, and the ability for participants to opt-out. Through the novel experimental approach and econometric analysis, we demonstrate that the decoy effect is not an artifact of hypothetical settings, the decoy effect is not driven by forced choice, and that decoys do little more than sway individuals at the point of indifference.

Contents

List of Tables	x
List of Figures	xii
1 The Truth Lies in Consequences: The Incentive Properties of Binary Choice Contingent Valuation Methods Without Expected Utility	1
1.1 Introduction	1
1.2 Individual choice setting assuming expected utility	5
1.3 Independence, and the Allais Paradox	6
1.4 Expected Utility Analysis Without the Independence Axiom	10
1.5 Single binary choice questions without the independence axiom	14
1.6 Discussion	18
2 Behavioral dynamics of tax compliance under an information services initiative	20
2.1 Introduction	20
2.2 Conceptual Framework	23
2.2.1 Basic economic theory	24
2.2.2 Insights from behavioral economics	25
2.3 Experimental Design	26
2.3.1 Decision setting	26
2.3.2 Treatments	30

2.3.3	Participants and procedures	31
2.4	Testable hypotheses	32
2.4.1	Basic economic hypotheses	32
2.4.2	Testable behavioral hypotheses	34
2.5	Results	35
2.5.1	Basic analysis of treatment effects	37
2.5.2	Behavioral Dynamics	41
2.6	Conclusion	45
3	How Relevant is Irrelevance?	
	Testing Independence with Increasingly Irrelevant Alternatives	47
3.1	Introduction	47
3.1.1	Related literature	49
3.2	Open questions	52
3.3	Conceptual framework and testable hypotheses	54
3.3.1	Willingness to pay	55
3.4	Experimental Design	56
3.4.1	Treatments	59
3.4.2	Product Categories	60
3.4.3	Participants and procedures	61
3.5	Econometric Method	64
3.5.1	Willingness to pay estimation	64
3.5.2	Finding the point of indifference with random prices	65
3.6	Results	67
3.6.1	Analysis of willingness to pay	70
3.7	Conclusion	74
	Bibliography	77

A	Appendix A	86
A.1	Proof of Proposition 1	86
A.2	Proof of Proposition 2	88
B	Appendix B	91
B.1	Theory of line-item reporting	91
B.2	Optimal evasion given tax policy parameters	93
C	Appendix C	94
C.1	Selected Experiment Screenshots	94
D	Appendix D	97
D.1	Experiment Instructions	97
E	Appendix E	99
E.1	Selected Experiment Screenshots	99
F	Appendix F	103
F.1	Example results from the conditional logit estimation	103
Vita		105

List of Tables

2.1	Experiment Parameters	28
2.2	Experiment Treatments	30
2.3	Data Description	36
2.4	Tax Reporting Model I	37
2.5	Information Services and expected under-reporting	38
2.6	Information Acquisition Model I	40
2.7	Tax Reporting Model II	42
2.8	Information Acquisition Model II	44
3.1	Choice set composition by treatment	57
3.2	Treatments by main effects	60
3.3	Treatment pairs by session	63
3.4	Hypothetical Forced Choice	68
3.5	Real forced choice	69
3.6	Hypothetical Opt-out	70
3.7	Real Opt-out	71
3.8	Differences in marginal willingness to pay across treatments	72
3.9	Differences in estimated target share across treatments	72
3.10	Willingness-to-pay: Parametric Estimation	73
3.11	Opt-out rates	74
3.12	Willingness-to-pay: Non-parametric Estimation	75

B.1 Evasion calculation by income	93
---	----

List of Figures

1.1	Expected utility indifference curves in the unit triangle	11
1.2	Risk preferences in the unit triangle	12
1.3	Expected utility indifference curves and the Common Consequence effect	12
1.4	Expected utility indifference curves and the Common Ratio effect . .	12
1.5	Indifference curves exhibiting betweenness	13
1.6	An indifferent voter's choice problem in the unit triangle.	14
1.7	Strict preference for the status-quo option over the proposal	15
1.8	Illustration of Proposition 4(ii), where $A \sim B$ when $\mu_i = 1$, and $C \succ D$ represent a decrease in μ_i , and therefore in order for $C \sim D$ it must be the case that $b_i^* > c_{ij}^*$	18
3.1	Illustration of experimental choice set conditions	54
C.1	Income earnings task	94
C.2	Treatment 3, Tax decision screen, information requested	95
C.3	Treatment 3, Tax decision screen, after information acquired	95
C.4	Audit selection process	96
C.5	Results screen	96
E.1	Experimental Instructions	100
E.2	Example product purchase scenario	101
E.3	Example results screen	102

Chapter 1

The Truth Lies in Consequences: The Incentive Properties of Binary Choice Contingent Valuation Methods Without Expected Utility

1.1 Introduction

A *referendum* is a direct vote in which constituents are asked to either accept or reject a particular proposal, and is a form of direct democracy. Constitutional amendments, laws, and specific government policies are often accepted or rejected based on the outcome of a referendum. Referenda are common to many state and local ballots, and are generally viewed as an acceptable mechanism for social choice. Under reasonable assumptions, a binding referendum is incentive compatible, i.e. the agent has an incentive to truthfully reveal his preference (Farquharson, 1969).

An *advisory referendum* is a non-binding referendum that is used to gauge public opinion on an important issue. The uses of advisory referenda vary. In some cases, the outcome is used to determine whether a binding referendum on the same proposal,

or an alternative proposal, should be held at a later date. In others, the outcome is used as an input to the policy process. The proposals considered through advisory referenda are diverse. As examples, the City of Chicago held an advisory referendum on whether to be a “nuclear-free zone”. In Vermont, there was a referendum on whether the General Assembly should consider enacting a lottery to generate revenue for the state.

In addition to advisory referenda held in general elections, economists commonly use survey-based advisory referenda for the purpose of estimating values associated with non-market goods, such as environmental quality. In this context, advisory referenda represent a particular stated preference elicitation mechanism. It is widely known that stated preference methods are the only approach for measuring non-use values, as well as potential policy outcomes that are beyond the scope of existing data. In this study, under both expected utility (EU) and a form of non-expected utility (NEU) theory, I investigate the incentive properties of advisory referenda when participants believe, or otherwise when it is explicitly known, that alternatives to the proposal they are asked to vote on are simultaneously being considered by the agency.

In recent theoretical work (Carson and Groves, 2007; Vossler et al., 2009), a set of sufficient conditions under which advisory referenda (in the survey context or otherwise) are incentive compatible: (i) participants care about the outcome of the proposal; (ii) the authority has the ability to enforce payment; (iii) the elicitation involves a yes or no vote on a single project; and (iv) the probability that the proposed project is implemented is weakly monotonically increasing with the proportion of yes votes. Carson and Groves (2007) classify settings where conditions (ii) and (iv) hold as consequential. Recent field research suggests that elicited willingness-to-pay (WTP) differs between respondents who believe the survey to be consequential and those who do not (Bulte et al., 2005; Nepal et al., 2009; Herriges et al., 2010). Laboratory experiments suggests that advisory referenda that arguably meet the above sufficiency conditions, including consequentiality, are demand revealing (Carson et al., 2006; Landry and List, 2007; Vossler and Evans, 2009). Further, a recent

external validity study by [Vossler and Watson \(2012\)](#) finds that the results of an advisory survey, conditional on consequentiality, match the outcome of a parallel, binding public referendum on the provision of local public goods.

When considering the debate regarding consequentiality, it is important to note that conditions (*ii* – *iv*) rely on participant beliefs that are at least somewhat outside the control of the survey researcher. So while the focus in the literature has been on examining the role of consequentiality, condition (*iii*) may be equally important. In actuality, it may be known that the agency is simultaneously considering multiple proposals. For example, although voters are asked about whether the city should purchase specific undeveloped lands for purpose of open space preservation, through town meetings and media voters may be informed of alternative proposals, such as those involving different land parcels. More generally, voters are left to their beliefs, and even in the absence of explicit information they may nevertheless believe that other proposals are under consideration. For example, a voter may view a proposal to cut greenhouse gases by 50% to be too ambitious and instead believe that her vote in an advisory referendum may help determine whether a more modest proposal (say, one that would lead to a 20% reduction) will be implemented.

The purpose of this study is to investigate the incentive properties of advisory referenda when violations of condition (*iii*) occur, in the form described above, using a particular form of NEU theory developed by [Machina \(1982\)](#) and [Dekel \(1986\)](#) to accommodate the typical behavior observed in related choice settings. Specifically, the preference representation used in this paper replaces the independence axiom assumed in EU theory with two less-restrictive assumptions: betweenness and fanning-out. Betweenness replaces the independence axiom and allows for context dependent risk attitudes. The fanning-out hypothesis then governs the precise way in which risk preferences change given the unique circumstances in which values are elicited.

Theoretically, an individual votes on an advisory referendum according to her certainty equivalent for the uncertain outcome of the proposal, which defines her WTP for the proposal. When the individual perceives there to be another proposal

for the same project, she votes according to her certainty equivalent for the uncertain outcome of the proposal she was asked to consider, *conditional* on the existence of the other proposal, and this defines her *threshold acceptable* cost for the proposal. The analysis leads to three main insights:

First, under NEU, an individual's risk attitude changes according to how much better or worse she perceives the specific proposal she is considering to be in relation to the set of all possible proposals. For example, the better an individual perceives the proposal she is considering to be in relation to the set of all possible proposals, the more risk averse she becomes. The individual votes according to the comparison between the proposal's cost and her threshold acceptable cost. Therefore, becoming more risk averse corresponds to a threshold acceptable cost which is now lower than her true WTP for the proposal.

Second, there is an interesting interaction between the "degree" of consequentiality and an individual's threshold acceptable cost. Depending on the circumstances, a decrease in consequentiality could potentially increase or decrease the likelihood that she votes in favor of a proposal, in relationship to a vote cast according to her true value.

Third, under EU theory, beliefs about consequentiality and the possibility that other proposals are being considered do not lead to a loss of incentive compatibility. This is an important consideration because EU lacks the flexibility and explanatory power to allow beliefs to change an individual's response to a specific proposal in this manner and still be able to accurately estimate WTP for that proposal.

The paper proceeds as follows: Section 1.2 develops a formal model for individual choice in the context of advisory referenda assuming EU preferences. Section 1.3 highlights the similarity between advisory referenda and the common consequence and common ratio effects, for which the particular form of NEU used in this paper was explicitly developed to explain. Section 1.4 explains how the NEU theory is different from EU theory. Section 1.5 formally develops the potential

consequences NEU preferences have on non-market values estimated from advisory referenda, and Section 1.6 closes with discussion.

1.2 Individual choice setting assuming expected utility

A government agency is considering what project to undertake from the finite set $A = \{a_1, \dots, a_n\}$. The alternatives in A are binary choice proposals. As part of the decision process, the agency must assess constituent values for the different elements of each proposal. Proposals are characterized by two variables, the specific project a_j , and the proposed cost c_{ij} of implementing the proposal. Constituents are simply asked whether they would vote for or against proposal a_j at an individual cost of c_{ij} .

From an individual's point of view, her vote on a proposal has only some probability μ_i of determining the outcome of the process. Implicitly, the choice given to her by the agency is between either implementing project a_j at cost c_{ij} or implementing nothing at zero cost. Taking the nature of the individual choice task into account, individual i believes that if she votes in favor of proposal (a_j, c_{ij}) then with probability μ_i the project a_j will be undertaken at a cost c_{ij} to her. By the same token, if she votes against the proposal then there is the same probability μ_i that no project will be undertaken. Because of this, μ_i measures individual i 's beliefs about her *consequentiality*.

The value of proposal a_{ij} may differ from individual to individual, and from the voter's point of view, at least, that value might be stochastic. To account for this, let $\tilde{x}_{ij}$ denote the random variable governing the monetary value of project j offered to voter i .

Individual i also forms beliefs about what will happen if she is not consequential, that is, if the agency decides to act on some proposal other than the one given to her. In that case the agency might choose some other project from the set A , or it

might choose to do nothing at all. The individual forms beliefs over the net benefits accruing to her, and denote the resulting random variable by $\tilde{s}_i$.

Individual i 's decision problem becomes a simple one. She votes in favor of the proposal if and only if

$$\mu_i(\tilde{x}_{ij} - c_{ij}) + (1 - \mu_i)\tilde{s}_i \succeq \mu_i(\tilde{0}) + (1 - \mu_i)\tilde{s}_i, \quad (1.1)$$

where $\tilde{0}$ denotes the degenerate random variable that pays zero with certainty.


The decision faced by the individual is composed of three separate sources of uncertainty: their value for the final proposal ($\tilde{x}_{ij}$), their value for the net benefits occurring from the agency choosing a different proposal ($\tilde{s}_i$), and the probability that the proposal she considered will be the final outcome from the proposal (μ_i). In such a setting the probability μ_i can be broken into two parts, the probability that the cost identified in the proposal being voted on is the one that will be used, and the probability that the individual under consideration is pivotal in the vote.

From this characterization, it is clear that under EU the individual's voting decision is independent of the value of μ_i and $\tilde{s}_i$ so long as $\mu_i > 0$. The voting decision is based on comparing the value for the proposal determined by the individual's certainty equivalent for the random variable $\tilde{x}_{ij}$ to the cost for the proposal c_{ij} .

1.3 Independence, and the Allais Paradox

The independence axiom states that for any three lotteries X , Y and Z where $X \succeq Y$, then for every $0 < \alpha < 1$ it must be the case that $\alpha X + (1 - \alpha)Z \succeq \alpha Y + (1 - \alpha)Z$. In other words, the choice between X and Y is independent of both α and Z .

If we consider the choice presented in (1.1), between voting “for” and voting “against” a proposal in an advisory referendum, an individual is choosing between



According to the independence axiom, the voter's choice is independent of both μ_i and $\tilde{s}_i$.

The example above highlights the reliance of the advisory referendum on the independence axiom in order to obtain truthful preference revelation from voters. People vote in favor of the proposal as long as the proposal has a higher value than the status-quo, regardless of the extent to which she thinks the agency will act on some proposal other than the one she was asked to consider.

Independence is a common assumption and is necessary for expected utility analysis. Unfortunately, systematic violations of the independence axiom are common as well. Two of the more well known violations of the independence axiom are the “Common Consequence” effect and the “Common Ratio” effect.

The common consequence effect is best highlighted in the following example, made famous by [Allais \(1953\)](#). An individual is given the choice between two lotteries:


- Choice A: 10% chance of \$500 million; 89% chance of \$100 million; 1% chance of nothing.

- Choice B: Certainty of receiving \$100 million.

In most instances, individuals choose lottery B. The same individual is then given the choice between two different lotteries:

- Choice A': 10% chance of \$500 million; 90% chance of nothing.
- Choice B': 11% chance of \$100 million; 89% chance of nothing.

Many choose lottery A' in this second choice set that had chosen B in the first choice set, which is a violation of the independence axiom. This is more clear through presenting both choice sets as



When $Z = \$0$ most prefer A (from the second choice set) and when $Z = \$1M$ most prefer B (from the first choice set). It is easy to see that this is a violation of independence, as someone that prefers B from the first choice set should prefer B' from the second.

The common ratio effect is another violation of the independence axiom also documented by [Allais \(1953\)](#). Again an individual chooses between two lotteries:


- Choice A: 80% chance of \$4000; 20% chance of nothing.
- Choice B: Certainty of receiving \$3000.

Here again, most choose lottery B. The same individual is then given the choice between two different lotteries:

- Choice A': 8% chance of \$4000; 92% chance of nothing.

- Choice B': 10% chance of \$3000; 90% chance of nothing.

Many choose lottery A' in this second choice set that had chosen B in the first choice set, which is again a violation of the independence axiom. Consider the following representation of both choice sets:



When $P = .10$ most prefer A (from the second choice set) and when $P = 1$ most prefer B (from the first choice set). Again, it is easy to see that this behavior violates independence, as someone that prefers B from the first choice set should prefer B' from the second.

Many other violations of the independence axiom exist, however we chose to highlight these two as they are immediately applicable to an advisory referendum. Changing the probability P in the example of the common ratio effect is identical to changing the probability μ_i that a vote is consequential in an advisory referendum. Similarly, changing the value of Z in the example that produced the common consequence effect is equivalent to changing the non-consequential outcome $\tilde{s}_i$.

These violations of the independence axiom have led to the development of utility representations that do not rely on the independence axiom. The direct connection between the common consequence and ratio effects and an advisory referendum, allow the same utility representations developed to explain behavior observed in common consequence/common ratio experiments to be applied to the theory of the advisory referendum.

1.4 Expected Utility Analysis Without the Independence Axiom

The common consequence and common ratio effects are widely viewed as the primary departures from expected utility. As such, they have led to the development of many different preference representations that seek to accommodate the behavior that expected utility would otherwise deem “paradoxical”.

In expected utility, the preference over lotteries each of the form $P = (x_1, p_1; \dots; x_n, p_n)$ that yield an outcome of x_i with probability p_i where $\sum p_i = 1$ can be represented with a *preference function* of the form

$$V_{EU}(P) = u(x_1)p_1 + u(x_2)p_2 + \dots + u(x_n)p_n \quad (1.2)$$

In this framework, an individual prefers some lottery $P^* = (x_1^*, p_1^*; \dots; x_n^*, p_n^*)$ over the lottery $P = (x_1, p_1; \dots; x_n, p_n)$ if and only if $V_{EU}(P^*) > V_{EU}(P)$ and is indifferent between them if and only if $V_{EU}(P^*) = V_{EU}(P)$, where $u(\cdot)$ is the individual’s von Neumann-Morgenstern utility function.

Expected utility has a great deal of flexibility in representing many aspects of attitudes towards risk and also has testable implications that hold regardless of the shape of the utility function $u(\cdot)$. These implications follow from $V_{EU}(\cdot)$, which takes a linear form $V_{EU}(P) = \sum u(x_i)p_i$ for some set of coefficients $\{u(x_i)\}$. Therefore, expected utility preferences can be described as being *linear in the probabilities*, a product of the independence axiom.

A graphical depiction helps to highlight the behavioral restrictions implied by the linearity in the probabilities hypothesis implied by expected utility. If attention is restricted to the set of all distributions (p_1, p_2, p_3) over the three outcomes $\{x_1, x_2, x_3\}$ where $x_1 < x_2 < x_3$, the set of distributions can be represented by points in the unit triangle, since $p_2 = 1 - p_1 - p_3$ (see Figure 1.1). Because x_3 is the most preferred




Figure 1.1: Expected utility indifference curves in the unit triangle

outcome, all movements to the northwest (which increase p_3 at the expense of p_1 and p_2) result in stochastically dominating lotteries and therefore are preferred.

The individual's indifference curves in Figure 1.2 are depicted as parallel straight lines and are given by the solutions to

$$p_1 u(x_1) + (1 - p_1 - p_3) u(x_2) + p_3 u(x_3) = \text{constant} \quad (1.3)$$

These indifference curves have a common slope of

$$[u(x_2) - u(x_1)] / [u(x_3) - u(x_1)] \quad (1.4)$$

which depends on the concavity of the individual's utility function. The more risk averse is the individual, the steeper are the individual's indifference curves, which is illustrated in Figure 1.2.

The dashed lines in Figure 1.2 are the iso-expected value lines which are the solutions to

$$x_1 p_1 + x_2 (1 - p_1 - p_3) + x_3 p_3 = \text{constant}. \quad (1.5)$$

Since northeast movements along these lines do not change the expected value of the distribution but do increase the probabilities of the best (x_3) and worst (x_1) outcomes at the expense of the middle outcome (x_2), they represent simple increases in risk. When $u(\cdot)$ is concave (risk averse), its indifference curves will have a steeper slope




Figure 1.2: Risk preferences in the unit triangle




Figure 1.3: Expected utility indifference curves and the Common Consequence effect

than the iso-expected value lines and such increases in risk move an individual from more preferred to less preferred indifference curves.

Figures 1.3 and 1.4 illustrate how the choices B and A' violate expected utility in reference to the common consequence and common ratio effects. In both cases, if B is preferred in the first choice set, it must be on a higher indifference curve. If




Figure 1.4: Expected utility indifference curves and the Common Ratio effect

the indifference curves are drawn parallel to one that satisfies this condition, it is not possible to place A' on a higher indifference curve than B in the second choice set.

In order to accommodate such behavior, an alternative representation to expected utility replaces independence with *betweenness*. Betweenness says that if an individual is indifferent between two lotteries, he should also be indifferent between those two lotteries and any probability mixture of those two lotteries. In other words, it should not matter if the choice between two indifferent lotteries is determined by the decision maker or by some random device. In terms of the triangle diagram, if an individual is indifferent between two points in the triangle, he should also be indifferent between all points lying on the straight line connecting the two points. The assumption of betweenness means that indifference curves in the triangle are straight, but not necessarily parallel, lines (see Figure 1.5).




Figure 1.5: Indifference curves exhibiting betweenness

The key to explaining the behavior seen in examples of the Allais paradox is the flexibility that results from the ability of risk preferences to change when different gambles are considered. The choice of B and A' from our example would then mean that those individuals are less risk averse over lotteries whose expected value is (relatively) smaller. Therefore, indifference curves “fan-out” across the unit triangle. The theory of fanning indifference curves was developed by [Machina \(1982\)](#) and is illustrated in figure 1.5.

1.5 Single binary choice questions without the independence axiom

Consider the decision of a voter i whose maximum willingness-to-pay (WTP) for $\tilde{x}_{ij}$ is denoted c_{ij}^* , where c_{ij}^* is the cost of $\tilde{x}_{ij}$ such that $[\tilde{x}_{ij} - c_{ij}^*] \sim \delta_0$. Suppose the voter decides whether to vote yes or vote no. If individual i 's vote is the one that determines the final project implemented by the town, she has the choice between two options she is indifferent to: getting δ_0 for certain by voting no, or getting $\tilde{x}_{ij}$ by voting yes and incurring a cost of c_{ij}^* .

The lottery q is one which has the outcome $(\tilde{x}_{ij} - c_{ij}^*)$ with probability p and δ_0 with probability $(1 - p)$. If the voter is indifferent between $(\tilde{x}_{ij} - c_{ij}^*)$ and δ_0 , the lottery q is located somewhere on the straight line between δ_0 and $(\tilde{x}_{ij} - c_{ij}^*)$. Assuming betweenness, the voter is therefore indifferent between the three lotteries: δ_0 , $(\tilde{x}_{ij} - c_{ij}^*)$ and q . Figure 1.6 shows this graphically using the probability triangle diagram.




Figure 1.6: An indifferent voter's choice problem in the unit triangle.

The straight line connecting $(\tilde{x}_{ij} - c_{ij}^*)$ and δ_0 divides the unit triangle into two regions: a better-than region consisting of all lotteries lying north-west of the q lottery indifference curve, and a worse-than region consisting of all lotteries lying south-east of the q lottery indifference curve.

Now consider the setting where each voters response is an input into the decision making process for some authority, whose informed decision will ultimately determine

the final outcome. Let the threshold acceptable cost (the voter's maximum WTP for proposal $\tilde{x}_{ij}$ in this scenario) be denoted b_i^* , where b_i^* is the cost that solves

$$\mu_i(\tilde{x}_{ij} - b_i^*) + (1 - \mu_i)\tilde{s}_i \sim \mu_i\delta_0 + (1 - \mu_i)\tilde{s}_i \quad (1.6)$$

Consider again the indifference condition $[\tilde{x}_{ij} - c_{ij}^*] \sim \delta_0$ where c_{ij}^* is i 's *true* maximum WTP. Truthful preference revelation therefore requires that $b_i^* = c_{ij}^*$ for all $\tilde{s}_i$ and $0 < \mu_i < 1$. In other words, truthful preference revelation requires that the maximum WTP for $\tilde{x}_{ij}$ not change if $\tilde{x}_{ij}$ is mixed with any $\tilde{s}_i$. This requirement is equivalent to the restrictions imposed by the independence axiom.

Now consider the following exercise: Suppose that preferences satisfy betweenness and are governed by the fanning-out hypothesis. Assuming that $\tilde{s}_i$ is an element of the better-than set of the unit triangle, the indifference curve associated with $\mu_i\delta_0 + (1 - \mu_i)\tilde{s}_i$ is steeper than the indifference curve between δ_0 and $[\tilde{x}_{ij} - c_{ij}^*]$. This implies that $\mu_i\delta_0 + (1 - \mu_i)\tilde{s}_i$ is strictly preferred to $\mu_i[\tilde{x}_{ij} - c_{ij}^*] + (1 - \mu_i)\tilde{s}_i$, which yields, according to the indifference condition (1.6), $b_i^* < c_{ij}^*$. Graphically, figure 1.7 illustrates the indifference curve associated with the *true* status-quo outcome δ_0 and the steeper indifference curve associated with the “advisory” status-quo outcome $\mu_i\delta_0 + (1 - \mu_i)\tilde{s}_i$, which is attributed to fanning-out. The dashed line in figure 1.7




Figure 1.7: Strict preference for the status-quo option over the proposal

is the assumed indifference curve under EU, and therefore the independence axiom. According to the fanning-out hypothesis, individuals become more risk averse as they become better-off, which translates to steeper indifference curves in the better-than

portion of the unit triangle. This increased aversion to risk translates into a lower certainty equivalent (i.e. the threshold acceptable cost b_i^*) associated with the risky prospect $\tilde{x}_{ij}$.

The results from the above exercise can be interpreted as follows: when the proposal considered by voter i is worse than the proposal voter i perceives will be implemented if her vote is inconsequential ($\tilde{s}_i$), then voter i will have a lower maximum WTP for $\tilde{x}_{ij}$ than in the scenario where her vote alone determines whether $\tilde{x}_{ij}$ is implemented or if the proposal is rejected in favor of the status-quo.

The three possible preference ordering scenarios for the perceived inconsequential outcome are: i) The voter is indifferent between ($\tilde{s}_i$) and δ_0 , ii) The voter prefers ($\tilde{s}_i$) to δ_0 , and iii) The voter prefers δ_0 to ($\tilde{s}_i$). The implications for each of these scenarios are formally expressed in Proposition 1.

Proposition 1.1. *The following are true if preferences are consistent with betweenness and satisfy the fanning-out hypothesis:*

- (i) *If voter i is indifferent between the perceived inconsequential outcome ($\tilde{s}_i$) and the status-quo (δ_0), then the threshold acceptable cost used to determine her response to an advisory referendum is equivalent to her true maximum WTP ($b_i^* = c_{ij}^*$).*
- (ii) *If voter i prefers the inconsequential outcome ($\tilde{s}_i$) to the status-quo (δ_0), then the threshold acceptable cost used to determine her response to an advisory referendum is lower than her true maximum WTP ($b_i^* < c_{ij}^*$), which reduces the likelihood that she votes in favor of proposal $\tilde{x}_{ij}$ at cost c_{ij} .*
- (iii) *If voter i prefers the status-quo (δ_0) to the inconsequential outcome ($\tilde{s}_i$), then the threshold acceptable cost used to determine her response to an advisory referendum is higher than her true maximum WTP ($b_i^* > c_{ij}^*$), which increases the likelihood that she votes against the proposal $\tilde{x}_{ij}$ at cost c_{ij} .*

for all $0 > \mu_i > 1$.

Proof. See Appendix A. □

Proposition 1 establishes instances where there is a possibility that an individual will vote according to a threshold acceptable cost which is different from their true maximum WTP. The relevance of this result is dependent on the individual's perception of the inconsequential outcome ($\tilde{s}_i$) being *sufficiently different*^{*} from the status-quo. Proposition 1 therefore highlights the potential for individual's to base their response to advisory referenda on values which are different from their true maximum WTP.

Proposition 2 outlines the changes which occur to the threshold acceptable cost b_i^* , given a change in the perception of consequentiality μ_i , for each state of the individual's perception of the inconsequential outcome ($\tilde{s}_i$).

Proposition 1.2. *If individual i 's perceived consequentiality is such that $\mu_i = 1$ (i.e. she believes that her vote alone will determine whether proposal $\tilde{x}_{ij}$ is implemented in favor of the status-quo, δ_0), then her threshold acceptable cost (b_i^*) for proposal $\tilde{x}_{ij}$ is equivalent to her true maximum WTP (c_{ij}^*), and if preferences are consistent with betweenness and satisfy the fanning-out hypothesis, the following are true for any decrease in perceived consequentiality:*

- (i) *If voter i is indifferent between the inconsequential outcome ($\tilde{s}_i$) and the status-quo (δ_0), her threshold acceptable cost (b_i^*) for proposal $\tilde{x}_{ij}$ remains equivalent to her true maximum WTP (c_{ij}^*) as her perceived consequentiality (μ_i) decreases.*
- (ii) *If voter i prefers the inconsequential outcome ($\tilde{s}_i$) to the status-quo (δ_0), her threshold acceptable cost (b_i^*) for proposal $\tilde{x}_{ij}$ increase away from her true maximum WTP (c_{ij}^*) as her perceived consequentiality (μ_i) decreases.*

^{*}The definition of “sufficiently different” in this context has yet to be defined and is the topic of further investigation.

(iii) If voter i prefers the status-quo (δ_0) to inconsequential outcome ($\tilde{s}_i$), her threshold acceptable cost (b_i^*) for proposal $\tilde{x}_{ij}$ decrease away from her true maximum WTP (c_{ij}^*) as her perceived consequentiality (μ_i) decreases.

Proof. See Appendix A. □




Figure 1.8: Illustration of Proposition 4(ii), where $A \sim B$ when $\mu_i = 1$, and $C \succ D$ represent a decrease in μ_i , and therefore in order for $C \sim D$ it must be the case that $b_i^* > c_{ij}^*$.

Proposition 2 bears a lot of similarity to Proposition 1, but rather than dealing with changes occurring from a shift in the value of the inconsequential outcome given a fixed perception of consequentiality, Proposition 2 highlights the changes occurring from a shift in perceptions of consequentiality given a fixed value of the inconsequential outcome.

Disentangling these two effects is important because it offers an explanation for all types of observed behavior. For example, imagine an instance where individuals who felt their responses were inconsequential but had estimated WTP values identical to those who felt their responses were consequential. It may seem as though consequentiality does not play a roll in value formation, however it may be the case that the value of the perceived inconsequential outcome ($\tilde{s}_i$) is not *sufficiently different* from status-quo.

1.6 Discussion

This paper is the first to extend the theory of advisory referenda to accommodate the possibility of systematic violations of the independence axiom. This framework

offers a great deal of flexibility and accommodates all types of observed behavior, while preserving many of the desirable properties of expected utility.

Two different sources of behavioral biases were investigated. The results of the analysis offer two methods to control for these biases. Past research efforts (Vossler et al., 2010; Herriges et al., 2010) have tried to control for survey respondents' perceptions of consequentiality through (ironically enough) inconsequential follow-up questions. This research suggests that, in conjunction with estimating respondents' perceptions of consequentiality, there may also be returns to estimating what respondents *actually* think is going to happen if the results of the survey are ignored.

The theoretical framework built in this paper not only explains potential sources of the deviations of behavior from EU theory, it offers new and potentially very useful ways to measure non-market value. Further investigation into the strength and consistency of the biases could potentially lead to the development of new elicitation techniques which utilize behavior as a source of insight rather than unwanted noise.

Future extensions include using the insights from this model to conduct further studies in the field to explicitly test the theory. These studies will prove extremely helpful in understanding incentive properties surrounding binary choice preference questions.

Chapter 2

Behavioral dynamics of tax compliance under an information services initiative

2.1 Introduction

To “encourage” correct tax reporting it is likely that enforcement effort, audits and penalties, will continue to be primary tools in the tax authority’s arsenal. This approach is based on the basic model of tax evasion which views the taxpayer as engaging in an evasion “gamble” in which the bad state of nature involves the taxpayer being audited and paying a penalty on evaded taxes.* However, many tax agencies are exploring complementary instruments including the provision of information and assistance to taxpayers. This revised paradigm recognizes that tax administrators have a role as facilitators and providers of services to taxpayer-citizens. And, it opens up the possibility that the audit and service approaches to enhancing tax reporting may be synergistic.† Further, the “service” paradigm for tax administration

*This approach derives from the classic “economics of crime” pioneered by [Becker \(1968\)](#) and applied to tax evasion by [Allingham and Sandmo \(1972\)](#).

†The value of the taxpayer service derives from the costs imposed on the taxpayer for noncompliance. For the payoff maximizing individual, absent enforcement effort, service that resolves

fits squarely with the perspective that emphasizes the role social norms play in tax compliance (Feld and Frey, 2002), and these link directly to the behavioral issues that arise in understanding the dynamic interaction between taxpayers and the tax authority.

Some basic effects of an information service program on tax reporting have been recently examined by Alm et al. (2010) and by Vossler and McKee (2012). Using an experimental design that shares some common features with experiments reported here, Alm et al. (2010) find that taxpayers respond positively to service programs. However, Alm et al. (2010) do not report on the dynamic effects of prior audits. In another similar experimental setting, Vossler and McKee (2012) find that subjects are less likely to file a tax return when their tax liability is uncertain. However, Vossler and McKee (2012) find that the provision of information services offsets the effect of not filing and that simply providing the service, even an imperfect service, increases the propensity to file, and increases the accuracy of the filing. We continue in this research direction by implementing a richer design that allows us to investigate dynamic behavioral effects of tax audits, as well as the effects of varying the quality and cost of the service. Experimental data are especially useful here since the experiment allows for control of institutional features and addresses the problem of being unable to observe the actions of each individual taxpayer. An audit may not correctly reveal the true tax liability in field data, however this is explicitly induced in the lab setting, and therefore the exact amount of evasion is known. Finally, even though service programs have undergone some changes in the past, there is not a full spectrum of such programs in existence and so such field data as may exist are incomplete.

Our research utilizes controlled laboratory experiments with human decision makers and salient financial incentives in order to test the effects of audits on taxpayer reporting and information acquisition. Within the laboratory, we induce the true tax liability uncertainty would have no value to the taxpayer. However, a taxpayer wishing to honestly report would value the information since it would enable such honesty.

tax liability (which is not known with certainty to participants), and then identify the effects of information services (to resolve all or some of the uncertainty) by systematically varying the setting across groups of players. Since audits are random in our design, we are also able to investigate the effects of prior audits and information acquisition on tax reporting over time. This design then permits investigation of behavioral dynamics in two dimensions: tax reporting and information acquisition, in addition to the interaction between the service program and the audit program.^{*}

Because the tax reporting and information acquisition decisions are observed over several decision periods, our design allows us to examine the dynamic effects of prior audits on both the taxpayers' reporting behavior, and the taxpayers' subsequent utilization of the information services. Because a taxpayer's true liability is uncertain, and taxpayers are penalized for noncompliance,[†] information which resolves reporting uncertainty is valuable to the taxpayer. A taxpayer who has made the decision to "cheat" on their taxes, however, would not be willing to pay for information which they plan to ignore. Including information services in our design therefore allow us to tease out the distinction between "loss repair" (Andreoni et al. (1998); Maciejovsky et al. (2007)) and the "bomb crater effect" (Mittone, 2006) as responses to prior audits.

The bomb crater effect is a form of the gambler's fallacy. It is used to describe the behavior of a taxpayer who under-reports her true tax liability immediately following an audit due to her inaccurate belief that the probability of a consecutive audit is lower than the true random audit probability. Loss repair is when, after being penalized for under-reporting, an individual tries to recover the "losses" she occurred from the audit process.

Existing literature (e.g., Kastlunger et al. (2009); Erard (1990); Alm et al. (1992a); Alm et al. (1992b)) focuses only on the effects of past audits on the tax reporting decision. Typical findings are that individuals report a lower tax liability following

^{*}Endogenous or systematic audit rules would make it difficult to undertake this investigation as the behavioral impacts of an audit outcome would be clouded by the institutional change.

[†]Both in under-reporting and over-reporting their true taxes. Under-reporting may result in a penalizing audit. Over-reporting is "penalized" in the sense of forgone income.

an audit and various motives have been suggested to explain this dynamic response. However, to our knowledge our experiments are the first to examining dynamic behavior when a taxpayer service program is operating in conjunction with an audit program.

Our results suggest that, in the presence of uncertain tax liability, the audit process effects the tax reporting decision and information acquisition. We find greater support for loss repair behavior than the bomb crater effect. Immediately following an audit, information acquisition rates decrease and under-reporting increases only for non-compliant taxpayers. Compliant taxpayers who were audited actually report higher tax liabilities on average in subsequent rounds, while their propensity to acquire information was unaffected. Without the availability of information services to resolve uncertainty, it would not be able to distinguish between those who wish to comply but under-reported due to the uncertainty of their true liability, and those who purposely “evade” their taxes. The addition of information services is therefore key to distinguishing between the bomb crater effect and loss repair.

2.2 Conceptual Framework

In order to cleanly identify important effects related to information services and associated behavioral dynamics, we consider a stylized setting that captures some fundamental features of the personal income tax system while abstracting away from much of its complexities. The setting we consider is one where the risk-neutral taxpayer makes a tax reporting decision - in particular chooses a tax credit to report - and then files a return to the tax authority. The true tax liability is uncertain, which makes an information service potentially valuable. To motivate compliance, the tax authority undertakes audits with probability p . Audits are completely random and independent of whether other persons are audited or the reported tax liability. If an audit occurs, it perfectly reveals any unpaid taxes. In addition to being liable for unpaid taxes, there is a constant per-dollar penalty $\beta > 0$ assessed on unpaid taxes.

No refund is given if taxes are over-paid, and in this sense an audit is never beneficial. The audit process is static in that only the current period tax return is scrutinized and there is no possibility of penalties for (yet undiscovered) past non-compliance nor does a violation lead to a higher future audit probability.

2.2.1 Basic economic theory

A risk-neutral expected-utility maximizer simply weighs the expected marginal benefits and marginal costs of tax under-reporting. In the special case where tax liability is certain, given the above audit process, the marginal expected costs associated with every dollar of tax under-reporting is constant and equals $p(\beta+1)$. As such, a corner solution of full compliance (i.e. truthful reporting) arises if $p(\beta+1) \geq 1$, and otherwise the taxpayer fully evades (i.e. reports the lowest tax liability possible). With uncertain tax liability, depending on the distribution of tax liability values the taxpayer, interior or corner solutions are possible.

A “helpful” information service is one which: (1) leads the taxpayer to an optimal choice that is more likely to be truthful, and (2) has value to the taxpayer. These conditions implicitly assume an interior solution in the absence of information. This is consistent with our experimental design and rules out situations where the information service has a null effect on tax reporting and zero value to the taxpayer. For example, consider the simple case where the taxpayer believes her liability is either \$1000 or \$2000 with equal probability. With expected costs of under-reporting sufficiently high, it would be optimal for her to report \$2000 as her liability. However, if an information service reveals that her *true* liability is \$1000, she is able to avoid paying too much in taxes. Not surprisingly, the value of the service increases with the “helpfulness” of the service. For example, an information service that reduces more uncertainty has more value. An ancillary implication relevant for our experiment is that taxpayers should be willing to pay more to acquire a more helpful information service.

2.2.2 Insights from behavioral economics

There are several behavioral responses to the audit process that have been documented in past experiments involving tax reporting decisions and simple random audit enforcement mechanisms. [Mittone \(2006\)](#) finds that, on average, tax compliance drops in the period immediately after an audit. Mittone labels this behavior as the “bomb crater” effect (BCE). Subjects behave as if the probability of being audited immediately following a period in which they were audited is significantly lower and therefore perceive the cost of evasion to be low. Mittone also finds that after several filing periods, compliance increases, which he argues is likely due to an increase in the perceived probability of an audit. Another behavioral response to the audit process is known as “loss repair” ([Andreoni et al. \(1998\)](#); [Maciejovsky et al. \(2007\)](#)). Loss repair is the notion that the penalties that are incurred during the audit process might induce subjects to “want to evade more in the future in an attempt to ‘get back’ at the tax agency” ([Andreoni et al. \(1998\)](#) pp. 844). Therefore, subjects experiencing audits and penalties may try to recover their ‘losses’ by engaging in tax evasion in future filings.

Information acquisition is unique to our experimental design, but to the extent the above behavioral motivations exist, one would expect related effects. In particular, if a taxpayer is motivated to under-report taxes in the period immediately following an audit, the value of the information service (and associated willingness to pay for it) should be significantly lower. Therefore, a result consistent with the BCE would find that information acquisition is lower in the period immediately following an audit. A result consistent with loss repair would find similar effects as the BCE for those who were penalized, but would find little effect on those that were audited and found compliant.

2.3 Experimental Design

2.3.1 Decision setting

Our experimental setting implements the fundamental elements of a voluntary reporting system. Participants earn income by performing a task and self-report their tax liability to a tax authority. Final tax liability is a function of earned income, the tax rate, and tax credits claimed. If an audit occurs unreported taxes are discovered. The audit process performs without error; if the individual has evaded taxes both the unpaid taxes and a penalty are collected.

A participant’s earnings for a decision period are her income, minus the taxes she reported, minus penalties, if applicable. Income is denominated in “lab dollars”. The overall earnings for the experiment are the sum of the lab dollars earned over all decision periods multiplied by a common (and known) lab to US dollar exchange rate. In each period of the experiment, participants earn income based upon their performance in a simple computerized task, in which they are required to sort numbers into the correct order. Those who finish the task the fastest earn the highest income of 1500 lab dollars for the period, those who finish in the middle of pack earn 1250 lab dollars, and the slowest earn 1000 lab dollars. Participants are presented information about the distribution of group earnings to ensure that they believe the relative nature of the earnings. The earnings task is the only source of interaction and payoff interdependence; this design implements a blind setting among the participants.

After earning income, participants are presented with a screen that informs them of their earnings and the tax policy parameters (tax rate, audit probability, and penalty rate).^{*} In each period, the participants decide whether to request

^{*}These are fixed throughout the experiment. Our experimental setting is very contextual and the presence of the income earning task provides, we argue, for the necessary degree of “parallelism” to the naturally occurring world that is crucial to the applicability of experimental results (Smith (1982); Plott (1987)). The experimental setting need not - and should not - attempt to capture all of the variation in the naturally occurring environment, but it should include the fundamental elements of the naturally occurring world for the results to be relevant in policy debates. In this regard, our experimental design uses tax language (which is presented via the subject interface), requires that the participants earn income in each period, and also requires that the participants

an information service (if one is available) and how much to claim in tax credits. Although other institutional details are embedded in the design (e.g. tax rate, taxable income, etc.), and in particular the tax form, the participant can only manipulate her tax liability through her credit reporting choice. As there are penalties for tax underreporting if audited, and foregone earnings associated with over-paying taxes, there is value to resolving any uncertainty regarding the tax credit. The expected tax credit is calculated according to the formula $1000 - 0.5 * (\text{earned income})$, such that the expected credits equal 500, 375 and 250 for the three income categories (1000, 1250, 1500). The amount of the credit is high relative to the initial tax liability so that the credit decision is financially salient. One important feature of our design is that, given the tax rate and credit formula, the expected after tax income is 1,000 lab dollars across all initial income levels.

The “true” credit amount is uncertain and is a random draw from a uniform distribution, defined as plus or minus 100% of the expected credit amount. The true credit amount is independent across decision periods and individuals. A participant’s true credit remains unknown unless she acquires information or is audited. Given this design, uncertainty - and, hence, the value of resolving it - increases with the expected credit (or, analogously, decreases in income). With uncertainty, prior to making a credit choice or acquiring information (if possible), each participant sees the supports of the uniform distribution that coincides with her income. If an information service is available, participants can acquire the information with the click of a button.*

disclose tax liabilities in the same manner as in the typical tax form. As in the naturally occurring setting, there is a time limit on the filing of income. A clock at the bottom of the screen reminds the participants of the time remaining, and there is a penalty for failing to file on time set equal in all sessions to 10 percent of taxes owed; also, the individual is automatically audited if he or she fails to file on time, so that the participant pays the non-compliance fine as well.

*Such information reduces the cognitive burden of computing tax liabilities. The issue of tax liability uncertainty differs from enforcement uncertainty. As [Alm et al. \(1992b\)](#) demonstrate, given a setting where taxes are not used to fund a public good, the tax authority may use enforcement uncertainty to increase compliance. Theory predicts that uncertain penalties increase compliance by risk-averse agents and this is borne out in the data from a set of experiments. [Alm and McKee \(2006\)](#) extend this and report on the compliance effects of informing the taxpayer their return will be audited with certainty.

The participants are informed of the audit probability and the penalty rate, and know these values with certainty. In all sessions, the tax rate is fixed at 50% of earned income, the audit probability is fixed at 30%, and the penalty rate is fixed at 300% of unpaid taxes. Our audit rate is much higher than actual full audit rates in the United States. However, the IRS conducts a range of audits, and for many types of audits the actual rates are quite high.* The penalty rate is consistent with penalties imposed by the IRS for tax underreporting. Enforcement effort is held constant since the effects of enforcement efforts have been widely investigated and we only need this effort to be salient in the current setting to give value to the information that resolves tax liability uncertainty.† Table 2.1 summarizes the key parameters of the experiment.

Table 2.1: Experiment Parameters

Parameter / variable	Value(s)
Earned Income	1000, 1250 or 1500 lab dollars
Audit Probability (p)	30%
Penalty Rate (β)	300% on unpaid taxes
Tax Rate	50% on Earned Income
Tax Credit	Expected value: 1000 - (0.5 x Earned Income) Range: +/- 100% of expected value

Participants are able to revise their credit decision prior to filing their return, and the tax form updates their tax liability as the claimed credit is revised. Thus, they can observe the potential changes in their reported tax liability for each potential reporting strategy they investigate. A timer at the bottom of the tax form counts down the remaining time. The participants are allowed 90 seconds to file and the

*While overall audit rates are quite low, among certain income and occupation classes they are more frequent. The oft-reported IRS audit rate (currently less than one percent) is somewhat of an understatement. This reported rate usually refers to full audits. In fact, the IRS conducts a wide range of audit-type activities, including line matching and requests for information, and these activities are much more frequent. For example, in 2005 only 1.2 million individual returns (or less than one percent of the 131 million individual returns filed) were actually audited. However, in that year the IRS sent 3.1 million “math error notices” and received from third parties nearly 1.5 billion “information returns”, which are used to verify items reported on individual income tax returns.

†See Alm et al. (1992a)

counter begins to flash when there are fifteen seconds remaining. Thus, the process in the lab mimics that by which a taxpayer may well conduct different calculations in the time prior to actually filing her taxes (whether he or she uses one of the available tax software programs or simply does the tax return by hand). If an information service is available, this can be requested at any time and does not change the total amount of time for a period.

The audit selection process is completely random and the participants face the same probability in each period independent of current and past reporting behavior and past audit outcomes. The random audit selection process is illustrated by the use of a “virtual” bingo cage that appears on the computer screen. A box with blue and white bingo balls appears on the screen following the tax filing. The ratio of blue to white balls reflects the audit probability. The balls begin to bounce around in the box, and after a brief interval a door opens at the top of the box. If a blue ball exits, the participant is audited; a white ball signifies no audit.

When an audit occurs, the true value of the credit is used to determine taxes owed. The individual’s declarations are examined. If the individual has under-reported her tax liability, she must make up for the difference as well as pay a penalty. If an individual has over-reported her tax liability no over payments are returned to the individual.* Tax revenues and any penalties paid are not redistributed to the participants in order to ensure that the participants focus on the individual income disclosure decision and not on any public good provision decision. After the tax return is filed and an audit (if any) is determined, the participant is shown one final screen that summarizes everything that happened during the period. After two practice periods to allow subjects to gain familiarity with the interface, the process just described is repeated for a total of 20 paid periods. To minimize potential end-of-game effects the number of periods is not disclosed prior to its realization.

*Certain errors on the part of the taxpayer may not be easily verified in the event of an audit. For example, failure to claim a deduction for a charitable contribution because the taxpayer was uncertain of the status (e.g., 501c(3) status) of the organization may not be observed by the tax agency even in the event of an audit.

2.3.2 Treatments

With the exception of the variation in earned income, which is again varies across subjects in a session as determined by a simple task at the beginning of each decision period, we employ a between-subjects design. The main treatment variables (varied across sessions only) are the presence/absence of an information service, the quality of the service if provided, and the cost of obtaining the information. These factors are held constant throughout a session. There are five basic treatments (see Table 2.2). The first (T1) is a treatment with certain tax liability, which we use as a baseline for comparison against uncertain information treatment. In this treatment, participants are automatically given information on their true credit and there is no notion of an information service. In the second treatment (T2), the individual’s tax credit is uncertain and there is no information service available. This establishes a second baseline for comparison. In the remaining three treatments, there is an information service available. The status quo in the information service treatments, i.e. if the information service is not utilized, is identical to the uncertainty baseline.

Table 2.2: Experiment Treatments

Tax Liability Uncertain	Service Provided?			
	No	One Source (Complete and Correct)	Two Simultaneous Sources (One Correct)	Two Sequential Sources (One Correct)
No	T1	N/A	N/A	N/A
Yes	T2	T3	T4	T5
		Price of Information: \$0, \$50, \$100	Price of Information: \$0, \$50, \$100	Price of Information: \$0, \$50, \$100

The “perfect” information service reveals the true credit with certainty (T3). Under the other two information service types, the service is imperfect in the sense that up to two possible credit amounts can be provided and each amount has a 50% chance of being correct. Specifically, under the “simultaneous” information service treatment (T4) the authority simultaneously provides two credit amounts, one of

which is the truth while the other is a decoy. With the “sequential” information service (T5), the participant can make up to two information requests and with each request is delivered one possible credit amount. If two requests are made, then the simultaneous and sequential services reveal the same information. However, the sequential information treatment leaves the possibility that only one credit amount is delivered, in which case it still has the same 50% chance of being the truth.

To assess the value of information services, we vary the cost to acquire information in the information service treatments (see Table 2.2). The three cost levels are 0, 50 and 100 lab dollars for the perfect and the simultaneous information settings. For the sequential setting, these costs are halved and assessed separately for the two sources.

2.3.3 Participants and procedures

The experiments were conducted at dedicated experimental laboratories at the University of Tennessee and Appalachian State University, which both utilized the same software and experimental protocol, and have similar computer networks. The participant pools included students and non-students (university staff, mostly).^{*} Students and non-students participated at separate times, and the lone difference in student versus non-student sessions is that the latter utilized a lower lab dollar to US dollar exchange rate (375 to 1 versus 750 to 1) in order to reflect the higher opportunity cost of participation. Recruiting was conducted using the Online Recruiting System for Experimental Economics (ORSEE) developed by Greiner (2004). Databases of potential participants were built using announcements sent via email to university students and staff. Registered individuals were contacted, via email, and were permitted to participate in only one tax experiment.[†] Only participants recruited specifically for a session were allowed to participate, and no participant had prior experience in this experimental setting. Methods adhere to all

^{*}An individual session included only students or non-student participants; they were not mixed in a session.

[†]Other experimental projects were ongoing at the time and participants may have participated in other types of experiments.

guidelines concerning the ethical treatment of human participants. Earnings averaged \$25 for student participants and \$45 for non-students. Sessions lasted between 60 and 90 minutes. A total of 730 participants took part in these sessions.

The experiment session proceeded in the following fashion. Each participant sits at a computer located in a cubicle, and is not allowed to communicate with other participants. The instructions are conveyed by a series of computer screens that the participants read at their own pace, with a printed summary sheet provided and read aloud by the experimenter. (Appendix C provides representative screen shots from the experiment and Appendix D provides instructions from one of the treatments.) Clarification questions are addressed after the participants have completed the instructions and two practice periods. The participants are informed that all decisions will be private; the experimenter is unable to observe the decisions, and the experimenter does not move about the room once the session starts to emphasize the fact that the experimenter is not observing the participants' compliance decisions. This reduces, to the extent possible, peer and experimenter effects that could affect the decisions of the participants. All actions that participants take are made on their computer. After the 20 paid decision periods, participants are asked to fill out a brief questionnaire, which collects basic demographics including information on tax reporting experience. Payments are made privately at the end of the session.

2.4 Testable hypotheses

2.4.1 Basic economic hypotheses

With our chosen audit probability and penalty rate, when the true credit is known with certainty it follows that the expected cost of under-reporting by one lab dollar equals $0.3(3+1)=1.2$ lab dollars such that it is optimal to report truthfully.* When the credit is uncertain, based on the assumed uniform distributions, the taxpayer

*We note that in [Alm et al. \(2010\)](#) the expected cost is much less than 1, and the optimal strategy in that experiment (all treatments) is a corner solution of maximal evasion.

will optimally evade through over-claiming the credit. The extent of the deviation from truthful reporting increases with the level of uncertainty. In expectation, all income levels receive the same after tax income, however uncertainty decreases with income. Because those with the lowest income have the widest uncertainty range, theory suggests* the highest relative amount of tax evasion for these individuals. Point predictions from the basic theory have that it is optimal to evade by 333 lab dollars for those with earned income of 1000, by 250 for those with an income of 1250, and by 167 for those with an income of 1500.

The decision of whether to request the information service(s) to resolve uncertainty (at least partially) is driven by the value of information. Theoretically, and quite intuitively, the taxpayer's willingness-to-pay (WTP) is increasing in the initial level of uncertainty as well as the accuracy of the information. In the context of the experimental design, those with lower incomes face a larger range of uncertainty and, *ceteris paribus*, have a higher WTP for information. Further, knowing the true credit is more valuable than receiving two possible amounts only one of which is correct. In terms of point predictions, since information has value, in all situations information should be requested when it is free. At the other extreme, in all situations no information should be requested at our highest cost amount of \$100 (or \$50 for one imperfect information source). At the middle cost amount, those at the lowest income level should request the information (imperfect or perfect), at the middle income level it is beneficial to request perfect information, and it is not beneficial for those with high income to request information.[†]

The basic economic hypotheses we evaluate with our experiment are:

- Hypothesis 1. The cost of the information service decrease the propensity to acquire information.
- Hypothesis 2. Greater uncertainty increases the propensity to acquire information (i.e. lower income levels have a higher propensity to acquire information).

*see Appendix B

†see Appendix B

- Hypothesis 3. Information services decrease the under-reporting of tax liabilities.

2.4.2 Testable behavioral hypotheses

The instructions and information provided to the experiment subjects is explicit about the fixed audit probability, the purely random selection process, and independence over periods. Therefore, economic theory would predict that the amount of tax credit claimed by subjects will be independent of their audit history. Given our experimental design, we can test for the BCE and loss repair effects in the absence of other confounds that may exist in naturally occurring settings (e.g. increased future audit probability or auditing past returns). These two effects can in particular be identified by comparing pre- and post-audit credit reporting decisions. A basis from which to distinguish between the two competing theories arises as there is predicted to be a difference between those who were audited and found to be compliant, and those who were audited and found to be in violation and were therefore penalized. If compliant taxpayers do not evade more in the period immediately following a period where they were audited, but penalized violators do, then those results would more favorably comport with loss repair than with BCE.

Given the immediate response (if any) to being audited, the persistence of the effect is also of interest. Theoretically, there should be no immediate response to being audited, and therefore its effects would not persist. The main testable (null) hypothesis related to behavioral conjectures are summarized below:

- Hypothesis 4. An audit has no immediate effects on the level of tax under-reporting or the propensity to acquire information.
- Hypothesis 5. If an audit has an immediate effect on the level of tax under-reporting or the propensity to acquire information, then an audit will not have a lasting effect on the level of tax under-reporting or the propensity to acquire information.

2.5 Results

In the analysis that follows, we first estimate linear regressions to provide a snapshot of the basic treatment effects regarding uncertainty and information services on the tax reporting and information acquisition decisions. These models are presented in Tables 2.4 and 2.6, respectively. Then, we add additional structure and variables to the models which allow us to focus on dynamic behavior. These models are presented in Tables 2.7 and 2.8. To control for possible heteroskedasticity and autocorrelation of unknown form in the regressions, we use robust standard errors with clustering at the participant-level. Further, heteroskedasticity and autocorrelation robust t and F statistics are used when evaluating hypotheses. Table 2.3 provides a description and summary statistics for key variables used in these models.

For the tax reporting regressions, the dependent variable is the level of *expected* under-reporting. As every dollar taken as a credit reduces taxes paid (pre-audit) by one dollar, under-reporting is calculated as the reported credit less the “true” credit. For cases of uncertainty where no information is acquired, the true credit used in the calculation is the midpoint of the range of possible actual credit amounts shown to the participant. When two possible credit amounts are acquired through the information service, the average of the two is the expected true credit. Finally, in the simultaneous information treatment where only one piece of information is acquired, the expected credit is the midpoint of the original uniform distribution and the single possible credit draw. For the information acquisition regressions, we use as the dependent variable a binary indicator variable where a value of one denotes acquisition of the service (i.e. we estimate linear probability models).*

Finally, for the tax reporting models we estimate the treatment effects and behavioral dynamics separately by specific experiment “conditions” as defined by

*These formulations are consistent with the theory, which is also from the perspective of the taxpayer. However, since the information services are unbiased, and given a large number of random credit draws are accumulated over participants and rounds, if we instead use the (ex post) actual level of tax evasion as the dependent variable this leads to trivial differences in results.

Table 2.3: Data Description

Variable Name	Description	Sample Mean	(std. dev.)
Expected Under-reporting	Difference between credit claimed and (expected) actual credit	167.98	(319.90)
Information Acquisition	=1 if information service acquired; =0 otherwise	0.58	(0.49)
Income	Income from the income earnings task. Takes on values of 1000, 1250, or 1500	1271.78	(197.13)
Cost	Cost of information service, in lab dollars. Takes on values of 0, 50 or 100	46.02	(41.58)
Penalized Last Period	=1 if subject was audited in the previous period and penalized; =0 otherwise	0.17	(0.37)
Compliant Last Period	=1 if subject was audited in the previous period and not penalized; =0 otherwise	0.13	(0.34)
Penalty Persistence	The inverse of the number of rounds since the last audit where a penalty was incurred; equals zero in period immediately after audit	0.11	(0.18)
Compliant Persistence	The inverse of the number of periods since the last audit that did not result in a penalty; equals zero in period immediately after audit	0.09	(0.17)
Subjective Probability	The number of times the subject has been audited in past periods divided by the number of past periods.	0.26	(0.14)

Note: the descriptive statistics for Cost and Information Acquisition are computed for Treatments 3 - 5 only.

treatment and information interactions. The first two, “Certainty Baseline” and “Uncertainty Baseline”, simply correspond with all observations from T1 and T2, respectively. The third, “No Information”, includes observations from information service treatments where information was not acquired. The remaining two conditions correspond to observations where information was acquired. “Perfect Information” is associated with T3, and “Imperfect Information” is associated with T4, and those in T5 who sequentially requested information from one or both sources.

2.5.1 Basic analysis of treatment effects

Our analysis first investigates the basic treatment effects identifiable through the experimental design. Tax Reporting Model I (Table 2.4) estimates the mean levels of tax evasion by experiment condition and income level. Information Acquisition Model I (Table 2.6) estimates the mean propensity to acquire information, by treatment, for each unique income and information cost combination.

Table 2.4: Tax Reporting Model I

Dependent Variable: (Expected) under-reporting					
Income Level	Experiment Condition				
	Certainty Baseline (N=1,620)	Uncertainty Baseline (N=1,520)	No Information (N=4,798)	Perfect Information (N=2,029)	Imperfect Information (N=4,493)
Income=1000	160.36*** (31.38)	219.28*** (29.61)	76.67*** (24.82)	46.79** (22.45)	43.31*** (16.01)
Income=1250	186.97*** (30.52)	260.38*** (27.79)	218.84*** (19.07)	72.20*** (20.09)	97.53*** (11.96)
Income=1500	257.32*** (32.32)	310.31*** (35.37)	363.36*** (21.00)	111.39*** (17.52)	141.08*** (13.89)
N=14,454					
R ² =0.29					
F=41.87					
Notes: *, ** and *** denotes estimates that are statistically different from zero at the 10% , 5% and 1% significance levels, respectively. Cluster-robust standard errors are in parentheses.					

One prominent effect from Tax Reporting Model I, is that the level of expected under-reporting is increasing in income. For the Certainty Baseline and Perfect Information conditions, this effect is not consistent with theory, which predicts zero expected under-reporting for all income levels. For the other conditions, subjects face uncertainty and theory predicts that expected under-reporting is decreasing in income, and in our design uncertainty is decreasing in income. Therefore, unobserved behavioral factors related to income appear to be important determinants in the tax reporting decisions. When there is uncertainty with respect to the actual tax liability,

these unobserved behavioral factors are strong enough to counteract the economic net benefits of reporting truthfully.*

The average difference in the level of expected under-reporting is statistically significant beyond the 5% level between any of the three information service conditions and the certainty baseline, uncertainty baseline or no information condition. Also evident from this model is that participants in the information service treatments who do not acquire information (i.e. the No Information subgroup) tend to have reasonably high levels of under-reporting. However, under-reporting for this subgroup is statistically different from, and overall lower than, the under-reporting in the certainty baseline ($F = 4.88; p < 0.01$) or the uncertainty baseline ($F = 6.20; p < 0.01$) for this subgroup.

Table 2.5 succinctly summarizes the effects of information services on under-reporting. It's clear from this table that those who receive information services have the lowest levels of under-reporting. In fact, those receiving information evade roughly 80%, 70% and 60% less, across the respective income levels, as compared to those in the uncertainty baseline.

Table 2.5: Information Services and expected under-reporting

Dependent Variable: (Expected) under-reporting				
	Certainty Baseline	Uncertainty Baseline	No Information	Received Information
Income = 1000	160.36*** (31.38)	219.28*** (29.61)	76.67*** (24.82)	44.35*** (13.08)
Income = 1250	186.97*** (30.52)	260.38*** (27.79)	218.84*** (19.07)	89.98*** (10.35)
Income = 1500	257.32*** (32.32)	310.31*** (35.37)	363.36*** (21.00)	131.29*** (11.01)

*This finding is consistent with earlier findings in a related experiment by (Vossler et al., 2010)

Overall, under-reporting among those who do not acquire information is similar to the levels seen by those in the uncertainty baseline. Coupled with the fact that those who receive information tend to report a tax liability closer to the truth, suggests that the availability of the information service increases truthful reporting.*

Based on our analysis of Tax Reporting Model I, we arrive at the following results:

- Result 1. Tax evasion increases with income.
- Result 2. Those who acquire information evade less than those who do not.

The analysis of Information Acquisition Model I reveals that information requests are increasing with income and decreasing with cost. Turning first to costs, if the information treatments are pooled, the information acquisition rate when the information cost is 100 is lower and statistically different from the information acquisition rate when the information cost is 50 ($F = 3.62, p < 0.01$). Similarly, when the information cost is 50, the information acquisition rate is lower and statistically different from the information acquisition rate when the information cost is 0 ($F = 13.72, p < 0.01$).

Now turning to the effects of income on information acquisition, the pooled difference between the information acquisition rate at the 1500 income level is higher and statistically different from the information acquisition rate at the 1250 income level ($F = 2.92, p < 0.01$). The overall difference in acquisition rates is also higher for the 1250 versus 1000 income groups, however this difference is not significant at conventional levels ($F = 1.52, p = 0.14$).

The negative effect of cost on information acquisition is consistent with theory, however the positive effect of the income level on information acquisition rate is not. Recall that all income levels have the same expected post-tax value of 1000, that the credit is a function of income, and that the uncertainty range (+/- 100% of

*We also estimated an alternative version of this model where effects are allowed to vary by treatment rather than by condition. The treatment-specific result posited here can be shown statistically based on this model.

Table 2.6: Information Acquisition Model I

Dependent Variable: (Information acquired=1; otherwise=0)			
Parameter Setting	Experiment Treatment		
	Perfect Info Available (T3)	Simultaneous Info Available (T4)	Sequential Info Available (T5)
Income = 1000; Cost = 0	0.76*** (0.05)	0.84*** (0.05)	0.89*** (0.03)
Income = 1000; Cost = 50	0.33*** (0.07)	0.42*** (0.06)	0.44*** (0.08)
Income = 1000; Cost = 100	0.20*** (0.06)	0.20*** (0.05)	0.40*** (0.07)
Income = 1250; Cost = 0	0.77*** (0.05)	0.83*** (0.04)	0.90*** (0.03)
Income = 1250; Cost = 50	0.39*** (0.06)	0.56*** (0.05)	0.45*** (0.06)
Income = 1250; Cost = 100	0.24*** (0.05)	0.30*** (0.05)	0.44*** (0.05)
Income = 1500; Cost = 0	0.90*** (0.03)	0.86*** (0.04)	0.90*** (0.04)
Income = 1500; Cost = 50	0.56*** (0.06)	0.60*** (0.06)	0.60*** (0.06)
Income = 1500; Cost = 100	0.30*** (0.06)	0.30*** (0.06)	0.41*** (0.07)
N=11,320			
R ² =0.68			
F=116.50			

Notes: *, ** and *** denotes estimates that are statistically different from zero at the 10% , 5% and 1% significance levels, respectively. Cluster-robust standard errors are in parentheses.

the expected “true” credit) is decreasing with income. Therefore, when participants are in the lowest income level of 1000, theoretically they should have a higher WTP for information. However, this effect is not evident from our analysis. One possible explanation may be that participants decide whether they can “afford” the information based entirely on how well they performed in the income earnings task.

Another possibility is that participants are motivated by relative earnings. When a participant earns a low income, she may be compelled to “keep up with the Joneses” by not paying for information. Overall, the analysis of Information Acquisition Model I has led to the following conclusion:

- Result 3. The propensity to acquire information increases with income and decreases with information cost.

2.5.2 Behavioral Dynamics

The results from Tax Reporting Model II and Information Acquisition Model II are summarized in Tables 2.7 and 2.8, respectively. The variables used in the analysis are constructed to test for post-audit behavioral dynamics. The models also control for subjective probabilities and the basic treatment effects identified previously. The variables “Penalized Last Period” and “Compliant Last Period” are binary variables that indicate whether the subject was audited in the previous period, and whether she either under-reported her tax liability (and therefore paid a penalty) or whether she was compliant (and therefore did not pay a penalty). The “Persistence” variables measure the lasting effects of these two different audit outcomes and are the inverse of the number of rounds that have passed since the most recent “Compliant” or “Penalizing” audit. To parse between the immediate and lasting effects of the audits, the “Persistence” variables are equal to zero in the round that immediately follows an audit.^{*} The “Subjective Probability” variable is constructed using the number of audits in prior rounds divided by the number of prior rounds.[†]

The basic treatment effects identified in the simpler tax reporting and information acquisition models continue to persist when dynamic behavioral controls enter the

^{*}These “Persistence” variables are deliberately constructed so that the (absolute) effect of a past audit declines over time. This is consistent with more general regression specifications that estimate the separate period-by-period effects of past audit outcomes.

[†]Information from the training rounds was used to avoid having to omit period 1 observations from the analysis.

Table 2.7: Tax Reporting Model II

Dependent Variable: (Expected) under-reporting

	Experiment Condition				
	Certainty Baseline (N=1,620)	Uncertainty Baseline (N=1,520)	No Information (N=4,798)	Perfect Information (N=2,029)	Imperfect Information (N=4,493)
Intercept	-109.34 (96.05)	83.91 (98.28)	-332.26*** (68.22)	-72.90 (64.54)	-89.05* (47.91)
Income (in 1000s)	0.22*** (0.07)	0.20*** (0.07)	0.50*** (0.04)	0.15*** (0.04)	0.18*** (0.04)
Cost			-0.61* (0.36)	-1.11*** (0.34)	-0.79*** (0.23)
Penalized Last Period	177.55*** (41.06)	130.70*** (42.47)	249.42*** (26.63)	177.31*** (33.19)	72.99*** (25.02)
Compliant Last Period	-45.00 (46.45)	16.26 (52.41)	-134.98*** (32.62)	44.26 (28.71)	-51.39** (25.20)
Penalty Persistence	296.30*** (85.23)	163.66* (89.49)	469.56*** (54.82)	365.60*** (73.03)	169.20*** (50.04)
Compliant Persistence	-273.92*** (102.37)	-143.84 (107.46)	-363.81*** (70.27)	-3.59 (60.81)	-142.17*** (55.72)
Subjective Probability	-6.07 (166.92)	-406.51** (166.92)	-420.39*** (89.69)	-245.61*** (86.42)	-92.30 (75.58)
N=14,454 R ² =0.39 F=37.21					

Notes: *, ** and *** denotes estimates that are statistically different from zero at the 10% , 5% and 1% significance levels, respectively. Cluster-robust standard errors are in parentheses.

models: under-reporting increases with income, and information acquisition increases with income and decreases with cost.

For all experimental conditions in Tax Reporting Model II there is a positive significant effect corresponding to being audited and penalized, which also has a positive and significant lasting effect. As an illustration, those who report taxes without information report on average 249.42 less in taxes after a penalizing audit,

and the persistence of this effect is rather strong. In the second period following a penalizing audit, under-reporting is $469.56/2$ or 234.80 and by the tenth period following a penalizing audit, its effect on under-reporting is $469.56/10$ or 47.00 lab dollars.

The effects of being audited for individuals who were compliant are largely insignificant. The only significant coefficients for “Compliant Last Period” are under the “No Information” and the “Imperfect Information” conditions. This is also true for “Compliant Persistence”, with the addition of the Certainty Baseline condition. In all situations where the compliance-related estimates are significant, their effects are negative and in the opposite direction as the penalty-related estimates. This result does not support the BCE, which predicts that audits would have positive and significant effects on under-reporting, regardless of whether or not the audit was penalizing. In each condition, “Compliant Last Period” is statistically different from “Penalized Last Period” and “Compliant Persistence” is statistically different from “Penalty Persistence”. The results of the Tax Reporting Model II analysis appear to favor Loss Repair over the BCE, however the results of Information Acquisition Model II are needed to make a stronger, more convincing argument for Loss Repair. For now, we can draw the following conclusions from analysis of Tax Reporting Model II:

- Result 4. Penalizing audits increase the under-reporting of tax liabilities immediately following an audit.
- Result 5. The increase in the under-reporting of tax liabilities from those that were penalized in an audit is persistent and significant.

The results of Information Acquisition Model II indicate a significant and persistent decrease in the propensity to acquire information immediately following a penalizing audit in all three information treatments. The decrease in the propensity to acquire information is behaviorally consistent with the findings from the tax reporting analysis given that, if people plan to under-report their true liability, information

Table 2.8: Information Acquisition Model II

Dependent Variable: (Information acquired=1; otherwise=0)			
Explanatory Variable	Experiment Treatment		
	Perfect Info Available (T3)	Simultaneous Info Available (T4)	Sequential Info Available (T5)
Intercept	0.28*** (0.11)	0.61*** (0.11)	0.70*** (0.11)
Income (in 1000s)	0.30*** (0.07)	0.21*** (0.08)	0.12 (0.08)
Information Cost	-0.55*** (0.06)	-0.56*** (0.05)	-0.49*** (0.05)
Penalized Last Period	-0.17*** (0.04)	-0.15*** (0.04)	-0.17*** (0.04)
Compliant Last Period	0.05 (0.04)	-0.02 (0.04)	-0.02 (0.04)
Penalty Persistence	-0.34*** (0.09)	-0.31*** (0.09)	-0.31*** (0.09)
Compliant Persistence	0.26*** (0.09)	0.02 (0.10)	0.01 (0.10)
Subjective Probability	0.51*** (0.16)	0.03 (0.18)	0.24 (0.16)
N=11,320			
R ² =0.68			
F=143.18			

Notes: * , ** and *** denotes estimates that are statistically different from zero at the 10% , 5% and 1% significance levels, respectively. Cluster-robust standard errors are in parentheses.

which resolves uncertainty regarding their true liability has less value. It may also be the case that some individuals did not request information as a way to justify under-reporting through plausible deniability (i.e. I didn't know I couldn't claim that). The coefficients on "Compliant Last Period" in Information Acquisition Model II are not significant, which is also true for "Compliant Persistence", except for T3. These results suggest that the propensity to acquire information is not significantly affected

by non-penalizing audits. This is further evidence in favor of Loss Repairing behavior. The behavioral findings from the analysis of Information Acquisition Model II are:

- Result 6. The propensity to acquire information is lower immediately following a penalizing audit.
- Result 7. The lower the propensity to acquire information after a penalizing audit is a persistent effect.

The conclusions drawn from Tax Reporting Model II and Information Acquisition Model II are complementary and provide two different sources of results which reject the BCE in favor of Loss Repair. Both models demonstrate persistent behavioral effects after a penalizing audit, suggesting that Loss Repair works through multiple channels, and that misperception about the random audit process is not the driving force behind post audit behavioral phenomena.

2.6 Conclusion

Our most basic finding is that, as predicted by economic theory, the provision of information, whatever the quality, significantly increases tax compliance. Although we have not investigated subject pool effects for these treatments, other work using data from similar experimental settings suggests that observed behavior is broadly consistent across pools (Alm and McKee, 2011).*

The most notable contribution of this research concerns the observed dynamic response to the audit process on tax reporting decisions and the propensity to acquire potentially valuable information. In particular, our findings suggest that current tax reporting is affected by prior experiences with the audit process in a way that is

*Further, as noted above, Alm and McKee (2011) demonstrate the external validity of the experimental setting through a series of comparisons with field data results. This effectively addresses the criticisms of some who have questioned the use of lab experiments in tax compliance research (see Gravell (2008) (commenting on Alm et al. (2009)); Cadsby et al. (2006)). Recall, for the current experiments we have conducted sessions at two institutions and with two pools (students and non-students) at each. Thus we have several ways the pool effects could be analyzed.

consistent with Loss Repair. The lower propensity to acquire costly information which resolves the uncertain tax liability is also consistent with Loss Repair. Therefore, taxpayers who wish to recover losses from a penalizing audit have no desire to be informed of their true tax liabilities. The results of both the tax reporting and information acquisition models build a strong case for loss repair using two different sources of identification.

As a potential policy option, our findings suggest that an information program that directly informs those who were audited, rather than waiting for the information request from the taxpayer, could be an effective method to increase tax compliance.

Research with the data reported in this paper is underway, which further investigates the decision to acquire information, and the factors affecting the propensity to acquire more information (i.e. a second information “draw” in the sequential information setting). With a complex tax system, taxpayers are predicted to respond positively to the provision of information services that reduce the costs of computing true tax liabilities. The results reported here demonstrate: first, when information services are provided the level of underreporting is lowered, and second, the aggregate underreporting levels are lower when information services are provided, even when only a fraction (in this case, 58%) of participants utilize the service. Our experimental setting does not incorporate the cost to the tax agency of providing information services, however the improved tax reporting behavior suggests there is potential for a positive return to providing this service. Finally, the response of participants to the cost of acquiring information was predictable. While the “costs” in the experimental setting were monetary, we would expect a similar response to higher costs even if they were in the form of higher transaction costs, such as a longer waiting time to receive assistance, which are also topics of our ongoing research.

Chapter 3

How Relevant is Irrelevance?

Testing Independence with

Increasingly Irrelevant Alternatives

3.1 Introduction

Rational choice theory in economics implies that one's preference between two options should not depend upon the presence or absence of any other option. That is, if you prefer Skittles to M&Ms, then this should be true whether you are choosing only between the two or otherwise. However, over the last three decades, research in economics, psychology, marketing and elsewhere - typically through the use of controlled laboratory experiments - has demonstrated that this hypothesis of rational choice theory sometimes fails. One particular type of violation, labeled as the "decoy effect", has been shown to occur when an irrelevant option with particular characteristics is thrown in the mix. Specifically, the irrelevant option or "decoy" is strictly less preferred to only one of the other options (the "target") and its presence makes the targeted option more appealing. Technically, the decoy is said to be asymmetrically dominated. In the candy example, it could be the case that some

generic brand of M&Ms is offered at a higher price than the M&Ms. The decoy effect would then work by driving the consumer to now prefer M&Ms over Skittles.

Given the above example, it is not surprising that there is ample evidence of businesses trying to take advantage of the decoy effect. A fast food restaurant may offer a cheeseburger for the same price as a hamburger, a computer company may offer a computer which is identical to the competition and one which has more memory for the same price, or a car manufacturer may continue to produce a line of mid-level vehicles with a poor sales record because having a mid-level vehicle attracts more people to the more expensive luxury vehicles that they produce. It is easy to construct examples where decoy goods could potentially be used to drive sales towards a target good, however the strength and effectiveness of the decoy effect is less apparent. Outside of business contexts, decoy effects have been linked to dating (Lee et al., 2008), political candidates (Pan et al., 1995), job candidates (Highhouse, 1996) and policy issues (Herne, 1997).

The above examples showcase the broad appeal of the subject, but also highlight one of the major challenges in identifying behavioral phenomena: the motivations underlying observed choices in natural-occurring settings are difficult to decipher. In most cases, although one might observe me buying Skittles when the generic M&Ms are not available, I am unlikely to be observed in an identical situation where the generic M&Ms are available for purchase. Or, if I am, there are other confounds that make interpreting my purchase behavior difficult; for example, I might buy (genuine) M&Ms in the second situation simply because I just happened to be in a particular mood for chocolate or was buying candy for a friend. To help facilitate identification, researchers in economics, psychology and other disciplines have often relied on controlled experiments.

In this study, we build upon the existing experimental literature on decoy effects through an innovative design that comes closer than previous work in replicating in the laboratory the fundamental features of a consumer choice setting. In particular, our design involves: (i) choices with financial consequences; (ii) actual, rather than

researcher-manufactured or hypothetical, goods, and *(iii)* the ability for participants to “opt out”, i.e. choose none of the goods. The design allows us to answer several open empirical questions, including whether the observed decoy effects from past studies are an artifact of hypothetical choice and whether decoys significantly change preferences - as measured through elicited willingness to pay (WTP) - or simply drive some to the targeted good when they are otherwise indifferent between the choice options in the absence of a decoy.

3.1.1 Related literature

In this section we briefly review the existing literature, with a focus on the common characteristics of previous experimental designs. We note that there is a great deal of debate over why certain individuals are influenced by the inclusion of the decoy in a choice set. Our intent is not to investigate or critique past explanations of the decoy effect. For a thorough discussion of proposed explanations, we point the interested reader to [Herne \(1996\)](#).

Hypothetical choice

The majority of choice experiments which have shown decoy effects ([Huber et al. \(1982\)](#); [Huber and Puto \(1983\)](#); [Ratneshwar et al. \(1987\)](#); [Simonson \(1989\)](#); [Simonson and Tversky \(1992\)](#); [Lehmann and Pan \(1994\)](#); [Heath and Chatterjee \(1995\)](#)) have used hypothetical goods to construct choice sets. The hypothetical goods used in many of these choice experiments are products which may not exist, and are described using only a few attributes (e.g. two cars, one which has a ride quality of 65 and fuel economy of 36 miles per gallon, and the other which has a ride quality of 95 and fuel economy of 21 miles per gallon). Other experiments have used monetary gambles ([Wedell, 1991](#)), political candidates ([Pan et al., 1995](#)), political issues ([Herne, 1997](#)), and job candidates ([Highhouse, 1996](#)), but these too have been hypothetical situations

in the sense that there were no consequences to choosing one option over another - financial or otherwise.

The use of real goods is much less common. Some of the relevant experiments using real goods are [Simonson and Tversky \(1992\)](#) who gave their subjects the good they chose if selected randomly and [Herne \(1999\)](#) who used real monetary gambles following the design of [Wedell \(1991\)](#). Even when choice sets are constructed using real goods, there are some potential problems that may arise. For example, gambles were picked in [Herne \(1999\)](#) to provide an unambiguous way to define dominance. However, [Simonson \(1989\)](#) have argued that when numerical values are used to describe alternatives, the increased complexity of the decision task could lead to increased decision errors.

It has also been shown that if individuals make choices based on a familiar reference, the lack of familiarity or a reference may lead to a new reference in which they make their decision based on the information available at the time their decision is made. This could lead to behavior that “defies common sense” ([Wedell, 1991](#)). Further evidence of familiarity affecting choice is given by [Ratneshwar et al. \(1987\)](#) who find that attraction effects are enhanced when attributes lack precise meaning for the subjects.

Forced choice

Forcing participants to choose among the available options is a common characteristic among decoy effect studies. In a similar decoy effect experiment, [Dhar and Simonson \(2003\)](#) find that forcing people who essentially would be out of the market for the goods to choose one of the goods can produce systematically biased results about relative preferences. Specifically, the results from [Dhar and Simonson \(2003\)](#) suggest that the opt-out option systematically affects some choice options more than others, and consequently forced-choice procedures may lead to incorrect conclusions.

Focus on choice probabilities at the point of aggregate indifference

Consistent within this line of literature is the apparent manipulation of the choice sets. Researchers have gone to great lengths to construct a choice set involving goods that - in aggregate - subjects were indifferent to. This involved either the manipulation of probabilities for lotteries (Herne (1999); Wedell (1991)) or the construction of choice sets using hypothetical goods (Bateman et al. (2008); Huber et al. (1982); Lehmann and Pan (1994); Ratneshwar et al. (1987); Simonson (1989)). Another potentially problematic area when using hypothetical goods is that subjects may find it difficult to make choices based on a marginal increase in dimension rather than a change in the total level of a particular dimension (Huber et al. (1982); Lehmann and Pan (1994); Ratneshwar et al. (1987); Simonson (1989)). An example would be a consumer faced with the following problem:

Which would you prefer?

Product	Quality Rating (out of 100)	Price
Beer A	45	\$ 1.50
Beer B	60	\$ 2.00

It is difficult to imagine someone who would not only know what utility would be gained from 15 more quality points, but also be able to place a monetary value on that increase. This example not only demonstrates the confusion one might have over numerical values and loss of familiarity, but also provides a way to show the ease with which this choice set might be manipulated. For example, if a pilot experiment revealed that 75% of subjects preferred Beer A and 25% of subjects preferred Beer B, the use of hypothetical products allows the experimenter to adjust the quality or price of the beer until the “desired” status quo of a 50/50 split between respondents is met. It could be argued that this even split is due to increased confusion among study participants, rather than aggregate indifference (Simonson, 1989). Existing designs which have aggregate indifference as the status quo reflect behavior at the

point of indifference, which is unlikely to reflect actual choice situations, and thus unclear what happens in other situations.

3.2 Open questions

The purpose of this research study is to use an innovative experimental design that overcomes some key limitations of previous studies while providing novel insight on the magnitude (specifically, the money equivalent) and variation (based on the proximity between the target and decoy) of the decoy effect.

First, most experiments in the past have involved hypothetical choice situations. There is ample evidence (List and Gallet, 2001) that decisions in such settings do not accurately reflect market behavior, as there is simply no incentive to truthfully reveal one’s preferences. In our design, we utilize both real and hypothetical choice settings and, based on existing literature, we expect to find differences across these choice settings. However, it remains an open question whether hypothetical bias distorts the magnitude of decoy effects.

Second, the option to “opt out” of buying any of the products one is presented with is absent from the design of most past experiments. It is unclear whether being “forced” to make a decision is a key determinant in finding a decoy effect. Further, we specifically chose one of the product categories (women’s shaving razors) for which an entire portion of the population (males) should be uninterested. Therefore, if being forced to buy a product is a significant driver of the decoy effect, given participants rather not choose any of the products, the women’s razors should show pronounced results.

Third, although almost all past studies find evidence of a decoy effect, existing designs limit the metrics available from which to gauge the magnitude of the effect. Typically, experiments have simply identified the share of study participants choosing option A over B with and without the decoy good. Even if the decoy switches 30 percent of consumers towards the target good, this does not tell us by “how much”

preferences have changed. As such, whether the decoy good fundamentally changes one's WTP for goods or simply helps to resolve indifference without any meaningful change in value remains to be answered. More informative is obtaining a money measure of the value of the two options with and without the decoy present. Clearly, the decoy effect could be thwarted by competitors if the decoy effect amounts to pennies.

Fourth, past experiments involve situations where participants are split in their preference of A or B. Situations where participants have "aggregate indifference" are situations where the decoy effect should be the strongest. Because most studies were designed to have aggregate indifference as a baseline, the question of the existence/strength of a decoy effect as we move away from this scenario remains an open question. Aggregate indifference may also mean that individuals are more likely to be indifferent. Given that aggregate indifference is not applicable to most real-world settings, our design uses a number of different products for which the study participants have different aggregate preferences.

Finally, the characteristics of a viable decoy have yet to be investigated. That is, whether certain characteristics of the decoy (relative price difference between the target good's price, lower quantity of the same good, lower quality of a similar good, or off brand of a similar brand name good) make for a stronger decoy effect. Our design varies how "close" the decoy is to the target option across the quality, quantity, brand, and the price dimensions.

Our results suggest that there is not a significant difference between hypothetical and real choice settings, but the decoy effect is an artifact of forced choice setting and that, if anything, there is a negative effect on consumer WTP for goods when a decoy is present. Our results also suggest that consumers forced to buy something they are not in the market for behave differently than those potentially in the market. We find little evidence that any of the characteristics of the decoy, beyond being irrelevant, matter in terms of the strength of the effect on WTP.




Figure 3.1: Illustration of experimental choice set conditions

The next section outlines the conceptual framework and testable hypotheses. It is followed by a brief discussion of our experimental design and then a results section. The paper concludes with discussion and extensions.

3.3 Conceptual framework and testable hypotheses

We will refer to the two option choice set as the “core” choice set, as this is our basis of comparison. Further, we refer to the two options in the core choice set as the target and competitor goods. We assume that the target and competitor goods are close substitutes with levels of attributes defined along two dimensions. Figure 3.1 is a graphic representation of a generic choice set defined along two dimensions (quality and price). It is important to note from the figure that neither the competitor nor the target dominate in both dimensions.

By construction, a decoy option is dominated by the target good in one dimension and dominates the competitor good in that same dimension. Referring to figure 3.1 it is clear that the decoy is asymmetrically dominated by the target good. That is, the decoy can be directly compared to the target good and deemed inferior (higher cost for equal quality), however the decoy cannot be directly compared to the competitor

because of the difference in both cost and quality. Any choice set that includes a decoy as one of the options will be referred to as a decoy choice set. Our experimental design allows us to test the following (null) hypotheses:

- Hypothesis 1. In a real, consequential choice setting, a decoy good does not effect the propensity to purchase a target good.
- Hypothesis 2. If the option to “buy nothing” is a choice available to consumers, a decoy good does not effect the propensity to purchase a target good.
- Hypothesis 3. Familiarity with a product category does not reduce the strength of the decoy effect.
- Hypothesis 4. WTP for a targeted or competitor good is unaffected by the presence of a decoy good.

3.3.1 Willingness to pay

Besides testing for whether or not the decoy effect persists when choices are made concerning real goods with real consequences, it is also useful to know how the decoy effect changes preferences. Past studies demonstrate the decoy effect through showing an aggregate preference for one good over another when a decoy is present where in the absence of a decoy there was aggregate indifference between the two goods. Of course, aggregate indifference likely means that many individuals themselves are (near) indifference. As such, it’s not clear whether the decoy fundamentally changes one’s WTP for goods or simply helps to resolve indifference without any meaningful change in value.

In order to obtain an accurate estimate of identify the change in WTP that may occur in the presence of a decoy, it is necessary to include an “opt-out” option in the choice set. Theoretically, this allows one measure differences in utility between the consumption of a good and the status quo (i.e. no good). Including a price attribute in the choice set allows one to estimate the marginal utility of income, which can then

be used to convert this utility difference into WTP. Of course, allowing for opt-out mimics real life purchasing situations and the decoy effect could strictly be driven by participants being forced to make a decision. When participants are forced to choose between two costly options, we can only measure differences in utility (and WTP) between the two goods.

The next section outlines the components of the experimental design necessary to test the above hypotheses.

3.4 Experimental Design

To identify key treatment effects of interest, the design varies: (a) presence/absence of decoy; (b) the good; (c) price of the goods; (d) presence/absence of a “buy nothing” option; and (e) whether the decisions are hypothetical or involve actual money.

Design Elements

Each experimental session is comprised of three separate elements: an earnings task, a product purchasing treatment, and a product purchasing treatment with real financial consequences. The unrelated experiment is used as an earnings task to mitigate “house money” effects* and is always conducted at the beginning of the experimental session.

The design utilizes six distinct choice sets, defined by the particular target, competitor, and decoy (if any) good used, which are presented as product purchase scenarios. The choice sets are summarized in Table 3.1. Investigating different products will help establish the robustness of the decoy effect.

*“House money” refers to a laboratory endowment which participants treat differently than money they brought in from outside the lab [Clark \(2002\)](#). House money has been shown to cause participants to make less self-interested (or more risky) choices in a number of experiments ([Harrison, 2007](#); [Ackert et al., 2006](#); [Thaler and Johnson, 1990](#)). To mitigate any effects an endowment may have on subjects’ propensity to consume, participants “worked” for their income through a timed data entry task. Participants knew how much money they earned prior to making a purchase decision and were instructed that this money was theirs to keep, less any money they choose spend on products.

Each treatment is comprised of 18 product purchase scenarios. In particular, the respondent faces each of the six choice sets three times, with the prices of the goods varying across the same incidences of the same choice set.* In each product purchase scenario, subjects are shown a picture, description, and a price for each of the two or three goods (depending on the treatment) from one of six choice sets.

Table 3.1: Choice set composition by treatment

Choice			Treatments			
Set		Good	1/5	2/6	3/7	4/8
1	Target	Duracell AA 10-pack	X	X	X	X
	Competitor	Energizer AA 10-pack	X	X	X	X
	Decoy (if present)	Duracell AA 2-pack		X		X
	Buy Nothing	Buy Nothing			X	X
2	Target	Duracell AA 10-pack	X	X	X	X
	Competitor	Energizer AA 10-pack	X	X	X	X
	Decoy (if present)	Duracell AA 6-pack		X		X
	Buy Nothing	Buy Nothing			X	X
3	Target	Bic Soleil 3-blade razor	X	X	X	X
	Competitor	Schick Quattro 4-blade razor	X	X	X	X
	Decoy (if present)	Store brand 3-blade razor		X		X
	Buy Nothing	Buy Nothing			X	X
4	Target	Schick Quattro 4-blade razor	X	X	X	X
	Competitor	Bic Soleil 3-blade razor	X	X	X	X
	Decoy (if present)	Store brand 4-blade razor		X		X
	Buy Nothing	Buy Nothing			X	X
5	Target	UT Pint Glass	X	X	X	X
	Competitor	UT Coffee Mug	X	X	X	X
	Decoy (if present)	Unbranded Pint Glass		X		X
	Buy Nothing	Buy Nothing			X	X
6	Target	UT Coffee Mug	X	X	X	X
	Competitor	UT Pint Glass	X	X	X	X
	Decoy (if present)	Unbranded Coffee Mug		X		X
	Buy Nothing	Buy Nothing			X	X

*The initial design exposed participants to only one real or one hypothetical setting consisting of 15 choice questions.

Determining Prices

Randomly varying the prices of the products allows for the identification of the participant's WTP. Prices for the target and competitor goods are determined through one of two processes which depend on the choice setting (hypothetical or real) and the product purchase scenario (1-18). For hypothetical choice settings, prices are determined as follows:

Scenario 1-6:

Target Price = draw from uniform distribution with supports [\$0,\$10]

Competitor Price = draw from uniform distribution with supports [\$0,\$10]

Scenario 7-18:

Target Price = draw from uniform distribution with supports [\$0,\$10]

Competitor Price = Target Price + draw from uniform distribution with supports [\$-1,\$1]

In the real choice setting, prices are determined by:

Scenario 1-6:

Target Price = draw from uniform distribution with supports [\$0,\$5]
60% of the time, and supports [\$0,\$10] 40% of the time

Competitor Price = draw from uniform distribution with supports [\$0,\$5]
60% of the time, and supports [\$0,\$10] 40% of the time

Scenario 7-18:

Target Price = draw from uniform distribution with supports [\$0,\$5]
60% of the time, and supports [\$0,\$10] 40% of the time

Competitor Price = Target Price + draw from uniform distribution with supports [\$-1,\$1]

The decoy price was not dependent on the setting or scenario, but only on the relative prices of the target and competitor good. The decoy price was determined as follows:

If Target Price > Competitor Price

Decoy Price = Target Price + draw from uniform distribution with supports [\$0,\$1]

If Target Price $\leq$ Competitor Price

Decoy Price = Target Price + draw from uniform distribution with supports [0,1], multiplied by the difference between the Target Price and the Competitor Price

Prices for the real and hypothetical scenarios were determined from pilot sessions. In particular, price distributions were chosen to identify the total WTP of each good while allowing a fair amount of choice situations where the non-core good can be argued to be a decoy. Note that our desire to estimate total WTP, which mandates we have sufficient variation in the price of the competitor good, precludes having unambiguously a decoy in each choice set. The pilots revealed substantial differences between hypothetical and real choice (between \$5 and \$8 in the hypothetical setting and \$0 and \$2 in the real), with no one choosing to purchase any of the goods at prices above \$3.04 for the real scenarios and less than 5% of hypothetical purchases taking place above \$10. Initial results from the pilot experiment also revealed that the battery brands were close substitutes and the razors were close substitutes; meaning that very few people chose the higher cost product when the price difference was above \$1 (12% in the hypothetical setting and 0% in the real setting).

3.4.1 Treatments

The main effects in the experiment are: decoy verses no decoy, forced verse opt-out (the option to select none of the products) and real verses hypothetical choice settings. This simple 2x2x2 design results in a total of 8 treatments. Each subject participates

in two treatments during a session, one real and one hypothetical.* The order in which subjects participate in the real and hypothetical treatment is varied between sessions. The 8 treatments are summarized in table 3.2.

Table 3.2: Treatments by main effects

Treatment	Setting	Choice	Products
T1	Hypothetical	Forced	No Decoy
T2	Hypothetical	Forced	Decoy
T3	Hypothetical	Opt-out	No Decoy
T4	Hypothetical	Opt-out	Decoy
T5	Real	Forced	No Decoy
T6	Real	Forced	Decoy
T7	Real	Opt-out	No Decoy
T8	Real	Opt-out	Decoy

The baseline treatments (T1 and T3) serve as a means to elicit relative preferences (T1) and measures of baseline WTP (T3) for the core options. In treatments T1 and T2, participants are “forced” to select one of the product from the choice set, and in treatments T3 and T4 participants have the option to opt-out of having to select a product from the choice set. Treatments T3 and T4 have the additional benefit of more closely mimicking an actual purchasing decision in a market setting compared to treatments T1 and T2 because of the option to purchase nothing.

Hypothetical bias, whereby participants overstate their true value for a good when responding to inconsequential value elicitation questions, is a well known problem throughout the valuation literature. It remains an open question, however, whether hypothetical bias will serve to distort the decoy effect. For this reason, T5 through T8 are the real choice setting representations of T1 through T4.

3.4.2 Product Categories

This study utilizes three different product categories to construct the choice sets in both the real and hypothetical settings: batteries, university branded drinking glasses,

*With the exception of the original four sessions which involved only one treatment.

and womens' shaving razors. The products used can be found in Table 3.1. Each product is a well known brand name and allows us to test the decoy effect in three different dimensions: a quantity decoy, a quality decoy, and a brand decoy.

We define the quantity decoy as a smaller amount of the target good. Batteries were chosen for the quantity decoy because of their advantageous packaging and because of their use in past experiments (Heath and Chatterjee, 1995). Both brands, Duracell and Energizer, have no apparent difference in quality. Irwin et al. (1993) states that preference reversals are most likely to occur when consumers are indifferent. For this reason it is also likely that batteries would be susceptible to decoy effects.

The quality decoy is defined as a generic "store" brand of the target good and it is assumed that the generic brand is perceived as having lower quality than the name-brand good. There are many different brands of razors which have similar attributes (i.e. number of blades, color, etc.) including store brands that imitate the packaging and design of name-brand razors. An added benefit of using women's razors is that it forces an entire population of participants (males) to make a decision concerning a product that is unintended for their use. This should highlight whether forcing participants to choose among alternative they should be uninterested in drives any part of the decoy effect.

Unlike the quality decoy (where we assume brand is a signal of quality), the brand decoy is defined as having all the same qualities as the target good, less the brand name. In the case of University branded drinking vessels, it is assumed that adding an officially licensed university logo to a coffee mug or pint glass does not change its quality, but rather it adds value specifically through branding.

3.4.3 Participants and procedures

Experiments were conducted in a dedicated experimental economics laboratory at the University of Tennessee. The participant pool consists of current University

of Tennessee undergraduate students. Recruiting was conducted using the Online Recruiting System for Experimental Economics (ORSEE) developed by Greiner (2004). Subjects were contacted, via email, and were permitted to participate in only one experimental session.* Only participants who had registered for a particular session were allowed to participate, and no participant had prior experience in this experimental setting. Methods adhere to all guidelines concerning the ethical treatment of human participants. Earnings ranged between \$10 and \$16 for the earnings task and averaged \$14. Sessions lasted between 25 and 40 minutes and a total of 382 participants took part in these sessions.

Sessions proceed in the following fashion. Each participant sits at a computer located in a cubicle, and is not allowed to communicate with other participants. The participants are informed that all decisions will be private and that the experimenter is unable to observe their decisions during the experiment. To emphasize the fact that the experimenter is not observing the participants decisions, and to minimize possible peer and experimenter effects that could affect the decisions of the participants, the experimenter remains seated in the same spot throughout the session. All of the participants' decisions are made on their computer.

Participants are informed that the session consists of three separate, short experiments. Participants are provided with the instructions for the unrelated earnings experiment and these instructions are read aloud by the experimenter. (Appendix A provides representative screenshots from the experiment and appendix C provides instructions from one of the treatments.)

Once the earnings task is complete, the experimenter regains the attention of the participants and an information screen regarding the first of two product purchasing experiments (either real or hypothetical) is displayed on participants' monitors. The choice setting and product purchase scenarios are described and then the experimenter passes around products representative of the actual products under consideration.

*Other experimental projects were ongoing at the time and participants may have participated in unrelated experimental sessions.

After the instructions are read and each participant has an opportunity to inspect each of the products, participants proceed through 18 product purchase scenarios at their own pace.

Table 3.3: Treatment pairs by session

Session	First Product Treatment	Second Product Treatment
1	T3	T7
2	T7	T1
3	T7	T4
4	T3	T6
5	T8	T3
6	T4	T5
7	T4	T8
8	T8	T2
9	T5	T3
10	T1	T5
11	T1	T8
12	T5	T2
13	T2	T7
14	T6	T1
15	T6	T4
16	T2	T6

In the experiments with real financial consequences, before participants start the product purchase scenarios they are informed that one of the scenarios will be randomly chosen by the computer to be financially binding and that they will actually receive the product they selected in that scenario and its purchase price would be subtracted from their earnings. Whether participants complete the hypothetical or real treatment first was varied between sessions and treatment pairs were based on optimal experimental design to minimize the correlation between treatment variables. Table 3.3 outlines the session specific treatment pairs.

To minimize any ordering effects, each purchase scenario is assigned a number (1-18) and the computer randomly selects each scenario without replacement until the participant has completed all 18 scenarios.

After completing both purchasing treatments, participants fill out a brief demographic questionnaire and are then paid privately at the end of the session and given any good they purchased in the consequential treatments.

3.5 Econometric Method

3.5.1 Willingness to pay estimation

The standard underlying analytical framework for estimating WTP is the random utility model. The indirect utility derived from alternative j in choice set k for individual i is given by the expression

$$V_{ijk} = \beta_t t_{ijk} + \beta \mathbf{x}_{ijk} + \theta \mathbf{a}_{ijk} + \eta_{ijk} \quad (3.1)$$

which is comprised of a deterministic component, which includes the vector of non-price product attributes ($\mathbf{a}_{ijk}$), a cost attribute (t_{ijk}) as well as a vector of individual specific characteristics ($\mathbf{x}_{ijk}$), and a random error term (η_{ijk}). The m -dimensional vector $\mathbf{a}_{ijk}$ describes the attributes associated with alternative j from the choice set k . The vector of coefficients β is assumed to be common across all participants and the stochastic i.i.d. error term η_{ijk} captures unobserved participant heterogeneity and is distributed Type I extreme value.

This specification is: (a) linear in the unknown parameters, (b) assumes the marginal utilities of each attribute do not vary across participants, and (c) assumes constant marginal utility for each attribute.

Characterizing utility as in (3.1) provides a convenient and fairly direct means to measure value using the *conditional logit* model (?) which is based on the marginal changes in attributes. Note that we obtain the following:

- The marginal utility of income is $-\beta_t$
- The marginal rate of substitution between attributes r and s is: $\frac{\partial V_{ijk}/\partial a_r}{\partial V_{ijk}/\partial a_s} = \frac{\beta_r}{\beta_s}$

- The marginal WTP for attribute r is: $\frac{\partial V_{ijk}/\partial a_r}{\partial V_{ijk}/\partial t_j} = -\frac{\beta_r}{\beta_t}$

Based on the assumptions of the error term in (3.1), we also know that the probability that respondent i chooses alternative r from choice set k is given by

$$P_{ijk} = \frac{\exp(a'_r \theta_r)}{\exp(\mathbf{a}'_{ijk} \boldsymbol{\theta})}. \quad (3.2)$$

3.5.2 Finding the point of indifference with random prices

In order to relate our results to those of past experiments, we need to determine from the choice experiment data the point of aggregate indifference. Since prices vary within and across products for each of the six product scenarios, this gives us the ability to estimate a price differential that leaves the participant indifferent between the available alternatives in the choice set. This *indifference price* is unique to this study, as this is the first study of its kind to estimate values for products while simultaneously investigating the decoy effect. Where past studies are able to make direct comparisons of the choices between the core choice sets and decoy choice sets, our approach will require the following additional steps:

Step 1: Estimate the price difference between the target and competitor good in the core choice set treatment that makes individuals indifferent between the two goods (the indifference price).

Step 2: Estimate the choice probability associated with the target good in the decoy treatment, evaluated at the indifference price from step 3.1.

To estimate the indifference price, we undertake a conditional logit analysis for each product category in the forced, core choice set treatments (T1 and T5). When only two choices are available in the choice set, the panel structure of the data requires that we assume one of the two options (in our case, the competitor) to be the “pseudo-status quo” option. In this sense, the conditional logit analysis normalizes the characteristics of the competitor good and the results are based on the differences

between the target good and the competitor good.* Finding the indifference price is straightforward once the distribution of WTP for the target good is estimated.

Rearranging (3.2), we solve

$$\frac{1}{1 + \exp(-x'_{ki}\beta)} = 0.5 \quad (3.3)$$

which yields the solution to the indifference price.

Not only is this necessary to make comparisons with past decoy effect research, the use of forced and opt-out treatments restricts the dimensions we have to make comparisons between treatments to marginal WTP. This approach enables a direct comparison of the two choice distributions at the estimated indifference price and the associated propensity to choose the target good.

Experimental controls

To control for possible order effects, decoy “closeness” effects, and wealth effects, our full econometric model accounts for all relevant interactions. Wealth effects are controlled through accounting for subject earnings from the income task. An example of the econometric model with the full set of control variables is given in Appendix D.

To control for possible order effects, we include variables for the session order, scenario order, and the orientation of the choice question on the screen. Session order is a dummy variable that equals one for the second product purchase treatment that subjects participate in. Scenario order is a discrete variable from one to 18 designating when a subject saw a particular product purchase scenario within a treatment. The orientation variable identifies the position of the choice question assigned to the target good within a particular product purchase scenario. The orientation variable takes

*The results from such an analysis, where only comparative differences are used, is useful in producing comparative differences in WTP for the core goods. However, forced choice data distorts estimates of WTP when a true zero-cost status-quo is not an option.

on the values 1-3 which represent the left most position on the screen (1), center position (2), and right most position (3).

Decoy closeness is accounted for along two different dimensions: price and order. The “Decoy Price Difference” variable is constructed using the difference between the decoy price and the target price. Because we have designated a decoy as an asymmetrically dominated good, we had to control for instances where the target, competitor, and decoy prices were not properly ordered to define the decoy as being asymmetrically dominated. These instances include when the target price is betwixt the (low cost) competitor good and the (high cost) decoy good.

From the full econometric model, we tested conducted joint tests of significance to eliminate variables that were causing mischief by adding noise to the estimation.

3.6 Results

The first step in the analysis investigates the basic treatment effects identifiable through the experimental design. The calculated indifference prices and the estimated proportion of target good purchases resulting from those indifference prices for each of the product categories is organized by the four main treatment effects. The results of the hypothetical forced (HF) choice treatments (T1 and T2) are summarized in table 3.4; real forced (RF) choice treatments (T5 and T6) are summarized in table 3.5; hypothetical opt-out (HO) option treatments (T3 and T4) are summarized in table 3.6; and finally real opt-out (RO) option treatments (T7 and T8) are summarized in table 3.7.

Across the four different choice settings, we see a statistically significant decoy effect in each product category for at least one of the goods. The decoy effect in the battery category with the quantity decoy is significant in all four choice settings, as is the 4-blade store brand razor quality decoy. The magnitude of the decoy effect, the shift of the proportion of consumers predicted to choose the target good, is consistent

Table 3.4: Hypothetical Forced Choice

Target	Decoy	Indifference Price	Target Chosen ^γ	Obs.
Batteries		\$0.33 (0.16)		2026
Duracell	2-pack of Duracell	\$1.37** (0.55)	0.80*** (0.08)	
Duracell	6-pack of Duracell	\$1.44*** (0.59)	0.82*** (0.08)	
Drinking Vessels		\$2.11 (0.49)		1804
UT Glass	Generic glass	\$4.84** (1.20)	0.81*** (0.08)	
UT Mug	Generic mug	\$0.54** (0.49)	0.69*** (0.06)	
Razors		\$1.29 (0.36)		1756
Bic Razor	3-blade store-brand	\$0.50* (0.30)	0.69*** (0.07)	
Schick Razor	4-blade store-brand	\$2.05 (0.76)	0.69* (0.14)	

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively. ^γEstimated proportion that target good chosen given the indifference price.

with the results of past experiments.* Table 3.9 summarizes the difference in the target share between treatment conditions.

Table 3.8 summarizes the difference in WTP between treatment conditions. Based on existing literature on hypothetical bias, we expect to see a difference in the marginal WTP between real and hypothetical treatments. This is not always the case, however. What we see in the forced treatments seems to support the existing literature on all but the decoy drinking vessel scenarios, the 3-blade decoy scenario (for men and the aggregate), and the baseline and 6-pack battery scenarios; and, in the forced settings, all the coefficients have the expected signs. Comparisons between

*See Heath and Chatterjee (1995) for a meta analysis.

Table 3.5: Real forced choice

Target	Decoy	Indifference Price	Target Chosen ^γ	Obs.
Batteries		\$0.05 (0.07)		1368
Duracell	2-pack Duracell	\$0.06 ^A (0.08)	1.00 ^A (0.00)	
Duracell	6-pack Duracell	\$0.66 (0.42)	0.93*** (0.11)	
Drinking Vessels		\$0.64 (0.22)		1210
UT Glass	Generic glass	\$3.00** (0.99)	0.93*** (0.07)	
UT Mug	Generic mug	\$0.52 (0.31)	0.53 (0.08)	
Razors		\$0.16 (0.10)		1166
Bic Razor	3-blade store-brand	\$0.18 (0.29)	0.70* (0.11)	
Schick Razor	4-blade store-brand	\$0.23 ^A (0.10)	1.00 ^A (0.00)	

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively. ^AOnly the target good was ever chosen in this scenario. ^γEstimated proportion that target good chosen given the indifference price.

real and hypothetical treatments when there is an opt-out option are less encouraging and are the topic of future investigation.*

We are able to infer from our analysis, the following two results:

- Result 1: There is a statistically significant decoy effect in treatments using real goods.

*The women's razors were intended to test familiarity as a driver of the decoy effect. However, the difference between males and females is not significant in the HF treatments ($\chi^2 = 1.92$, $p = 0.59$), the RF treatments ($\chi^2 = 4.30$, $p = 0.23$), the HO treatments ($\chi^2 = 3.52$, $p = 0.47$), or the RO treatments ($\chi^2 = 1.49$, $p = 0.83$).

Table 3.6: Hypothetical Opt-out

Target	Decoy	Indifference Price	Target Chosen ^γ	Obs.
Batteries		\$0.51 (0.37)		3858
Duracell	2-pack Duracell	\$1.96*** (0.63)	0.68** (0.09)	
Duracell	6-pack Duracell	\$2.57*** (0.63)	0.75*** (0.08)	
Drinking Vessels		\$2.82 (0.66)		2739
UT Glass	Generic glass	\$3.40 (1.39)	0.56 (0.10)	
UT Mug	Generic mug	\$0.10** (0.88)	0.79*** (0.08)	
Razors		\$0.67 (0.70)		2466
Bic Razor	3-blade store-brand	\$1.17 (1.60)	0.66 (0.10)	
Schick Razor	4-blade store-brand	\$4.02* (2.00)	0.76*** (0.08)	

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively. ^AOnly the target good was ever chosen in this scenario. ^γEstimated proportion that target good chosen given the indifference price.

- Result 2: There is a statistically significant decoy effect in treatments where participants have the option to opt-out of buying a product.

3.6.1 Analysis of willingness to pay

Estimates of WTP for the different product categories are found in Table 3.10. Based on existing literature on hypothetical bias, we expect to see a difference in the marginal WTP between real and hypothetical treatments, and indeed this is what we observe.

Table 3.7: Real Opt-out

Target	Decoy	Indifference Price	Target Chosen ^γ	Obs.
Batteries		\$0.60 (0.21)		3648
Duracell	2-pack Duracell	\$0.89*** (0.42)	0.86*** (0.04)	
Duracell	6-pack Duracell	\$0.07** (0.27)	0.70*** (0.07)	
Drinking Vessels		\$1.80 (0.33)		3255
UT Glass	Generic glass	\$1.58 ^A (0.46)	1.00 ^A (0.00)	
UT Mug	Generic mug	\$0.66* (0.50)	0.74*** (0.08)	
Razors		\$1.35 (0.44)		3075
Bic Razor	3-blade store-brand	\$1.98*** (1.20)	0.98*** (0.02)	
Schick Razor	4-blade store-brand	\$0.23 ^A (0.80)	1.00 ^A (0.00)	

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively. ^γEstimated proportion that target good chosen given the indifference price.

From our analysis, in nearly every scenario, hypothetical or real, when a decoy effect was shown to affect choice, there was a negative effect on WTP. Other than the Duracell batteries, for which there was an increase in WTP in the hypothetical setting, the only statistically significant differences in WTP are negative. Comparing the differences in WTP, or the marginal willingness to pay across treatments, it seems as though the story lies in these differences and not the absolute level of WTP.*

*Table 3.8 shows marginal WTP is not statistically different between the HF and HO treatments. The relationship between marginal WTP and choice is not immediately apparent and remains a topic for future investigation.

Table 3.8: Differences in marginal willingness to pay across treatments

Product	Decoy	Forced Hypothetical - Forced Real	Forced Hypothetical - Opt-out Hypothetical	Forced Hypothetical - Opt-out Real	Forced Real - Opt-out Hypothetical	Forced Real - Opt-out Real	Opt-out Real - Opt-out Hypothetical
Batteries		0.28	-0.18	0.93***	-0.46	0.65***	-1.11***
	2-pack	1.31***	-0.59	0.48	-1.90***	-0.83**	-1.07
	6-pack	0.78	-1.13	1.37**	-1.91***	0.59	-2.50***
Drinking Vessels		1.47***	-0.71	0.31	-2.18***	-1.16***	1.02
	Glass	1.84	1.44	3.26***	-0.40	1.42	-1.82
	Mug	0.02	0.44	-0.12	0.42	-0.14	0.56
Razors		1.13***	0.62	-0.06	-0.51	-1.19***	0.68
	3-blade	0.32	-0.67	-1.48	-0.99	-1.80	0.81
	4-blade	1.82**	-1.97	1.82*	-3.79**	0.00	-3.79*

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively.

Table 3.9: Differences in estimated target share across treatments

Product	Decoy	Forced Hypothetical - Forced Real	Forced Hypothetical - Opt-out Hypothetical	Forced Hypothetical - Opt-out Real	Forced Real - Opt-out Hypothetical	Forced Real - Opt-out Real	Opt-out Real - Opt-out Hypothetical
Batteries	2-pack	-0.20**	0.12	-0.06	0.32***	0.14***	0.18*
	6-pack	-0.11	0.07	0.12	0.18	0.23*	-0.05
Drinking Vessels	Glass	-0.12	0.25**	-0.19***	0.37***	-0.07	0.44***
	Mug	0.16	-0.10	-0.05	-0.26**	-0.21*	-0.05
Razors	3-blade	-0.01	0.03	-0.29***	0.04	-0.28***	0.32***
	4-blade	-0.31**	-0.07	-0.31***	0.24***	0.00	0.24***

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively.

Table 3.10: Willingness-to-pay: Parametric Estimation

Decoy	None	Duracell		None	Energizer	
		2-pack	6-pack		2-pack	6-pack
Hypothetical N = 3858	7.49 (0.50)	8.71* ^A (0.78)	9.25** (0.80)	6.98 (0.50)	6.75 ^A (0.67)	6.68 (0.79)
Real N = 3648	0.61 (0.19)	0.84 (0.23)	0.83 (0.27)	1.22 (0.22)	-0.06*** (0.38)	0.76* (0.25)
Decoy	None	UT Glass		None	UT Mug	
		Glass	Mug		Glass	Mug
Hypothetical N = 2739	8.08 (0.71)	7.04 (0.90)	6.17** (0.98)	5.26 (0.72)	3.64 (1.40)	6.27 (0.91)
Real N = 3255	1.85 (0.25)	1.37 ^A (0.37)	0.76* (0.45)	0.06 (0.27)	-0.21 ^A (0.37)	0.09 (0.40)
Decoy	None	Schick		None	Bic	
		3-blade	4-blade		3-blade	4-blade
Hypothetical N = 2466	5.09 (0.75)	1.78** (1.30)	2.82** (1.00)	4.41 (0.80)	2.95 (1.20)	-1.19*** (2.10)
Real N = 3075	0.77 (0.29)	-2.84*** (1.30)	-1.26*** ^A (0.65)	-0.58 (0.37)	-0.86 (0.51)	-1.02 ^A (0.63)

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively.

- Result 3: Decoy goods seem to have a negative (if any) impact on willingness to pay.

The driving force behind Result 3 may be the overwhelming difference in opt-out rates among participants, particularly in the real choice settings. Table 3.11 illustrates this point. In order to test if this result is an artifact of the parametric estimation technique which allows for negative WTP values, we also estimate WTP non-parametrically. The results, however, are very similar to those found in the parametric estimation. Table 3.12 shows a similar pattern of WTP estimates which are lower in almost every significant case.

Table 3.11: Opt-out rates

Decoy	Batteries		
	None	2-pack	6-pack
Hypothetical	32%	36%	34%
Obs.	875	215	201
Real	78%	80%	76%
Obs.	692	261	268
Decoy	Drinking Vessels		
	None	Glass	Mug
Hypothetical	36%	51%	39%
Obs.	513	203	200
Real	70%	79%	78%
Obs.	556	271	263
Decoy	Razors		
	None	3-blade	4-blade
Hypothetical	50%	62%	64%
Obs.	439	219	171
Real	86%	87%	90%
Obs.	522	287	225

3.7 Conclusion

This study built upon the existing experimental literature on decoy effects through an innovative design which preserved the fundamental features of a consumer choice setting in the laboratory. Our design involved choices with financial consequences, real goods, and the ability for participants to opt out. Through our novel experimental approach and econometric analysis, we were able to demonstrate that: (1) the decoy effect is not an artifact of hypothetical settings; (2) the decoy effect is not driven by forced choice; and finally (3) due to the lack of a meaningful and significant change in WTP, that is an increase in the WTP for the target good or decrease in the WTP for the competitor good, our analysis suggests that decoys do little more than sway individuals at the point of indifference.

Table 3.12: Willingness-to-pay: Non-parametric Estimation

Decoy	None	Duracell		None	Energizer	
		2-pack	6-pack		2-pack	6-pack
Hypothetical	\$ 9.50 (0.36)	\$ 9.64 (1.01)	\$ 10.16 (1.17)	\$ 8.25 (0.28)	\$ 9.30 (1.01)	\$ 10.82** (1.32)
Obs.	523	143	149	519	122	136
Real	\$ 2.72 (0.21)	\$ 2.35 (0.31)	\$ 2.56 (0.37)	\$ 2.09 (0.18)	\$ 1.93 (0.26)	\$ 1.94 (0.29)
Obs.	419	171	178	403	170	170
Decoy	None	UT Glass		None	UT Mug	
		Glass	Mug		Glass	Mug
Hypothetical	\$ 11.35 (0.74)	\$ 8.39*** (0.61)	\$ 9.22** (0.76)	\$ 9.89 (0.67)	\$ 6.67*** (0.5)	\$ 8.83 (0.82)
Obs.	327	165	146	271	111	124
Real	\$ 2.93 (0.14)	\$ 3.12 (0.2)	\$ 3.26* (0.18)	\$ 3.00 (0.15)	\$ 2.01*** (0.18)	\$ 1.94*** (0.21)
Obs.	342	194	187	350	150	145
Decoy	None	Bic		None	Schick	
		3-blade	4-blade		3-blade	4-blade
Hypothetical	\$ 8.85 (0.51)	\$ 6.38*** (0.47)	\$ 6.57*** (0.65)	\$ 9.21 (0.55)	\$ 7.61** (0.51)	\$ 8.52 (0.6)
Obs.	258	153	77	278	194	115
Real	\$ 2.70 (0.14)	\$ 1.76*** (0.14)	\$ 1.68*** (0.21)	\$ 2.74 (0.14)	\$ 1.79*** (0.14)	\$ 2.60 (0.16)
Obs.	353	187	127	352	208	152

Note: Cluster robust standard errors are in parentheses. *, **, and *** denote that parameter is statistically different from status-quo at the 10%, 5%, and 1% significance levels, respectively.

Although we estimate minimal effects on WTP, the long-run effects are still uncertain. If a decoy drives somebody to buy a product, she would gain familiarity with the product that over time can strengthen her (relative and absolute) preferences for it. Even if the decoy is eliminated, perhaps through price changes or product changes, a person may nevertheless continue to purchase the product. As such, a

competitor may still lose market share over time, even if they change their marketing strategy to combat what was a decoy effect.

Aside from the potential long-run effects of preference formation, it may be the case that there are goods for which the WTP difference or purchase probability is more or less pronounced. For products such as less familiar goods, high-valued goods, different types of good (durable versus nondurable), there can be considerable differences, which remains an open question.

We plan to further investigate the decoy effect using the data set obtained from this study. These extensions include investigating the demographic information obtained from the post-experiment questionnaire for any possible effects, and utilizing a mixed logit framework to relax the assumptions implicitly made by the conditional logit model (most notably the IIA assumption).

Bibliography

Bibliography

- Ackert, L., Charupat, N., Church, B., and Deaves, R. (2006). An experimental examination of the house money effect in a multi-period setting. *Experimental Economics*, 9(1):5–16. [56](#)
- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine. *Econometrica: Journal of the Econometric Society*, pages 503–546. [7](#), [8](#)
- Allingham, M. and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of public economics*, 1(3-4):323–338. [20](#)
- Alm, J., Cherry, T., Jones, M., and McKee, M. (2009). Encouraging filing via tax credits and social safety nets. *The IRS Research Bulletin: Proceedings of the 2008 IRS Research Conference*, pages 43–57. [45](#)
- Alm, J., Cherry, T., Jones, M., and McKee, M. (2010). Taxpayer information assistance services and tax compliance behavior. *Journal of Economic Psychology*, 31(4):577–586. [21](#), [32](#)
- Alm, J., Jackson, B., and McKee, M. (1992a). Estimating the determinants of taxpayer compliance with experimental data. *National Tax Journal*, 45(1):107–114. [22](#), [28](#)
- Alm, J., Jackson, B., and McKee, M. (1992b). Institutional uncertainty and taxpayer compliance. *The American Economic Review*, pages 1018–1026. [22](#), [27](#)

- Alm, J. and McKee, M. (2006). Audit certainty and taxpayer compliance. *National Tax Journal*, 59(4):801–816. [27](#)
- Alm, J., K. B. and McKee, M. (2011). On the external validity of tax compliance experiments. *prepared for presentation at 2011 IRS Research Conference, Washington, DC*. [45](#)
- Andreoni, J., Erard, B., and Feinstein, J. (1998). Tax compliance. *Journal of economic literature*, 36(2):818–860. [22](#), [25](#)
- Bateman, I., Munro, A., and Poe, G. (2008). Decoy effects in choice experiments and contingent valuation: asymmetric dominance. *Land Economics*, 84(1):115. [51](#)
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217. [20](#)
- Bulte, E., Gerking, S., List, J., and De Zeeuw, A. (2005). The effect of varying the causes of environmental problems on stated wtp values: evidence from a field study. *Journal of Environmental Economics and Management*, 49(2):330–342. [2](#)
- Cadsby, C., Maynes, E., and Trivedi, V. (2006). Tax compliance and obedience to authority at home and in the lab: A new experimental approach. *Experimental economics*, 9(4):343–359. [45](#)
- Carson, R. and Groves, T. (2007). Incentive and informational properties of preference questions. *Environmental and Resource Economics*, 37(1):181–210. [2](#)
- Carson, R., Groves, T., and List, J. (2006). Probabilistic influence and supplemental benefits: a field test of the two key assumptions behind using stated preferences. *Unpublished manuscript*. [2](#)
- Clark, J. (2002). House money effects in public good experiments. *Experimental Economics*, 5(3):223–231. [56](#)

- Dekel, E. (1986). An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom* 1. *Journal of Economic Theory*, 40(2):304–318. [3](#)
- Dhar, R. and Simonson, I. (2003). The effect of forced choice on choice. *Journal of Marketing Research*, pages 146–160. [50](#)
- Erard, B. (1990). *The influence of tax audits on reporting behavior*. [22](#)
- Farquharson, R. (1969). *Theory of voting*. Blackwell Oxford. [1](#)
- Feld, L. and Frey, B. (2002). Trust breeds trust: How taxpayers are treated. *Economics of Governance*, 3(2):87–99. [21](#)
- Gravell, J. (2008). Comments. *The IRS Research Bulletin: Proceedings of the 2008 IRS Research Conference*, pages 59–60. [45](#)
- Greiner, B. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics*, 10(23):63–104. [31](#), [62](#)
- Harrison, G. (2007). House money effects in public good experiments: Comment. *Experimental Economics*, 10(4):429–437. [56](#)
- Heath, T. and Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, pages 268–284. [49](#), [61](#), [68](#)
- Herne, K. (1996). The role of decoys in choice: a review of research on context dependent preferences. *Risk Decision and Policy # 160*;, 1(1):105–119. [49](#)
- Herne, K. (1997). Decoy alternatives in policy choices: Asymmetric domination and compromise effects. *European Journal of Political Economy*, 13(3):575–589. [48](#), [49](#)

- Herne, K. (1999). The effects of decoy gambles on individual choice. *Experimental Economics*, 2(1):31–40. [50](#), [51](#)
- Herriges, J., Kling, C., Liu, C., and Tobias, J. (2010). What are the consequences of consequentiality? *Journal of Environmental Economics and Management*, 59(1):67–81. [2](#), [19](#)
- Hey, J. and Strazzera, E. (1989). Estimation of indifference curves in the marschak-machina triangle a direct test of the fanning out hypothesis. *Journal of Behavioral Decision Making*, 2(4):239–260.
- Highhouse, S. (1996). Context-dependent selection: The effects of decoy and phantom job candidates* 1. *Organizational Behavior and Human Decision Processes*, 65(1):68–76. [48](#), [49](#)
- Huber, J., Payne, J., and Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, pages 90–98. [49](#), [51](#)
- Huber, J. and Puto, C. (1983). Market boundaries and product choice: Illustrating attraction and substitution effects. *Journal of Consumer Research*, pages 31–44. [49](#)
- Irwin, J., Slovic, P., Lichtenstein, S., and McClelland, G. (1993). Preference reversals and the measurement of environmental values. *Journal of Risk and Uncertainty*, 6(1):5–18. [61](#)
- Kastlunger, B., Kirchler, E., Mittone, L., and Pitters, J. (2009). Sequences of audits, tax compliance, and taxpaying strategies. *Journal of Economic Psychology*, 30(3):405–418. [22](#)
- Landry, C. and List, J. (2007). Using ex ante approaches to obtain credible signals for value in contingent markets: Evidence from the field. *American Journal of Agricultural Economics*, 89(2):420–429. [2](#)

- Lee, L., Loewenstein, G., Ariely, D., Hong, J., and Young, J. (2008). If i'm not hot, are you hot or not? physical-attractiveness evaluations and dating preferences as a function of one's own attractiveness. *Psychological Science*, 19(7):669–677. [48](#)
- Lehmann, D. and Pan, Y. (1994). Context effects, new brand entry, and consideration sets. *Journal of Marketing Research*, pages 364–374. [49](#), [51](#)
- List, J. and Gallet, C. (2001). What experimental protocol influence disparities between actual and hypothetical stated values? *Environmental and Resource Economics*, 20(3):241–254. [52](#)
- Machina, M. (1982). "Expected Utility" Analysis without the Independence Axiom. *Econometrica: Journal of the Econometric Society*, 50(2):277–323. [3](#), [13](#)
- Machina, M. (1987). Choice under uncertainty: Problems solved and unsolved. *The Journal of Economic Perspectives*, 1(1):121–154.
- Maciejovsky, B., Kirchler, E., and Schwarzenberger, H. (2007). Misperception of chance and loss repair: On the dynamics of tax compliance. *Journal of Economic Psychology*, 28(6):678–691. [22](#), [25](#)
- Mitami, Y. and Flores, N. (2011). Public goods referenda without perfectly correlated prices and quantities. *European Economic Association & Econometric Society, Oslo, Aug.*
- Mittone, L. (2006). Dynamic behaviour in tax evasion: An experimental approach. *Journal of Socio-Economics*, 35(5):813–835. [22](#), [25](#)
- Neilson, W. (1993). An expected utility-user's guide to nonexpected utility experiments. *Eastern Economic Journal*, 19(3):257–274.
- Nepal, M., Berrens, R., and Bohara, A. (2009). Assessing perceived consequentiality: Evidence from a contingent valuation survey on global climate change. *International Journal of Ecological Economics and Statistics*, 14(P09):14–29. [2](#)

- Pan, Y., O'Curry, S., and Pitts, R. (1995). The attraction effect and political choice in two elections. *Journal of Consumer Psychology*, 4(1):pp. 85–101. [48](#), [49](#)
- Plott, C. (1987). Dimensions of parallelism: Some policy applications of experimental methods. *Laboratory experimentation in economics: Six points of view*, pages 193–219. [26](#)
- Ratneshwar, S., Shocker, A., and Stewart, D. (1987). Toward understanding the attraction effect: The implications of product stimulus meaningfulness and familiarity. *Journal of Consumer Research*, pages 520–533. [49](#), [50](#), [51](#)
- Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research*, pages 158–174. [49](#), [50](#), [51](#)
- Simonson, I. and Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*. [49](#), [50](#)
- Smith, V. (1982). Microeconomic systems as an experimental science. *The American Economic Review*, 72(5):923–955. [26](#)
- Thaler, R. and Johnson, E. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management science*, pages 643–660. [56](#)
- Vossler, C., Doyon, M., and Rondeau, D. (forthcoming). Truth in consequentiality: Theory and field evidence on discrete choice experiments. *American Economic Journal: Microeconomics*. [2](#), [19](#)
- Vossler, C. and Evans, M. (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. [2](#)
- Vossler, C. and McKee, M. (2012). Efficient tax reporting: The effects of taxpayer information services. *Working Paper*. [21](#), [91](#)

- Vossler, C., McKee, M., and Jones, M. (2010). The impact of taxpayer information services on tax reporting and tax filing. *presented at the 2010 meetings of the Southern Economics Association, Atlanta, GA.* [38](#)
- Vossler, C. and Watson, S. (2012). The role of consequentiality in the external validation of stated preference methods through public referenda. *Working paper. Department of Economics, The University of Tennessee.* [3](#)
- Wedell, D. (1991). Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(4):767. [49](#), [50](#), [51](#)

Appendix

Appendix A

Appendix A

A.1 Proof of Proposition 1

Proof. (i): If the voter is indifferent between the status-quo and the perceived inconsequential outcome,

$$\delta_0 \sim \tilde{s}_i$$

then using the definition of betweenness, she is also indifferent between any mixture of the status-quo and the perceived inconsequential outcome,

$$\delta_0 \sim (\beta\delta_0 + (1 - \beta)\tilde{s}_i) \sim \tilde{s}_i \tag{A.1}$$

for all $\beta \in [0, 1]$.

Defining the maximum WTP for proposal $\tilde{x}_{ij}$ as the cost c_{ij}^* which solves

$$[\tilde{x}_{ij} - c_{ij}^*] \sim \delta_0,$$

assuming betweenness, she is also indifferent between the proposal $\tilde{x}_{ij}$ at cost c_{ij}^* and the perceived inconsequential outcome,

$$[\tilde{x}_{ij} - c_{ij}^*] \sim \tilde{s}_i.$$

Again, by the definition of betweenness, voter i is also indifferent between any mixture of the proposal $\tilde{x}_{ij}$ at cost c_{ij}^* and the perceived inconsequential outcome

$$[\tilde{x} - c_{ij}^*] \sim \beta[\tilde{x} - c_{ij}^*] + (1 - \beta)\tilde{s}_i \sim \tilde{s}_i \quad (\text{A.2})$$

for all $\beta \in [0, 1]$. Therefore, if we define the threshold acceptable cost b_i^* as the cost that solves:

$$(\beta[\tilde{x} - b_i^*] + (1 - \beta)\tilde{s}_i) \sim (\beta\delta_0 + (1 - \beta)\tilde{s}_i). \quad (\text{A.3})$$

Condition A.1 established that

$$(\beta\delta_0 + (1 - \beta)\tilde{s}_i) \sim \tilde{s}_i$$

for all $\beta \in [0, 1]$. Therefore, condition A.3 can be expressed as

$$(\beta[\tilde{x} - b_i^*] + (1 - \beta)\tilde{s}_i) \sim \tilde{s}_i$$

which, by condition A.2, implies

$$(\beta[\tilde{x} - b_i^*] + (1 - \beta)\tilde{s}_i) \sim [\tilde{x}_{ij} - c_{ij}^*] \quad (\text{A.4})$$

and therefore $b_i^* = c_{ij}^*$ for all $\beta \in [0, 1]$.

(ii) and (iii) from (i): If voter i is indifferent between the status-quo and the perceived inconsequential outcome $\tilde{s}_{i1}$, and the proposal $\tilde{x}_{ij}$ at cost c_{ij}^* and the perceived inconsequential outcome $\tilde{s}_{i1}$ such that:

$$\delta_0 \sim \tilde{s}_{i1} \sim [\tilde{x}_{ij} - c_{ij}^*],$$

it must be the case that

$$\beta\delta_0 + (1 - \beta)\tilde{s}_{i1} \sim \beta[\tilde{x}_{ij} - c_{ij}^*] + (1 - \beta)\tilde{s}_{i1}.$$

Assume now that there is a perceived inconsequential outcome $\tilde{s}_{i2}$ which satisfies

$$\tilde{s}_{i2} \succ \tilde{s}_{i1} \quad (\tilde{s}_{i2} \prec \tilde{s}_{i1}).$$

From Machina (1982, Hypothesis 5(iii)), we know that

$$\beta\delta_m + (1 - \beta)\tilde{s}_{i2} \sim \beta[\tilde{x}_{ij} - b_i^*] + (1 - \beta)\tilde{s}_{i2}$$

if and only if

$$\delta_m > \delta_0 \quad (\delta_m < \delta_0), \tag{A.5}$$

where δ_m is the certainty equivalent of $[\tilde{x}_{ij} - b_i^*]$ and δ_0 is the certainty equivalent of $[\tilde{x}_{ij} - c_{ij}^*]$. Condition A.5 therefore requires that the threshold acceptable cost b_i^* be lower (higher) than voter i 's maximum WTP:

$$b_i^* < c_{ij}^* \quad (b_i^* > c_{ij}^*)$$

□

A.2 Proof of Proposition 2

Proof. (i): Individual i 's maximum WTP for proposal $\tilde{x}_{ij}$ is defined as the cost c_{ij}^* which solves:

$$[\tilde{x}_{ij} - c_{ij}^*] \sim \delta_0.$$

If the voter is indifferent between the status-quo and the perceived inconsequential outcome such that

$$\delta_0 \sim \tilde{s}_i$$

then from condition (A.4), it was show that

$$(\beta[\tilde{x} - b_i^*] + (1 - \beta)\tilde{s}_i) \sim [\tilde{x}_{ij} - c_{ij}^*]$$

and therefore that $b_i^* = c_{ij}^*$ for all $\beta \in [0, 1]$.

(ii): Individual i 's maximum WTP for proposal $\tilde{x}_{ij}$ is defined as the cost c_{ij}^* which solves:

$$[\tilde{x}_{ij} - c_{ij}^*] \sim \delta_0.$$

Assume $\bar{x}$ and $\tilde{s}_i$ are such that

$$\bar{x} \succ \delta_0 \succ \tilde{s}_i. \tag{A.6}$$

By solvability, there must be some value of q which solves

$$q\bar{x} + (1 - q)\tilde{s}_i \sim \delta_0$$

which, by the definition of betweenness, implies

$$q\bar{x} + (1 - q)\tilde{s}_i \sim [\tilde{x}_{ij} - c_{ij}^*]. \tag{A.7}$$

Now define any decrease in consequentiality as an increase in $r \in (0, 1)$ such that more probability mass is shifted to the inconsequential outcome $\tilde{s}_i$. By construction, the status-quo δ_0 is preferred to the inconsequential outcome $\tilde{s}_i$, and therefore, following

Machina (1982, Hypothesis 5(v)),

$$qr\bar{x} + (1 - qr)\tilde{s}_i \succ r\delta_0 + (1 - r)\tilde{s}_i$$

which from condition (A.7) implies

$$(r[\tilde{x}_{ij} - c_{ij}^*] + (1 - r)\tilde{s}_i) \succ (r\delta_0 + (1 - r)\tilde{s}_i). \quad (\text{A.8})$$

The threshold acceptable cost is defined as the cost b_i^* which solves

$$(r[\tilde{x} - b_i^*] + (1 - r)\tilde{s}_i) \sim (r\delta_0 + (1 - r)\tilde{s}_i) \quad (\text{A.9})$$

which from condition (A.8) implies

$$(r[\tilde{x} - c_{ij}^*] + (1 - r)\tilde{s}_i) \succ (r[\tilde{x}_{ij} - b_i^*] + (1 - r)\tilde{s}_i). \quad (\text{A.10})$$

Truthful preference revelation requires that

$$(r[\tilde{x} - c_{ij}^*] + (1 - r)\tilde{s}_i) \sim (r[\tilde{x}_{ij} - b_i^*] + (1 - r)\tilde{s}_i) \quad (\text{A.11})$$

and from Proposition 3(iii), it must be the case that that $b_i^* > c_i$ in order to maintain condition (A.2).

(iii):The same as (ii), only working in the opposite direction. □

Appendix B

Appendix B

B.1 Theory of line-item reporting

The theory presented here is part of a more thorough presentation made by [Vossler and McKee \(2012\)](#), with adaptations made to fit the needs of this experimental setting. Our decision setting is characterized as follows. A risk-neutral taxpayer chooses whether to file, and if filing is her choice what to report on one or more “line items” (an entry which the taxpayer has discretion over what to report) on the tax form. We assume that the taxpayer considers directly the tax liability associated with her line item reports which allows us to generally characterize the optimal decision regardless of whether the line item is associated with a credit, deduction, reported income, or otherwise.

Audits occur with probability p and are completely random and independent of whether other persons are audited or the reported tax liability. Audits on tax returns perfectly reveal unpaid taxes separately for each line item on the tax form. In addition to being liable for unpaid taxes upon audit, there is a constant per-unit penalty $\beta > 0$ assessed on unpaid taxes.

The actual tax liability on one or more line items is uncertain, and there may be an information service available to partially or fully resolve the uncertainty. Let x_l^0 denote the actual tax liability associated with line item l .

From the perspective of the taxpayer, tax liability is a random variable x_l with distribution function $F(x_l)$, which is assumed to have positive density $f(x_l)$ on the interval $[a_l, b_l]$.

for each line item on the tax form the taxpayer chooses a tax liability to report, denoted R_l . The optimal reporting problem is then one of choosing a tax liability R_l in order to minimize expected costs:

$$\min_R \sum_l \left\{ R_l + p \left\{ (\beta + 1) \int_{R_l}^{b_l} (x_l - R_l) f(x_l) dx_l \right\} \right\} \quad (\text{B.1})$$

The optimal reporting choice for a particular line item, R_l^* , is implicitly defined by

$$1 = p(\beta + 1) \int_{R_l^*}^{b_l} f(x_l) dx_l \quad (\text{B.2})$$

for every l . The interpretation is that the taxpayer minimizes cost by equating the marginal cost of taxes reported with the expected marginal cost of the audit. The first-order necessary conditions can instead be written as

$$F(R_l^*) = 1 - \frac{1}{p(\beta + 1)}. \quad (\text{B.3})$$

An interior solution exists for R_l^* the interval $[a_l, b_l]$ if $\frac{1}{p(\beta+1)} < 1$. Otherwise, there is a corner solution $R_l^* = a_l$, i.e. the taxpayer engages in maximum tax evasion.

It is possible in general for the optimal reported liability to be under, over or equal to the true liability. For instance, even if $E[x_l] = x_l^0$ (i.e. beliefs are unbiased) there is the potential value to over-report in expectation as it decreases the probability (and expected cost) of being found to have underreported.

When liability is certain, it is not possible to have over-reporting as optimal, as paying too much tax provides no benefit regardless of whether an audit occurs. Instead, under certainty, the solution is to fully comply when $\frac{1}{p(\beta+1)} < 1$, and to engage in maximum evasion when $\frac{1}{p(\beta+1)} > 1$. Thus, uncertainty in the former case - if anything - leads the taxpayer *away* from the truth. In the latter case, uncertainty has no effect as the taxpayer will be at the corner solution of maximum evasion regardless.

B.2 Optimal evasion given tax policy parameters

Focusing first on the optimal reporting decision, with the experiment parameters ($p = 0.3; \beta = 3$) we have that $\frac{1}{p(\beta+1)} < 1$. Thus, when liability is uncertain, the optimal reporting decision is defined by equation (B.2). With certainty or receipt of the perfect information service, it is optimal to report the truth. For both the deduction and credit decision, based on the experiment parameters it is optimal to under-report the true tax liability (i.e. over-claim credits). To see this, with uncertainty and the uniform distribution employed, based on equation (B.2), the solution to the cost minimization problem is

$$R_l^* = \frac{a_l - b_l}{p(\beta + 1)} + b_l, \quad (\text{B.4})$$

which was used to construct table B.1.

Table B.1: Evasion calculation by income

Income Level (a_l)	Lower-bound on liability (a_l)	Optimal Credit R_l^*	Expected Credit	Optimal under-reporting
\$1000	\$1000	\$833	\$500	\$333
\$1250	\$750	\$625	\$375	\$250
\$1500	\$500	\$417	\$250	\$167

Appendix C

Appendix C

C.1 Selected Experiment Screenshots




Figure C.1: Income earnings task

Subject 3 **Tax Form** **Round 1**

Group Earnings Summary

\$LAB

RED bar is your earnings

Your Earnings Summary

	Lower Actual	Upper Actual
Income Earned	1,500	1,500
Initial Taxes	750	750
Credits (actual is amount allowed)	0	500
Information Request		
Final Taxes	750	250

? indicates the actual value is a whole number within the range shown.

Original Information
Information Request

Department Of Treasury **Individual Income Tax Return**

Income

1 Income earned 1500

2 Initial taxes 750

Tax Credits

3 Credit 0

7 8 9
4 5 6
1 2 3
0 Clear

Taxes

4 Final taxes 750

Tax Policy

Tax Rate (tax = rate x taxable income) 50%

Enforcement Policy

Penalty For Late Filing 10%

Penalty Rate (penalty = rate x unpaid taxes) 300%

Audit Chance

The likelihood of an audit 50%

Time Remaining (Seconds) = 50

File

Figure C.2: Treatment 3, Tax decision screen, information requested

Subject 3 **Tax Form** **Round 1**

Group Earnings Summary

\$LAB

RED bar is your earnings

Your Earnings Summary

	Lower Actual	Upper Actual
Income Earned	1,500	1,500
Initial Taxes	750	750
Credits (actual is amount allowed)	0	500
Information Request	201	201
Final Taxes	549	549

? indicates the actual value is a whole number within the range shown.

Original Information
Information Request

Department Of Treasury **Individual Income Tax Return**

Income

1 Income earned 1500

2 Initial taxes 750

Tax Credits

3 Credit 0

7 8 9
4 5 6
1 2 3
0 Clear

Taxes

4 Final taxes 750

Tax Policy

Tax Rate (tax = rate x taxable income) 50%

Enforcement Policy

Penalty For Late Filing 10%

Penalty Rate (penalty = rate x unpaid taxes) 300%

Audit Chance

The likelihood of an audit 30%

Time Remaining (Seconds) = 55

File

Figure C.3: Treatment 3, Tax decision screen, after information acquired




Figure C.4: Audit selection process

Subject 3

Round Summary

Round 1

Audit Results: AUDITED

Cumulative Wealth \$901

Your Earnings Summary	Actual	Filed	Audited	Final
Income Earned	1,500	1,500	1,500	1,500
Initial Taxes	750	750	750	750
Credits (actual is amount allowed)	201	201	201	201
Final Taxes	549	549	549	549
Earnings After Taxes	951	\$951	\$951	\$951
Penalty (under reporting taxes owed)			0	0
Penalty for late filing			0	0
Cost of Information Service One			50	50
Earnings =		\$901		

Tax Policy

Tax Rate

50%

(tax = rate x taxable income)

Audit Chance

The likelihood of an audit

30%

Enforcement Policy

Penalty For Late Filing

10%

Penalty Rate

300%

(penalty = rate x unpaid taxes)

Time Remaining (Seconds) =

28

Continue

Figure C.5: Results screen

Appendix D

Appendix D

D.1 Experiment Instructions

Experiment Overview

- You will be participating in a market simulation that lasts several decision “rounds”. In each round, you first play an earnings game and then face a tax reporting decision.
- In the earnings game, you sort the numbers 1 through 9. Your Income earned is determined by how fast you sort the numbers relative to others. The participant in your group with the fastest time receives the highest Income earned.
- In the tax reporting stage, you fill out and file a tax form. How much you earn from the tax reporting decision depends on how much you claim in Tax Credit and whether or not you are audited. Note that the on-screen instructions do not specify the tax policy parameters (e.g. tax rate, penalty rate, etc.), but those specified below will be in effect for this experiment.
- Each round is completely independent from the others, which means your decisions in one round in no way affect the outcome of any other round.

How your earnings are determined each round

- On the tax form, your Initial Taxes will be calculated automatically. This amount is determined by multiplying your Income earned by a tax rate of 50%.
- You decide how much to claim in Tax Credit on the tax form. Each dollar you claim in credits reduces your Final taxes by one dollar. This amount is subtracted from the Initial Taxes to determine your Final Taxes. If Final Taxes is a negative number, this reflects a tax refund.
- You will be shown a range of tax credits (this range is highlighted in white on the left side of the decision screen), which depends on your Income earned. Each amount within the range has an equal chance of being your actual tax credit, which is the highest amount you can claim without possible penalty. You can choose to claim any amount between 0 and 1000.
- You have an information service available to you at a cost of \$50. By clicking on the “Request Information” button you will know the *exact* amount of your actual tax credit.
- You have a 30% chance of being audited. Audits are determined completely at random and do not depend on how much you or anyone else claims in tax credits.
- If you are not audited, your earnings for the round are your Income earned minus Final taxes.
- If you are audited, but claimed *less* than or *equal* to the actual tax credit, your earnings for the round are your Income earned minus Final taxes. Know that if you under-reported the credit you will not receive additional money through the audit process.
- If you are audited, and claimed *more* than the actual tax credit, you pay back the extra tax credit you claimed and also pay a penalty.
 - The penalty is equal to 300% multiplied by the amount of *extra* tax credit you claimed. Thus, if you claimed an extra \$100 your penalty is $\$100 \times 300\%$ or \$300.
 - Your earnings for the round are then Income earned *minus* Final taxes *minus* the extra tax credit you claimed *minus* the penalty.

Appendix E

Appendix E

E.1 Selected Experiment Screenshots





Figure E.1: Experimental Instructions

Scenario 1

Please choose one of the following options:



Purchase a 10 pack of Energizer
AA batteries at a price of \$5.04

Choose



Purchase a 10 pack of Duracell
AA batteries at a price of \$3.58

Choose



Purchase a 2 pack of Duracell
AA batteries at a price of \$4.21

Choose

I do not want to buy any of the products.

Choose


Figure E.2: Example product purchase scenario

Results

The computer randomly chose to make your decision from Scenario 4 binding.

You will receive the product shown below in addition to your earnings from the first experiment.

The cost of this item, **\$5.72**, will be deducted from your earnings.



Earnings from first experiment:		\$10.00
Cost of item:	-	\$5.72
Cash:		\$4.28

Please proceed to finish the questionnaire

OK

Figure E.3: Example results screen

Appendix F

Appendix F

F.1 Example results from the conditional logit estimation

Table F1: Conditional logit estimates for Treatment 4- Batteries

Variable	Coeff	Std.Err.	z	P>z
Price	-0.55	0.04	-12.29	0.00
Duracell	2.61	1.11	2.35	0.02
Energizer	3.44	0.96	3.59	0.00
Duracell:CTD#:2-pack decoy	-2.69	0.52	-5.21	0.00
Duracell:CTD#:6-pack decoy	-2.70	0.47	-5.75	0.00
Duracell:2-pack decoy	-1.54	2.96	-0.52	0.60
Duracell:6-pack decoy	2.61	3.15	0.83	0.41
Energizer:2-pack decoy	-2.45	2.93	-0.83	0.40
Energizer:6-pack decoy	1.34	3.06	0.44	0.66
Duracell:Decoy Price Difference:2-pack decoy	-0.34	0.22	-1.56	0.12
Duracell:Decoy Price Difference:6-pack decoy	-0.31	0.22	-1.44	0.15
Energizer:Decoy Price Difference:2-pack decoy	0.06	0.23	0.28	0.78
Energizer:Decoy Price Difference:6-pack decoy	-0.08	0.33	-0.24	0.81
Duracell:Earnings	-0.02	0.07	-0.25	0.80
Duracell:Earnings:2-pack decoy	0.10	0.20	0.49	0.63
Duracell:Earnings:6-pack decoy	-0.10	0.23	-0.44	0.66
Energizer:Earnings	-0.10	0.06	-1.57	0.12
Energizer:Earnings:2-pack decoy	0.16	0.19	0.85	0.40
Energizer:Earnings:6-pack decoy	-0.10	0.21	-0.46	0.65
Duracell:Session Order	1.86	0.39	4.78	0.00
Duracell:Session Order:2-pack decoy	-0.65	0.73	-0.88	0.38
Duracell:Session Order:6-pack decoy	-1.12	0.74	-1.51	0.13
Energizer:Session Order	2.09	0.37	5.68	0.00
Energizer:Session Order:2-pack decoy	-1.16	0.66	-1.77	0.08
Energizer:Session Order:6-pack decoy	-1.26	0.77	-1.64	0.10
Duracell:Scenario Order	-0.08	0.02	-3.55	0.00
Duracell:Scenario Order:2-pack decoy	0.08	0.05	1.73	0.08
Duracell:Scenario Order:6-pack decoy	0.06	0.05	1.13	0.26
Energizer:Scenario Order	-0.04	0.02	-1.73	0.08
Energizer:Scenario Order:2-pack decoy	-0.01	0.04	-0.23	0.82
Energizer:Scenario Order:6-pack decoy	0.03	0.04	0.58	0.56
Duracell:Picture Position	0.23	0.20	1.14	0.26
Duracell:Picture Position:2-pack decoy	0.22	0.38	0.58	0.56
Duracell:Picture Position:6-pack decoy	-0.15	0.34	-0.46	0.65
Energizer:Picture Position	0.07	0.21	0.32	0.75
Energizer:Picture Position:2-pack decoy	0.27	0.33	0.84	0.40
Energizer:Picture Position:6-pack decoy	0.04	0.38	0.09	0.93

Vita

Caleb Andrew Siladke was born on October 6, 1985 in Traverse City, Michigan. He graduated from Traverse City West Senior High School in 2004 and proceeded to Northern Michigan University, where he received his B.S. with a major in economics and a minor in mathematics. Caleb then entered the Graduate School at the University of Tennessee, where in August of 2010 he obtained his M.A. in economics. Caleb then continued at the University of Tennessee in pursuit of his Ph.D. in economics.