

Essays on Resource Allocation Efficiency and Behavior

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Julianna Marie Butler
August 2014

Copyright © 2014 by Julianna Marie Butler
All rights reserved.

DEDICATION

This dissertation is dedicated to my family, Mom, Dad, Chris and Ben, for their unending support and love, and to my dear friends, Nick, Xiaowen, Michael, Ahiteme, Jens, Jill, Steve, Becky, Emily, Jing and Beth for their fellowship, encouragement and inspiration during graduate school.

ACKNOWLEDGEMENTS

I would like to express my appreciation and gratitude to Dr. Christian Vossler for his advice and guidance both in writing this dissertation and on the job market; without his support, this project would not have been possible. Further, I would like to thank my other committee members, Dr. Michael Price, Dr. Scott Gilpatric and Dr. Christopher Clark for their help during the research process and for everything they taught me during graduate school. Special thanks to Dr. Jonathan Alevy, Dr. William Neilson and Dr. Jens Schubert for their help with the first chapter. Further, I would like to express appreciation to participants at the ESA North American Conference in November 2012 and the faculty in the Department of Economics at the University of Tennessee, Knoxville for helpful comments. Finally, I would also like to acknowledge funding from the Bureau of Land Management and the Department of Economics Research Incentive Fund at the University of Tennessee, Knoxville.

ABSTRACT

This dissertation is comprised of three papers in the field of microeconomics. The first examines bidder's choice auctions using both field and laboratory experiments. The field experiments demonstrate that traditional bidder's choice auction theory does not always hold; the laboratory experiments subsequently isolate several characteristics of this auction format to explain why. We find that while price revelation does not impact the revenue superiority of the auction mechanism, multi-good demand significantly reduces the revenue premium. Intuitively, risk aversion plays less of a role when bidders have the opportunity to win multiple goods. The second chapter is theoretical and presents a dynamic Markov labor market tournament in which the manager does not have the ability to incentivize agents using money. Instead, the manager can use task assignment to reward and punish agents who are in and out of favor with him. This situation frequently characterizes public organizations such as schools and government agencies. The prize of the tournament is the difference between groups in the present value of the agent's expected utility. We show that when the manager must delegate a certain number of tasks and when agents' cost of contractible effort is a convex function, the manager can incentivize optimal non-contractible effort by agents. However, the total cost to the manager is higher than if the manager was able to use monetary incentives. The third chapter is an experimental paper that elicits consumer willingness to pay for food products labelled "natural". The "natural" label is not regulated in the United States; however, several manufacturers are currently under lawsuit for selling "natural"-labelled food that contains genetically modified ingredients. This study uses an incentive-compatible mechanism and a survey to connect consumers' beliefs to the premium that they associate with the "natural" label. Primarily, we find that consumers who believe "natural" means "no genetically modified organisms" (42% of our sample) are willing to pay a premium for "natural" food, whereas consumers who do not have this belief actually exhibit a negative premium. The overall effect is near zero, although the identified heterogeneity suggests that "natural" labels are potentially misleading.

TABLE OF CONTENTS

INTRODUCTION	1
CHAPTER 1: Multi-Good Demand in Bidder’s Choice Auctions	4
References	39
Appendix A: Theoretical Work	40
Appendix B: Tables	41
Appendix C: Figures	48
Appendix D: Instructions	51
CHAPTER 2: A Dynamic Markov Tournament Model of Task Assignment and Up-and-Down Competition for Status	58
References	82
Appendix A: Theoretical Work	84
Appendix B: Tables	87
CHAPTER 3: Willingness to Pay for Goods with Unregulated and Potentially Misleading Labels: the Case of “Natural”-Labelled Groceries	88
References	110
Appendix A: Tables	112
Appendix B: Figures	117
Appendix C: Instructions	119
Appendix D: Survey	134
CONCLUSION	142
VITA	143

LIST OF TABLES

Table 1.1: Average Revenues	41
Table 1.2: Average Revenues (Using Only Markets with Variation in Most Preferred Good)	41
Table 1.3: Average Bids in GBG Auctions and 3 rd Phase of BC Auctions	41
Table 1.4: Average Revenues in GBG Auctions and 3 rd Phase of BC Auctions	41
Table 1.5: Average Bids for Most Preferred Good in GBG and Phase 1 of BC	42
Table 1.6: Change in Bids for BC Auctions when Most Preferred Still Available in Next Phase	42
Table 1.7: Change in Bids for BC when Most Preferred Still Available by Rank in Initial Period	42
Table 1.8: Effect of Demographics and Personality Measures on Bids	43
Table 1.9: Experimental Design (Laboratory)	44
Table 1.10: Average Revenues	44
Table 1.11: Average Bids by Treatment	44
Table 1.12: BC Auctions: Average Values and Bids	45
Table 1.13: Good-by-Good (Benchmark) Auctions: Average Values and Bids	46
Table 1.14: Estimates of Bidding Behavior	47
Table 2.1: Transition Probabilities	87
Table 3.1: Session Schedule	112
Table 3.2: Variable Definitions and Means	113
Table 3.3: Between-Subject Analysis Using Part 1 Bids	114
Table 3.4: Treatment Effects Using Part 1 Data	115
Table 3.5: Treatment Effects Using Part 2 Data	116

LIST OF FIGURES

Figure 1.1: Proportions of Participants in Each Risk Preference Group	48
Figure 1.2: Average Rankings for Each Good in Each Auction Format	48
Figure 1.3: Proportions of Participants Who Ranked Each Good as Their Most Preferred	49
Figure 1.4: Average Revenues	49
Figure 1.5: Risk Elicitation Results	50
Figure 1.6: Average Revenues Including Theoretical GBG Revenues	50
Figure 3.1: Potato Chip Pictures and Labels	117
Figure 3.2: Toothpaste Pictures and Labels	118

INTRODUCTION

This dissertation is comprised of three papers in the field of microeconomics. The auction mechanism examined in the first chapter has been used in a variety of settings including the sale of real estate, antiques and customized phone numbers. The paper focuses on the design of the mechanism: when it theoretically raises higher revenue than other mechanisms and under what conditions the theory predicts behavior in reality. The second chapter also focuses on mechanism design. A dynamic tournament is applied to a labor market setting that is commonly found in public organizations. The theoretical results describe conditions under which a manager can achieve optimal effort from his employees. Finally, the third chapter answers a pressing empirical question regarding consumer willingness to pay for a “natural” label on food. Several pending lawsuits partially motivate the paper. Although the three chapters are not directly related and each represents a separate, stand-alone paper, the tools used throughout this dissertation are similar. For instance, laboratory experiments are used in both the first and third chapters, and microeconomic theory is used in both the first and second chapters. This introduction will provide an overview of each chapter including its motivation, the methods used, and the primary result(s).

Bidder’s choice (or “right-to-choose”) auctions are of particular interest to parties who wish to sell multiple similar goods. Economic theory has shown that this type of auction, where the high bidder wins the right to choose one good from among the available goods, results in higher revenue under risk aversion than traditional good-by-good auctions. Most theoretical and experimental work focuses on bidder’s choice auctions where bidders have value for only one of the available goods. Chapter 1 presents a field experiment and a lab experiment that allow for price revelation and multi-good demand, which are typically found in bidder’s choice auctions used for the sale of condos, antiques, customized telephone numbers and other groups of similar goods. We find that while price revelation does not have a significant effect on revenue, multi-good demand mutes the theoretical revenue superiority the bidder’s choice mechanism. This is consistent with the notion that the perceived risk of losing one’s most preferred good is softened when there is a chance to win multiple goods. This result implies that bidder’s choice auctions

are preferred in settings where each bidder is likely to strongly prefer one good over the others, though this need not be the same good for every bidder. Further, this work demonstrates the complementarities of the field and laboratory settings to answer questions which are not clearly resolved using only one setting.

In the second chapter, my coauthors and I develop a dynamic Markov model to capture the incentives in indefinitely-repeated tournaments in labor market settings where agents compete both to “move up” as well as to avoid a “move down”. Such settings naturally arise regardless of whether explicit performance incentives or an organizational hierarchy exist. We show that when monetary incentives are available, the dynamic tournament approaches the first-best outcome, but we also allow for the possibility that the principal’s only available incentive mechanism is the assignment of undesirable tasks to agents who are out-of-favor. Inability to change salaries or demote workers is common for public organizations, such as government agencies and schools. For instance, a school principal may not be able to monetarily reward or sanction teachers based on performance, but typically has discretion within the labor contract to vary class assignments and resources such as teacher’s aides. We model agents as being either in or out of favor with the principal in any given period; those who are out of favor are assigned more undesirable tasks. The prize of the tournament is the difference between groups (in favor and out of favor) in the present value of the agent’s expected utility. We assume that agents’ effort cost of completing contractible tasks is such that these costs are minimized by assigning equally burdensome tasks to all agents. Therefore the principal can motivate non-contractible effort through differential task assignment, but this entails an efficiency cost. The model demonstrates that employers may seek flexibility to vary task assignments in labor contracts not only to adapt to changing circumstances, but also to enable them to motivate non-contractible effort when agents’ compensation is fixed.

Food labeling has been widely studied, especially in the context of consumer willingness to pay for features that are considered healthy, such as organic content. Additionally, labels are highly regulated by the government; for instance, the phrase “low fat” cannot be used for foods with more than 3 grams of fat per serving. Especially for labels indicating low environmental impact, most of the theoretical literature acknowledges that there is some level of fraud in the

market for regulated labels, but the effects of an unregulated phrase on consumer demand are unclear empirically. In Chapter 3, an incentive-compatible approach is used to elicit willingness to pay (WTP) for grocery items with and without “natural” labels, several of which have genetically modified ingredients. Several pending lawsuits regarding genetically modified ingredients in food labeled as “natural” partially motivate our paper. Primarily, we find that consumers who believe “natural” means “no genetically modified organisms” (42% of our sample) are willing to pay a premium for “natural” food, whereas consumers who do not have this belief actually exhibit a negative premium. The overall effect is near zero, although the identified heterogeneity suggests that “natural” labels are potentially misleading and further that there is potential for firms to exploit uninformed consumers.

CHAPTER 1

Multi-Good Demand in Bidder's Choice Auctions

Abstract

Bidder's choice (or "right-to-choose") auctions are of particular interest to parties who wish to sell multiple similar goods. Economic theory has shown that, under risk aversion, this type of auction, where the high bidder wins the right to choose one good from among the available goods, results in higher revenue than traditional good-by-good auctions. Most theoretical and experimental work focuses on bidder's choice auctions where bidders have value for only one of the available goods. This paper presents a field experiment and a lab experiment that allow for price revelation and multi-good demand, which are typically found in bidder's choice auctions used for the sale of condos, antiques, customized telephone numbers and other groups of similar goods. We find that while revealing winning prices does not have a significant effect on revenue, multi-good demand mutes the theoretical revenue superiority the bidder's choice mechanism. This is consistent with the notion that the perceived risk of losing one's most preferred good is softened when there is a chance to win multiple goods. This result implies that bidder's choice auctions are preferred in settings where each bidder is likely to strongly prefer one good over the others, though this need not be the same good for every bidder. Further, this work demonstrates the complementarities of the field and laboratory to answer questions which are not clearly resolved using only one setting.

I. Introduction

In a study of condominium sales, Ashenfelter and Genesove (1992) reported results that they argued “should surprise most economists”: they found that the prices of condominium units in New Jersey varied significantly depending on the way they were sold. Units sold at auction were valued more highly than those sold through bilateral negotiation. The specific auction institution in use was a “bidder’s choice” (also known as a “right-to-choose”) auction in which the winner, rather than receiving a specific condominium, earned the right to choose their preferred unit from among those remaining¹. Bidder’s choice auctions are also commonly used in the sale of customized telephone numbers, antiques, bank branches following mergers, and other sequential sales of multiple similar goods. Theorists and experimentalists have explored the issue raised by Ashenfelter and Genesove (1992) considering more generally the allocation of multiple heterogeneous goods to a pool of bidders. The research has typically studied sequential or good-by-good auctions rather than bilateral bargaining in order to create a clean counterfactual to understand the performance of the bidder’s choice institution (Burguet 2007; Goeree, Plott and Wooders 2004; Eliaz, Offerman and Schotter 2008).

Importantly, Burguet (2007) demonstrates that in theory, the bidder’s choice auction raises higher revenue than a simple sequential auction when bidders are risk averse. The mechanism can “thicken markets” by creating competition across goods that are evaluated independently of each other in the sequential setting². Existing studies generally provide support for Burguet’s theory, though the role of risk aversion is somewhat unclear (Goeree, Plott and Wooders 2004; Eliaz, Offerman and Schotter 2008). Further, these studies only allow each

¹ According to the National Association of Realtors, a “bidder’s choice” auction is: “a method of sale whereby the successful high bidder wins the right to choose a property from a grouping of similar or like-kind properties. After the high bidder’s selection, the property is deleted from the group, and the second round of bidding commences, with the high bidder in round two choosing a property, which is then deleted from the group...” This process continues until all goods have been sold.

² Intuitively, the possibility that one’s preferred good will be chosen early makes the value of the later auctions less certain. Risk-averse buyers therefore are willing to pay a premium to secure their favored good in an early round.

bidder to have positive value for one good (single-good demand)³. Our study makes novel contributions by studying the bidder's choice auction in field settings, where Ashenfelter and Genesove's questions arose, while retaining the controls introduced by laboratory studies. Specifically, we address several previously unexplored characteristics of this auction format, including multiple values per bidder, as described below.

First, we compare bidder's choice and sequential auctions with a variety of consumer goods using a diverse pool of subjects in Reno, Nevada. The research design is that of a "framed field experiment" that incorporates relevant elements of traditional field studies in economics, including real goods and a diverse population of consumers, with appropriate experimental controls (Harrison and List 2004). The field setting differs from the laboratory environment in several ways, most critically in that an individual's bids are based on "homegrown" values for actual goods rather than induced values. We find that the bidder's choice mechanism fails in this environment, in contrast with theory and previous experimental work. Consequently, the second part of our research uses a lab experiment to focus on two elements of bidder's choice auctions (multi-good demand and price revelation) that are not yet fully understood. We find that while price revelation does not have any significant effect, multi-good demand significantly mutes of revenue superiority of the bidder's choice mechanism, which explains the results of the field experiment relative to other studies.

Multi-good demand in the bidder's choice setting has many applications. For instance, winning condominium bidders frequently choose a condo that they plan to rent out and continue bidding for the remaining condos, demonstrating that they have values for multiple units. Harstad (2009) presents anecdotal evidence of winning bidders in art auctions choosing an artwork and remaining to bid for further rights to choose. He also discovered that mergers or acquisitions in the banking industry, where branches are sold via bidder's choice auctions, result

³ For example, Eliaz et al. (2008) perform an experimental test of the theory where the high bidder drops out after he has won in the first phase of the auction, and the remaining bidders place bids for the right to choose from among the remaining goods. However, there are many instances in the field where it is the norm for winning bidders to remain in the auction.

in more branches sold than the number of purchasers. This indicates that some purchasers must attempt to purchase several branches in different locations, implying multi-good demand⁴.

The U.S. Bureau of Land Management considered using bidder's choice auctions to sell wild horses in 2009, which originally inspired our field experiment. In this experiment, bidders presumably had positive values for all three of the goods (an iPod package, a hiking equipment package, and a wine package), bidders continued to bid in all phases of the auction, and most bidders were risk averse. In addition, the winning price in each phase was revealed, similar to sales in real estate or antiques via bidder's choice auction. The results indicate that the theoretical revenue superiority of the bidder's choice mechanism under risk aversion may be overstated. However, since personal values for the objects were private, it is impossible to know for sure how the bidders updated their beliefs regarding values in each round. Further, it is possible that the effects of risk aversion may not be as pronounced when bidders have a chance of obtaining surplus from multiple goods than when they only have positive value for one good.

Therefore, the next stage of our research, the lab experiment, aims to bridge the gap between theory and the field. We are able to isolate the effects of multi-good demand and price revelation. Specifically, in one set of treatments, the participants are notified of the winning bid after each phase, while in other treatments, they are only told who won and which object was selected. Further, the level of competition is also varied: in one set of treatments, all participants bid in every round regardless of whether they have already won, and in another set of treatments, only two participants have a positive value for each good.

As previously stated, behavior in bidder's choice auctions is of particular interest to parties who need to sell several similar objects, and could also be applied in natural resource settings where rights to land, greenhouse gas emissions, etc., need to be allocated efficiently. Developing an understanding of how bidders behave under different conditions is important in designing auctions settings, and this research aims to aid this development. We find that price revelation does not have an effect on revenues, but the revenue premium is significantly higher

⁴ Additionally, it is easy to imagine that a firm may have positive value for more than one phone number or web address. For instance, Comcast owns both 1-800-COMCAST and 1-800-XFINITY. Google owns various misspellings of "google.com" to ensure users can reach their search engine, even if they are typing too quickly.

under single-good demand than multi-good demand⁵. This is consistent with the notion that the perceived risk of losing one's most preferred good is muted when there is a chance to win multiple goods. Further, this result implies that bidder's choice auctions should be used in settings where each bidder is likely to strongly prefer one of the goods over the others, though this need not be the same good for every bidder.

The paper proceeds as follows. Section II provides a detailed review of the literature and emphasizes the contribution of this research. Section III outlines the theoretical background. Sections IV and V detail the field experiment and the laboratory experiment, respectively, including design choices and results. Finally, a discussion is provided in Section VI. References follow.

II. Literature Review

2.1 Theoretical Research

Ashenfelter and Genesove (1992) inspired much of the work on pooled auctions and bidder's choice auctions when they observed declining prices in successive rounds of condominium sales⁶. The authors hypothesized that bidders who were aware that waiting may lead to a lower price were risk averse; the bidders did not want to shade their bids as theory would predict because they were afraid of losing the condo. Shortly after, Gale and Hausch

⁵ The revenue premium is calculated as the revenue from a bidder's choice auction minus the revenue from a corresponding standard good-by-good auction.

⁶ While this paper focuses on bidder's choice auctions, research on similar pooled auctions are worthy of mentioning. Menezes and Monteiro (1998) show that a simultaneous pooled auction also yields the same revenue as a standard sequential auction, though their model uses a first-price rather than a second-price standard auction (in contrast to the related literature). Salmon and Iachini (2007) provide an experimental analysis of pooled auctions and find that pooled auctions yield substantially higher revenues than ascending auctions. They find that this increase in revenue is not due risk or loss aversion, and they instead provide an attentional bias hypothesis where bidders overweight the surplus from winning their most preferred good as opposed to a lesser preferred good. These pooled auctions, while similar to bidder's choice auctions, are different in that bidders are not aware of the "remaining" goods when they place their bids; the second-highest bidder wins the second right-to-choose, but he is not aware of this outcome when he places his bid.

(1994) compared the theoretical revenue of a bidder's choice auction to that of a standard sequential auction and found that the former is larger. However, their model is limited to a two-bidder, two-good case.

Burguet (2005 and 2007) formalized the theory of behavior for the bidder's choice auction most frequently used today as a basis for lab and field experiments. He uses a simple illustration with two bidders and then extends the model to include any number of additional bidders⁷. The results clearly show that the bidder's choice auction should raise more revenue than the standard sequential good-by-good auction when bidders are risk averse. Burguet also shows that bidder's choice auctions are efficient and concealing information as to which goods have been selected allows the seller to achieve higher revenue.

Interestingly, Burguet touches briefly on the topic of taste diversity (less than perfect substitutability among the goods). This is analogous to what is later referred to as non-persistent competition (single-good demand) in the literature; the extreme case is having positive value for only one good and zero values for the other goods. He shows that greater taste diversity may increase seller revenues even if it reduces buyers' willingness to pay for some objects. While this is not one of Burguet's main conclusions, our research confirms this result. To our knowledge, this dimension of taste diversity has not previously been tested in a lab or field setting and this constitutes one of the main contributions of our paper. The variance of values over goods for each bidder has been overlooked in much of the literature, yet this dimension is an important gap between the theory and the field, as described in the introduction.

Harstad (2010) points out that persistent competition (multi-good demand) is the norm in many instances of bidder's choice auctions. He builds a theoretical model where bidders have positive values for multiple goods and demonstrates that, under risk neutrality, the distribution of equilibrium revenue from a standard good-by-good auction is a mean-preserving spread of the distribution of revenue from a bidder's choice auction. His discussion comments on the risk averse seller's preference for a bidder's choice auction, but does not comment on the risk preferences of the bidders in the base model. Harstad extends the model to include cases where

⁷ See Section III for more detail.

bidders may believe that one good is valued over the other goods by the majority; he refers to this good as “the usual favorite”. This modification may lead a winning bidder to choose a good other than his most preferred in order to reduce competition in future rounds. While we do not model this specifically (values in our design are drawn randomly), a few of our bidders do choose goods that are not their most preferred. This may indicate some gravitation toward a “usual favorite” belief, or may simply be a mistake by these bidders.

2.2 Experimental Research in the Laboratory

Experimental work on bidder’s choice auctions is a little sparse, though several recent papers have reported successful results. Goeree et al. (2004) find that bidder’s choice auctions raise more revenue than standard sequential ascending auctions under risk aversion. The authors are able to compare observed bids with theoretically predicted bids to estimate a common risk aversion parameter: on average, their bidders have the utility function: $u(x) = x^{0.39}$. Goeree et al. recognize the value of testing multi-good demand, though they phrase the idea in a slightly different way in their concluding remarks: “One extension is to consider bidders who value more than one item. It is an open question whether the revenue superiority of the ascending bidder’s choice auction extends to richer valuation structures where the simultaneous ascending auction has proven to perform well.” This paper aims to fill this gap in the existing literature.

Our experimental design closely follows that of Eliaz et al. (2008), who demonstrate that bidder’s choice auctions raise higher revenue than the theoretically optimal auction and show how withholding some information or restricting quantity can benefit sellers. However, they argue that risk aversion may not be the only factor contributing to aggressive bidding in bidder’s choice auctions. They incorporate a “no information” treatment where bidders do not know which good has been selected as they bid in each phase (similar to a pooled auction)⁸. In this treatment, risk averse bidders are expected to bid below the risk-neutral bid (instead of above the risk-neutral bid as predicted in a regular bidder’s choice treatment) due to the fact that they

⁸ This “no information” treatment is not to be confused with the “no information” treatment in our experiment; Eliaz et al. refer to no information on goods selected, whereas we refer to no information on prices.

essentially face a lottery over which good they will win. The authors find that bidders in the “no information” treatment nonetheless raise their bids, implying that a different behavioral phenomenon may be at work. By calculating the equilibrium bids for many different hypothetical numbers of participants, the authors are able to show that bidders behave as if they are competing with many subjects. The intuition is that bidders perceive the competition to be bigger than it actually is; they distort the probability their good will be taken in each phase and do not realize that they only need to compete with the one other person who values their same good.

This interesting result may also explain our results in part; however, we would expect this behavior to be extremely muted, if not non-existent, in our experiment. In our multi-good demand treatments, bidders are, in fact, in competition with every other bidder in their group (all bidders have values for every good). Therefore, for this bias to hold, bidders would have to believe that they are in competition with people who do not exist, which seems unlikely. Further, if this bias was solely responsible for the revenue superiority of the bidder’s choice mechanism, we would not expect a significant difference between the bidder’s choice auctions and the benchmark auctions for multi-good demand. However, we do find that the revenue premium is significant, even though it is diminished greatly from the single-good demand case. This paper, however, abstracts from analyzing this in detail and rather focuses on establishing the differences in revenue between single and multi-good demand⁹. The other main finding of Eliaz et al. – that quantity restriction may increase seller surplus by allowing the seller to keep one good without losing revenue – presumably would hold for our multi-good treatments as well, though we do not test this explicitly.

⁹ All of the treatments conducted by Eliaz et al. assume bidders only have positive value for one good. Bidders drop out of the auction if they have already won or if their good has already been selected.

2.3 Experimental Research in the Field

To our knowledge, there exists only one prior field study on bidder's choice auctions¹⁰. Alevy et al. (2010) find support for the Burguet's original theory by using water volumes that differed by reservoir source and time of availability as goods. Farmers in Chile bid for the water volumes in two treatments: a standard sequential auction and a bidder's choice auction. Arguably, the farmers had strong preferences for specific goods (volumes of water at a specific time and place) – the authors state, “bids decline substantially for the less preferred goods in both auction institutions, reinforcing the finding of heterogeneity in preferences” – suggesting that the model more closely resembles a single-good demand situation than multi-good demand.

In this study, on the other hand, the three goods auctioned in the field experiment (an iPod, hiking equipment, and fine wine) presumably have some substantial value to every bidder. In fact, the three goods had almost exactly the same retail value so it is reasonable to assume that the bidders had similar values for the three goods (or that at least that the variance of values was less than the variance of values in the Alevy 2010 field work). Our field experiment thus resembles a situation where competition persists more so than previous work. This, coupled by the fact that we do not find support for the original theory in the field (despite clearly risk averse participants)¹¹, suggests that the revenue superiority of bidder's choice auctions does not hold (or is muted) by multi-good demand. Our subsequent lab experiment, therefore, seeks to further develop and provide evidence toward this hypothesis.

¹⁰ By field study, we are referring to experiments that did not occur in an experimental laboratory. In the field study discussed here, participants' values for the goods were private and homegrown (i.e. the values were not induced) and participants bid and paid with their own money.

¹¹ In our field experiment, bidder's choice auctions and standard good-by-good auctions yield approximately the same revenue.

III. Theoretical Background

3.1 Single-good Demand

Burguet (2007) illustrates the intuition of the revenue superiority of the bidder's choice auction with this simple example. Two bidders each have unit demand for one of two goods. Each bidder is equally likely to prefer either good – the bidders prefer the same good with probability one half and prefer different goods with probability one half. The payoff to each bidder for winning their preferred good is one (1) minus the price and the payoff for winning the non-preferred good is zero (0) minus the price. The price paid is equal to the second-highest bid. The goods are auctioned off in two phases where both bidders place bids for the right to choose their preferred good in phase one. The bidder who wins in the first phase of the auction does not participate in phase two. In the second phase, the remaining bidder has a fifty percent chance of “winning” her preferred good and a fifty percent chance that the remaining good is her non-preferred good. Expected utility in the second phase is given by the right-hand side of the following equation.

$$u(1 - R) = \frac{1}{2}u(1) + \frac{1}{2}u(0)$$

In the first phase, bidders will not be willing to pay more than R , which will make them indifferent between the two phases. Normalizing $u(1) = 1$ and $u(0) = 0$ allows one to easily see that $R = 1/2$ for a risk neutral bidder and $R > 1/2$ for a risk averse bidder with a concave utility function. This bid of $R > 1/2$ represents the seller's revenue for a bidder's choice auction with risk averse bidders. A standard second-price good-by-good sequential auction, on the other hand, would yield revenue equal to one half regardless of the bidders' risk preferences; if both bidders prefer the same good, the seller gets one (1), and if the bidders prefer different goods, the seller gets zero (0). Therefore, under risk aversion, the bidder's choice auction raises more revenue than the standard good-by-good auction. Intuitively, a bidder in a bidder's choice auction faces a tradeoff between paying more and greater risk that her preferred good is not still

available in the second phase. A risk averse bidder is willing to bid higher in the first phase to avoid the risk of losing her preferred good.

Eliaz et al. (2008) extend this theory to account for infinite bidders and infinite goods. As previously stated, they also add a dimension of quantity restriction. However, the bid function that these authors derive applies only when each bidder values exactly one good (single-good demand)¹². Since our experiments involve both single-good and multiple-good demand, we continue with Burguet's example by modifying it to account for bidders with values for multiple goods.

3.2 Multi-good Demand

Now we turn to cases involving multi-good demand, where bidders do not drop out because they have positive values for multiple goods. Now, both bidders have positive value for both goods: one (1) for the preferred good and α for the less-preferred good where $0 < \alpha < 1$. Again, there is a fifty percent chance that the two bidders prefer the same good and a fifty percent chance that they prefer different goods. Expected utility in the second phase is now represented by the right-hand side of the following equation.

$$u(1 - X) = \frac{1}{2}u(1 - \alpha) + \frac{1}{2}u(0)$$

To elucidate, let us refer to the winning bidder in the first phase as “Bidder A” and the other bidder as “Bidder B”. If the bidders prefer *different* goods, Bidder A chooses her preferred good in the first phase and bids up to α in the second phase. Bidder B bids up to 1 in the second phase since her preferred good is still available, but pays a price of α (the second-highest bid). If the bidders prefer *the same* good, both bidders have values of α for the remaining good in the second phase. Since the second phase is essentially a second-price good-by-good auction for the

¹² The theoretical bid function derived by Eliaz et al. (2008) also assumes risk-neutrality of the bidders. Therefore, the bids that we observe are higher than what would be predicted by their theory. In their experiment, the authors use a benchmark treatment (a second-price sealed-bid good-by-good auction), where risk preferences have no effect, as a comparison for the bidder's choice mechanism. We follow suit and use this same benchmark in our experiments as well.

remaining good, both bidders will bid up to their value. Consequently, Bidder B receives a payoff of zero: either she doesn't win or she wins and pays the second-highest bid (α)¹³.

In the first phase, a bidder will bid X which makes her indifferent between winning in the first phase and facing the lottery in the second phase. This is represented in the left-hand side of the above equation. As in Burguet's example, X will be larger under risk aversion than risk neutrality; the bidder will be willing to give up some surplus in order to secure her preferred good in the first phase and avoid the lottery. In comparing X in this example to R from Burguet's original model of single-good demand, it is clear that $X > R$; a bidder will bid higher in the first phase under multi-good demand in this model. However, this is simply the result of increasing demand for the goods (the good that had zero value to each bidder in the single-good demand case has positive value in the multi-good demand case)¹⁴.

The more interesting result answers the question, "which case (single versus multi-good demand) raises more revenue *above the benchmark*?"¹⁵ First, we will show that the bidder's choice mechanism raises the same revenue (in expectation) as the benchmark good-by-good auction under risk neutrality for multi-good demand¹⁶. (Recall that under non-persistent competition, Burguet showed that risk neutrality yields expected revenue equal to $1/2$ in both the bidder's choice auction and the benchmark.) Under multi-good demand, risk neutrality leads to an expected revenue equal to $X + \alpha$ in both auction formats.

In the benchmark, the seller's revenue is equal to $1 + \alpha$ if the two bidders prefer the same good and $\alpha + \alpha$ if the two bidders prefer different goods. Therefore, expected revenue is:

¹³ Assume, here, that there exists a tie-breaking rule given that the bidders are completely identical.

¹⁴ Note that both of these simple models use one (1) as the value of the preferred good. Since this value is constant and the value for the non-preferred good has increased, demand has increased. We choose to model the values in this way for simplicity of comparison. Similarly, the values for both single and multi-good demand treatments are drawn from the same support in our experiment. This reflects the choice that a seller may face when choosing an auction mechanism: given fixed bidder values, is it better to use the bidder's choice auction when competition persists or does not persist? Or, similarly, should I use a bidder's choice auction or a standard good-by-good auction when I believe that bidders have low variance in their values? What about when they have high variance?

¹⁵ This is what we refer to as the revenue premium of the right-to-choose mechanism: the right-to-choose revenue minus the good-by-good revenue. We are essentially asking which of the following is greater: (the risk averse R – the risk neutral R), or (the risk averse X – the risk neutral X)?

¹⁶ Recall that this result was shown more formally by Harstad (2010).

$$GBG \text{ Revenue } (MG) = \frac{1}{2}(1 + \alpha) + \frac{1}{2}(\alpha + \alpha) = \frac{1}{2} + \frac{3}{2}\alpha.$$

In the bidder's choice format, the seller's revenue is equal to the price paid in the first phase plus the price paid in the second phase: $X + \alpha$. Since we are considering a risk-neutral bidder, we can normalize $u(x) = x$ so that the bidder's first phase bid function becomes the following equation.

$$1 - X = \frac{1}{2}(1 - \alpha) + \frac{1}{2}(0)$$

Solving this for X yields $X = (1/2) + (1/2)\alpha$. Adding α for the price paid in the second phase provides the expected seller's revenue for the bidder's choice format.

$$RTC \text{ Revenue } (MG) = \frac{1}{2} + \frac{1}{2}\alpha + \alpha = \frac{1}{2} + \frac{3}{2}\alpha.$$

Now we show that the revenue premium of the bidder's choice format is higher under multi-good demand than single-good demand simply by comparing the variance of the second-phase lotteries. Risk aversion should cause a bidder to be willing to give up some surplus in the first phase to avoid facing a lottery in the second phase. Given a particular level of risk aversion, should a bidder raise her first-phase bid (over the risk neutral bid) more under single or multi-good demand? It is simple to show that the variance of the second-phase lottery (the right-hand side of the bid function) under single-good demand is $1/4$ and the variance of the second-phase lottery under multi-good demand is $(1/4) + (1/8)\alpha^2 - (1/4)\alpha$. Since $0 < \alpha < 1$, the former variance must be larger than the latter¹⁷. Intuitively, a bidder is more afraid of losing her most preferred good in the first phase when she does not have a chance at positive surplus in later rounds.

PROPOSITION 1: The risk averse bidder should raise her first-phase bid higher (over the corresponding risk-neutral bid) under single-good demand than multi-good demand because the variance of the alternative lottery in the second phase is greater.

¹⁷ For a formal proof, please see Appendix A.

We predict, therefore, that the revenue premium will be higher in single-good demand treatments than multi-good demand treatments. This result would be consistent with Burguet's aforementioned theory of increased seller revenue with increased taste diversity. Further, we predict that bidding behavior will reflect risk preferences; a more risk averse bidder will bid higher, given their value. Finally, we predict that a greater variance in values for a bidder in a multi-good demand treatment will lead to that bidder submitting a higher bid, given that the bidder is risk averse.

We realize, however, that risk aversion may not be the only force driving possible results. Recall one of the secondary results of Eliaz et al. (2008) mentioned earlier. The authors calculate equilibrium bids for different numbers of hypothetical competitors given the random value draws in their experiment. They find that bidders behave as if they are competing with five other bidders, when in reality, they are only competing with the one other bidder who values the same good. This bias could conceivably affect behavior in multi-good demand auctions as well. However, bidders would have to perceive that the competition encompasses more participants than were actually participating (for the bias to work in the same direction). Under multi-good demand, bidders do compete with all five other participants in their group; in order for this bias to cause raised bids, bidders would have to believe they were competing with more than five other people, even though there are only six people in each group.

It could be the case that the bias exists under single-good demand, but does not exist under multi-good demand – i.e. the bias is the reason for the revenue superiority of the bidder's choice format when bidders only value one good, but multi-good demand eliminates this bias because bidders do compete with everyone in their group. If this is the case, we would expect that revenue under multi-good demand would be the same for both the bidder's choice auction and the benchmark. Our aim is to provide explanations for varying field results and this behavioral bias could very well be at work in the field.

Another seemingly feasible behavioral bias could result from bidders only using their highest value (value for their most preferred good) to determine their bid in multi-good demand treatments as long as their most preferred good is still available. Previous literature has shown

that experimental subjects in pooled auctions may weight their most preferred outcome more heavily than less-preferred outcomes due to an attentional bias (Salmon and Iachini 2007)¹⁸. If this is the case, we would predict that single and multi-good demand treatments would yield the same or very similar revenue premiums¹⁹.

3.3 Information

The results of two field experiments, ours and Alevy et al. (2010), vary in their support of theory: our field experiment does not support the revenue superiority of the bidder's choice mechanism while the 2010 field work does. We can identify two major differences between these field experiments and previous laboratory experiments: information revelation and single versus multi-good demand²⁰. In order to attempt to explain the contradictory results, we must vary each of these attributes individually and simultaneously, which results in a 2x2x2 design. Our lab experimental design is explained further in the following section.

Price information may allow bidders to update their beliefs regarding the bounds of the value distribution. The distribution is always uniform over the support $[1, 100]$, but public prices will allow some expectation of the realization of these values. This, in turn, may alter a bidder's expectation of the probability of winning in subsequent phases based on the updated order statistic. We predict that treatments where prices are revealed after each phase may exhibit different results in the second and third phases than treatments where information is withheld (the first phase should be unaffected since prices are not revealed until after the first phase is complete).

¹⁸ See literature review.

¹⁹ If the attentional bias was complete (i.e. the bidders only used their most preferred good in determining their bid), then we would expect the revenue premium to be exactly the same for single and multi-good demand (given identical values for the most preferred good). If the attentional bias only caused bidders to underweight outcomes associated with their less preferred goods by some proportion, we would expect the revenue premium to be larger than in the absence of the bias. This is due to the fact that the existence of less-preferred outcomes in our design actually decreases the revenue premium (bids are not raised as high above the risk-neutral bid).

²⁰ Recall that the variance in bidders' values (which can also be thought of as the persistence of competition, the level of taste diversity, or the level of heterogeneity in preferences) presumably differed between the two field experiments.

As a preview, our field results indicate that in some cases, bidders who placed the second-highest bid in the first phase actually decreased their bid in the second phase. These bidders probably thought that they would win in the second phase after finding out that their bid set the price in the first phase. They decreased their bid in attempt to gain extra surplus, but did not take into account that it is theoretically optimal for every bidder to raise their bid in each subsequent phase (given that their most preferred good is still available) until submitting a bid equal to their value in the final phase²¹. We predict that this behavior of decreased bids in the second phase may occur in our lab experiment as well. To test this, we execute additional experimental sessions so that each of our other four treatments (bidder’s choice and standard good-by-good, for each single and multi-good demand) are executed both with and without information revelation.

IV. Field Experiment

4.1 Experimental Design

To study the bidder’s choice institution in the field, we conducted 30 auction markets in the spring and fall of 2008. Subjects were randomly assigned to 16 markets in which the bidder’s choice (hereafter “BC”) institution was implemented and 14 markets in which the standard sequential good-by-good auction (hereafter “GBG”) institution was implemented, using a between-subjects design. A total of 155 subjects participated in the study. With the exception of three BC and two GBG markets which had six bidders, the markets contained five bidders each.

Subjects were recruited broadly from the Reno population with outreach to the community taking place through flyers and announcements in local stores and through community organizations. Subjects were also recruited from existing databases of non-students who had participated in previous field experiments, and from University of Nevada Reno staff.

²¹ This optimal path of shading bids less and less in each phase is supported by the bid function derived by Eliaz et al. (2008).

Sessions were held on both the north and south side of Reno in accessible locations as well as on the University of Nevada Reno campus. Statistical tests indicate that participants do not significantly differ across the BC and GBG treatments in demographic characteristics that include gender, age, education, and income, or in personality traits.

To cleanly observe the impact of the auction institution, the goods for sale in the BC and GBG auction settings were identical. The goods, or more appropriately bundles, consisted of (i) hiking equipment that included a backpack, water filtration device, and first aid kit, (ii) an Apple iPod and speaker system, and (iii) three bottles of high quality wines. Each bundle had a retail value of approximately \$250. Within each market the goods were sold in three auction phases. In each phase, a single good was allocated to the highest bidder using a second-price rule. In the BC institution, the good sold was the right to choose from the remaining bundles, which varied with the auction phase and market history. In the GBG institution, the good for sale was announced prior to the auction. The order in which goods were sold in the GBG auctions was determined randomly prior to the first auction phase.

The auctions were hand-run, with bidding cards for three phases distributed to participants at the start of the session. In all sessions, treatment specific instructions on the bidding process were distributed to participants and read aloud by the experimenters. An example of allocation through the second-price rule was discussed in detail. After reading the instructions, but before submitting bids, subjects had the opportunity to visually inspect the goods²².

In addition to the auction, each session included a risk elicitation exercise, and a short survey. The risk elicitation closely followed the protocol developed by Holt and Laury (2002) and consisted of a series of 10 binary choices, each between a safe and risky lottery. The payoffs were \$200 and \$160 for the safe lottery, and \$385 and \$10 for the risky lottery. The probability of gaining the higher payout increased from 10% to 100% across the ten choices as is standard with this protocol. In this implementation subjects were paid with a one-third probability, with the outcome determined independently across subjects after the questionnaire was completed. To

²² The full instructions are available upon request.

determine payoffs, experimental monitors would (i) roll a 10-sided die to pick one of the questions for potential payment and (ii) roll a 6-sided die to determine if subjects were paid based on their response to the selected question. Subjects were paid the outcome of their choice if a 1 or 2 resulted from the die roll and received nothing otherwise.

The final element of each session was the completion of a survey which included (i) the collection of demographic data, (ii) the elicitation of personality traits, and (iii) the cognitive reflection test (CRT), which contains three questions intended to measure impulsivity and intelligence (Frederick 2005). Subjects received \$2 for each of the CRT questions answered correctly. A series of 40 questions contained in the International Personality Item Pool (IPIP) were used to measure the traits of assertiveness, sociability, performance motivation, risk-taking, confidence, beliefs about intelligence, and efficacy. The personality items were measured using a five-point Likert scale.

4.2 Field Results

Figure 1.1²³ illustrates that most of our participants are risk averse; it shows the proportion of participants in each risk preference group. T-tests show that the proportions are statistically different between auction types for the risk-loving group (p -value = 0.04), but not for the risk-neutral or risk averse groups (p -values = 0.87 and 0.15 respectively). Since risk posture is irrelevant for GBG auctions and BC participants are highly risk averse, we expect BC theory to hold.

Each participant was asked to rank the 3 goods from “Most Preferred” (a ranking of 1) to “Least Preferred” (a ranking of 3). Figure 1.2 demonstrates that while the iPod package was preferred over the wine and hiking packages, the preferences are very similar between BC and GBG auctions. However, we do find a significant difference in preferences between treatments for the hiking and iPod packages (p -values = 0.02 and 0.01 respectively). The wine package preferences are not statistically different between treatments (p -value = 0.81). Since revenues are

²³ All figures are in Appendix C.

driven by those who have the highest values for each good, we also examine the proportions of participants who ranked each good as their most preferred (illustrated in Figure 1.3). The proportions are not statistically different between treatments for the hiking or wine packages (p -values = 0.09 and 0.11 respectively), but are statistically different for the iPod package (p -value = 0.00). Again, however, we do find that the ordering of the goods is the same across treatments. We acknowledge that a difference in preferences could affect our results in part; however, we show later in the section that it is not the primary cause for our main findings.

Since the majority of our auction participants are risk averse, BC theory tells us to expect higher revenues from the BC auctions than the GBG auctions. However, we do not find a significant difference in revenues between the two types of auctions (t -test p -value= 0.61). Table 1.1²⁴ includes average revenues for each phase and average market revenues. The average revenues in the GBG markets are not statistically different between phases, ruling out any order or wealth effects.

In case a lack of variation in preferences is driving the result, we temporarily eliminate any markets where all participants' most preferred good was the same. One GBG auction and four BC auctions meet this condition. The remaining markets' revenues are reflected in Table 1.2. BC auctions yield higher revenue than GBG auctions, but the difference is not significant (p -value = 0.78). Although this result is more in line with theory, we expected BC revenues to be significantly higher than GBG revenues.

RESULT 1: In contrast to theoretical prediction, the bidder's choice auction does not raise higher revenue than the good-by-good auction in the field experiment.

To further explore why our results are not in line with theory, we examine bidding behavior. We would expect that the third phase of BC auctions would have the same result as GBG auctions since both are second-price auctions for one good. The average bids for GBG

²⁴ All tables are in Appendix B.

auctions and Phase 3 of BC auctions are compared in Table 1.3. The iPod package never made it to the third phase in the BC auctions. We cannot reject that the average bid in GBG auctions is different from the average bid in the 3rd phase of BC auctions for wine (t-test p-value= 0.17, Mann-Whitney p-value= 0.30). However, we do reject the null for hiking at the 5% level: the average bid in GBG auctions is significantly higher than the average bid in the 3rd phase of BC auctions (t-test p-value= 0.01, Mann-Whitney p-value= 0.06). What might have happened during the course of the BC auctions to cause bidders to bid less in the 3rd phase than in GBG auctions for the same good? One major difference in the format of BC auctions relative to GBG auctions is that 1st and 2nd placed bidders in BC auctions get feedback.

We hypothesize that bidders in BC auctions may have changed their values for the goods over time. Past lab experiments have shown that BC auctions result in higher revenues, but the authors assume that the value a participant has for a good at the start of the auction remains the same throughout; participants are assigned a value and this cannot change during the experiment. It is easily possible, however, that participants in our field experiment update their values based on others' bids. A participant may see the winner in the first or second phase choose a good other than the one they believed was most valuable. Since all participants are aware of the second-highest bid (the amount the winner pays), the participant may believe he made a mistake judging the value of the good.

In addition to the monetary values of the goods, the participant may update his belief on the *relative* values of the goods. For instance, if a participant sees the iPod and the wine packages get chosen in the first and second phases, he may update his belief on the popularity of hiking relative to listening to music and drinking wine. If he finds that hiking is not as popular an activity as he originally believed, he may not bid as high for the good. Finally, a participant may get discouraged in an BC auction by watching others winning over him and choosing their favorite goods. In a GBG market, on the other hand, a participant may not be discouraged by losing in an auction for his least favorite good; he knows he was not really trying to compete with the other participants for that good.

We also analyze the revenues in GBG auctions versus the last phase of BC auctions and the findings are in line with the comparison of the bids (see Table 1.4). We cannot reject the null hypothesis that there is a difference in revenues for wine (t-test p-value= 0.88, Mann-Whitney p-value= 0.93). We can reject the null for hiking, but only at the 10% significance level when using a t-test (t-test p-value= 0.08, Mann-Whitney p-value= 0.04). We further investigate behavior by analyzing bids in more detail.

Next, we examine the first phase of BC auctions in comparison with GBG auctions. We would expect that bidders would shade their bids in the first phase of BC auctions and, in fact, they do (see Table 1.5). Bids are compared between GBG and Phase 1 of BC by declared most preferred good. In other words, the first column includes the average bid in Phase 1 of BC for participants who preferred the hiking package and the average bid in GBG “Hiking” auctions for participants who preferred the hiking package. The average bid for the most preferred good in GBG auctions is not significantly higher than in the first phase of BC auctions for the wine and hiking packages. We cannot reject the null in t-tests or Mann-Whitney tests at the 5% or 10% levels. However, the difference is significant for the iPod at the 10% level (t-test p-value= 0.08, Mann-Whitney p-value= 0.08). Note that current theory does not specify *how much* bidders will shade during the first phase of an BC auction, but just that they will decrease their bid from their true value. This is what we observe.

RESULT 2: Although some anomalies exist, bidding behavior is generally in line with theory when comparing 1st and 3rd phase bids between BC and GBG in the field experiment.

We further analyze bidding behavior by calculating the change in bids for each participant over time in BC auctions. Table 1.6 displays the change in bids for individuals whose most preferred good is still available in the next phase. For instance, the first cell displays the average change in bids between Phase 1 and Phase 2 for participants whose most preferred good is the hiking package if the hiking package is still available (the hiking package was not chosen

by the first winner). We find mixed support for the theory here; some bidders decreased their bids and some increased. We would have expected all bidders to increase their bids if their preferred good was still available.

To delve deeper, we examine how bidders change their bids by initial rank. Table 1.7 summarizes the results. Each bidder was ranked from highest bid (rank=1) to lowest bid (rank=5 or 6 depending on number of participants in the market) in each phase. We find that bidders who were ranked 2nd lower their bids on average for the next phase. Bidders who were ranked 2nd knew their rank because the price paid by the winner was announced. Bidders who were ranked greater than 2nd did not know their rank, but did know that they were not ranked first or second. It turns out that bidders who were ranked 2nd were responsible for drops in bids; bidders ranked 3, 4 or 5 increased their bids on average.

This contradicts the traditional theory; bids should continue to increase in each phase when the bidder's most preferred good is still available. In phases prior to the last phase, bidders should shade their bids just enough so that they are indifferent between winning and facing the lottery that occurs in the last phase. As phases progress, this shading should become less and less, assuming the most preferred good is still available. (In the last phase, which is essentially a good-by-good auction for the remaining good, bidders should bid their value.) This result leads to speculation over whether bidders may have updated their expectations of the goods' values as new information, such as the first good selected, was revealed. For instance, one subject may believe that the Wine package would be the most popular (and therefore possibly easier to resell), but is surprised when the winner in the first phase chooses the iPod. The subject now lowers his private value for the Wine, even though his bid should theoretically increase in the second phase. This updating cannot be controlled in the field setting since we do not observe private values. Consequently, it is difficult to conclude why exactly the BC institution did not raise higher revenues than GBG auctions. This difficulty motivates the second portion of this research, the laboratory experiment, which will be discussed in later sections.

RESULT 3: Bidders who are ranked 2nd in the BC auction in Phase 1 decrease their bids in Phase 2; this is definitely a contradiction of theory.

Next, we explore how demographics and personality measures affect bidding behavior using regression analysis. We find that when a bidder's most preferred good *or* second most preferred good is still available in an BC auction, he increases his bid, as expected. Bidders also increase their bids significantly in GBG auctions for their most preferred goods. Table 1.8 provides these regression results. In addition, we find that sociability²⁵ is a significant negative predictor of BC bids whereas confidence has a significantly positive impact on GBG bids.

As expected, an indicator for a risk-averse individual is significant in predicting BC bids but not GBG bids. However, the direction of impact is not in line with theory; risk-averse individuals should bid more than risk neutral or risk seeking individuals because they do not want to risk losing their preferred good. Instead, we find that risk-averse individuals bid less. Rather than having aversion to losing their preferred good, our bidders are averse to paying too much for a good. According to theory, BC auctions produce higher revenues only if bidders are risk averse. The negative coefficient on our risk aversion indicator could partially explain why we see equivalent revenues between the two auction types.

We also examine how bidders behave in comparison with theory by constructing an indicator for circumstances where bids should have increased in BC auctions. If a participant's most preferred good is still available, they should always increase their bid in the next phase. They should also increase their bid if their least favorite of the remaining goods gets chosen. For instance, if the order (by preference) of goods taken for a participant is "1, 3, 2" (most preferred good taken in 1st phase, least preferred good taken in 2nd phase, 2nd most preferred good taken in 3rd phase), the participant should increase his bid from the 2nd phase to the 3rd phase. If the preference order of goods taken is "2 1 3", the participant should increase his bid from the 1st phase to the 2nd phase, but theory is silent on what he should do from the 2nd phase to the 3rd

²⁵ Sociability and confidence are personality measures derived from the International Personality Item Pool (IPIP) questions mentioned previously.

phase (it depends on how much he has been shading and the difference in value to him of the goods).

We find that in more than half of cases (54%) where theory predicts a bid increase in BC auctions, participants actually decrease their bids or leave their bids unchanged. This result alone demonstrates that the theoretical BC bidding strategy is not followed in the field, regardless of where BC versus GBG preferences or revenues stand. One could argue that participants do not understand the optimal BC bidding strategy. However, we already saw that BC bidders are, in fact, shading their bids across the board (for all goods) during the first phase. It seems quite unlikely that this could be a coincidence; instead, participants probably underestimated how much other bidders would shade or made emotional decisions based on the outcomes of each phase. This caused theory to break down during the second and third phases.

Table 1.8 reports the effects of demographics and personality measures on an indicator of whether a participant increased their bid. A Probit model is used and the sample is restricted to instances where theory says the participants' bids should have increased. An indicator for above-average education is significant but in the opposite direction from what one might expect; participants are less likely to increase their bid when they should if they are educated. This may be due to over-analyzing the bidding strategy. The results of an IQ quiz are more intuitive; smarter participants are more likely to increase their bids when they should. Confidence, efficacy and sociability are also significant predictors of how likely a participant is to follow theory. Notice that our indicator for risk aversion still has a negative coefficient, though not significant here.

As previously discussed, it is difficult to determine exactly why theory broke down in this field study. As is the nature of field studies, private values are unknown; further, any updating of private values are also unknown. Contradictory to theory, risk aversion does not appear to play a role, as it is insignificant in our econometric analysis regardless of how the parameter is defined.²⁶ Further, we observe bidding paths that are illogical, unless participants'

²⁶ Using the results of our Holt and Laury style risk elicitation, we tried defining the risk parameter in several different ways, including an indicator for very risk averse subjects, an indicator for mildly or very risk averse subjects, a numeric variable which reports how many "safe" choices the subject chose, an indicator for subjects

private values changed throughout the auction. These inconsistencies motivate the second stage of our research: the laboratory experiment. In the lab, we are able to control values by inducing them. Further, we can isolate possible confounding points which have not yet been explicitly tested in the literature. Specifically, we isolate price revelation (providing information on the prices of winning goods, which may have led to some of the aforementioned behavioral biases), and persistent competition (subjects having value for multiple goods). As discussed in the second section, previous work has focused on bidder’s choice auctions where participants dropped out after their only preferred good had been chosen. In our field study, however, it is quite plausible that participants had values for all three packages. In the next section, we illustrate theoretically why this may cause the bidder’s choice mechanism to be less superior than previously thought.

V. Lab Experiment

5.1 Experimental Design

Our experimental approach closely follows Eliaz et al. (2008), though our dimensions of variation differ to focus on the effects of information and multi-good demand. In one set of treatments, bidders are informed of the winning price (the second-highest bid) in each phase (referred to as “I” treatments), whereas other treatments do not provide this information (referred to as “NI” treatments). In another set of treatments, bidders draw random values for three goods in each round (multi-good demand – “MG”), while in single-good demand treatments (“SG”), bidders only draw a random value for one good in each round. All of these treatments are tested using a bidder’s choice, or right-to-choose, auction (“BC”) and a standard sequential good-by-good auction (“GBG”). This 2x2x2 design yields a total of eight treatments.

Values were drawn from a uniform distribution over the support [1, 100]. For SG treatments, three sets of preferred goods (values) were drawn ex ante and used repeatedly for

who chose more safe choices than average, an indicator for subjects who reported consistent risk preferences (did not switch back and forth between the risky choices and the safe choices), etc.

different groups. For MG treatments, we varied the order of the sets of values for different groups to control for ordering effects. Consistent with Eliaz et al. (2008), we did this to ensure that differences in revenue are attributable to differences in behavior rather than differences in the vectors of random variables generated.

Eight sessions (one for each treatment) were executed at the University of Tennessee, Knoxville (UTK) and eight sessions (one for each treatment) were executed at the University of Alaska, Anchorage (UAA) in April 2012. The experimental laboratories at the two universities have similar recruiting procedures, attracting undergraduate students from a variety of disciplines. Four groups of six bidders each occupied the laboratory for each of the sessions held in Tennessee. Due to capacity constraints, two or three groups of six bidders participated in each of the eight treatments in Alaska. There were a total of 324 participants. The sessions lasted about 70 minutes and most participants earned between 15 and 40 U.S. dollars in total²⁷. Table 1.9 displays the experimental design.

The experimental sessions proceeded as follows. First, the subjects were asked to participate in a risk elicitation similar to the one popularized by Holt and Laury (2002). Instructions for the risk elicitation were read aloud while subjects followed along with on-screen instructions. The computer program then allowed the subjects to make 10 risk decisions, one of which would be selected at random and paid out at the end of the session²⁸. Next, the subjects were given hard copies of the auction instructions and asked to read along while the instructions were read aloud²⁹. The subjects were encouraged to ask clarifying questions before the experiment began. The subjects were randomly assigned into groups of 6 and were unaware of the identities of the other 5 participants in their group.

The experiment consisted of a practice round, followed by 10 paid rounds. In each round, there were 3 phases. In BC treatments, each phase was an auction for a right to choose³⁰. In GBG

²⁷ This total includes earnings from all 10 auction rounds plus earnings from a risk elicitation.

²⁸ We elicited the risk preferences of the subjects prior to the experiment to ensure that their responses were not affected by their experiences of wins and losses during the experiment. We did not reveal the results of the risk elicitation until the end of the session to avoid any endowment effects.

²⁹ The instructions and screen shots from the experiment are available upon request.

³⁰ Note, however, that the last phase was identical to a good-by-good auction for the remaining good.

treatments, each phase was a standard auction for one of the goods (the goods were labeled “A”, “B” and “C”). Bidders were instructed to submit bids ranging from zero to their value. In GBG auctions, the highest possible bid was a subject’s value for the good being auctioned. In BC treatments, the highest possible bid was a subject’s highest value (in SG treatments, the subject’s highest value was also their only value). We did not allow the subjects to overbid in order to decrease the effect of cognitive mistakes. While learning effects have been the focus of other experimental work, we are focused on cleanly identifying the effects of multi-good demand and information. This arrangement also minimized bankruptcy. A small percentage of subjects in BC treatments did still go bankrupt due to choosing the wrong good³¹. In these cases, if the subject did not recover from the loss, we paid the subject a show-up fee. Eliminating these subjects from the data does not significantly change the results.

After each phase, all bidders were informed whether they won. In BC treatments, all bidders were also informed of the good that was chosen by the winner in their group. In “I” (information) treatments, all bidders were informed of the price of the good sold in their group (the second-highest bid). No subject was ever told the name or ID number of any other subjects in their group so they could not infer that one particular person won more or less often. In SG treatments, subjects whose preferred good sold in phase 1 or 2 “dropped out” of the auction; they faced a screen that read, “Please wait while the other members of your group bid in Phase ...”. In the GBG SG treatments (good-by-good auctions where bidders only have positive value for one good), subjects only bid in phases when their preferred good was auctioned. For instance, a subject who preferred good B faced a screen in the first phase that said, “Please wait while good A is auctioned.” In MG treatments, all subjects participated in every phase.

³¹ For example, suppose a subject’s values are 80, 60 and 40, the subject bids 51 and wins. For simplicity, further suppose that the second-highest bid (price) is 50. If he chooses the good that he values at 40, he receives negative 10 tokens. If he chooses the good that he values at 60, he receives positive 10 tokens, but this is inferior to choosing the good that he values at 80, which yields positive 30 tokens. Choosing any good other than the most preferred is illogical given our experimental design. However, as mentioned earlier, Harstad (2010) theorizes that a bidder may choose a good that is not his most preferred if he believes the good is a “usual favorite”. The bidder’s motive is to eliminate the most popular good in hopes of obtaining his most preferred good in a later round for a low price. In our experiment, all values are drawn randomly from a uniform distribution so there is not a “usual favorite”. However, it is possible that a few subjects misunderstood the implications of random drawings and wanted to test Harstad’s theory as a potential strategy. As stated, this behavior only occurred in a few cases and does not appear to affect our overall results.

Earnings consisted of the subject's value (randomly drawn) minus the price they paid for the good (the second-highest bid). All values and prices were expressed in tokens. In SG treatments, each 8 tokens equaled one dollar; in MG treatments, each 4 tokens was equal to a dollar. The different exchange rates were based on the fact that equilibrium earnings must be less in MG treatments³². Earnings were totaled over the 10 rounds. In the final stage of the experiment, the risk elicitation results were revealed and each subject's total earnings were calculated. Subjects were paid in cash and in private.

5.2 Laboratory Results

First, we simply compare the average revenues (sum of prices paid) for each BC treatment to its GBG counterpart. These values are displayed in Table 1.10. BC revenues are significantly higher than GBG revenues, as expected. T-tests reveal no significant revenue differences between any information treatment and its "no information" ("NI") counterpart (i.e. no significant difference between BC – I – SG and BC – NI – SG or between GBG – I – MG and GBG – NI – MG, etc.). Hence, for the following comparisons, we pool information conditions. The difference between single and multi-good demand is significant, as expected (t-test p-values are 0.000). This result is a consequence of the fact that more values in MG treatments than SG treatments necessarily decreases the gap between the highest and second-highest values in any phase, thereby raising prices. This also decreases bidder surplus and is the reason for the differing exchange rates between the two treatment groups.

The interesting result lies in the difference in the differences between BC and GBG revenues by type of demand. Put differently, the revenue premium of the bidder's choice mechanism (compared to the benchmark GBG auction) is larger when bidders only have demand

³² In SG treatments, there were 6 random values drawn per group per round. In MG treatments, there were 18 random values drawn per group per round. Due to the increased number of draws, the spread between the highest and second-highest values at any given time in MG treatments was, on average, much less than in an SG treatment. (Recall from the theory that the revenue in MG auctions must be higher than SG auctions, given that the same value for the preferred good. The idea is the same here: values were drawn from the same support for both MG and SG auctions, but since more values were drawn in MG auctions, the demand for each good is essentially higher.) The specific exchange rates of 8 and 4 tokens per dollar were determined based on earnings in a pilot session held prior to the experiment.

for one good. We compare the BC—GBG difference in the top two rows in Table 10 (SG treatments) to the BC—GBG difference in the bottom two rows (MG treatments). While the BC revenue premium is statistically significant in both cases (p -values = 0.000), the difference is significantly larger for SG treatments. To see this visually, refer to Figure 1.4. The difference in the differences is statistically significant (48.28, (SE = 4.22) for SG and 18.29 (SE = 2.53) for MG). This is consistent with our theoretical prediction that risk averse bidders do not raise their bids as high (over the risk-neutral bid) under multi-good demand as under single-good demand.

The risk elicitation demonstrates that the majority of participants are risk averse; the results are displayed in Figure 1.5. A risk neutral subject would have switched from the lottery to the certainty equivalent at either Row 8 or Row 9 (the two darker shaded columns in the figure). However, we see many subjects choosing the certainty equivalent prior to Row 8, indicating risk aversion³³.

RESULT 4: While there is no treatment effect associated with the information treatment, there is a significant consequence of multi-good demand. There is a revenue premium for the bidder's choice mechanism regardless of single-good or multi-good demand, but the premium for multi-good demand is significantly smaller. This is consistent with theory.

Due to the fact that bids in multi-good demand treatments are necessarily higher than bids in single-good demand treatments (a result of the larger number of values drawn), we can only compare bidding behavior between the two treatments by using the benchmark GBG auction as a baseline for each treatment, as we have done to analyze revenues. Table 1.11 compares the first phase BC bid for each type of competition alongside the GBG bid for the subject's most preferred good (i.e. if the subject's most preferred good is "A", then this is his first phase bid; if the subject's most preferred good is "B", then this is his second phase bid, etc.). This comparison allows us to analyze each bidder's bid based on his highest value. The first phase BC bids are

³³ The complete instructions for the risk elicitation and screen shots displaying the ten choices faced by participants are available upon request.

significantly less than the GBG benchmark for both single and multi-good demand (p-values are 0.047 and 0.001 respectively). This demonstrates that the general bidder's choice model holds for our experiment: bidders shade their bids from their values in the first phase.

The difference in the BC and GBG bids is larger for multi-good demand (3.18) than single-good demand (2.75), though the difference in the differences is not significant. Recall that the theoretical risk neutral bid is higher for MG than SG given the same value for the most-preferred good (from the theoretical section: $X > R$). On the other hand, risk aversion causes less of an upward force on bids in MG than SG. Hence, we cannot draw any conclusions from this difference and difference comparison.

Table 1.12 provides the BC bids for single and multi-good demand. The data used includes bids where a subject's most preferred good is still available. When bidders only have value for one good (SG treatments), this includes all bids until a bidder drops out³⁴. This allows us to see how the bidding path for the most preferred good progresses. Notice that the average highest value decreases from Phase 1 to Phase 2 and again from Phase 2 to Phase 3. This reflects the fact that some bidders drop out of the auction because they either win or their most preferred good is taken by another bidder.

The progression of bids (as a percentage of value) in multi-good demand treatments is significantly flatter than in single-good demand treatments (Phase 1 p-value = 0.00, Phase 2 p-value = 0.04, Phase 3 p-value = 0.04). Again, there are two effects at work: the theoretical bid as a proportion of value is higher (than SG), but the effects of risk aversion should be muted. Overall, we observe that subjects bid very high percentages of their values in multi-good demand; as we have already seen, this ultimately leads to higher revenue than the multi-good benchmark. Thus, subjects must be bidding higher than the risk neutral bid. Average bids in the third phase are slightly below average values for both SG and MG treatments. We find that this is the case in the benchmark GBG auctions as well; we turn to this analysis next.

³⁴ In SG treatments, a bidder drops out when his most preferred good is no longer available. In MG treatments, a bidder continues to bid when his most preferred good is no longer available; at this point, his bid cannot be higher than his second-highest value. For comparison's sake, we limit the MG observations to bidders whose most preferred good is available in the phase of interest.

RESULT 5: Bidding behavior in the lab experiment is generally in line with theory. The bid paths (calculated as bid/value for subjects whose most preferred good is still available) significantly differ between single and multi-good demand.

Table 1.13 provides the theoretical bids for the benchmark treatments (which are the values) alongside the actual bids. Subjects consistently bid a few tokens below their values in the benchmark treatments. Interestingly, this shading is more pronounced under single-good demand than multi-good demand. The average percentage of value that was bid is significantly higher in MG treatments than SG treatments in each phase (p-value = 0.000 for Phase 1, p-value = 0.000 for Phase 2, p-value = 0.004 for Phase 3)³⁵. Multi-good demand has an effect on bidding in the benchmark treatments as well as the bidder's choice treatments. However, this effect is purely behavioral; there is no theoretical foundation for bidding less than value in a second-price good-by-good auction.

In light of this behavioral discovery, we retest revenues to reflect optimal behavior in the benchmark treatments. That is, we compare bidder's choice revenues to the theoretical benchmark revenue; if bidders had behaved according to theory in GBG auctions, would our main result hold? Figure 1.6 displays this new information. We find that our revenue result is, in fact, robust to GBG behavioral biases. The difference between the bidder's choice (BC) revenue and the benchmark (GBG) revenue is significant for both single and multi-good demand. Additionally, the difference in these differences is also significant; the revenue premium is significantly reduced when bidders value more than one good.

In the final part of our analysis, we use econometric methods to demonstrate how risk preferences, variance over values and individual characteristics affect bidding behavior. Our model focuses on the first and second phase BC bids where a subject's most preferred good is available and GBG bids for the most preferred good, allowing us to compare outcomes between

³⁵ The actual difference in tokens between PC and NPC is also statistically significant for the first two phases (3.72 is a significantly larger difference than 1.31, etc.). We tend to focus on percentage differences, however, because the risk neutral theoretical predictions provide bids in terms of proportion of value.

persistent and non-persistent competition. We know that no subject has a value of zero for their most preferred good, so a bid of zero is never optimal in our model; we do not need to control for left-hand censoring. It is also never optimal for a subject to bid more than 100% of his value so the fact that the bids are capped at 100% should not cause a bias. Therefore, we use simple linear models with errors clustered at the individual bidder level to analyze bidding behavior. Alevy et al. (2010) use a Tobit model to allow for corner solutions where a subject's optimal bid is zero. In their model, however, they incorporate bids for all available goods, which include goods that may have no value to the bidder. Our use of induced values in a laboratory experiment, on the other hand, allows us to focus on Phase 1 bids (R and X from the theory section), which drive revenues.

Table 1.14 presents the OLS estimates. The baseline for the model is the benchmark GBG auction. We use subjects' bids as a percentage of subjects' highest values as the dependent variable. Since theory predicts a bidder's BC bid should be a fraction of their value, this approach allows us to explore how various characteristics affect bidding behavior. We find that the bidder's choice mechanism lowers bids as compared to the benchmark by about 12% of the subject's value for single-good demand and about 13% for multi-good demand (during the first phase).

We also find that subjects increase their bids significantly as the experiment progresses, but at a decreasing rate (the coefficient on "Round" is positive and significant, while the coefficient on "Round²" is negative and significant). As expected (though contrary to the field study), subjects raise their bids in the second phase if their preferred good is still available; the interaction term for the BC treatment and the second phase has a positive, significant coefficient for single-good demand³⁶. This interaction term is not significant for multi-good demand. This is probably due to the fact that the optimal bidding path is flatter for multi-good demand than single-good demand (as long as the most preferred good is available), as seen in Table 1.12. Thus, the effect is more subtle.

³⁶ We do not analyze phase differences for the benchmark because: (i) theoretically, drawing the highest value for good "B" versus good "A" or good "C" should not affect bids and (ii) we see from Table 5 that, in fact, bids in each benchmark treatment are not significantly different between phases.

Additionally, we find that grade point average has a positive and slightly significant effect on the percentage of value bid under multi-good demand. Contrary to expectation, an indicator for the variance of a subject's values in multi-good demand is not significant. This suggests that subjects may exhibit a threshold bias. It may be the case that bidders perceive *any* positive values for their lesser preferred goods to be the same (versus zero values in SG treatments); i.e. the subject considers that he has a chance to earn surplus in later rounds, but does not consider what that surplus may be. If this is the case, the effects of risk aversion would be muted.

Finally, we also control for risk preferences. We do not expect that risk preference should have an effect in the benchmark, so the fact that this variable is not statistically significant is not a surprise. However, we also interact risk preference with BC; this variable is not significant as well, contrary to expectation. It is possible that the risk elicitation we used was too coarse of a measure to pick up on the differences in risk preference which might affect bidding behavior. The other possibility is that the behavioral bias proposed by Eliaz et al. (2008) could theoretically bias bids, though we argue in the discussion that this is extremely unlikely³⁷.

VI. Discussion

Bidder's choice auctions have been shown to yield higher revenue than simple good-by-good auctions. Theoretically, this is a result of risk aversion, but Eliaz et al. (2008) find evidence to support the bidder's choice premium is partially the result of a behavioral bias, causing subjects to believe they are competing with more people than they actually are. In SG scenarios, this bias is plausible and could theoretically be responsible for the BC revenue premium. However, this bias almost certainly would not affect bidding behavior in MG scenarios because bidders would need to believe that they were competing with people that do not actually exist.

³⁷ We also test for endogeneity of risk preference and find that it is not endogenous.

Since we do find a revenue premium in our MG treatments, we conclude that the premium is the result of risk aversion, not a behavioral bias.

While the field experiment did not reveal a revenue premium, the laboratory experiment results show that the revenue premium of the bidder's choice mechanism is significantly greater under single-good than multi-good demand. Bidding behavior is generally consistent with the theory and we find that price revelation does not have a significant effect. This suggests that multi-good demand, not information, is probably the reason that our field study finds results that are at odds with previous work.

In conclusion, we find that multi-good demand mutes the revenue superiority of the bidder's choice institution, consistent with the notion that the perceived risk of losing one's most preferred good is softened when there is a chance to win multiple goods. This result implies that bidder's choice auctions should be used in settings where each bidder is likely to strongly prefer one of the goods over the others, though this need not be the same good for every bidder. This conclusion is consistent with Burguet's (2005) result that greater taste diversity leads to greater revenue. In addition, the results explain why our field experiment finds contrasting results to a previous field study conducted by Alevy et al (2010): our field environment is arguably a case of multi-good demand, which mutes the revenue superiority of the mechanism, while the greater taste diversity (closer to single-good demand) that characterizes Alevy et al. 2010 leads the authors to find substantial support for theory.

Future work may include additional field or lab experiments to cleanly distinguish behavioral biases from risk aversion; a more finely tuned and detailed risk elicitation than is typically used may be helpful since the differences in bids may be very small. Interdependent preferences for goods may also be an interesting extension. For example, a prospective landlord who wins an auction for a condo on the seventh floor of a building may subsequently increase his value for another condo on the same floor (i.e. the landlord's preferences for goods depend on the good(s) he has already acquired). This type of preference structure has implications for broadcast spectrum auctions and plausibly many other applications.

References

- Alevy, Jonathan E., Oscar Cristi, and Oscar Melo. "Right-to-Choose Auctions: A Field Study of Water Markets in the Limari Valley of Chile." *Agricultural and Resource Economics Review* 39 (2), 213-226. 2010.
- Alevy, Jonathan E., Jeffrey Englin, Tom Harris, Man-Keun Kim, Li Yiheng, Michael Price, and Jens Schubert. "BLM Wild Horse and Burro Policy: Auction Design and Feasibility Study." Report to US Bureau of Land Management. 2009.
- Ashenfelter, Orley and David Genesove. "Testing for Price Anomalies in Real-Estate Auctions." *The American Economic Review* 82 (2), 501-505. 1992.
- Burguet, Roberto. "The Condominium Problem; Auctions for Substitutes." *Review of Economic Design* 9, 73-90. 2005.
- Burguet, Roberto. "Right to Choose in Oral Auctions." *Economics Letters* 95, 167-173. 2007.
- Eliaz, Kfir, Theo Offerman, and Andrew Schotter. "Creating Competition out of Thin Air: An Experimental Study of Right-to-Choose Auctions." *Games and Economic Behavior* 62, 383-416. 2008.
- Gale, Ian L. and Donald B. Hausch. "Bottom-Fishing and Declining Prices in Sequential Auctions." *Games and Economic Behavior* 7, 318-331. 1994.
- Goeree, Jacob K., Charles R. Plott, and John Wooders. "Bidder's Choice Auctions: Raising Revenues Through the Right to Choose." *Journal of the European Economic Association* 2 (2-3), 504-515. 2004.
- Harstad, Ronald. "Auctioning the Right to Choose when Competition Persists." *Decision Analysis*, Special Michael Rothkopf Memorial Issue 7, 78-85. 2010.
- Holt, Charles A. and Susan K. Laury. "Risk Aversion and Incentive Effects." *American Economic Review* 92 (5), 1644-1655. 2002.
- Menezes, Flavio M. and Paulo K. Monteiro. "Simultaneous Pooled Auctions." *Journal of Real Estate Finance and Economics* 17 (3), 219-232. 1998.
- Salmon, Timothy and Michael Iachini. "Continuous Ascending versus Pooled Multiple Unit Auctions." *Games and Economic Behavior* 61, 67-85. 2007.

Appendix A

Theoretical Work

Proof of Proposition 1:

Here, we show that the variance of the second-phase lottery faced by bidders in the simple 2-bidder, 2-good case is larger under single-good demand than multi-good demand.

We use the following formula for the variance of a lottery, where A and B are payoffs and $\Pr(A)$ is the probability that outcome A occurs and $\Pr(B)$ is the probability that outcome B occurs.

$$Var = (A - EV)^2 \Pr(A) + (B - EV)^2 \Pr(B)$$

Under single-good demand, the variance of the second-phase lottery equals:

$$Var_{SG} = \left(1 - \frac{1}{2}\right)^2 \left(\frac{1}{2}\right) + \left(0 - \frac{1}{2}\right)^2 \left(\frac{1}{2}\right) = \frac{1}{8} + \frac{1}{8} = \frac{1}{4}$$

Under multi-good demand, the variance of the second-phase lottery equals:

$$\begin{aligned} Var_{MG} &= \left(1 - \alpha - \frac{1}{2} + \frac{1}{2}\alpha\right)^2 \left(\frac{1}{2}\right) + \left(0 - \frac{1}{2}\right)^2 \left(\frac{1}{2}\right) = \left(\frac{1}{2} - \frac{1}{2}\alpha\right)^2 \left(\frac{1}{2}\right) + \frac{1}{8} \\ Var_{MG} &= \left(\frac{1}{4} + \frac{1}{4}\alpha^2 - \frac{1}{4}\alpha - \frac{1}{4}\alpha\right) \left(\frac{1}{2}\right) + \frac{1}{8} = \frac{1}{4} + \frac{1}{8}\alpha^2 - \frac{1}{4}\alpha \end{aligned}$$

Since $0 < \alpha < 1$, it must be the case that $\frac{1}{8}\alpha^2 - \frac{1}{4}\alpha < 0$. Therefore, $Var_{SG} > Var_{MG}$.

Appendix B

Tables

Table 1.1: Average Revenues

	Phase 1	Phase 2	Phase 3	Market
BC	78.55 (23.08)	64.69 (26.12)	50.16 (18.30)	193.40 (53.49)
GBG	74.07 (29.96)	55.51 (29.66)	75.27 (40.19)	204.85 (67.96)

Table 1.2: Average Revenues (Using Only Markets with Variation in Most Preferred Good)

	Phase 1	Phase 2	Phase 3	Market
BC	82.67 (23.93)	71.59 (26.62)	51.26 (11.29)	205.52 (53.87)
GBG	71.69 (29.78)	49.77 (21.33)	77.21 (41.14)	198.68 (66.53)

Table 1.3: Average Bids in GBG Auctions and 3rd Phase of BC Auctions

	Hiking	Wine
BC (3 rd Phase)	34.46 (20.62)	36.18 (23.47)
GBG	50.75 (41.40)	48.08 (40.65)

Table 1.4: Average Revenues in GBG Auctions and 3rd Phase of BC Auctions

	Hiking	Wine
BC (3 rd Phase)	45.92 (11.13)	59.49 (28.07)
GBG	58.93 (21.64)	62.84 (43.97)

Table 1.5: Average Bids for Most Preferred Good in GBG and Phase 1 of BC

	Hiking	Wine	iPod
BC (in phase 1 by 1 st preferred good)	54.22 (38.85)	72.06 (38.24)	63.65 (49.27)
GBG (for 1 st preferred)	82.67 (53.45)	80.72 (49.77)	85.00 (61.86)

Table 1.6: Change in Bids for BC Auctions when Most Preferred Still Available in Next Phase

	Hiking	Wine	iPod	All
Change from Phase 1 to Phase 2	5.84 (10.21)	2.38 (17.44)	2.96 (19.30)	3.29 (17.45)
Change from Phase 2 to Phase 3	2.83 (3.71)	-23.50 (21.76)	(.) (.)	-7.70 (18.72)
Both Phase Changes	4.33 (7.49)	-6.25 (22.03)	2.96 (19.30)	1.00 (18.09)

Table 1.7: Change in Bids for BC when Most Preferred Still Available by Rank in Initial Period

	Rank = 2	Rank = 3	Rank = 4	Rank = 5
Change from Phase 1 to Phase 2	-8.53 (22.60)	5.73 (13.77)	13.20 (25.95)	5.13 (10.63)
Change from Phase 2 to Phase 3	-4.00 (12.73)	2.00 (5.20)	-31.00 (.)	2.00 (2.65)
Both Phase Changes	-7.63 (20.47)	5.07 (12.63)	5.84 (29.40)	4.27 (9.09)

Table 1.8: Effect of Demographics and Personality Measures on Bids

	Dep. Var = BC Bids	Dep. Var = GBG Bids	Dep. Var = Indicator for Increased Bid
Most Preferred Wine	41.03** (8.73)	49.49** (9.75)	--
Most Preferred Hike	42.73** (14.55)	53.28** (15.29)	--
Most Preferred Ipod	36.76** (7.03)	49.72** (10.44)	--
Second Preferred Wine	14.71** (6.39)	6.49 (8.78)	--
Second Preferred Hike	14.12** (5.84)	17.80** (8.19)	--
Second Preferred Ipod	5.35 (14.03)	31.83** (9.16)	--
Wine	3.82 (8.03)	-4.11 (7.13)	--
Ipod	(omitted)	14.61** (7.66)	--
Risk	-13.64** (6.48)	-7.91 (9.12)	-0.30 (0.52)
Educ Above Average	3.22 (6.78)	7.64 (7.48)	-1.69** (0.50)
Income Above Average	6.45 (6.56)	-3.27 (9.25)	0.74 (0.45)
Iq Quiz	4.05 (3.87)	1.11 (3.37)	0.58** (0.25)
Assert	2.76 (1.89)	-2.39 (1.64)	-0.01 (0.09)
Motivation	0.07 (1.21)	-1.42 (1.56)	0.07 (0.07)
Confidence	-1.29 (1.51)	4.41** (1.99)	-0.23** (0.11)
Efficacy	0.59 (1.74)	-1.95 (1.77)	-0.17* (0.09)
Social	-1.95* (1.14)	-0.88 (1.17)	0.15** (0.07)
constant	25.91** (11.40)	35.95** (12.83)	0.99 (0.79)

One asterisk (*) indicates statistical significance at the 10% level and two asterisks (**) indicates significance at the 5% level. Standard errors are in parentheses. Robust standard errors are reported for regressions in the 1st and 2nd columns. The last column reports the results of a Probit model where the dependent variable indicates an increased bid from the prior period. The sample for this model is restricted to instances where theory suggests bids should increase.

Table 1.9: Experimental Design (Laboratory)

Treatment	Number of UTK groups	Number of UAA groups	Subjects per group	Total subjects
BC – NI – SG	4	3	6	42
BC – I – SG	4	2	6	36
BC – NI – MG	4	3	6	42
BC – I – MG	4	3	6	42
GBG – NI – SG	4	2	6	36
GBG – I – SG	4	3	6	42
GBG – NI – MG	4	3	6	42
GBG – I – MG	4	3	6	42
Total	32	22	--	324

Table 1.10: Average Revenues

BC Treatment	Average Revenue (Std. Dev.)	GBG Treatment	Average Revenue (Std. Dev.)
BC – NI – SG	152.41 (36.63)	GBG – NI – SG	105.68 (27.52)
BC – I – SG	146.03 (37.15)	GBG – I – SG	97.34 (33.18)
BC – NI – MG	229.96 (20.22)	GBG – NI – MG	211.37 (22.64)
BC – I – MG	232.20 (19.14)	GBG – I – MG	214.20 (22.60)

Table 1.11: Average Bids by Treatment

	BC – 1 st Phase (Std. Err.)	GBG – Most Preferred (Std. Err.)	Difference (Std. Err.)
Single-good Demand	48.94 (0.96)	51.69 (0.99)	2.75 (1.38)
Multi-good Demand	71.06 (0.73)	74.24 (0.64)	3.18 (0.97)

Table 1.12: BC Auctions: Average Values and Bids

	Average Highest Value (Std. Dev.)	Average Actual Bid (Std. Dev.)	Average Bid / Value (Std. Dev.)
SG: Phase 1	54.78 (27.59)	48.94 (26.85)	0.892 (0.20)
SG: Phase 2	50.78 (26.62)	46.29 (25.38)	0.916 (0.17)
SG: Phase 3	42.35 (25.96)	39.42 (24.54)	0.934 (0.14)
MG: Phase 1	75.83 (17.40)	71.06 (21.17)	0.935 (0.17)
MG: Phase 2	72.96 (17.29)	68.49 (20.87)	0.938 (0.18)
MG: Phase 3	69.53 (16.76)	66.55 (18.26)	0.959 (0.13)

The Average Highest Value column provides the average highest value remaining in the phase in question; i.e. the value for subjects whose most preferred good is still available.

Table 1.13: Good-by-Good (Benchmark) Auctions: Average Values and Bids

	Average Value for Most Preferred (Std. Dev.)	Average Bid for Most Preferred (Std. Dev.)	Average Value – Bid (Std. Dev.)	Average Bid / Value (Std. Dev.)
SG: Phase 1	54.82 (30.15)	51.10 (30.24)	3.72 (11.20)	0.927 (0.19)
SG: Phase 2	54.38 (22.56)	51.55 (23.09)	2.83 (6.81)	0.932 (0.17)
SG: Phase 3	55.14 (29.50)	52.42 (29.46)	2.72 (7.38)	0.933 (0.17)
MG: Phase 1	79.65 (14.60)	78.34 (15.02)	1.31 (4.57)	0.984 (0.06)
MG: Phase 2	70.38 (19.50)	69.09 (20.50)	1.30 (6.82)	0.980 (0.10)
MG: Phase 3	77.84 (16.18)	75.62 (18.63)	2.22 (9.93)	0.972 (0.13)

The Average Value for Most Preferred is the mean value for the most preferred good. Since Phase 1 was always an auction for good “A” in the benchmark treatments, the first cell in this table provides the average value for subjects who preferred good “A” in SG treatments. The fourth row down in the first column provides the average highest value for subjects whose most preferred good was “A”. This table does not include bids for lesser preferred goods in MG (for purposes of a clean comparison between SG and MG). Note that average values are also average theoretical bids in the GBG auctions.

Table 1.14: Estimates of Bidding Behavior

	Model 1: Single-good Demand	Model 2: Multi-good Demand
	Dependent Variable = Bid / Value	
BC	-.119** (.055)	-.130** (.053)
Round	.024*** (.006)	.014*** (.004)
Round ²	-.001*** (.000)	-.001*** (.000)
BC x Phase 2	.024*** (.008)	.002 (.009)
Risk Preference	-.004 (.005)	-.006 (.005)
BC x Risk	.012 (.008)	.013 (.008)
Variance Indicator	--	.004 (.010)
Age	.001 (.001)	-.002 (.001)
GPA	.012 (.009)	.017* (.009)
Gender	.004 (.018)	-.004 (.014)
Constant	.808*** (.060)	.950*** (.043)
N	2080	2220
F	6.20	4.52
R ²	.036	.047

This table includes bids for the first phase of BC auctions and the second phase of BC auctions when the subject's most preferred good is still available. The table includes bids for GBG auctions for each subject's most preferred good. One asterisk, two asterisks, and three asterisks indicate significance at the 10%, 5%, and 1% levels, respectively.

Appendix C

Figures




Figure 1.1: Proportions of Participants in Each Risk Preference Group




Figure 1.2: Average Rankings for Each Good in Each Auction Format (1 = Most Preferred)




Figure 1.3: Proportions of Participants Who Ranked Each Good as Their Most Preferred




Figure 1.4: Average Revenues

The standard errors for SG treatments are: 2.71 for GBG and 3.23 for BC; the standard error for the difference is 4.22. The standard errors for MG are: 1.91 for GBG and 1.66 for BC; the standard error for the difference is 2.53.




Figure 1.5: Risk Elicitation Results




Figure 1.6: Average Revenues Including Theoretical GBG Revenues

The standard errors for SG treatments are: 2.50 for Theoretical GBG and 3.23 for BC; the standard error for the difference is 2.66. The standard errors for MG are: 1.87 for Theoretical GBG and 1.66 for BC; the standard error for the difference is 1.57.

Appendix D

Instructions

This appendix includes the field instructions for the Bidder's choice Auction and the laboratory instructions for the Bidder's choice Auction with Single-good Demand and Full Information. Instructions for other treatments are available upon request.

Field Instructions

Welcome to Jonesie's Auctions. You have the opportunity today to bid in an auction where we will be selling the three bundles of goods displayed on the table in front of you. We will provide you an opportunity to examine each of the items before the bidding begins. We ask that you do not talk with any of the other participants during the session. If you have a question at any time during the session, please raise your hand and a monitor will come to your seat and answer it in private.

Description of the available goods

Good 1: I-Pod and Speakers

- 2 GB I-Pod Nano with 500 song capacity
- JBL On Stage Micro portable music dock for I-Pod

Good 2: Hiking Equipment and Backpack

- REI Ridgeline backpack
- REI Hiker First Aid Kit
- Katadyn Hiker Microfilter

Good 3: Riedel Wine Glasses and Wine

- Set of 4 Riedel Chardonnay Glasses
- One bottle of 2006 Laird Family Estate Carneros Chardonnay
- Set of 4 Riedel Pinot Noir Glasses
- One bottle of DuNah Vineyards Russian River Valley Pinot Noir
- Set of 4 Riedel Cabernet/Merlot Glasses
- One bottle of 2004 Chappallet Napa Valley Cabernet Sauvignon

There are five bidders in this auction which will consist of three phases. Rather than sell the goods one by one, we will sell 'rights to choose' one by one. If in any phase you win one of the rights to choose, you will be able to choose which of the goods remaining at that time you want. To be more precise, in each phase a 'right to choose' is sold to the highest bidder. In the first phase, all five bidders will submit a bid for the first right to choose. The highest of these five bidders wins the first right to choose and selects the good that he or she prefers. At the end of the first phase, every bidder will be informed whether they won the first right to choose and which good was selected by the winning bidder.

Once the winning bidder from the first phase has selected their preferred item, the second phase starts. In the second phase all bidders will submit a new bid for the second right to choose. The highest of these bids wins the second right to choose and selects the good that he or she prefers from amongst the two remaining items. At the end of the second phase, every bidder will be informed whether they won the second right to choose and which good was selected by the winning bidder. In the third and final phase, all bidders will submit a new bid for the remaining item. The highest bidder in the third phase will win the final item.

Auction Rules:

In each phase, you are asked to submit a bid indicating the maximum amount you are willing to pay to acquire the 'right to choose' your most preferred item from the set of available items. Bids may be submitted in intervals as fine as one cent although there is no restriction on the amounts that you can bid. If you do not place a bid, it will be counted as a bid of zero dollars. Once I have received bids from all five bidders, I will order them from highest to lowest to determine the winner in that phase. The price that the winner in each phase pays depends on the bids of the other participants in the market. To be precise, in each phase the individual that submits the highest bid will be awarded the "right to choose" their preferred item for a price equal to the second highest bid submitted for that phase. If you do not submit the highest bid, you will not win the 'right to choose' in that phase and will not be asked to pay anything.

If two (or more) individuals submit the same high bid, then one of these bidders will be randomly selected and awarded the "right to choose" for that phase. In such an instance, the winner pays a price equal to their own bid amount.

Example

If the bids in the first phase are ranked highest to lowest as follows:

\$A (bid from bidder A)
\$D (bid from bidder D)
\$E (bid from bidder E)
\$B (bid from bidder B)
\$C (bid from bidder C)

Bidder A would win the ‘right to choose’ his most preferred item from the set of three available items and would pay a price equal the amount of the bid submitted by bidder D.

After Bidder A selected his most preferred item, the bidding process would be repeated with everyone submitting a bid for the ‘right to choose’ their most preferred of the remaining two items. If the bids in the second phase are ranked highest to lowest as follows:

\$E (bid from bidder E)
\$C (bid from bidder C)
\$F (bid from bidder F)
\$B (bid from bidder B)
\$A (bid from bidder A)

Bidder E would win the ‘right to choose’ his most preferred item from the set of two available items and would pay a price equal the amount of the bid submitted by bidder C.

Once Bidder E selected her most preferred good, the bidding process would be repeated one final time with bidders submitting a bid for the final item.

Example

Before you submit your actual bids, I would like you to work through an example. Consider an auction where the following bids were submitted in the first phase. We want you to determine who will win the auction and how much they will pay to obtain the good.

Bidder 1's First Bid = 1103¥

Bidder 2's First Bid = 850¥

Bidder 3's First Bid = 1200¥

Bidder 4's First Bid = 250¥

Bidder 5's First Bid = 475¥

Take the two highest bids and order them from highest to lowest:

Highest Bid _____ 2nd Highest Bid _____

Now, determine which bidder has won the first 'right to choose' and the amount that he or she will have to pay. Fill in those numbers here.

Winning Bidder _____ Amount Paid _____

To assure that you understand how this auction mechanism operates, I will check your work after you complete this example. Please raise your hand once you have completed the example.

Final Transaction:

The winners in each phase will be required to pay me (cash or check) for the items that they have selected at the end of the session. Once I have received payment, the respective item will be awarded to the winning bidder.

I understand that you may not have anticipated the need to bring cash or your checkbook with you for this experiment. In the case that you do not have the necessary cash (or a check) to pay for the items, we will provide you with a stamped envelope in which to mail the payment. Upon receipt of your cash or check, I will send you the items that you won. All postage will be paid by Jonesie's Auctions for items mailed to the winners.

Note that while this is a real auction for the items displayed on the table in front of you, I plan to use data on the bids in this auction for economic research. I guarantee to sell all three of the items to the winners of this five-bidder auction, whatever the final auction prices turn out to be. Your bids represent binding commitments to purchase the items you win at the prices specified by the auction outcomes.

Good luck – we now invite you to spend a few minutes examining the goods on the table at the front of the room. Once you have examined the items, please return to your seats. Once everyone has been seated, we will ask you to write your bid for the first phase on the sheet provided.

Lab Instructions

Welcome to this experiment on economic decision-making! This experiment consists of 10 rounds plus 1 practice round. At the start of the session, you will be randomly assigned to a group of 6 people and you will remain in this same group for all ten rounds. Importantly, you will not know the identity of the other five participants in your group and the other participants in your group will not know your identity. You will earn tokens in the experiment by purchasing a good you value in a market. At the end of the experiment your tokens will be exchanged for dollars. Each 8 tokens is equal to 1 dollar. Your total earnings in the experiment will equal the sum of your earnings in all 10 rounds.

Values of the Goods

In each group, 3 goods will be available for sale in each round: good A, good B, and good C. Each participant will have a positive value for only one of the goods in each round. Values for this good are randomly determined and will lie between 1 and 100 tokens. That is each number between 1 and 100 is equally likely to be assigned as your value. The other goods have no value (=0 tokens) to the participant. Each participant will receive a different value for his or her preferred good. The value of this good for each participant does not depend on the values of the preferred goods for the other participants. You will have the opportunity to earn money by purchasing your preferred good at a price less than your assigned value.

At the start of each round, you will be informed of which good you prefer and how much you value it. You will not know the preferred goods or the values of the other participants and the other participants will not know your preferred good or value. Among the 5 other participants in

your group, there will be one other participant who prefers the same good as you do and his or her value is also determined randomly from the interval between 1 and 100.

Which good a participant prefers changes (randomly) from round to round. This implies that the person who prefers the same good as you will also change from round to round. Each participant will receive a new value for the preferred good in each round. The value for a preferred good in one round does not depend on the value for the preferred good in any other round.

Sale of the Goods

Rather than sell the goods one by one, the market will sell “rights to choose” one by one. If in any phase you win one of the rights to choose, you will be able to choose which of remaining goods you wish to purchase. To be more precise, in each phase a right to choose is sold to the highest bidder. In the first phase, all six bidders will submit a bid for the first right to choose. The highest of these six bidders wins the first right to choose and selects the good that he or she prefers. At the end of the first phase, every bidder will be informed whether they won the first right to choose and which good was selected by the winning bidder.

Once the winning bidder from the first phase has selected their preferred item, the second phase starts. In the second phase the remaining bidders (whose preferred goods are still unsold) will submit new bids for the second right to choose. The highest bidder wins the second right to choose and selects the good that he or she prefers from amongst the two remaining goods. At the end of the second phase, every bidder will be informed whether they won the second right to choose and which good was selected by the winning bidder. In the third and final phase, the remaining bidders (whose preferred good is still unsold) will submit new bids for the remaining good. The highest bidder in the third phase will win the final good. This process will be repeated in each of the ten rounds.

Prices of the Goods

In each phase, you will be asked to submit a bid indicating the maximum amount you are willing to pay to acquire the “right to choose” your most preferred good from the set of available goods. You may submit any number up to your value for your most preferred good. The price that the winner in each phase pays depends on the bids of the other participants in the market. To be precise, in each phase, the individual that submits the highest bid will be awarded the right to choose their preferred good for a price equal to the second-highest bid submitted for that phase. The profit to the bidder from winning will be equal to his or her value minus the price he or she pays, so $\text{profit} = (\text{value} - \text{price})$. At the end of each phase, the price paid by the winning bidder

will be announced to all six members of the group. If you do not submit the highest bid, you will not win the right to choose in that phase and you will not pay anything.

If two (or more) individuals submit the same high bid, then one of these bidders will be randomly selected and awarded the right to choose for that phase. In such an instance, the winner pays a price equal to their own bid amount.

Example

Suppose the bids in the first phase are ranked highest to lowest as follows:

\$A (bid from bidder A)
\$D (bid from bidder D)
\$E (bid from bidder E)
\$B (bid from bidder B)
\$C (bid from bidder C)
\$F (bid from bidder F)

Bidder A would win the right to choose his most preferred good from the set of three available goods and would pay a price equal the amount of the bid submitted by bidder D. After Bidder A selects his most preferred good, the bidding process would be repeated with the remaining bidders (whose preferred goods are still unsold) submitting a bid for the right to choose their most preferred of the remaining two items.

Final Payout

At the end of the experiment, your total tokens earned will be displayed on your screen. You will be asked to fill out a short, anonymous survey and then you will be paid in private. If you have any questions at any time, please raise your hand.

CHAPTER 2

A Dynamic Markov Tournament Model of Task Assignment and Up-and-Down Competition for Status

Abstract

We develop a dynamic Markov model to capture the incentives in indefinitely-repeated tournaments in labor market settings where agents compete both to “move up” as well as to avoid a “move down”. Such settings naturally arise regardless of whether explicit performance incentives or an organizational hierarchy exist. We show that when monetary incentives are available the dynamic tournament approaches the first-best outcome, but we also allow for the possibility that the principal’s only available incentive mechanism is the assignment of undesirable tasks to agents who are out-of-favor. Inability to change salaries or demote workers is common for public organizations, such as government agencies and schools. For instance, a school principal may not be able to monetarily reward or sanction teachers based on performance, but typically has discretion within the labor contract to vary class assignments and resources such as teacher’s aides. We model agents as being either in or out of favor with the principal in any given period; those who are out of favor are assigned more undesirable tasks. The prize of the tournament is the difference between groups (in favor and out of favor) in the present value of the agent’s expected utility. We assume that agents’ effort cost of completing contractible tasks is such that these costs are minimized by assigning equally burdensome tasks to all agents. Therefore the principal can motivate non-contractible effort through differential task assignment, but this entails an efficiency cost. The model demonstrates that employers may seek flexibility to vary task assignments in labor contracts not only to adapt to changing circumstances, but also to enable them to motivate non-contractible effort when agents’ compensation is fixed.

I. Introduction

Many situations are characterized by agents that compete for rank or status in an ongoing, indefinite contest. Agents who outperform their peers may “move up”, while those who underperform “move down”. The advantages of moving up as well as avoiding a move down engender competition. In the workplace employees may work hard not only to compete for promotions but also to avoid demotion, just as many sports leagues worldwide, such as the English football league, employ a system of promotion and relegation in which a fixed number of the lowest performing teams in the top league are demoted at the end of a season while the highest performing teams in the second-tier league are promoted. Even in a fairly “flat” (i.e. non-hierarchical) organization such as a school, which may have little room for inducing effort or performance through competition for rank or financial rewards, there nevertheless can be significant consequences associated with being in or out of favor with the principal, such as class assignments and the allocation of such scarce resources as teacher aides.

We develop an indefinite dynamic Markov tournament model with competition for status within the organization. An agent’s utility payoff each period depends on his status, and each period a fixed number of agents will be moved up and moved down in the organization. We explicitly model the possibility that high status may be rewarded non-monetarily through the principal’s discretion of task assignment. For example, in a school, teachers who are in favor with the principal may be assigned more desirable classes to teach or provided additional resources such as teacher’s aides. We show that when high status can be rewarded monetarily and discounting is negligible this type of infinitely repeated dynamic tournament can function as an efficient mechanism, inducing non-contractible effort without paying rents to agents, so that first-best effort is obtained. We then show how differential assignment of contractible tasks can similarly be employed to motivate non-contractible effort. However, if the agents’ cost of contractible effort is a convex function, employing differential task assignment entails an efficiency cost. The outcome is therefore second-best as the principal faces a tradeoff between

implementing incentives to induce optimal non-contractible effort and obtaining completion of necessary contractible tasks at least cost.

By demonstrating how competition for status which determines task assignments within an organization can be employed as an incentive mechanism, our model offers a new perspective on the value to an employer of flexibility over job assignments within labor contracts. Of course a firm or other principal/employer is likely to value such flexibility for many reasons, such as being better able to adapt to changing technology or market conditions. But our model illustrates that contractual flexibility that gives a manager significant discretion over employees' task assignments yields an important motivational tool that can elicit greater effort for a workforce. In this context, it is not surprising that unions may resist such flexibility, or demand compensation for it, in labor negotiations¹. Workers will recognize that if they accept a contract that permits greater discretion in their assignments, this can compel them to exert more costly effort in the competition for status that will be ongoing within the organization. Indeed, such negotiations have occurred between UPS and the Teamsters union; a clause in one of their contracts reads: "Job reassignments will be on an as-needed basis only, due to reduction or transfer of the work. Seniority will be recognized in all job reassignments" (Teamster Local 150: UPS Contract Updates). Further, the union states that one of their goals is "stronger language that strengthens the rights of... workers to bid on overtime and job assignments" (Teamsters for a Democratic Union: UPS Contract Scorecard).

Our theory represents a departure from existing dynamic models in the tournament literature which, beginning with Rosen's (1986) seminal work, have focused on elimination tournaments in which competition is "up or out" and the game has a finite number of periods

¹ Anecdotal evidence supports this: "Unions typically direct their job-description efforts toward setting defined boundaries for positions, usually wanting to define the work that employees can perform within specific job classifications" (Joinson 2001). It is also possible to think of task assignment in a broader sense that includes schedule flexibility; a particularly undesirable task could be one that needs to be completed at night or on a holiday (e.g. teaching an evening class). Zeytinolgu (2005) states, "... traditional union preference [includes] regularity of work and/or skepticism regarding flexible scheduling, which they tend to view as a risk for losing control to employers." Further, a recent contract negotiated by the teacher's union in New York highlights how the city hopes to achieve higher quality in education by providing flexibility to schools in terms of work rules and length of school day (New York Times Editorial Board 2014). This could be explained in part by the principals' ability to incentivize more effort from teachers when they have more control.

(because a hierarchy is inherently finite)². Additionally, the present work significantly complements work on public organizations (e.g. Gersbach and Keil, 2005; Heinrich and Marschke, 2010; Perry and Porter, 1982), which emphasize that a manager's ability to use monetary incentives in a public organization (such as a government agency or school) is limited. One potential non-monetary incentive is public recognition. According to Heinrich and Marschke, there is some evidence that employees in the public sector may be particularly motivated by public recognition relative to monetary compensation. Our model can be applied to this context in that certain tasks or projects may be more likely to result in public recognition and are therefore more desirable. To the best of our knowledge this literature has not considered how a manager's use of task assignments can be employed to generate competition for status and thereby motivate performance and this is therefore one of our primary contributions³.

The relevance of our model for government agencies is potentially quite large. Several pertinent aspects of government culture are characterized by Wilson (1989) in his text on bureaucracy. First, he makes clear that managers in government have very limited ability to alter pay or give promotions, but managers can "give people attractive or miserable job assignments" (p. 156) as an alternative. Further, government job descriptions are so abstract that it is very difficult for managers to justify bonuses even in the few circumstances when they do have the ability to use money⁴. Wilson also points out that in certain branches, including military, the goal is not only to use job assignment as a reward mechanism, but also to provide equal opportunities

² To the best of our knowledge, with the exception of work by Liu and Neilson (2009) and Gilpatric, Vossler and Liu (2013), both the indefiniteness of the game as well as the nature of the movement of players (i.e. both up and down) has not been considered in previous tournament models. An important feature of our model relative to Liu and Neilson (2009) is that we assume agents compete in separate tournaments according to their organizational level, rather than in a single tournament. This simplifies the model dramatically (importantly, there is an analytical solution), and significantly increases the amount of effort induced by competition. Just as crucial, given that inducements for effort generally differ across levels, the single-tournament becomes a competition that only those at a particular level are likely to win. This characterization does not appear to fit the settings we endeavor to model very adequately.

³ Gersbach and Keil (2005) characterize a public organization where the principal does not have direct monetary means to incentivize agents, similar to the current paper. However, the model developed by Gersbach and Keil focuses on incentivizing agents to reveal unproductive tasks in their departments, not put forth effort *per se*.

⁴ Although fairly rare, sometimes managers in government are allowed to allocate a pot of bonus money among their employees. Wilson cites an example from a report by the Weatherhead School of Management, where "agency heads often gave small bonuses to many people rather than large ones to a few", partially due to the fact that they were unable to justify bonuses in organizations with "vague and complex goals".

to all officers by rotating them, even though it is disadvantageous in terms of having experts in a particular area. This demonstrates that task assignments are not static in many agencies; organizations may have reasons outside our model to want to keep task assignments in flux. This could suggest that the efficiency cost of using our mechanism may not be too great. Finally, government agencies “are often prepared to accept less money with greater control than more money with less control,” according to Wilson (p. 179). “This is because of the high priority they attach to autonomy, or turf.”⁵ This mentality favors incentive mechanisms which rely on operational control, like task assignments, instead of money.

On the other hand, our model is also potentially useful for corporations. If an organization transitions from a period of rapid growth to being more “mature” and stable in size, it may move from having lots of opportunity for promotion (and thus more traditional tournament rewards) to having a much larger role played by task assignments. Similarly, an organization that has money for bonuses or raises during good times may institute a pay freeze or eliminate the bonus pool during a downturn, thus forcing the organization to rely on non-monetary incentives. It is important to note that the model presented in this paper is not meant to be an alternative to other tournament schemes, but rather, we model a distinct setting which commonly arises in public organizations and may arise in other organizations as well.

The model developed here builds on the framework of Gilpatric, Vossler and Liu (2013), which uses a dynamic Markov tournament to show that regulators benefit from placing firms in a tournament where they compete to avoid being targeted for future audits. The theory and supporting laboratory experiment demonstrate that the tournament setup achieves significant leverage over a simple random audit mechanism⁶. The dynamic tournament model that the

⁵ Wilson states that political support is key for government bureaus and “ideally, a government bureau would like to be the only organization in town curing cancer and would like to have no limitations on how it goes about achieving that cure. The typical bureau is in a much less happy state of affairs... it must [operate] under the watchful and critical eyes of countless subcommittees, interest groups, and journalists... all else being equal, big budgets are better than small. But all else is not equal.” He goes on to cite many instances of government agencies which have chosen smaller budgets and less responsibilities to maintain control over how to operate.

⁶ The authors’ theory involves firms being placed in two groups: a high-probability audit group and a low-probability audit group. Those in the targeted group compete to be transitioned to the non-targeted group, and those in the non-targeted group compete to avoid being moved to the targeted group via an indefinite Markov chain model. The regulator incorporates both peer-evaluation and the firm’s own compliance history to shape

authors use is derived from the indefinite Markov chain models employed in the dynamic regulatory enforcement literature (e.g. Landsberger and Meilijson, 1982; Greensberg, 1984; Harrington, 1988; Harford and Harrington, 1991). We use a similar tournament framework, but apply it to the labor market setting which has key differences, such as the need to satisfy agents' participation constraint.

In our model, agents are either in favor or out of favor with the principal in any given period. The organization is limited to two groups: the high group (in favor) and the low group (out of favor). A tournament occurs within each group in each period. Agents are required to complete contractible tasks and choose effort that increases output or performance in a non-contractible dimension. Non-contractible effort is unobservable to the principal. The agent(s) with the highest output or performance in the low group in each period are subsequently promoted to the high group, and the agent(s) with the lowest output in the high group in each period are subsequently demoted to be in the low group. Agents are paid a fixed fee in each period. The effort cost of completing assigned tasks in the contractible dimension is higher for those agents in the low group if differential task assignment is employed as an incentive. The prize of the tournament is the difference between groups in the present value of the agent's expected utility. Unlike other tournament models, in our model strong incentives arise from competition for status within an organization without a significant hierarchy.

An important characteristic of our model employing differential task assignment to reward high status is that in this setting more intense competition is desirable because it reduces the difference in task assignments required to motivate any level of effort. Therefore an increase in the variance of the random shocks that impact agents' performance (which dulls the competitive incentive) is costly to the principal. This finding lies in contrast to the standard Lazear and Rosen (1981) tournament framework where greater variance in random shocks can be offset by an increased prize spread to maintain the effort incentive at no cost to the principal (assuming agents are risk neutral).

enforcement effort. Laboratory experiments broadly confirm the theoretical comparative statics: targeting leads to significant enforcement leverage and the dynamic tournament exhibits strong audit cost, audit probability and transition effects.

The paper proceeds as follows. First, we present a dynamic Markov tournament model with task incentives. Next, we compare this model to a general model of a dynamic Markov tournament, demonstrating the inefficiency that results from the task model. Several propositions follow. A discussion section concludes.

II. A Dynamic Markov Tournament with Task Incentives

Our model consists of a firm or other organization whose objective is to maximize the value of expected output net of labor costs. There are two dimensions of labor effort, one which is contractible and one which is non-contractible. The contractible labor consists of tasks like administrative work, classes which must be taught, or spots which must be filled in an assembly line. The completion of these tasks is observable by the manager and verifiable by a third party and, therefore, employees who do not fulfill the task requirements can be denied compensation. Every period there are a certain number of contractible tasks that must be completed. On the other hand, non-contractible effort cannot be directly observed by the manager or verified by a third party. This effort contributes to the production of a valuable output, but just as in a standard tournament model, the output is subject to random shocks. For example, a school principal must ensure all classes are taught each semester. Simply being present and meeting certain basic requirements of teaching is a contractible task. However, the output (student learning) also depends on teacher effort in a manner that is non-contractible and the principal will therefore desire to motivate this non-contractible effort. We model how the principal can use a dynamic tournament to induce non-contractible effort from agents even if all agents receive equal payment every period.

While the manager or principal seeks to maximize net benefit, he faces three constraints. The first is that all of the contractible tasks must be completed, as mentioned above. Second, the agents must agree to enter into the contract and thus participate in the tournament. Agents are assumed to have an outside option and the present value of their expected utility from the

tournament must be greater than their reservation utility. Third, it must be incentive compatible for the agents to provide the level of non-contractible effort the principal seeks, which in our context requires that the desired effort is the symmetric Nash equilibrium of the dynamic tournament. In this section we will demonstrate that the manager can achieve this result when high status can be rewarded with less required contractible effort. However, we will see later that this method of incentivizing agents is costly and thus, the manager chooses to incentivize less than optimal effort from agents.

Our framework consists of N risk-neutral⁷ agents who are sorted into two groups: G_1 and G_2 . The agents are better off when they are in favor with the manager (group G_1). The game is infinite with discount factor δ . In each period, each agent chooses non-contractible effort μ at cost $c(\mu)$, where $c' > 0$, $c'' > 0$. The agent's output from non-contractible effort is given by y , which is the sum of effort and a random component: $y = \mu + \varepsilon$ (the random component is an *i.i.d.* draw from the distribution H across agents and periods). Further, we utilize the following additional notation:

f_1, f_2	fixed payment to an agent in a period, conditional on group assignment
n_1, n_2	number of agents in G_1 and G_2
m_1, m_2	number of agents in G_1 and G_2 selected for tournament (i.e. agents for whom the manager will review their work in a given period)
ρ_1, ρ_2	probability of being selected for tournament (i.e. probability that the agent's work will be reviewed by the manager in a given period) for agents in G_1 and G_2
τ	number of agents in each period transitioned from G_1 to G_2 and vice versa
A	value of output y (common to both groups)
e_1, e_2	contractible effort required of each agent in G_1 and G_2 in each period
E	total contractible effort required to complete all necessary tasks in each period such that $n_1 e_1 + n_2 e_2 = E$

⁷ Both Lazear and Rosen (1981) and Nalebuff and Stiglitz (1983) discuss risk averse contestants. As with other incentive mechanisms, risk aversion among agents implies only a second-best outcome can be achieved with a tournament because motivating effort is traded-off against agents' exposure to risk. This result applies in the present context as well.

$z(e)$ cost of contractible effort for the agents, where $z' > 0$, $z'' > 0$

In each period, a separate tournament occurs in each group. Specifically, $m_1 \leq n_1$ agents from G_1 (in-favor group) are selected for tournament (i.e. the manager reviews their work; this allows for the possibility that the manager may not be able to inspect the work of every agent in every period) and the τ agents with the lowest output are placed in G_2 the following period, while the other agents remain in G_1 . Similarly, $m_2 \leq n_2$ agents from G_2 (out-of-favor group) are selected for tournament and the τ agents with the highest output are placed in G_1 the following period, while the other agents remain in G_2 ($\tau < m_1$ and $\tau < m_2$). Agents choose non-contractible effort before being selected for the tournament.

The manager can motivate agents using a spread of undesirable tasks between the high (in favor) group and the low (out of favor) group. In this section, we assume that $f_1 = f_2 = \gamma$ and $e_1 < e_2$. The manager must pay all agents the same fixed wage each period, but the manager can assign tasks such that the prize of the tournament is the assignment of less contractible tasks. Therefore, the payoff to an agent in G_1 (in-favor group) in the task differentiation model is $\pi_{1,it} = \gamma - c(\mu_{it}) - z(e_1)$, and analogously for an agent in G_2 . Thus, the payoff to agents in G_1 is higher than the payoff to agents in G_2 . Let the probability that an agent in G_1 who is selected for the tournament ranks among the bottom τ agents (and therefore gets transitioned to G_2) be represented by $Q_i(\mu_i, \mu_{-i})$ and the probability that an agent in G_2 who is selected for tournament ranks among the top τ agents (and therefore gets transitioned to G_1) be represented by $R_j(\mu_j, \mu_{-j})$. In each period, two tournaments take place; they differ in that the contest in the high group is a competition to avoid ranking at the bottom while the contest in the low group is a competition to rank at the top, but this is not consequential because the two tournaments are otherwise completely symmetric. Both tournaments are standard symmetric rank-order tournaments as developed by Lazear and Rosen (1981) and Nalebuff and Stiglitz (1983).

Applying a result from Nalebuff and Stiglitz, the probability that an agent in G_1 who chooses effort μ_i when the other agents in G_1 choose μ_{-i} ranks in exactly the k th position up from the bottom (e.g. $k=1$ denotes ranking last) is the following.

$$Q_{ik}(\mu_i, \mu_{-i}) = \int \frac{(m_1 - 1)!}{(m_1 - k)! (k - 1)!} h(\varepsilon_i) (H(\varepsilon_i + \mu_i - \mu_{-i}))^{k-1} (1 - H(\varepsilon_i + \mu_i - \mu_{-i}))^{m_1 - k} d\varepsilon_i$$

The probability that i ranks among the bottom τ is then $Q_i(\mu_i, \mu_{-i}) = \sum_{k=1}^{\tau} Q_{ik}(\mu_i, \mu_{-i})$. For identifying the symmetric equilibrium of the tournament, we require the marginal effect of effort on the probability of ranking among the bottom τ be evaluated when $\mu_i = \mu_{-i}$. The marginal effect on the probability of ranking in position k is given below.

$$\begin{aligned} \frac{\partial Q_{ik}(\mu_i, \mu_{-i})}{\partial \mu_i} \Big|_{\mu_i = \mu_{-i}} &= \int \frac{(m_1 - 1)!}{(m_1 - k)! (k - 1)!} (h(\varepsilon_i))^2 \left\{ (1 - H(\varepsilon_i))^{m_1 - k - 1} (H(\varepsilon_i))^{k-2} \right\} * \\ &\quad \{ (k - 1)(1 - H(\varepsilon_i)) - (m_1 - k)H(\varepsilon_i) \} d\varepsilon_i \end{aligned}$$

The effect of effort in symmetric equilibrium on the probability of ranking among the bottom τ is then: $\frac{\partial Q_i(\mu_i, \mu_{-i})}{\partial \mu_i} \Big|_{\mu_i = \mu_{-i}} = \sum_{k=1}^{\tau} \frac{\partial Q_{ik}(\mu_i, \mu_{-i})}{\partial \mu_i} \Big|_{\mu_i = \mu_{-i}}$.

The G_2 tournament is directly analogous except that effort increases the probability of an agent's ranking among the top τ in the group. Because the random component of output is drawn from the same distribution H regardless of group, it follows that $\frac{\partial Q_i(\mu_i, \mu_{-i})}{\partial \mu_i} \Big|_{\mu_i = \mu_{-i}} = -\frac{\partial R_j(\mu_j, \mu_{-j})}{\partial \mu_j} \Big|_{\mu_j = \mu_{-j}}$ if $m_1 = m_2$.

Agents are assumed to be risk neutral and maximize the expected present value of their utility. The dynamic game follows a Markov chain process with a transition matrix representing the probabilities for an agent to transition (or not transition) out of his current group. The matrix is given in Table 2.1 in Appendix B. For example, the bottom middle cell represents the probability that agent who is in G_2 in period t will transition to G_1 in period $t + 1$.

Let V_{1t} be the expected present value to an agent of being in group 1 at time t (and analogously for group 2). Then we have:

$$V_{1t} = \pi_{1t} + \delta(1 - \rho_{1t}Q_{it})V_{1,t+1} + \delta\rho_{1t}Q_{it}V_{2,t+1}$$

$$V_{2t} = \pi_{2t} + \delta \rho_{2t} R_{jt} V_{1,t+1} + \delta (1 - \rho_{2t} R_{jt}) V_{2,t+1}$$

The expected present value of utility is the sum of utility in the current period and the discounted expected present value of utility starting from the next period, accounting for the probabilities associated with the two possible states the agent may experience in the following period. The agents maximize these expected present value utilities at any given point in time. Applying the ergodic theorem for Markov chains, the optimal strategy for an agent is stationary, i.e. conditioned only on an agent's current state (group), not on the period in the game (Kohlas, 1982; Harrington, 1988). Stationarity allows us to drop the time subscript and we impose symmetric behavior on all agents. Thus, we obtain the following first order conditions.

$$G_1: \frac{\partial \pi_i}{\partial \mu_i} = \delta (V_1 - V_2) \rho_1 \frac{\partial Q_i}{\partial \mu_i} |_{\mu_i = \mu_{-i}} \quad (1)$$

$$G_2: \frac{\partial \pi_j}{\partial \mu_j} = -\delta (V_1 - V_2) \rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}} \quad (2)$$

where:

$$(V_1 - V_2) = \frac{(\pi_1 - \pi_2)}{1 - \delta \left(1 - \left(\frac{\tau}{n_1} \right) - \left(\frac{\tau}{n_2} \right) \right)} \quad (3)$$

This set of three equations implicitly defines the equilibrium of the dynamic game entailing symmetric behavior by agents (all agents follow identical strategies conditional on their group)⁸. Note that an agent maximizing his stage-game payoff would exert no effort in this model. The equations above show the incentive arising from the dynamic game which depends on the value of $(V_1 - V_2)$. This difference is the prize at stake in both tournaments (among the agents in G_1 and among the agents in G_2), and the magnitude of the difference depends on two things: the difference in contractible effort required, $z(e_2) - z(e_1)$; and the equilibrium

⁸ The existence of pure strategy equilibrium in any rank order tournament requires sufficient variance of the random component of agents' output. This is required to make the equilibrium effort satisfy general incentive compatibility such that the effort identified by the marginal optimality conditions is not dominated by "opting out" of competition and choosing zero effort. Nalebuff and Stiglitz (1983) discuss this detail. As is standard, we assume this condition is met.

transition probabilities, which determine the “stickiness” of the states (in or out of favor)⁹. The lower the transition probabilities, the more valuable it is to be in G_1 (in-favor group) rather than G_2 . This result was identified by Gilpatric, Vossler and Liu (2013).

Recall that the principal does not have monetary prizes or other means of rewarding the agents in the “in-favor” group with the exception of task assignment. This situation is common in public organizations such as schools or government agencies. Additionally, managers in other sectors may have limitations on their ability to reward employees monetarily due to civil service rules, union contracts, or other constraints. Nevertheless, most managers do have some sort of discretion over how to allocate assignments among employees. In fact, task delegation is commonly an important aspect of the manager’s responsibility. In our framework, the monetary compensation is fixed and contractible task effort is allowed to vary between groups. The manager can reward agents with high status by assigning them less undesirable tasks. We will show that the manager or principal is able to induce non-contractible effort from agents using contractible task assignment. Consequently, labor contracts that provide the manager flexibility in contractible task assignments serve the purpose of motivating non-contractible effort in addition to other benefits that flexibility may provide (e.g. the ability to adjust to a changing environment). However, the use of task assignment as a motivational tool does cause an efficiency loss compared to the general model.

2.1 Optimizing the Dynamic Tournament Labor Contract

The manager’s problem is to maximize profit from the agents’ effort subject to the constraints; specifically, the profit is equal to the value of expected output minus total cost each period. Recall that the incentive compatibility constraint is satisfied if the desired effort is incentivized by the tournament. Further, the participation constraint will be slack for agents in

⁹ The first order conditions defining effort in each group (equations 1 and 2) show that effort is increasing in $(V_1 - V_2)$, which is effectively the prize spread. Also, we have $(\pi_1 - \pi_2) = z(e_2) - z(e_1) + [c(\mu_i) - c(\mu_j)]$; therefore, by looking at equation 3, we can see $(V_1 - V_2)$ increases with $z(e_2) - z(e_1)$. The denominator of the r.h.s. of (3) is clearly increasing with the number of players transitioned each period, τ , which therefore decreases the prize spread $(V_1 - V_2)$ and reduces equilibrium effort.

the high group when it holds for agents in the low group. We assume the outside option for an agent yields constant utility per period of $\underline{u}$. Then the relevant constraint is $V_2 \geq \frac{\underline{u}}{1-\delta}$, that is, the present value of the equilibrium payoff stream for an agent in G_2 (out-of-favor group) equals the present value of the utility stream from opting out of the tournament. Therefore, the manager's problem is as follows.

$$\max_{\gamma} NA\mu - (n_1\gamma + n_2\gamma) \quad s. t. V_2 \geq \frac{\underline{u}}{1-\delta}$$

The optimal non-contractible effort level, denoted μ^* , is defined by $c'(\mu^*) = A$. Given this, it is optimal to equate the effort incentive across groups which can be achieved by setting $n_1 = n_2$ and $m_1 = m_2$, which also implies $\rho_1 = \rho_2$. Optimal effort for agents in G_1 will occur if $\frac{\partial \pi_i}{\partial \mu_i} = -c'(\mu^*) = -A$, and analogously for agents in G_2 . Using equation 2, this is also: $\delta(V_1 - V_2)\rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j=\mu-j} = A$. Rearranging, we get:

$$(V_1 - V_2) = \frac{A}{\delta \rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j=\mu-j}} \quad (4)$$

When agents in both groups choose a common effort level (here we are showing they both choose the optimal effort), then $(\pi_1 - \pi_2) = z(e_2) - z(e_1)$. We can now solve for the required contractible effort spread that achieves optimal effort by substituting equation 3 into equation 4 and noting that $1 - \left(\frac{\tau}{n_1}\right) - \left(\frac{\tau}{n_2}\right) = 1 - \frac{4\tau}{N}$ when $n_1 = n_2$.

$$z(e_2) - z(e_1) = \frac{c'(\mu)}{\delta \rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j=\mu-j}} \left(1 - \delta \left(1 - \frac{4\tau}{N}\right)\right) \quad (5)$$

Now, we solve for the participation constraint, which will be slack for agents in the high group when it holds for agents in the low group. The participation constraint is:

$$\gamma = \underline{u} + c(\mu) + z(e_2) - \frac{2\tau}{N} * \frac{c'(\mu)}{\rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \quad (6)$$

This constraint shows that ensuring agents do not opt out of the tournament when they are out of favor requires that the per-period fixed payment, γ , be sufficient to compensate them for foregoing their outside option and for the cost of contractible and non-contractible effort, less the present value of expected future rents that the agent expects to obtain from future periods of being in favor.

Finally, the total cost¹⁰ of compensation in each period is $TC_T = \gamma N$. We can substitute equation 6 into total cost to get:

$$TC_T = N\{\underline{u} + c(\mu) + z(e_2)\} - 2\tau * \frac{c'(\mu)}{\rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \quad (7)$$

We can rearrange equation 5 to get an expression for the last term in 7:

$$2\tau * \frac{c'(\mu)}{\rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} = N \left\{ \frac{1}{2} z(e_2) - \frac{1}{2} z(e_1) - \frac{1}{2} (1 - \delta) \left(\frac{c'(\mu)}{\delta \rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \right) \right\}$$

Substituting this into the cost function, we have:

$$TC_T = N\{\underline{u} + c(\mu) + z(e_2)\} - N \left\{ \frac{1}{2} z(e_2) - \frac{1}{2} z(e_1) - \frac{1}{2} (1 - \delta) \left(\frac{c'(\mu)}{\delta \rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \right) \right\}$$

In the limit, as discounting becomes negligible, we find that

$$\lim_{\delta \rightarrow 1} TC_T = N \left\{ \underline{u} + c(\mu) + \frac{z(e_1) + z(e_2)}{2} \right\}$$

¹⁰ Note that the “T” subscript on total cost is for “Task Model”; this is needed to differentiate from total cost in the general model presented in the next section.

The principal chooses a non-contractible effort level to induce, $\hat{\mu}$, and maximizes the expected difference between total benefit and total cost. (Recall that output, y , is the sum of non-contractible effort, μ , and an error term, ε ; in expectation, y is simply μ .) The principal's maximization problem is:

$$\max_{\hat{\mu}}(TB - TC) = NA\hat{\mu} - N \left\{ \underline{u} + c(\mu) + \frac{z(e_1) + z(e_2)}{2} \right\}$$

The first order condition yields:

$$A = c'(\hat{\mu}) + \frac{\partial \left[\frac{z(e_1) + z(e_2)}{2} \right]}{\partial \hat{\mu}}$$

Note that the average cost of contractible effort assigned increases as non-contractible effort increases. That is: $\partial \left[\frac{z(e_1) + z(e_2)}{2} \right] / \partial \hat{\mu} > 0$. (In order to induce higher non-contractible effort from agents, the principal must increase the spread between e_1 and e_2 .) This causes the average cost of completing these tasks to increase because $z'' > 0$. Thus, $A > c'(\hat{\mu})$.

III. Comparison to a General Dynamic Markov Tournament

In this section, we present a general model of a dynamic Markov tournament for comparison purposes. In the general model, the manager can use monetary incentives to motivate agents. We now assume that $f_1 > f_2$. Further, we make the assumption that the manager evenly distributes the contractible tasks among all N agents (i.e. $e_1 = e_2 = E/N$). Note that the least-cost way for the manager to get all E contractible tasks completed is to distribute them evenly due to $z'' > 0$. Therefore, the payoff to an agent in G_1 (in-favor group) in the general model is $\pi_{1,it} = f_1 - c(\mu_{it}) - z(E/N)$, and analogously for an agent in G_2 .

When agents in both groups choose a common effort level, then $(\pi_1 - \pi_2) = (f_1 - f_2)$. For a given non-contractible effort level, μ , that the firm wishes to motivate, equation 5 from the task model now becomes:

$$(f_1 - f_2) = \frac{c'(\mu)}{\delta \rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \left(1 - \delta \left(1 - \frac{4\tau}{N} \right) \right) \quad (5b)$$

This is the fixed payoff spread that achieves optimal effort. Similarly, the participation constraint for this model is also a function of contractible effort.

$$f_2 = \underline{u} + c(\mu^*) + z \left(\frac{E}{N} \right) - \frac{2\tau}{N} * \frac{A}{\rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \quad (6b)$$

This constraint shows that ensuring agents do not opt out of the tournament when they are out of favor requires that the per-period fixed payoff f_2 be sufficient to compensate them for forgoing their outside option and for the cost of contractible and non-contractible effort, less the present value of expected future rents that the agent expects to obtain from future periods of being in favor.

Finally, we can find the total cost of compensation paid each period¹¹, which is $TC_G = n_1 f_1 + n_2 f_2 = N f_2 + n_1 (f_1 - f_2)$. Substituting in equations 5b and 6b, we get:

$$TC_G = N \left\{ \underline{u} + c(\mu^*) + z \left(\frac{E}{N} \right) - \frac{2\tau}{N} * \frac{A}{\rho \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \right\} + \frac{N}{2} \left\{ \frac{A}{\delta \rho_2 \frac{\partial R_j}{\partial \mu_j} |_{\mu_j = \mu_{-j}}} \left(1 - \delta \left(1 - \frac{4\tau}{N} \right) \right) \right\} \quad (7b)$$

In the limit, as discounting becomes negligible, we find that

¹¹ Note that the “G” subscript on total cost is for “General Model”; this is needed to differentiate from total cost in the task model.

$$\lim_{\delta \rightarrow 1} TC_G = N \left\{ \underline{u} + c(\mu^*) + z\left(\frac{E}{N}\right) \right\}$$

The principal chooses a non-contractible effort level to induce, $\hat{\mu}$, and maximizes the expected difference between total benefit and total cost¹². (Recall that output, y , is the sum of non-contractible effort, μ , and an error term, ε ; in expectation, y is simply μ .) In the general model, we will show that the principal chooses to induce μ^* .

$$\max_{\hat{\mu}} (TB - TC) = NA\hat{\mu} - N \left\{ \underline{u} + c(\hat{\mu}) + z\left(\frac{E}{N}\right) \right\}$$

The first order condition can be solved to show $c'(\hat{\mu}) = A$; therefore, $\hat{\mu} = \mu^*$.

PROPOSITION 1: As discounting becomes negligible (δ approaches 1) the general dynamic tournament yields optimal effort, μ^* , such that $c'(\mu^*) = A$.

3.1 Comparison of Models

Finally, we can compare the total cost from the general model to the total cost from the task model. Due to $z'' > 0$, $[z(e_1) + z(e_2)]/2 > z[(e_1 + e_2)/2]$. Recall that $n_1 = n_2 = (1/2)N$ and $n_1e_1 + n_2e_2 = E$. Thus, $z[(e_1 + e_2)/2] = z[E/N]$ and:

$$\frac{z(e_1) + z(e_2)}{2} > z\left(\frac{E}{N}\right)$$

¹² This could be approached in different ways; for instance, by considering a dynamic Markov tournament employed by competitive firms, analogous to the analysis of Lazear and Rosen (1981). In that case, competition for labor bids up agents' payoffs until the value of expected output equals costs each period. Note that in this case the participation constraint does not bind; rather, the fixed payments must maximize agents' present value expected utility subject to the zero profit constraint. Of course the expected utility of agents differs depending on their group, but as δ approaches 1 the solutions converge. Alternatively, we can suppose that the first-period fixed payoff to agents randomly placed in G_1 is $f_2 - (V_1 - V_2)$ in order to equate the expected utility of agents across groups at the start of the game. In this case the same f_1, f_2 pair maximize both V_2 and the present value of expected utility for agents in G_1 at the start of the game. In any case, the result is the same.

Therefore, $TC_T > TC_G$. Intuitively, the manager in the task differentiation model faces a tradeoff between implementing the appropriate incentives to induce optimal non-contractible effort and obtaining completion of the contractible tasks at least cost.

Recall that the principal chooses a non-contractible effort level to induce, $\hat{\mu}$, and maximizes the expected difference between total benefit and total cost. Further recall from the general model that the principal chooses to induce μ^* . The principal maximizes:

$$\max_{\hat{\mu}}(TB - TC) = NA\hat{\mu} - N\left\{\underline{u} + c(\hat{\mu}) + z\left(\frac{E}{N}\right)\right\}$$

The first order condition can be solved to show $c'(\hat{\mu}) = A$; therefore, $\hat{\mu} = \mu^*$. In the task differentiation model, on the other hand, the principal's maximization problem is:

$$\max_{\hat{\mu}}(TB - TC) = NA\hat{\mu} - N\left\{\underline{u} + c(\hat{\mu}) + \frac{z(e_1) + z(e_2)}{2}\right\}$$

The first order condition yields:

$$A = c'(\hat{\mu}) + \frac{\partial \left[\frac{z(e_1) + z(e_2)}{2} \right]}{\partial \hat{\mu}}$$

Note that the average cost of contractible effort assigned increases as non-contractible effort increases, that is, $\partial \left[\frac{z(e_1) + z(e_2)}{2} \right] / \partial \hat{\mu} > 0$. In order to induce higher non-contractible effort from agents, the principal must increase the spread between e_1 and e_2 . This causes the average cost of completing these tasks to increase because $z'' > 0$. We can conclude that $\hat{\mu}_G > \hat{\mu}_T$, or $\mu^* > \hat{\mu}_T$. Constrained by the lack of monetary incentives, the principal chooses to induce less non-contractible effort in the task differentiation model because he faces an additional cost of increasing the payoff spread.

PROPOSITION 2: When monetary incentives are not available, the manager can use the assignment of tasks to motivate optimal effort. However, an inefficiency is created: non-contractible effort will be less than μ^* , i.e. the effort incentivized is “second-best”.

IV. Extensions

4.1 Error Variance

An important result of Lazear and Rosen (1981) is that the prize spread required for a given effort level, $\hat{\mu}$, increases with the error variance. In the Lazear and Rosen (1981) framework, and in the general dynamic tournament model presented here, increasing the spread of payoffs is not costly (i.e. the tournament manager does not need to pay extra compensation to agents as a result of increasing $(f_1 - f_2)$, holding total payment, $(f_1 + f_2)$, constant). However, in the task differentiation model, increasing the spread between $z(e_2)$ and $z(e_1)$ is costly due to the convexity of the cost of contractible effort function and causes inefficiency. O’Keeffe, Viscusi and Zeckhauser (1984) argue that the error variance in a tournament can reflect a variety of phenomena such as uncertainty in environmental factors or the precision with which a principal monitors his agents. If it is the case that monitoring precision partially explains the error, then in the task differentiation model, the manager’s ability to precisely evaluate his employees is important for efficiency¹³. This result is in contrast to the standard tournament literature.

¹³ Note that the existence of pure strategy equilibrium in any rank order tournament requires sufficient variance of the random component of agents’ output. This is required to make the equilibrium effort satisfy general incentive compatibility such that the effort identified by the marginal optimality conditions is not dominated by “opting out” of competition and choosing zero effort. Nalebuff and Stiglitz (1983) discuss this detail. As is standard, we assume this condition is met.

PROPOSITION 3: In the task differentiation model, an increase in the variance of ε increases the spread of payoffs, $z(e_2) - z(e_1)$, required to motivate any effort level. This is costly for the principal and decreases the elicited effort, $\hat{\mu}$.

The task model we present suggests that constraining managers to use non-monetary incentives is sub-optimal. Thus, an inefficiency is created in many public organizations, such as government agencies and schools. However, this inefficiency may not be as bad as it seems. Inefficiency in tournament settings is common in other circumstances as well. A different line of literature which examines “limited liability” in tournaments (e.g. Krakel and Schottner, 2012) has a similar effect in terms of error variance, but for a different reason. Under limited liability, increasing the prize spread is costly because workers cannot earn negative payoffs; consequently, the prize for the loser is bounded at zero. Increasing the prize spread beyond this point causes the total prize payout to increase. Thus, an increase in the error variance – which requires an increase in the prize spread to maintain a given effort level – results in inefficiency. As previously stated, this result parallels our findings, but the source of the inefficiency differs. In our model, prizes are not bounded, but the cost of contractible effort is a convex function. Therefore, the inefficiency results from not assigning the contractible tasks evenly among the employees.

4.2 Transition Probability

Similarly to the error variance, the number of agents transitioned between groups in each period, τ , also has an effect on the manager’s cost. Recall equation 3:

$$(V_1 - V_2) = \frac{(\pi_1 - \pi_2)}{1 - \delta \left(1 - \left(\frac{\tau}{n_1} \right) - \left(\frac{\tau}{n_2} \right) \right)} \quad (3)$$

It is easy to see that increasing τ decreases the spread in the present value of utility, $(V_1 - V_2)$, ceteris paribus, and thus decreases equilibrium effort. In order to maintain a given effort level, $\hat{\mu}$, the manager will have to increase the spread of the fixed payoffs, $(f_1 - f_2)$ (or $z(e_2) - z(e_1)$).

In the general base model, increasing the spread of payoffs has no consequence in terms of the cost to the manager. In the task differentiation model, however, increasing $z(e_2) - z(e_1)$ is costly and increases the inefficiency of the model. Therefore, it is optimal for the manager to keep the number of agents transitioned as small as possible.

PROPOSITION 4: In the task differentiation model, the optimal number of agents transitioned from G_1 to G_2 and vice versa is $\tau = 1$.

V. Conclusion

Tournaments are frequently used to model labor market settings due to the hierarchal nature of employment; typically, there is a clear potential for promotion and, sometimes, demotion. We demonstrate that a tournament can also be used in non-hierarchal situations where many employees are at a similar rank, promotion does not play a significant role in employees' motivation, and immediate monetary incentives are not readily available (e.g. employees in a government agency, teachers in a public school or associate professors in a university). While the manager in such an organization may not be able to set up a monetary incentive program, he can still induce optimal non-contractible effort by appropriately assigning contractible tasks. The payoff spread is determined by the difference in contractible tasks for in-favor employees and out-of-favor employees. Thus, the prize spread is the difference (between the in-favor group and the out-of-favor group) in the present value of the employees' expected utility. However, it is more costly to the manager to induce optimal effort in this framework compared to a tournament with monetary incentives. Therefore, suboptimal effort results and the solution is "second-best".

Since increasing the payoff spread is costly in this task differentiation model, an increase in the error variance or an increase in the number of agents transitioned between groups contributes to the inefficiency. In both cases, the incentives underlying the tournament are

dulled, requiring the manager to increase the payoff spread to maintain a given effort level. A different line of literature which examines “limited liability” in tournaments (e.g. Krakel and Schottner, 2012) has a similar effect in terms of error variance, but for a different reason. As previously stated, this puts our model’s inefficiency into context; inefficiencies can arise for other reasons and thus, the task delegation setup we present is not necessarily a worse situation than using monetary incentives.

A possible extension involves different values of output between the two groups. For example, a principal may value non-contractible effort from teachers more highly in honors classes (a possible task “reward” for the in-favor group) because the parents of honors students are more demanding, or he may value non-contractible effort more highly in lower-level classes (a possible task “punishment” for the out-of-favor group) due minimum standard testing that the school must pass¹⁴. Regardless of which output is valued more highly (that of the in-favor agents or the out-of-favor agents), the model can be solved to show that the manager will induce different levels of non-contractible effort to reflect the difference in output values. In the general model, the manager will induce effort exactly such that the marginal cost of effort in each group is equal to the value of output for each group. In the task model, on the other hand, the non-contractible effort induced will depend on several factors, including the proportion of agents who are selected for tournament in each period. This suggests that the manager can influence the effort induced to appropriately reflect output values by varying the intensity with which he evaluates agents in each group. This extension is plausible in that it reflects a problem that managers are likely to face and further work in this area could lead to interesting results.

Our framework has important roots in the literature on public organizations. Several papers have noted that monetary incentives are rarely used in the public sector in comparison to the private sector. This may be partially due to the fact that there are insufficient funds for large bonuses and partially due to the unpopularity of rewarding employees in public service with cash payments (Heinrich and Marschke 2010). Further, it is often the nature of public service,

¹⁴ Similarly, in a government organization, the manager could place a high value on non-contractible effort put into a report that gets publicized even if this report is tedious to produce and is therefore the task “punishment” for the out-of-favor group.

especially in the case of education, that the hierarchy of the organization is fairly flat, at least for certain groups of employees. This does not mean, however, that managers have no opportunities to incentivize employees using a tournament. On the contrary, managers are frequently responsible for assigning courses, administrative tasks and other contractible responsibilities. While some programs may be very transparent, such as an “employee of the month” scheme, most of the managers’ decisions are probably based on which employees are “in favor” and which employees are “out of favor” in the current time period. Further, employees are likely to fall out of favor, or be “promoted” to being in favor, with the manager from time to time based on past performance. A dynamic Markov tournament is therefore a mechanism that may enable an organization without explicit performance incentives or hierarchy to achieve efficient labor outcomes.

Further, our model offers a new perspective on the value to managers of flexibility over job assignments. Of course flexibility is important for a firm or any organization in that it allows adaptation to changing technology or market conditions without having to re-negotiate labor contracts¹⁵. However, we show that there may be an additional advantage of flexibility in that it can be used as an incentive mechanism. A manager who can use contractible task assignments to reward agents in favor and punish those out of favor has the opportunity to motivate non-contractible effort by engaging employees in a competitive tournament for status.

¹⁵ Wright and Snell (1998) note that flexibility in workforce can be important for a firm’s ability to meet the needs of a dynamic environment.

References

- Altmann, S., A. Falk, and M. Wibral. "Promotions and Incentives: The Case of Multistage Elimination Tournaments." *Journal of Labor Economics* 30-1 (2012).
- Casas-Arce, P. and F. Martinez-Jerez. "Relative Performance Compensation, Contests, and Dynamic Incentives." *Management Science* 55-8 (2009).
- Dewatripont, M. "The Economics of Career Concerns, Part I." *Review of Economic Studies* (1999a).
- Dewatripont, M. "The Economics of Career Concerns, Part II." *Review of Economic Studies* (1999b).
- Editorial Board. "A Teachers' Contract for New York." New York Times (2014). Accessed July 2014: <http://www.nytimes.com/2014/05/03/opinion/a-teachers-contract-for-new-york.html>.
- Falk, A., E. Fehr, and D. Huffman. "The Power and Limits of Tournament Incentives." Working Paper (2008).
- Fershtman, C., and U. Gneezy. "The Tradeoff between Performance and Quitting in High Power Tournaments." *Journal of the European Economic Association* 9-2 (2011).
- Gersbach, H. and M. Keil. "Productivity Improvements in Public Organizations." *The Economic Journal* 115 (2005).
- Gibbons, R. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *Journal of Political Economy* (1992).
- Gilpatric, S., C. Vossler, and L. Liu. "Using Competition to Stimulate Regulatory Compliance: A Tournament-Based Dynamic Targeting Mechanism." Working Paper (2012).
- Harrington, W. "Enforcement Leverage when Penalties are Restricted." *Journal of Public Economics* 37 (1988).
- Heinrich, C. and G. Marschke. "Incentives and Their Dynamics in Public Sector Performance Management Systems." *Journal of Policy Analysis and Management* 29-1 (2010).
- Holmstrom, B. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* (1982).
- Joinson, Carla. "Refocusing Job Descriptions." HR Magazine 46-1 (2001).

- Kohlas, J. "Stochastic Methods of Operations Research." *Cambridge University Press, New York* (1982).
- Krakerl, M. and A. Schottner. "Internal Labor Markets and Worker Rents." *Journal of Economic Behavior and Organization* 84-2 (2012).
- Lazear, E. and S. Rosen. "Rank Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89 (1981).
- Nalebuff, B. and J. Stiglitz. "Prizes and Incentives: Towards a General Theory of Compensation and Competition." *The Bell Journal of Economics* 14 (1983).
- O'Keefe, M., W.K. Viscusi, and R. Zeckhauser. "Economic Contests: Comparative Reward Schemes." *Journal of Labor Economics* 2-1 (1984).
- Perry, J. and L. Porter. "Factors Affecting the Context for Motivation in Public Organizations." *Academy of Management Review* 7-1 (1982).
- Teamsters for a Democratic Union: UPS Contract Scorecard. Accessed July 2014:
<http://www.tdu.org/news/ups-contract-scorecard>.
- Teamster Local 150: UPS Contract Updates. Accessed July 2014:
http://www.teamsters150.org/index.cfm?zone=/docs/UPS_Contract_Updates.htm.
- Wilson, J. Q. Bureaucracy: What Government Agencies Do and Why They Do It. Basic Books, Inc. (1989).
- Wright, P. M. and S. A. Snell. "Toward a Unifying Framework for Exploring Fit and Flexibility in Strategic Human Resource Management." *The Academy of Management Review* 23-4 (1998).
- Zeytinolgu, I. U. Flexibility in Workplaces: Effects on Workers, Work Environment and the Unions. Geneva: IIRA/ILO (2005).

Appendix A

Theoretical Work

Proof of Proposition 1:

The optimization problem is given in (8')

$$\max_{f_1} V_2 = \frac{1}{1-\delta} \left\{ 2A\mu - f_1 - c(\mu) + \frac{4\delta\tau}{N} \left[\frac{f_1 - A\mu}{1-\delta\left(1-\frac{4\tau}{N}\right)} \right] \right\}$$

The first order condition for a maximum is

$$2A \frac{\partial \mu}{\partial f_1} - 1 - c'(\mu) \frac{\partial \mu}{\partial f_1} + \frac{4\delta\tau}{N} \left[\frac{1-A \frac{\partial \mu}{\partial f_1}}{1-\delta\left(1-\frac{4\tau}{N}\right)} \right] = 0.$$

As δ approaches 1 this becomes

$$\lim_{\delta \rightarrow 1} 2A \frac{\partial \mu}{\partial f_1} - 1 - c'(\mu) \frac{\partial \mu}{\partial f_1} + \frac{4\delta\tau}{N} \left[\frac{1-A \frac{\partial \mu}{\partial f_1}}{1-\delta\left(1-\frac{4\tau}{N}\right)} \right] = 2A \frac{\partial \mu}{\partial f_1} - 1 - c'(\mu) \frac{\partial \mu}{\partial f_1} + \frac{4\tau}{N} \left[\frac{1-A \frac{\partial \mu}{\partial f_1}}{\frac{4\tau}{N}} \right] = 0,$$

Which simplifies to $\frac{\partial \mu}{\partial f_1} [A - c'(\mu)] = 0$.

Proof of Proposition 2:

The firm's optimization problem is:

$$\max_{\hat{\mu}} (TB - TC) = NA\hat{\mu} - N \left\{ \underline{u} + c(\mu^*) + \frac{z(e_1) + z(e_2)}{2} \right\}$$

The first order condition yields:

$$A = c'(\hat{\mu}) + \frac{\partial \left[\frac{z(e_1) + z(e_2)}{2} \right]}{\partial \hat{\mu}}$$

In order to induce higher non-contractible effort from agents, the principal must increase the spread between e_1 and e_2 ; this causes the average cost of contractible effort to increase because $z'' > 0$. Thus, $\partial \left[\frac{z(e_1) + z(e_2)}{2} \right] / \partial \hat{\mu} > 0$ and $c'(\hat{\mu}) < A$.

Proof of Proposition 3:

Recall that $R_j(\mu_j, \mu_{-j})$ is the probability that an agent in G_2 who is selected for tournament ranks among the top τ agents and therefore gets transitioned to G_1 . This probability is a function of the error term, ε , as is $\frac{\partial R_j(\mu_j, \mu_{-j})}{\partial \mu_j}$. An increase in non-contractible effort increases the probability of ranking in the top τ agents; thus, $\frac{\partial R_j(\mu_j, \mu_{-j})}{\partial \mu_j} > 0$.

$$\begin{aligned} \frac{\partial R_{jk}(\mu_j, \mu_{-j})}{\partial \mu_j} \Big|_{\mu_j = \mu_{-j}} &= - \int \frac{(m_1 - 1)!}{(m_1 - j)! (k - 1)!} (h(\varepsilon_i))^2 \left\{ (1 - H(\varepsilon_i))^{m_1 - k - 1} (H(\varepsilon_i))^{k - 2} \right\} * \\ &\quad \{ (k - 1)(1 - H(\varepsilon_i)) - (m_1 - k)H(\varepsilon_i) \} d\varepsilon_i \quad \text{if } m_1 = m_2 \end{aligned}$$

When the error variance increases, then $\frac{\partial R_j(\mu_j, \mu_{-j})}{\partial \mu_j}$ decreases. Recall equation 5:

$$z(e_2) - z(e_1) = \frac{A}{\delta \rho_2 \frac{\partial R_j}{\partial \mu_j} \Big|_{\mu_j = \mu_{-j}}} \left(1 - \delta \left(1 - \frac{4\tau}{N} \right) \right)$$

As the denominator of this equation decreases, the prize spread, $z(e_2) - z(e_1)$, required to achieve any effort level increases.

Proof of Proposition 4:

The difference in expected utilities, $(V_1 - V_2)$, is the effective prize spread of the tournament. Recall equation 3:

$$(V_1 - V_2) = \frac{(\pi_1 - \pi_2)}{1 - \delta \left(1 - \left(\frac{\tau}{n_1} \right) - \left(\frac{\tau}{n_2} \right) \right)}$$

Clearly, $\partial(V_1 - V_2) / \partial \tau > 0$. An increase in the prize spread required for a given effort level further requires an increase in the payoff spread, $z(e_2) - z(e_1)$. Recall the total cost in the task differentiation model:

$$\lim_{\delta \rightarrow 1} TC_T = N \left\{ \underline{u} + c(\mu) + \frac{z(e_1) + z(e_2)}{2} \right\} = N \left\{ \underline{u} + c(\mu) + \frac{[z(e_1) + z(e_2) - z(e_1) + z(e_1)]}{2} \right\}$$

$$\lim_{\delta \rightarrow 1} TC_T = N \left\{ \underline{u} + c(\mu) + z(e_1) + \frac{[z(e_2) - z(e_1)]}{2} \right\}$$

Clearly, $\partial TC_T / \partial [z(e_2) - z(e_1)] > 0$.

Thus, $\partial TC_T / \partial \tau > 0$ and the manager is best off minimizing τ .

Appendix B

Tables

Table 2.1: Transition Probabilities

From Group ↓ To Group →	$\mathbf{G}_1$	$\mathbf{G}_2$
$\mathbf{G}_1$	$1 - \rho_1 Q_i$	$\rho_1 Q_i$
$\mathbf{G}_2$	$\rho_2 R_j$	$1 - \rho_2 R_j$

CHAPTER 3

Willingness to Pay for Goods with Unregulated and Potentially Misleading Labels: the Case of “Natural”-Labelled Groceries

Abstract

Food labeling has been widely studied, especially in the context of consumer willingness to pay for features that are considered healthy, such as organic content. In this study, we provide insight to the demand effects for an unregulated phrase found on many labels: “natural”. A plethora of currently pending lawsuits regarding this phrase demonstrates that research is needed to better understand consumer misconceptions. In an experimental setting, we use an incentive-compatible approach to elicit the willingness to pay of grocery shoppers for “natural”-labelled food products, several of which contain genetically modified organisms (GMOs). We find, on average, that there is an overall null effect of the “natural” label. However, when the sample is segregated based on the belief that “natural” means GMO-free, there is a positive “natural” premium for those who hold the belief and a negative premium for those who do not hold the belief. Additionally, we find evidence of framing effects which suggest that between-subject analysis is more reliable for this type of research than within-subject analysis. Our results have implications in both the public policy and legal arenas.

I. Introduction

The phrases “all natural”, “natural”, “100% natural”, etc. have been under fire recently. Consumers claim that the phrases are misleading when the associated products contain artificial ingredients, preservatives, and/or genetically modified organisms (GMOs). Many lawsuits have been considered and several are currently pending¹. These lawsuits have included popular products such as Wesson cooking oil, Campbell’s soup, Kix cereal, Truvia sweetener, SunChips, Tostitos chips, Goldfish crackers, Ben and Jerry’s ice cream, Naked juice and many others. Although several judges have asked the U.S. Food and Drug Administration to officially define and regulate the phrases (Frankel 2013), the FDA has prioritized other projects, leaving the decision of whether to award damages to the courts. In a recent email to USA Today, the FDA stated, “Defining ‘natural’ represents additional challenges when food has been processed and is no longer the product of the earth. Additionally, there are differing perspectives on how specific such a label should be” (Weise 2014). The FDA initially decided not to define the phrase in the 1990s; since then, manufacturers have been able to use the phrase free of any regulation². However, recent turmoil has caused some manufacturers to remove the “natural” label from their products³. In particular, the issue of GMOs in food labelled as “natural” has received intense

¹ Several lawsuits have been dismissed by judges and others have been settled out of court; so far, none have gone to trial, but there are still many cases waiting “in the pipeline” (Smith 2014). In addition to lawsuits specifically regarding the “natural” label (see Allen et al. 2013, Frankel 2013, and Smith 2014), there are a plethora of GMO labelling lawsuits, some of which also address “natural”; as of July 2013, 37 GMO labeling cases had been introduced in 21 U.S. states and several legislators had introduced a new potential federal law called the “Genetically Engineered Food Right to Know Act”, which, if passed, would require the labeling of all GM foods (see the Center for Food Safety and Cummins 2014).

² Although the “natural” term is not actively regulated, the FDA has occasionally sent warning letters to manufacturers who used a clearly synthetic ingredient in their “natural” product. However, these letters often go ignored according to a report by The Center for Science in the Public Interest (Silverglade and Heller 2010). The same report explicitly points out many natural claims that may be considered “deceptive” but have been allowed to remain in the marketplace.

³ In 2009, 30% of food products and 45% of beverages introduced in the U.S. were reportedly labelled “natural”; in 2013, only 22% of food products and 34% of beverages held the label (Smith 2014).

media coverage as the U.S. remains one of the only developed countries to not specifically require GMO labelling⁴.

To help inform the legal and policy debates, this study uses economics experiments with seasoned grocery shopper participants to determine whether and why there are demand effects generated by the use of “natural” labels. Using an incentive-compatible approach to elicit willingness to pay (WTP) for several grocery items with and without “natural” labels, we find that consumers who believe “natural” means “no genetically modified organisms” (42% of our sample) are willing to pay a premium for “natural” food, whereas consumers who do not have this belief actually exhibit a negative premium. The overall effect is near zero, although the identified heterogeneity suggests that “natural” labels are potentially misleading and further that there is potential for firms to exploit uninformed consumers.

Food labeling has been widely studied, especially in the context of consumer WTP for features that are considered healthy, such as organic certification. Additionally, most labels are highly regulated by the government; for instance, the phrase “low fat” cannot be used for foods with more than 3 grams of fat per serving (serving sizes are also regulated). Especially for labels indicating low environmental impact, most of the theoretical literature acknowledges that there is some level of fraud in the market for regulated labels, but the effects of an unregulated phrase on consumer demand are unclear empirically. Importantly, food labeling has increasingly become an environmental issue rather than simply a consumer health preference issue, especially in the case of genetically modified ingredients.

Due to pressure from activists and the media, many European retailers have decided not to carry GMO-containing foods. Conversely, GMOs are widely used in the U.S. One estimate by the Center for Food Safety indicated that possibly more than 70% of processed foods in U.S. supermarkets contain GMOs – specifically 85% of corn and 91% of soybeans, which are both very common ingredients in processed foods. While some consumers and organizations are concerned about GMOs in food, others are more optimistic. GM foods have the potential to reduce costs and reduce the need for chemical pesticides. Cost reductions could allow a

⁴ According the Center for Food Safety, 64 countries have mandatory labeling of GM foods; the U.S. is one of the few remaining developed countries without GM label regulations.

considerable increase in food production, helping poor countries to address hunger concerns. Further, genetic enhancement technology for plants and animals may be able to create healthier versions of food with added vitamins and less saturated fat. The U.S. FDA states that GM foods are, in general, just “as nutritious as foods from comparable traditionally bred plants”. On the other hand, the Center for Food Safety (a non-profit organization) advocates for GM food labeling laws. In addition, the environmental effects of GMOs are largely unknown. Introduction of these new species into existing ecosystems could potentially cause irreversible damage. Regulating the location and spread of these species could be very costly, if not impossible.

As awareness and demand for GMO-free products begins to heat up the U.S., research is needed to determine how U.S. consumers respond to GM ingredients. In particular, research regarding GMOs and the “natural” label is, to our knowledge, non-existent; yet, the plethora of pending lawsuits demonstrates that the issue requires attention. The Center for Science in the Public Interest urged the FDA and USDA to define “natural” and strictly enforce standards, while others feel that the natural label should be eliminated altogether, especially since the term “organic” is highly regulated and already serves the purpose of signaling fresh, non-GMO products to consumers (Plumer 2014). Existing literature largely focuses on the European Union, where the debate over GM food started more than a decade ago. Prior work has also focused specifically on willingness to pay for food with GM-specific labels⁵. This paper differs in several important ways.

We present an experimental elicitation of willingness to pay for food with and without “natural” labels. Our hybrid approach answers two main questions: a survey addresses consumer beliefs surrounding the phrase “natural” (e.g. Do consumers believe that the foods labeled as “natural” contain GM ingredients?) and an auction-style experiment elicits willingness to pay for “natural” food. Further, we explore other dimensions, such as beliefs surrounding health and

⁵ Lusk et al. (2005) provide a useful meta-analysis of 25 GM food valuation studies that report a total of 57 valuations for GM food. They find that, on average, consumers place a 23% premium on non-GM food. Notably, European consumers placed a significantly higher premium on GM food than U.S. consumers, and studies that used student samples yielded similar estimates to those that used non-students. The authors also found that hypothetical surveys yielded higher premiums than non-hypothetical studies, studies using a willingness to accept measure had higher premiums than those using a willingness to pay measure, and GM oil products were more acceptable to consumers than GM meat products. Further, GM products that provided enhanced nutrition compared to their non-GM counterparts significantly decreased the premiums associated with non-GM food.

environmental risks associated with GMOs. The proper counterfactual for the natural label is unclear given the current literature so we additionally include two baseline treatments. Finally, we explore how the WTP for the “natural” label differs on several dimensions, including whether the product is used in cooking or is consumed as is, whether the product is edible or simply used as a toiletry (e.g. toothpaste), and whether the product is typically given to children. In total, we use six product pairs, where one product in the pair is “natural” and the other product in the pair is not advertised as such.

Our procedures closely follow Huffman et al. (2003)⁶. Adults (age 18 and older) were paid to participate in a survey (including questions on demographics, beliefs regarding “natural” label regulation, and opinions of GMOs) and an incentive-compatible purchase procedure for several food items. Each “natural”-labelled food item had a “non-natural” counterpart. Although the brands were different, the products were almost identical; therefore, the generic labels presented to the participants were identical except for the “natural” indication. The Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) was used to elicit WTP for each variety⁷. Additional details on recruiting and procedures are provided later.

This design allowed us to explore how factors such as natural flavors and colors, GMO content, and the mere existence of the “meaningless” label, can be attributed to a potential price premium for “natural” food. It is possible that some consumers are aware that the phrase “natural” is not regulated, but are still willing to pay a premium for foods with “natural” labels because the label signals a higher likelihood of attempted health or environmental safety concern by the manufacturer. This last point is documented in the eco-label literature by Cason and Gangadharan (2002). These authors show that, in a laboratory setting, “cheap talk” by

⁶ Several papers by Huffman et al. (Huffman 2003; Huffman, Rousu, Shogren, and Tegene 2003; Huffman and Tegene 2002) used an experiment to elicit willingness to pay for GM foods under various information conditions. Participants bid on three food items: vegetable oil, tortilla chips and russet potatoes. The food labels were generic and only included information on the weight of the package, the expiration date, and the GM content depending on treatment. Participants were also given information on GM foods (including general information and scientific, human, financial and environmental impacts) from a variety of sources. As expected, the authors find that the anti-GM perspective increases the WTP gap between GM and non-GM foods, the pro-GM perspective decreases the WTP gap between GM and non-GM foods, and the third-party information acts as a moderator when provided in addition to the other sources. Overall, WTP for GM-labeled food was 14% less than food with a plain label.

⁷ Huffman et al. actually use an *n*th-price auction to elicit WTP rather than the BDM mechanism. The techniques are similar and the BDM approach is used by Noussair et al. (2004). The details are discussed in design section.

manufacturers regarding product quality is not enough to maintain an efficient outcome, but it is enough for the manufacturers to command a price premium for their inferior, fraudulently-labeled products. This is due to the fact that the majority of the quality claims in the market are true; thus, consumers are willing to pay a premium even under uncertainty.

Noussair et al. (2004) conducted a similar study regarding GMOs. These authors allowed participants in France to taste-test several varieties of biscuits before sequentially revealing information regarding GM content, information on GMOs in general, and finally, the brand names of the biscuits. The authors used a thorough training process to explain the purchase procedure: they allowed subjects to discuss the optimal bidding strategy and transitioned subjects to the purchase procedure using both induced value and real-product training rounds⁸. The results indicated that many participants were willing to purchase the GM varieties if the prices were sufficiently low. This is a surprise given that French surveys indicated an overwhelming opposition to GM foods; at the time, opposition to GM foods was very high in European countries⁹.

Our study differs in that it focuses on a potentially misleading label, not GM content per se. In the U.S., the lack of regulation regarding informative labels is the emphasis of recent lawsuits. Therefore, we are analyzing consumer knowledge and perception of foods and labels as well as price premiums that may be associated with different interpretations of the phrase “natural”. Primarily, we find that consumers who believe that “natural” means “no genetically modified organisms” are willing to pay a premium for “natural” food, whereas consumers who do not explicitly report this belief actually exhibit a negative premium. This result holds for a variety of different products. Further, explicitly pointing out the lack of a “natural” label for standard products significantly affects the premium; transparency could be important from a marketing standpoint. Finally, we find evidence of framing effects which suggest that between-

⁸ We use the same incentive-compatible purchase procedure and use similar training exercises.

⁹ Further, the 1% GM content threshold was treated differently than the 0.1% threshold, and the 0.1% threshold was treated differently from GM-free. The information revelation raised the price of the GM-free product; it also raised the prices of some of the GM-containing products, but not enough to overcome the drop in price from when GM content was revealed. Finally, revealing the brand names of the products raised prices for most of the products.

subject analysis is more reliable for this type of research than within-subject analysis. Section II outlines our experimental design, Section III presents the results and Section IV concludes.

II. Experimental Design

The design includes two equally important parts: a purchase procedure and a survey. We used the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) to sell 12 grocery products (6 product pairs) to 164 adult grocery shoppers in Knoxville, TN. Each product pair included a “natural”-labelled version and a standard version (a product that does not say “natural” on the label). The two versions were not the same brand, but were otherwise very similar. We used an eclectic variety of products to broaden the scope of our results: potato chips, peanut butter, crackers, cooking oil, cereal and toothpaste. Subjects viewed generic labels of the products and placed bids. The generic labels were identical except the natural version said “natural”, “all-natural”, or “100% natural”, depending on how the actual natural product was advertised. Prior to bidding, subjects participated in several training exercises to familiarize them with the BDM. After bidding, the subjects filled out a survey which elicited demographic information and beliefs and opinions regarding the “natural” label and GMOs. Table 3.1 depicts the sequence of events for each session. In total, 8 sessions were conducted at the University of Tennessee, Knoxville Experimental Laboratory in June 2014.

Our treatments varied on two dimensions. First and foremost, we (randomly) assigned subjects to see either the natural labels or the standard labels first. Subjects bid in two blocks: Part 1 and Part 2. Half the subjects saw natural labels in Part 1 and standard labels in Part 2, and the other half saw standard labels in Part 1 and natural labels in Part 2. We also randomized the order of the six products within each Part. Secondly, we varied whether we explicitly told subjects that the standard products were not advertised as natural. This provided two baselines: standard labels without any explanation and (the same) standard labels, accompanied by a statement that indicated the associated products do not say “natural” on the labels. The existing

literature does not provide clear guidance on which baseline is appropriate for our situation, which is why we include both. We discuss this in detail later.

2.1 Purchase Procedure

The purchase procedure we used was the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964). Similar to a second-price auction, the dominant strategy for subjects is to bid their maximum willingness to pay for the item being sold, regardless of risk preference. Unlike an auction, however, there can be multiple winners, thus decreasing the potential for non-demand revealing bids due to competition considerations¹⁰. In short, subjects place a bid for the item for sale. Next, a random price is determined by rolling dice. The distribution of prices is uniform over a predetermined interval¹¹. For instance, we used an interval of \$0.00 - \$5.99 for our main experiment; to achieve any price within the interval with equal probability, we asked a volunteer to roll a six-sided die with numbers 0-5 for the dollars, and then roll a ten-sided die with numbers 0-9 twice (once for each cents place)¹².

Participants were not aware of the interval when placing their bids. Instead, they were told that the interval would be capped at the maximum expected willingness to pay based on previous research. We pre-determined the intervals based on the retail prices of the products and the results from our pilot session. Finally, the participants' bids were compared to the random price. If a participant's bid was higher than or equal to the random price, they purchased the item

¹⁰ Noussair et al. (2004) also use this procedure. All participants have a positive probability of winning; thus, the frequency of insincere bidding is reduced.

¹¹ We considered programming our experiment using zTree software, as is typical for many experiments. However, we decided to use a traditional pen-and-paper setting, rather than a computerized mechanism. Participants were instructed to simply write their bids on paper. We asked for a volunteer to determine the sale prices by rolling dice. Given that the subjects were not college students, and therefore unfamiliar with the computerized laboratory setting, we believe that the transparency of the randomization process was important. While this pen-and-paper setting required additional data input time, the experiment was simple enough that the benefits of transparency outweighed the costs of data entry.

¹² For the training rounds, we used different intervals (e.g. \$1.00 - \$4.99 for a training round where the good for sale was \$3.00). We used a polyhedral dice set (which includes dice with a variety of sides and numbers) to achieve the intervals we needed. Participants were told that the intervals would not all be the same. Further, they were made aware of the interval used for the round only after they placed their bids for the round.

and paid the random price. If a participant's bid was lower than the random price, they did not purchase the item and did not pay anything.

2.2 Training

Subjects participated in several training exercises prior to the actual experiment. The purpose of the training was to ensure that subjects understood the best bidding strategy for the BDM. Prior work has differed on how to train subjects; some experimenters tell subjects the best strategy and explain it using examples (e.g. Plott and Zeiler 2005), while others provide a series of induced value rounds and let subjects figure out and discuss the strategy on their own (e.g. Noussair et al. 2004). We used a hybrid approach. First, we described the purchase procedure to the subjects and had them answer a set of practice example problems, assuming that the item for sale is a \$5.00 bill (e.g. "Suppose you bid \$3.20 and the random price determined by the volunteer is \$4.50. Would you purchase the \$5.00 bill? If so, what price would you pay?"). They were paid for correct answers. Subsequently, we discussed the optimal strategy with the participants and allowed them to ask questions.

Next, we implemented two training rounds where the item for sale was an amount of money (\$7.00 and \$3.00). For each of these training rounds, subjects placed a bid and then we played out the purchase procedure (rolling the dice and determining earnings) several times to allow for learning situations (random prices above and below the value of the sale item) and to allow the subjects to earn cash. For instance, if a subject bid \$6.90 in the first training round and the random price for the first trial was \$5.60, then the subject earned \$1.40 for that trial (\$7.00 - \$5.60). Rarely, subjects earned negative money for a trial if they bid higher than the value of the item. Most subjects bid exactly \$3.00 (the optimal bid) or within 10% of \$3.00 by the time we got to the second training round¹³.

¹³ As further proof (although possibly biased proof) that subjects understood the instructions, more than 90% of subjects responded "4" or "5" to the question, "Did you understand the instructions for the experiment today? Please rate your understanding on a scale from 1 to 5" where "5" was "I understood very well".

Finally, as the last part of the training, we sold an actual grocery item. The purpose of this training round was to allow subjects to see that they actually receive the items they purchase (we physically handed them the item if they purchased it) and to allow them to see that the optimal strategy is the same, regardless of whether the item for sale is an amount of money or a grocery product. There is precedent for the importance of this type of training in the literature: Noussair et al. (2004) use both induced value rounds and a training round with an actual product. In our experiment, we instructed the subjects to choose beforehand whether they would like to bid on a Snickers bar or a granola bar¹⁴. They placed their bids and then a random price was determined by rolling dice from a predetermined price interval. We immediately distributed the items to the subjects who purchased. Subjects were told ahead of time that any purchases they made would be subtracted from their earnings. We have no interest in the bids for this round in terms of our research question; this round was used purely as a training exercise.

2.3 Procedures, Products and Labels

After the training, we proceeded to the actual experiment where subjects placed bids on the grocery items. We used a total of 12 grocery items: a natural-labelled variety and a standard variety (without a natural label) for each of six products: canola cooking oil, kettle-cooked potato chips, creamy peanut butter, frosted bite-size wheat cereal, wheat crackers, and mint-flavored toothpaste. We used an eclectic set of products for several reasons. First, we wanted to use a product that is currently under lawsuit for using the “natural” label; Wesson cooking oil fits this requirement. Additionally, several varieties of chips have been involved in lawsuits in the past. Second, we wanted to use a non-edible product and the toothpaste fits this requirement.

Third, we wanted to include a product that is typically given to children; both the peanut butter and the frosted bite-size wheat cereal could fall under this condition. According to a report by The Center for Science in the Public Interest (Silverglade and Heller 2010), “products claiming to be natural, particularly those aimed at parents of young children, have a competitive

¹⁴ The granola bars were gluten-free and nut free. This provided an option for participants with popular dietary restrictions.

edge in the marketplace”. We hypothesize that the natural premium for products typically given to children may be different from the premium for other products – we also asked about the number of children in the household in our survey. Fourth, we thought it was important to have products that could be consumed immediately or soon after the experiment concluded (as opposed to products that are used in cooking, like cooking oil). The potato chips and the wheat crackers fit this need.

Fifth, and possibly most importantly, it was necessary to have some products that typically contain genetically modified organisms (GMOs). Although the subjects did not know whether the particular product in our experiment contained GMOs, they may have had some knowledge regarding whether the type of product has been associated with GMOs. Major brands of cooking oil, potato chips and wheat crackers have all been publicly associated with GMOs. For instance, several organizations publish lists to guide consumers on how to eat “non-GMO”. Lays potato chips (all varieties), Ruffles potato chips (all varieties), and Ritz crackers are some of the brand names on the “do not buy” list.

Subjects placed bids in two Parts: a “natural” Part and a “standard” Part¹⁵. We randomized which Part came first. Also, the order of products was randomized within each Part. This design allows us to examine both between- and within-subject effects. For the initial between-subject analyses, we only use the Part 1 bids (for half of the subjects, this is the natural products and for the other half this is the standard products). Subjects were informed that one round would be randomly chosen from each Part to be binding, but that the binding rounds would not be determined until after all bids had been placed. A volunteer rolled a six-sided die to determine the binding rounds at the end of the experiment. The random prices were then determined (also via a volunteer rolling dice) only for the binding rounds. After filling out the survey, subjects were paid in cash (in private) and were given the goods (if any) that they purchased. The bid packets and surveys were matched only by ID number to preserve anonymity.

¹⁵ We did not use these “natural” and “standard” names in the experiment, but we use them here to refer to labels with “natural” and labels without “natural”.

The labels for the products were generic and included information such as the type of product, the size, the flavor, etc. We chose pairs of products that were almost identical except for the natural label. For instance, both the natural and standard versions of the peanut butter were “creamy”, were 16 ounce containers, and had 0 grams trans-fat per serving. Both the natural and standard versions of the toothpaste were “clean mint” flavor, were American Dental Association Accepted, contained fluoride, were labelled “whitening”, and were approximately 4 ounce tubes¹⁶. The generic labels for both products in each pair were *identical* with the exception of the natural version saying “natural”. In other words, if there was a slight difference on the labels of the two products, we did not include that difference on the generic label¹⁷. Please refer to Figures 3.1 and 3.2 in Appendix B for examples of our labels and pictures of the actual products used. (Note that participants only viewed the generic labels when bidding, not the actual products.) Participants were informed that all of the information on the generic labels was completely accurate to the actual product label.

Finally, we also included potential brand names on our labels (this was identical for both the natural and standard labels). We did this to help the subjects to understand which type of product was for sale (e.g. a “wheat cracker” could take various forms, but most people know what type of product to picture when they are informed that the product is similar to a Ritz cracker). The brands listed were not necessarily the brands of the actual products, but all of the actual products were major brands (not store brands).

It is unclear what the proper counterfactual is for the “natural” label. To the best of our knowledge, there are no existing products that say “Unnatural” or “Not Natural” on their labels. Therefore, our standard labels simply did not say anything about natural. However, this does not necessarily mean that the product is not “natural” in the sense that it may or may not reflect the

¹⁶ We were able to find some of the product pairs in the exact same size containers. However, if the containers were similar, but not exactly the same, we included approximate sizes on the generic label. We always used an approximate size that was smaller than or equal to the size of the actual product. This way, participants would never pay more than they intended for a product, given its size.

¹⁷ For instance, the natural crackers said “Entertainment” on the actual product label, but the standard variety did not say “Entertainment”; even though the two products are extremely similar, the brand we chose for the standard variety did not advertise their product as a party cracker, or a cracker typically used “for entertaining”. Thus, we simply did not include this detail on our generic labels; our labels were identical with the exception of “natural” on the natural variety.

beliefs that subjects have regarding the natural label. For instance, a product that is not labelled “natural” doesn’t necessarily contain artificial flavors or colors. (Many of our subjects indicated in the survey that they believe “natural” means “no artificial flavors” and “no artificial colors”). Therefore, we included additional treatments in which we explicitly informed subjects that the products included in a particular Part *do not* say “natural” on their labels¹⁸.

The literature does not provide clear guidance on this issue. Huffman et al. (2003) examine labels which indicate that a product does or does not contain GMOs (i.e. “This product contains GMOs” or “This product does not contain GMOs” or no statement regarding GMOs). However, regulations regarding this type of labelling (“This product contains GMOs”) exist in Europe and are under consideration in some states in the U.S. Conversely, it seems highly unlikely that a regulation will ever exist which requires manufacturers to state “This product is not natural”, especially since the label is completely unregulated now¹⁹. Therefore, we implement what we call “explicit” and “implied” treatments so that we can use both counterfactuals. In the “explicit” treatments, a piece of paper with subject-specific instructions informed subjects that the products which followed did not say “natural” on the actual labels. In the “implied” treatments, the subjects just viewed the standard labels without any indication of whether or not the actual products were “natural”.

This allowed us to vary the transparency that subjects face. In reality, this level of transparency varies at grocery stores. For instance, in the case of cooking oil, both the “natural”-labeled variety and the variety without a “natural” label literally sit adjacent to one another on the shelf in most major grocery stores – it is very obvious that a “natural” variety exists. On the other hand, we had to go to two different stores to find “natural” and standard varieties of the wheat crackers which were otherwise identical. The likelihood that consumers notice whether a product says “natural” (and therefore have the opportunity to pay a premium for the natural variety) may depend on how obvious it is that the “natural” alternative exists.

¹⁸ Since different subjects saw the standard labels at different times (some saw the natural labels first and some saw the standard labels first), we did not state this verbally. Instead, we included a piece of paper following the main instructions for the Part that had information specific to the subject. Subjects were told to read this page silently before bidding.

¹⁹ We expect that if the FDA does define the term “natural”, it will probably be defined similarly to phrases like “low fat”; the label is not required, but if the manufacturer does use it, their product must meet certain conditions.

2.4 Survey

The survey included questions regarding demographic information, general questions about grocery shopping behavior, and questions eliciting knowledge and beliefs about “natural” foods and foods containing GMOs. For instance, one important question read: “What do you think the phrases ‘natural’ or ‘all natural’ actually mean when printed on a food label? (Please check all that apply.)” with options such as “No artificial flavors”, “No artificial colors”, “No pesticides”, “No dyes”, “Limited processing”, “Environmentally-friendly”, “No genetically modified ingredients”, and “Organic”. Of course, in actuality, “natural” is not regulated and therefore has no real meaning, but many participants checked at least a few options from our list of 10 options. Other participants checked “None of the above” and still others wrote in things like, “Only a few ingredients”.

Other questions asked participants to rate their agreement with various statements such as “I am concerned that food products containing genetically modified organisms (GMOs) pose a health risk”, “I am concerned that the production of products containing genetically modified organisms (GMOs) poses an environmental risk”, and “Foods labelled ‘natural’ or ‘all natural’ are healthier than foods without a ‘natural’ or ‘all natural’ label.” For these questions, subjects rated their agreement on a scale from 1 to 5.

2.5 Participants

In total, 164 non-student grocery shoppers living in the Knoxville, TN area participated in the experiment. In an effort to match the demographics of U.S. grocery shoppers, we use a disproportionate number of females and people who indicated that they shop on a regular basis for multiple-person households. Katsaras et al. (2001) report that 76% of grocery shoppers are women; our sample is 78% women. A survey by the Time Use Institute²⁰ reported that the average age of grocery shoppers is 47 and the median income is \$50,525; the average age of our

²⁰ The Time Use Institute: “Grocery Shopping: Who, Where and When” 2008. The data used is from The American Time Use Survey, sponsored by the U.S. Dept. of Labor and fielded by the U.S. Census Bureau.

sample is 46.8 and the median household income is within the range of \$40,000-\$60,000 (subjects only reported their household income range).

While many experiments use college students as participants, this demographic group is not entirely appropriate for our purposes. As Huffman et al. (2003) point out, the share of college-age grocery shoppers is below the share of college-age individuals in the population. We pre-screened members of the Knoxville community based on demographics; the vast majority of our participants are employed full-time. We paid our recruits \$10 for showing up, \$15 for completing the survey, and additional money from the practice calculations and training rounds. Subjects earned about \$35 on average²¹ to come to the University of Tennessee to participate (about 90% of subjects earned between \$30 and \$40). Each session lasted about 70 minutes.

III. Results

3.1 Between-Subject Analysis Using Part 1 Bids

The experimental design allows for identification of labelling effects using multiple comparisons²². First, we analyze only the Part 1 bids in order to achieve a clean, between-subject comparison. Approximately half of our subjects bid on products with natural labels first and the other half bid on products with standard labels first, allowing such a comparison. The results (displayed in Table 3.3) indicate that there is a statistically significant overall impact of the natural label for those who believed that “natural” means “no genetically modified organisms”. We further break this down by product and see that the effect is significant for cereal, crackers and oil, weakly significant for peanut butter, and insignificant for chips and toothpaste. The willingness to pay for peanut butter could be affected by the fact that some “natural”-labelled peanut butters have separation of oils, requiring them to be stirred; some subjects could be

²¹ The average earnings *before* the subjects had the opportunity to purchase grocery items was about \$35.

Presumably, subjects only purchased goods at prices that were favorable to them, so the actual average benefit to a participant was probably higher than \$35.

²² Table 3.2 provides variable descriptions and means.

willing to pay more for varieties with additives, which keep the peanut butter spreadable. It is also plausible that “natural” doesn’t have as strong an effect for non-edible items, such as toothpaste. Further, it could be the case that when people purchase chips, their aim is not to be healthy. Chips are generally snack foods and some buyers’ first priority with snacks could be taste; the “natural” label may lead them to believe that these “healthier” versions are less tasty.

Table 3.4 displays several treatment effects of interest based on the model in Table 3.3. We see that the “natural” premium for subjects who believe that “natural” means no GMOs is significant for the “explicit” treatment, but not for the “implicit” treatment. This suggests that when the lack of a “natural” label is specifically pointed out, subjects are willing to pay less for the standard products. Thus, it was worthwhile to include both baselines in our design. While the experimental setting is not necessarily representative of typical shopping experiences, this result implies that when consumers are more readily confronted by the difference in labels (e.g. when the two products sit adjacent to one another on the shelf), they are more likely to pay a premium for the “natural” product. During our research, we noticed several products for which this is the case: the “natural”-labelled and standard varieties of cooking oil, toothpaste, ice cream, mustard, hamburger seasoning and chips are all located next to one another in the grocery store. For other products, such as crackers, juice and cereal, shoppers have to visit either the “natural” foods section or another store to find “natural” varieties. For the “explicit” treatment, the significant premiums for subjects who believe “natural” means no GMOs range from \$0.47 (chips) to \$0.87 (cooking oil). The products had retail prices in the range of \$2.00 to \$5.00 so these premiums are substantial²³.

Interestingly, the “natural” label itself (for subjects who did not believe that “natural” means no GMOs) actually had a negative premium for most products, though this is only statistically significant for the “implied” treatment²⁴. Again, we hypothesize that some subjects

²³ The approximate retail prices are as follows (retail prices varied slightly from store to store depending on sales and promotions): peanut butter, \$3 - \$4; crackers, \$3 - \$4; cereal, \$4 - \$4.75; cooking oil, \$3 - \$3.50; potato chips, \$2.50 - \$3.50; toothpaste, \$2.50 - \$4.

²⁴ The “explicit” treatment raised the premium enough that it is not significant, but it is still negative.

may have believed that “natural” means “healthier” and that healthy products are less tasty²⁵. The last two rows in Table 3.4 provide the weighted effect for the “explicit” and “implicit” treatments. Approximately 42% of our sample believed that “natural” means no GMOs; thus, the no GMO effect is weighted by this percentage. These weighted effects are insignificant for most products so we conclude that there is a null overall effect of the “natural” label. However, clearly the subgroup who held the belief behaved quite differently. This could imply that manufacturers may benefit from holding a portfolio of products to capture rents from different segments of the market.

Although we do not report the results here, we also execute this model with additional demographic and shopping behavior variables. However, these variables are not very consequential and we justify removing them with statistical tests. The only statistically significant factors are gender and the number of trips taken each week to purchase groceries. Subjects who take one additional trip each week are, on average, willing to pay \$0.13 more in the “All Goods” model. Perhaps this is because these subjects are more familiar with the retail prices of different types of products and are thus less uncertain about how much they are willing to pay. We also explore the effects of other beliefs regarding “natural”, such as “no artificial content” and vague claims such as “limited processing” or “environmentally-friendly”, but do not find significant effects. Instead, it appears that “no genetically modified organisms” is the most important factor contributing to the “natural” premium among our variables of interest²⁶.

²⁵ In fact, we do ask about “natural” and “healthy” in our survey. Some subjects did indicate that they believe “foods labelled natural are healthier than foods without a natural label.” However, we cannot prove that subjects believed healthier foods are less tasty; we simply observe a negative premium and suggest this as a possibility.

²⁶ Responses to our Likert-scale survey questions are also explored as possible determinants of the “natural” premium in a within-subject analysis. For instance, we include subjects’ Likert scale responses to questions like, “Please rate your agreement with the following statement on a scale from 1 to 5. ‘I am concerned that the production of products containing genetically modified organisms (GMOs) poses an environmental risk’”, where 5 represents “I completely agree”. However, the results are mostly insignificant. Further, the within-subject analysis is possibly compromised due to framing effects biasing the Part 2 bids. Details are discussed later.

3.2 Part 2 Bids and Framing Effects

Next, we turn to the Part 2 bids. Note that this analysis is not as clean as using Part 1 bids. When the subjects placed bids in Part 2, they were probably aware of our research question – i.e. if they saw the standard labels first and natural labels second, or vice versa, they could infer by Part 2 that we were testing their willingness to pay for the natural label. Indeed, we do find quite a difference between Part 1 and Part 2. We replicated the model used in the Part 1 analysis but restricted it to Part 2 bids. The associated treatment effects are displayed in Table 3.5. Contrary to the Part 1 analysis, there is no longer preference heterogeneity conditional on beliefs that “natural” means GMO-free. However, there is a significant positive effect of the natural label itself. It is more pronounced in the “explicit” treatment, but is also statistically significant in the “implied” treatment for cereal and toothpaste.

We hypothesize that framing or experimenter effects play a large role in Part 2. At this point, subjects probably realized that the study was focused on the natural label. Further, subjects may have felt that the experimenters were looking for a “natural” premium (e.g. “the labels have changed between Part 1 and Part 2; the experimenters must be expecting me to behave differently”). This framing or experimenter effect could have outweighed the effect of the subjects’ own beliefs regarding the label; thus, the premium cannot be attributed to the “no GMO” belief. This finding has important design implications for future work. Experiments and surveys should be designed so that a clean, between-subject analysis can be extracted. While within-subject calculations may be illustrative in some cases, researchers should exercise caution when interpreting data.

3.3 Beliefs Regarding “Natural”

It is apparent from Table 3.2 that the majority of our subjects do believe that “natural” has some meaning; in addition to the means for the specific beliefs, which range from 15% of the sample to 70% of the sample, the average number of characteristics reported as beliefs was 4.24. This is consistent with a recent article published in USA Today, which reports that “two-thirds of

Americans think the word “natural” on the label of a packaged or processed food means it contains no artificial ingredients, pesticides or genetically engineered organisms... under federal labeling rules, the word natural means absolutely nothing²⁷” (Weise 2014).

We examine beliefs a little further to determine how they differ based on demographics. Not surprisingly, subjects who are more highly educated and subjects with young children report knowing significantly more about GMOs. Older subjects are more likely to believe that GMOs pose risks to human health or the environment; they are also more likely to believe foods labelled “natural” are healthier than foods without a natural label. Higher income, a higher level of educational attainment and being female all reduce the probability that subjects believe “natural” labels indicate a vague claim such as “limited processing” or “environmentally-friendly”. From a marketing standpoint, the heterogeneity in beliefs could mean that manufacturers might benefit from holding portfolios of products containing both “natural” and standard versions. However, demographics cannot predict the likelihood of believing “natural” means no GMOs, which is the belief that actually leads to a premium.

We conclude that, on average, there is a null effect of the “natural” label. However, when the sample is segregated based on the belief that “natural” means no GMOs, there is a positive “natural” premium for those who hold the belief and a negative premium for those who do not hold the belief. A policy to regulate phrases containing the word “natural” could possibly be welfare-enhancing for both those who believe “natural” means GMO-free as well as those without this belief for whom the phrase “natural” fosters negative associations.

IV. Conclusion

To our knowledge, this study is the first to examine the commonly used “natural” food label, which is currently unregulated. We execute an experiment to examine consumer response

²⁷ The article reports the results of a recently released survey by the magazine, Consumer Reports.

to the label using a variety of grocery products. A hybrid approach is used to both elicit willingness to pay (WTP) for foods with and without the label and elicit subjects' knowledge and beliefs surrounding the label and genetically modified organisms (GMOs). An incentive-compatible procedure is used to sell 12 products (6 product pairs). Each product pair consists of a "natural"-labeled variety and a variety without a "natural" label. A survey is used to better understand subjects' beliefs regarding the meaning of the label and issues related to GMOs, including health and environmental risks.

We find that the premium associated with the "natural" label is highly sensitive to the belief of the buyer regarding GMO content. Buyers who believe that "natural" means "no genetically modified organisms" are willing to pay a premium for a "natural" product, while buyers without the belief instead have a negative premium. We break this down by product in our between-subject analysis and find that the effect is significant for most products used in our experiment. Further, we find that explicitly pointing out the lack of a "natural" label on standard products decreases WTP relative to only plain labels. Finally, framing effects significantly change the results for the Part 2 labels, which suggests that future work should not rely solely on within-subject analysis.

Several pending lawsuits regarding GMOs in foods labeled as "natural" partially motivate our paper. The "natural" label is not regulated in the U.S. Further, labelling of GMOs in food products is not required in the U.S., even though it is required in almost every other developed country. Thus, the pending lawsuits will set an important precedent. The results of this study suggest that there is some merit to the plaintiffs' complaints; many of our subjects believed that the "natural" label has some meaning. Since the label is unregulated, "natural" could be a misleading claim. However, only the "no GMO" belief for the "natural" label is actually associated with a WTP premium. Forty-two percent of our sample held this belief, and, on average for the sample, the effect is zero. Further, subjects arrived at our laboratory with homegrown beliefs and we cannot necessarily rule out that there could be unobservable effects correlated with the "no GMO" belief. It could be case that consumers who believe "natural" means "no GMOs" are also those who exhibit a positive premium due to unobserved factors correlated with this belief. We can say that there is a correlation between those who believe

“natural” means “no GMOs” and those who are willing to pay a premium for a product labelled as “natural”. Regulation of the “natural” label could help to resolve the information asymmetry.

If policymakers do decide to regulate the “natural” label, there are several directions in which policy could move forward. One option is to define the “natural” label, requiring manufacturers to meet certain conditions before using the label on their products. According to our study, the conditions could include: no genetically modified material, no artificial flavors or colorings, no pesticides and no dyes, as many of our subjects believed that this is what “natural” indicates to them. Another option is to get rid of the “natural” label, banning it from food products. The “organic” label is regulated and already conveys many of the characteristics that consumers look for in the “natural” label. The final option, and possibly the most politically feasible option, is to require disclosure of GMO content on food labels. Non-profit organizations, such as the Center for Food Safety, have been advocating for GMO labelling for years. Additionally, most other developed countries have GMO labelling laws. Although this option would leave the “natural” label ambiguous, it would clarify the issue for consumers who may currently believe that “natural” means “no genetically modified organisms”. Since we find a premium for subjects with this belief, the policy would be welfare-enhancing.

Future work could include a wider range of products to determine whether these results hold for different characteristics; for instance, shampoo or soap (many of which are labelled “natural”) could be used. Additionally, one could compare the premiums associated with a “non-GMO” label (or an “organic” label) to the “natural” label by using these different labelling schemes in a controlled experiment. Finally, as lawsuits involving “natural” and regulations involving GMO labelling are currently evolving in the United States, one could take advantage of future policy changes as natural experiments to determine whether and under what circumstances consumers are willing to pay premiums for health and environmental-related labels.

References

- Allen, D. M., Coutroulis, C. S., and Ciotti, R. L. "Square Pegs into Round Holes: Class Certification in GMO Food Products Cases." *Bloomberg Law* (2013). Accessed June 2014: <http://www.bna.com/square-pegs-into-round-holes-class-certification-in-gmo-food-products-cases/>.
- Becker, G., M. DeGroot, and J. Marschak. "Measuring Utility by a Single-Response Sequential Method." *Behavioral Science* 9 (1964).
- Cason, T. and L. Gangadharan. "Environmental Labeling and Incomplete Consumer Information in Laboratory Markets." *Journal of Environmental Economics and Management* 43 (2002).
- Center for Food Safety. Accessed June 2014: <http://www.centerforfoodsafety.org/>.
- Cummins, Ronnie. "GMO and the 'Natural' Food Fight: The Treacherous Terrain of Food Labeling." Global Research Centre for Research on Globalization (2014). Accessed June 2014: <http://www.globalresearch.ca/gmo-and-the-natural-food-fight-the-treacherous-terrain-of-food-labeling/5363798>.
- Frankel, Alison. "Labeling Genetically Modified Food: Regulation via Litigation is Back." *Reuters* (2013). Accessed June 2014: <http://blogs.reuters.com/alison-frankel/2013/10/16/labeling-genetically-modified-food-regulation-via-litigation-is-back/>.
- Huffman, W. "Consumers' Acceptance of (and Resistance to) Genetically Modified Foods in High-Income Countries: Effects of Labels and Information in an Uncertain Environment." *American Journal of Agricultural Economics* 85 (2003).
- Huffman, W., M. Rousu, J. Shogren, and A. Tegene. "The Public Good Value of Information from Agribusinesses on Genetically Modified Foods." *American Journal of Agricultural Economics* 85 (2003).
- Katsaras, N., Wolfson, P., Kinsey, J., and Senauer, B. "Data Mining: A Segmentation Analysis of U.S. Grocery Shoppers." The Retail Food Industry Center, University of Minnesota Working paper 1-01 (2001).
- Lusk, J., Jamal, M., Kurlander, L., Roucan, M., and Taulman, L. "A Meta-Analysis of Genetically Modified Food Valuation Studies." *Journal of Agricultural and Resource Economics* 30-1 (2005).
- Noussair, C., S. Robin, and B. Ruffieux. "Do Consumers Really Refuse to Buy Genetically Modified Food?" *The Economic Journal* 114 (2004).

- Plumer, B. “‘All-natural’ Labels on Food are Meaningless. Let's Get Rid of Them”. Vox Conversations (2014). Accessed June 2014: <http://www.vox.com/2014/6/23/5827850/those-all-natural-labels-on-food-are-totally-meaningless>.
- Rousu, M., W. Huffman, J. Shogren, and A. Tegene. “Effects and Value of Verifiable Information in a Controversial Market: Evidence from Lab Auctions of Genetically Modified Food.” *Economic Inquiry* 45-3 (2007).
- Silverglade, B. and Heller, I. R. “Food Labeling Chaos: The Case for Reform”. Center for Science in the Public Interest, Washington, DC (2010).
- Smith, J. P. “More Lawsuits over ‘Natural’”. Sound Consumer (2014). Accessed June 2014: <http://www.pccnaturalmarkets.com/sc/1403/more-lawsuits-over-natural.html>.
- Weise, E. “66% of Consumers Wrongly Think ‘Natural’ Means Something”. USA Today (2014). Accessed June 2014: <http://www.wtsp.com/story/money/2014/06/17/natural-food-labels-no-meaning/10674755/>.

Appendix A

Tables

Table 3.1: Session Schedule

Step 1	Introduction and read training instructions
Step 2	Allow subjects to answer practice calculations; moderators grade calculations as subjects finish, and record earning on the Record Sheet
Step 3	Present answers to practice calculations to the group, discuss optimal bidding strategy, and answer questions
Step 4	Training Round 1: Read instructions, allow subjects to place bids on the item (\$7), determine random prices by having volunteers roll dice (4 trials), and allow subjects to calculate their earnings
Step 5	Training Round 2: Read instructions, allow subjects to place bids on the item (\$3), determine random prices by having volunteers roll dice (4 trials), and allow subjects to calculate their earnings
Step 6	Training Round 3: Read instructions, allow subjects to choose to bid on the granola bar or the candy bar (display both), allow subjects to place bids, determine a random price by having a volunteer roll dice, allow subjects to record costs on Record Sheet if they made a purchase, and give items to subjects who made purchases
Step 7	Main Experiment, Part 1: Read instructions, allow subjects to place bids
Step 8	Main Experiment, Part 2: Read instructions, allow subjects to place bids
Step 9	Determine one binding round from each Part by having a volunteer roll a 6-sided die, determine random prices for each binding round by having a volunteer roll dice, and allow subjects to record costs on Record Sheet if they made a purchase(s)
Step 10	Allow subjects to fill out survey
Step 11	Pay subjects and give them grocery products if they made a purchase(s)

Table 3.2: Variable Definitions and Means

Variable Name	Definition	Mean
Age	= subject's age	46.75
Female	= 1 if subject is female	0.78
Primary	= 1 if subject is the primary grocery shopper for the household	0.88
Trips	= number of trips subject takes to the grocery store each week	1.86
Natural	= 1 if bid was for a natural-labelled product; = 0 for standard product	
Standard	= 1 if bid was for a standard product; = 0 for natural product	
Explicit	= 1 if subject was explicitly informed that the product was not labelled natural	
Natural First	= 1 if the subject bid on the natural products first	
Belief: No Art Flavor	= 1 if subject reported that he believed "natural" meant "no artificial flavors", "limited processing", etc.; 0 otherwise	0.70
Belief: No Art Color		0.68
Belief: No Art Preservatives		0.68
Belief: No GMO		0.42
Belief: No Pesticides		0.41
Belief: No Dyes		0.38
Belief: Limited Processing		0.33
Belief: Higher Quality		0.25
Belief: Environment-Friendly		0.16
Belief: Organic		0.15
Total Beliefs	= total number of characteristics that the subject reported as his beliefs regarding the "natural" label (min = 0, max = 10)	4.24
Healthy	= the number subject indicated on a Likert scale from 1-5 to indicate agreement with these statements: "Foods labelled natural are healthier than foods without a natural label", "I am concerned that food products containing GMOs pose a health risk", "I am concerned that we as a society do not fully understand the impacts of GMOs", etc. Please see survey questions 22, 24, 25, 26 and 28 for exact wording.	3.18
GMO Know		2.71
GMO Health Risk		3.58
GMO Env Risk		3.55
GMO Concern		4.16
GMO Pay	= 1 if subject indicated that he is willing to pay extra for food that does not contain GMOs; 0 otherwise	0.62

Table 3.3: Between-Subject Analysis Using Part 1 Bids

Dependent variable: Bid, in \$							
Variable	All Goods	Peanut Butter	Cereal	Crackers	Chips	Oil	Tooth-paste
Natural × Belief: No GMO	0.54** (0.19)	0.43* (0.25)	0.80** (0.25)	0.58** (0.20)	0.29 (0.21)	0.89** (0.27)	0.29 (0.21)
Natural	-0.43** (0.19)	-0.01 (0.28)	-0.47* (0.25)	-0.54** (0.20)	-0.60** (0.22)	-0.50* (0.27)	-0.47** (0.22)
Standard × Belief: No GMO	-0.18 (0.21)	0.08 (0.29)	-0.08 (0.28)	-0.18 (0.21)	-0.34 (0.24)	-0.26 (0.28)	-0.29 (0.22)
Standard × Explicit	-0.25 (0.21)	-0.02 (0.29)	-0.15 (0.27)	-0.33 (0.21)	-0.44* (0.24)	-0.21 (0.27)	-0.36 (0.22)
Female	0.31** (0.15)	0.22 (0.21)	0.12 (0.21)	0.31* (0.16)	0.38** (0.17)	0.62** (0.22)	0.23 (0.17)
Trips per Week	0.13** (0.06)	0.07 (0.07)	0.17** (0.07)	0.09 (0.06)	0.17** (0.06)	0.19** (0.08)	0.11* (0.06)
constant	1.70** (0.23)	1.90** (0.36)	1.91** (0.28)	1.63** (0.24)	1.52** (0.25)	1.59** (0.30)	1.61** (0.24)
N	971	164	162	157	163	162	163
R ²	0.08	0.03	0.10	0.11	0.12	0.15	0.07

Notes: * and ** denote coefficient is statistically different from zero at the 10% and 5% significance levels, respectively. Cluster-robust standard errors in parentheses.

Table 3.4: Treatment Effects Using Part 1 Data

Hypothesis	All Goods	Peanut Butter	Cereal	Crackers	Chips	Oil	Tooth-paste
Natural No GMO = Standard No GMO (Explicit)	0.55** (0.24)	0.36 (0.35)	0.57* (0.32)	0.55** (0.26)	0.47* (0.28)	0.87** (0.33)	0.47* (0.26)
Natural No GMO = Standard No GMO (Implied)	0.29 (0.26)	0.34 (0.32)	0.42 (0.34)	0.21 (0.24)	0.03 (0.30)	0.65* (0.36)	0.12 (0.27)
Natural w/o Belief = Standard w/o Belief (Explicit)	-0.17 (0.20)	0.01 (0.29)	-0.31 (0.28)	-0.21 (0.21)	-0.16 (0.22)	-0.29 (0.27)	-0.11 (0.24)
Natural w/o Belief = Standard w/o Belief (Implied)	-0.43** (0.19)	-0.01 (0.28)	-0.47* (0.25)	-0.54** (0.20)	-0.60** (0.22)	-0.50* (0.27)	-0.47** (0.22)
Natural Weighted = Standard Weighted (Explicit)	0.13 (0.16)	0.16 (0.25)	0.06 (0.23)	0.11 (0.18)	0.10 (0.19)	0.20 (0.22)	0.14 (0.19)
Natural Weighted = Standard Weighted (Implied)	-0.12 (0.17)	0.14 (0.23)	-0.09 (0.22)	-0.22 (0.17)	-0.34* (0.20)	-0.01 (0.24)	-0.23 (0.19)

Notes: * and ** denote statistical significance at the 10% and 5% levels, respectively.

Table 3.5: Treatment Effects Using Part 2 Data

Hypothesis	All Goods	Peanut Butter	Cereal	Crackers	Chips	Oil	Tooth-paste
Natural No GMO = Standard No GMO (Explicit)	0.24 (0.26)	0.50 (0.31)	0.32 (0.31)	0.38 (0.26)	0.37 (0.30)	-0.22 (0.43)	0.12 (0.24)
Natural No GMO = Standard No GMO (Implied)	-0.33 (0.26)	-0.16 (0.33)	-0.34 (0.34)	-0.30 (0.26)	-0.19 (0.31)	-0.71* (0.37)	-0.27 (0.26)
Natural w/o Belief = Standard w/o Belief (Explicit)	0.94** (0.22)	1.00** (0.31)	1.21** (0.28)	0.92** (0.24)	0.89** (0.22)	0.77** (0.32)	0.88** (0.23)
Natural w/o Belief = Standard w/o Belief (Implied)	0.37** (0.18)	0.33 (0.26)	0.55** (0.25)	0.24 (0.19)	0.32 (0.21)	0.28 (0.28)	0.49** (0.21)
Natural Weighted = Standard Weighted (Explicit)	0.65** (0.18)	0.79** (0.24)	0.83** (0.22)	0.69** (0.19)	0.67** (0.20)	0.36 (0.30)	0.56** (0.18)
Natural Weighted = Standard Weighted (Implied)	0.07 (0.16)	0.12 (0.22)	0.17 (0.22)	0.01 (0.16)	0.10 (0.19)	-0.14 (0.24)	0.17 (0.18)

Notes: * and ** denote statistical significance at the 10% and 5% levels, respectively.

Appendix B

Figures

	Potato Chips – Natural Type: Kettle Cooked Flavor: Original (Plain) Size: 8.5 oz bag Brand: A major brand such as Lays, Utz, Herr's, Wise, or similar Product is labelled as "All Natural"
	Potato Chips Type: Kettle Cooked Flavor: Original (Plain) Size: 8.5 oz bag Brand: A major brand such as Lays, Utz, Herr's, Wise, or similar

Figure 3.1: Potato Chip Pictures and Labels

	
Toothpaste – Natural Flavor: Clean Mint Features: Fluoride Toothpaste, Clinically Proven, Cavity Protection, Whitening, American Dental Association (ADA) Accepted Size: approximately 4 oz tube Brand: A major brand such as Crest, Colgate, Aquafresh, Arm and Hammer, or similar Product is labelled as “Natural”	Toothpaste Flavor: Clean Mint Features: Fluoride Toothpaste, Clinically Proven, Cavity Protection, Whitening, American Dental Association (ADA) Accepted Size: approximately 4 oz tube Brand: A major brand such as Crest, Colgate, Aquafresh, Arm and Hammer, or similar

Figure 3.2: Toothpaste Pictures and Labels

Appendix C

Instructions

The training instructions and the Main Experiment Part 1 “explicit” instructions where the standard products came first are presented here. Other versions of the instructions are available upon request.

Introduction and Training

Welcome to the University of Tennessee Experimental Economics Laboratory. Your participation in today’s study is appreciated. If you have a question at any time, please raise your hand. Please refrain from verbally reacting to events that occur during the experiment. This is very important. Your decisions in today’s study are private to you. We ask that you do not communicate with other study participants. We also ask that, after the session, you do not discuss details with others who have registered to potentially participate in a future session of this study.

The decisions you make (including your answers to survey questions) will be linked only with your station ID, which was (randomly) chosen by you when you entered the lab. Your decisions will not be associated with your name or other identifying information. Your name will not be linked in any way to the results of the study. For accounting purposes, we will ask you to fill out a receipt at the end of the session. We do not keep the receipts; they are submitted to the UT Treasurer’s office.

In the main parts of the experiment, you will be asked to place bids to buy common grocery items. Although you will not earn money in these parts, you will have the opportunity to actually buy some of these items and take them home with you today. The purchase procedure will be new to you. We will first go through a series of training materials that will familiarize you with the purchase procedure. The good news is that you will have the opportunity to earn cash money during this training.

You will be paid in cash and in private at the end of the session. You will receive a show-up fee of \$10 for participating, an additional \$15 for completing a survey, and additional money from the training rounds. If you purchase one or more grocery items, the amount you pay for the product(s) will be subtracted from your earnings. **We will proceed through the written materials together. Please do not flip forward until instructed.**

General Procedures

In each decision “round” of this study, the moderator will offer an item for sale. In the initial training rounds, the item for sale will be an amount of money. In the other rounds, the item for sale will be a grocery product. In each round you place a bid to buy the item for sale.

We will use the following purchase procedure in all rounds:

1. You will place a bid on the item. You will not know the price prior to bidding.
2. The price of the item will be randomly drawn. A volunteer will be asked to roll dice to determine this price. The random price will be the same for all participants.
3. **If your bid is equal to or higher than the random price, you buy the item and pay the random price (not your bid!). If your bid is lower than the random price, you do not buy the item.**

Here are some possible scenarios based on the purchase procedure:

- You bid \$2. The random price is drawn to be \$1.50. Since your bid is equal to or higher than the random price, you buy the item at a price of \$1.50.
- You bid \$5. The random price is drawn to be \$5. Since your bid is equal to or higher than the random price, you buy the item at a price of \$5.
- You bid \$3. The random price is drawn to be \$3.50. Since your bid is lower than the random price, you do not buy the item.

It is important to point out some aspects of the procedure. First, different from auctions, you are not bidding against other players. The bids of other players do not impact whether you buy an item. We have large quantities of each item. If, for example, everyone bids an amount higher than the random price, each person will pay the random price and each person will receive the item. Second, different from some auctions, if you buy something, the price is not equal to your bid. Instead, you pay the randomly selected price.

Third, your bid sets the highest price for which you agree to buy the good. For example, if you bid \$6.25, this means that you agree to buy the item as long as the price is something less than or equal to \$6.25. Your bid of \$6.25 guarantees that you do not buy the item at prices above \$6.25.

Before bidding you should ask yourself “what is the highest price I am willing to pay for the item?” It is in your best interest to place a bid equal to this highest price.

“What If” Scenarios

To help you understand the procedures, we ask that you consider a number of “what if” scenarios. Here is the good news: you will be paid \$1 for each scenario you answer correctly. There is a bonus question, and you will be paid \$2 for a correct answer to this.

The item for sale in these scenarios is a \$5 bill. Remember: If your bid is equal to or higher than the random price, you buy the item and pay the random price (not your bid!). If your bid is lower than the random price, you do not buy.

1. Suppose you **bid \$2.50**. Then, a volunteer draws a **random price of \$4.00**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____
2. Suppose you **bid \$3.12**. Then, a volunteer draws a **random price of \$6.37**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____
3. Suppose you **bid \$5.00**. Then, a volunteer draws a **random price of \$4.25**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____
4. Suppose you **bid \$5.00**. Then, a volunteer draws a **random price of \$6.56**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____
5. Suppose you **bid \$7.16**. Then, a volunteer draws a **random price of \$4.12**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____
6. Suppose you **bid \$8.00**. Then, a volunteer draws a **random price of \$6.50**.
Based on the procedure we described would you purchase the \$5 bill? Yes No
If you answered “Yes”, what price would you pay? \$ _____

Bonus question. Given the purchase procedure, how much should you bid for the \$5 bill? Keep in mind that it is in your best interest to place a bid equal to the highest price you’re willing to pay.
You should bid: \$ _____ . _____

Please raise your hand when you are ready to have your calculations checked.

Training Round 1

In this training round you will have the opportunity to earn money.

The item for sale is \$7.

Your task in this round is to place a bid to buy the \$7.

After everyone has indicated their bid, a volunteer will roll dice to determine the random price. Although you will not know the price range before you bid, know that three dice will be rolled. The first will determine the dollars and the other two will determine the cents.

If your bid is equal to or higher than the random price, you will receive the \$7 and pay the random price. You will thus earn an amount equal to: \$7 *minus* the random price. If you make a purchase at a price that is higher than \$7, you will in fact have negative earnings (lose money).

If your bid is less than the random price, you will not receive the \$7. You will not pay the random price. You will earn \$0.

For training purposes, we will play out the purchase procedure (rolling the dice and calculating earnings) several times. However, you will only bid once. **You will not be able to change your bid after the random prices are determined.**

Please determine your bid at this time and write it here: \$ _____ . _____

Trial 1

Random price: \$ _____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes.** You bought the \$7.
Your earnings are equal to: \$7 – _____ = \$ _____ . _____
(random price)

☐ **No.** You did not buy the \$7. Your earnings for this trial are \$0.

Record your earnings on your Record Sheet.

Training Round 1—Continued

Trial 2

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes.** You bought the \$7.

Your earnings are equal to: $\$7 - \text{_____} = \$ \text{_____} . \text{_____}$
(random price)

☐ **No.** You did not buy the \$7. Your earning are for this trial are \$0.

Record your earnings on your Record Sheet.

Trial 3

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes.** You bought the \$7.

Your earnings are equal to: $\$7 - \text{_____} = \$ \text{_____} . \text{_____}$
(random price)

☐ **No.** You did not buy the \$7. Your earning are for this trial are \$0.

Record your earnings on your Record Sheet.

Trial 4

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes.** You bought the \$7.

Your earnings are equal to: $\$7 - \text{_____} = \$ \text{_____} . \text{_____}$
(random price)

☐ **No.** You did not buy the \$7. Your earning are for this trial are \$0.

Record your earnings on your Record Sheet.

Training Round 2

The item for sale is \$3. Your task in this round is to place a bid to buy the \$3.

After everyone has indicated their bid, a volunteer will roll dice to determine the random price. Although you will not know the price range before you bid, know that three dice will be rolled. The first will determine the dollars and the other two will determine the cents.

If your bid is equal to or higher than the random price, you will receive the \$3 and pay the random price. You will thus earn an amount equal to: \$3 *minus* the random price. If you make a purchase at a price that is higher than \$3, you will in fact have negative earnings (lose money). If your bid is less than the random price, you will not receive the \$3. You will not pay the random price. You will earn \$0.

Please determine your bid at this time and write it here: \$ _____ . _____

Trial 1

Random price: \$ _____ . _____

Is your bid equal to or higher than the random price? (check the box below)

- ☐ **Yes.** You bought the \$3. Your earnings are: \$3 – random price = \$ _____ . _____
- ☐ **No.** You did not buy the \$3. Your earning are for this trial are \$0.
-

Trial 2

Random price: \$ _____ . _____

Is your bid equal to or higher than the random price? (check the box below)

- ☐ **Yes.** You bought the \$3. Your earnings are: \$3 – random price = \$ _____ . _____
- ☐ **No.** You did not buy the \$3. Your earning are for this trial are \$0.
-

Trial 3

Random price: \$ _____ . _____

Is your bid equal to or higher than the random price? (check the box below)

- ☐ **Yes.** You bought the \$3. Your earnings are: \$3 – random price = \$ _____ . _____
- ☐ **No.** You did not buy the \$3. Your earning are for this trial are \$0.
-

Trial 4

Random price: \$ _____ . _____

Is your bid equal to or higher than the random price? (check the box below)

- ☐ **Yes.** You bought the \$3. Your earnings are: \$3 – random price = \$ _____ . _____
- ☐ **No.** You did not buy the \$3. Your earning are for this trial are \$0.

Training Round 3

You will not have the ability to earn money in this round. Instead, you will have the opportunity to buy a real grocery product. If you buy something, you will pay for it out of your earnings.

As before, if your bid is equal to or higher than the random price, you will buy the item and pay the random price. Otherwise, you will not buy the item.

The previous training rounds emphasize that it is in your best interest to place a bid equal to the highest price you are willing to pay for the item. By doing so, you will only purchase the item at prices you are willing to pay. You will not purchase the item at prices you are not willing to pay.

If you instead bid lower than this highest price you are willing to pay, you risk not purchasing the item at prices favorable to you.

If you instead bid more than this highest price you are willing to pay, you risk purchasing the item at prices that are not favorable to you.

The moderator will now show you two items to choose from: a candy bar and a granola bar. Please choose which item you would like to bid on and check the appropriate box below:

☐ Candy Bar

☐ Granola Bar

The random price will be an amount between \$0.00 and the maximum expected bid (based on previous research). A volunteer will roll dice to determine the price, and each price within the range will be equally likely.

Please determine your bid at this time. **You will not be able to change your bid after the random price is determined.**

Your bid: \$_____ . _____

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

- ☐ **Yes.** You will receive the item you bid on. The price will be subtracted from your earnings as a cost.
- ☐ **No.** You did not buy anything. No money will be subtracted from your earnings.

Record your cost for this round on your Record Sheet.

Main Experiment – Part 1

As in the last training round, you will not have the ability to earn money in the following decision rounds. Instead, you will have the opportunity to buy real grocery products. If you buy something, you will pay for it out of your earnings.

In this Part of the experiment, you will place bids for a variety of grocery items. You will not see the actual items, but you will be provided with an information label that describes the main characteristics of the item such as the type of item, and its size, flavor, etc. The information on the labels is accurate: the actual label on the real grocery item does reflect the information that you see in information label. In this Part, you will be asked to place bids for 6 different grocery items.

We have all of the grocery items in the lab today. All of the items have been recently purchased from popular grocery stores in Knoxville.

The procedure for purchasing items is the same as before. Your bid will be compared to a random price. You will purchase the good only if your bid is equal to or higher than the price. The random price will be a randomly drawn number between \$0.00 and the maximum expected bid (based on previous research). A volunteer will roll dice to determine the price, and each price within the range will be equally likely. The range of prices will not be the same as in the last training round.

As before, it is in your best interest to place a bid equal to the highest price you are willing to pay for the item. By doing so, you will only purchase the item at prices you are willing to pay. You will not purchase the item at prices you are not willing to pay.

If you instead bid lower than the highest price you are willing to pay, you risk not purchasing the item at prices favorable to you. If you instead bid more than this highest price you are willing to pay, you risk purchasing the item at prices that are not favorable to you.

Only one round from this Part will be implemented for real. We will have a volunteer roll a six-sided die to determine which round this is. Since you will not know which round will be selected prior to making any decisions, it is in your best interest to take each decision seriously as if it will determine an actual purchase.

On the next page you will see important information that is specific to you. Please read this information, and then proceed to place your bids. You have six minutes to do so.

Please Read:

The grocery items in this section are not advertised as “Natural”, “All Natural”, or “100% Natural” on their labels.

Please proceed to the labels now.

Part 1, Round 1

Potato Chips

Type: Kettle Cooked

Flavor: Original (Plain)

Size: 8.5 oz bag

Brand: A major brand such as Lays, Utz, Herr's, Wise, or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If “Yes”, you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If “No”, you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Part 1, Round 2

Toothpaste

Flavor: Clean Mint

Features: Fluoride Toothpaste, Clinically Proven, Cavity Protection,
Whitening, American Dental Association (ADA) Accepted

Size: approximately 4 oz tube

Brand: A major brand such as Crest, Colgate, Aquafresh, Arm and
Hammer, or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If “Yes”, you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If “No”, you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Part 1, Round 3

Wheat Crackers

Features: Round with small holes

Size: 8 oz box

Brand: A major brand such as Ritz, Toasteds, Carr's, Milton's,
Breton or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If “Yes”, you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If “No”, you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Part 1, Round 4

Peanut Butter

Type: Creamy

Features: 0 grams trans fat per serving

Size: 16 oz (1 lb) jar

Brand: A major brand such as Jif, Skippy, Smucker's, Planters, or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If "Yes", you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If "No", you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Part 1, Round 5

Cooking Oil

Type: Canola Oil

Features: 0 grams trans fat per serving, cholesterol free, 1 gram saturated fat per serving, 0 mg sodium per serving, 0 grams sugar per serving

Size: 48 fl oz (1.5 quart) bottle

Brand: A major brand such as Crisco, Wesson, Mazola, or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If “Yes”, you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If “No”, you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Part 1, Round 6

Wheat Cereal

Type: Frosted

Features: Bite-size, square-shaped cereal

Size: approximately 20 oz box

Brand: A major brand such as Kellogg's, Post, General Mills,
Nestle, or similar

Please read the label carefully before determining your bid. Since the purchase procedure is the same as before, remember that it will be in your best interest to place a bid equal to the highest price you are willing to pay for the item.

Your bid (in dollars): \$_____ . _____

The volunteer will now determine the random price. Please write the random price here:

Random price: \$_____ . _____

Is your bid equal to or higher than the random price? (Circle one.) Yes No

If "Yes", you bought the item (and will receive it at the end of the session)

Your cost is the random price, which is: \$ _____ . _____

If "No", you did not buy the item. Your cost is \$0.00.

On your record sheet, please indicate whether you made a purchase. If so, please record the cost of the item as well as indicate what you bought. Otherwise, simply record a cost of \$0.

Appendix D

Survey

Please answer the following questions. **Please write/check clearly.** Your responses will not be connected to your name or other identifying information.

1. What is your age? _____
2. What is your gender? (Please check one.)
 - ☐ Male
 - ☐ Female
3. What is the highest level of education you have attained? (Please check one.)
 - ☐ High School Diploma
 - ☐ Associate's Degree
 - ☐ Bachelor's Degree
 - ☐ Master's Degree
 - ☐ Ph.D. Degree
 - ☐ Other Advanced Certification/Professional Studies
4. Were you a student during the 2013-2014 academic year? (Please check one.)
 - ☐ Yes, full-time
 - ☐ Yes, part-time
 - ☐ No
5. Which of the following best describes your current employment status? (Please check one.)
 - ☐ Employed, full-time
 - ☐ Employed, part-time
 - ☐ Self-employed
 - ☐ Unemployed or retired
6. What is your marital status? (Please check one.)
 - ☐ Single
 - ☐ Married
 - ☐ Widowed
 - ☐ Divorced

7. Are you the primary grocery shopper in your household? (Please check one.)
- ☐ Yes
 - ☐ No
8. In 2013, what was your annual household income, before taxes? (Please check one.)
- ☐ \$5,000 or less
 - ☐ \$5,000 - \$10,000
 - ☐ \$10,001 - \$20,000
 - ☐ \$20,001 - \$40,000
 - ☐ \$40,001 - \$60,000
 - ☐ \$60,001 - \$80,000
 - ☐ \$80,001 - \$100,000
 - ☐ \$100,001 - \$120,000
 - ☐ \$120,001 - \$140,000
 - ☐ More than \$140,000
9. How many adults live in your household? _____
10. How many children live in your household? _____
11. Do you have any children living in your household under the age of 10? (Please check one.)
- ☐ Yes
 - ☐ No
12. How much does your household typically spend each week on groceries? (Please check one.)
- ☐ \$30 or less
 - ☐ \$31 - \$50
 - ☐ \$51 - \$100
 - ☐ \$101 - \$150
 - ☐ \$151 - \$200
 - ☐ \$201 - \$250
 - ☐ \$251 - \$300
 - ☐ \$301 - \$350
 - ☐ More than \$350

13. How many trips do you typically take to buy groceries each week? (Please check one.)

- ☐ Less than 1
- ☐ 1
- ☐ 2
- ☐ 3
- ☐ 4
- ☐ 5
- ☐ More than 5

14. Where have you shopped for groceries during the last month? (Please check all that apply.)

- ☐ Kroger
- ☐ Food City
- ☐ Ingles
- ☐ Publix
- ☐ Earth Fare
- ☐ Fresh Market
- ☐ Target
- ☐ Walmart
- ☐ Sam's Club
- ☐ Costco
- ☐ CVS / Walgreens / other convenience store
- ☐ Local farmer's markets
- ☐ Amazon or other online grocery retailer
- ☐ Other: _____

15. How many meals do you typically eat out (at a restaurant or fast food establishment) per week? (Please check one.)

- ☐ 0
- ☐ 1-5
- ☐ 6-10
- ☐ 11-15
- ☐ More than 15

16. When grocery shopping, do you specifically look for products that meet certain dietary restrictions? (If yes, please check all that apply. If no, please leave blank.)

- ☐ Gluten-free
- ☐ Lactose-free
- ☐ Vegetarian
- ☐ Vegan
- ☐ Low carb
- ☐ Low fat
- ☐ No nuts
- ☐ No soy
- ☐ No fish
- ☐ Other: _____

17. How important are each of the following to you when you purchase food? (Please check one box per row.)

	Not Important	Somewhat Important	Very Important
Nutrition Information			
Ingredients			
Country where food was produced			
Brand			
Package size			
Price			
Health-related labelling such as “gluten-free”			
Environmental-related labelling such as “Rainforest Alliance Certified”			
Other:			

18. When grocery shopping, how often do you notice labels that say “**natural**” or “**all natural**”? (Please check one.)

- ☐ Almost always or always
- ☐ Frequently
- ☐ Occasionally
- ☐ Almost never or never

19. When grocery shopping and when such an option is available, how often do you purchase foods that are labelled “**natural**” or “**all natural**”? (Please check one.)

- ☐ Almost always or always
- ☐ Frequently
- ☐ Occasionally
- ☐ Almost never or never

20. What do you think the phrases “**natural**” or “**all natural**” actually mean when printed on a food label? (Please check all that apply.)

- ☐ No artificial flavors
- ☐ No artificial colors
- ☐ No artificial preservatives
- ☐ No genetically modified ingredients
- ☐ No pesticides
- ☐ No dyes
- ☐ Limited processing
- ☐ Higher quality ingredients
- ☐ Environmentally-friendly
- ☐ Organic
- ☐ Other: _____
- ☐ None of these

21. What do you think the phrases “**natural**” or “**all natural**” should mean when printed on a food label? (Please check all that apply.)

- ☐ No artificial flavors
- ☐ No artificial colors
- ☐ No artificial preservatives
- ☐ No genetically modified ingredients
- ☐ No pesticides
- ☐ No dyes
- ☐ Limited processing
- ☐ Higher quality ingredients
- ☐ Environmentally-friendly
- ☐ Organic
- ☐ Other: _____
- ☐ None of these

22. Please rate your agreement with the following statement on a scale from 1 to 5. (Circle one number.)

“Foods labelled **“natural”** or **“all natural”** are healthier than foods without a **“natural”** or **“all natural”** label.”

I completely disagree				I completely agree
1	2	3	4	5

23. When you see a food label that says **“organic”**, do you trust that the product is truly organic? Please rate your level of trust on a scale from 1 to 5. (Circle one number.)

I do not trust the label at all				I completely trust the label
1	2	3	4	5

24. Please rate your knowledge of **genetically modified organisms (GMOs)** on a scale from 1 to 5. (Circle one number.) If you had never heard of **GMOs** prior to this question, please circle “1”.

I know nothing about GMOs				I know a lot about GMOs
1	2	3	4	5

25. Please rate your agreement with the following statement on a scale from 1 to 5. (Circle one number.)

“I am concerned that food products containing **genetically modified organisms (GMOs)** pose a health risk.”

I completely disagree				I completely agree
1	2	3	4	5

26. Please rate your agreement with the following statement on a scale from 1 to 5. (Circle one number.)

“I am concerned that the production of products containing **genetically modified organisms (GMOs)** poses an environmental risk.”

I completely disagree				I completely agree
1	2	3	4	5

27. When you see a food label that says “**Non-GMO**” or “**Non-GMO Project Verified**”, do you trust that the product is free of **genetically modified organisms (GMOs)**? Please rate your level of trust on a scale from 1 to 5. (Circle one number.)

I do not trust the label at all				I completely trust the label
1	2	3	4	5

28. Please rate your agreement with the following statement on a scale from 1 to 5. (Circle one number.)

“I am concerned that we as a society do not fully understand the impacts of **genetically modified organisms (GMOs)**.”

I completely disagree				I completely agree
1	2	3	4	5

29. Are you willing to pay extra for a food item that does not contain **genetically modified organisms (GMOs)**? (Please check one.)

- ☐ Yes
☐ No

30. Did you feel that you were well-compensated for your participation in this experiment? Please rate your satisfaction with the compensation on a scale from 1 to 5. (Circle one number.)

I was compensated very poorly				I was compensated very well
1	2	3	4	5

31. Did you understand the instructions for the experiment today? Please rate your understanding on a scale from 1 to 5. (Circle one number.)

I understood very poorly				I understood very well
1	2	3	4	5

CONCLUSION

This dissertation entitled, “Essays in Resource Allocation Efficiency and Behavior,” is comprised of three papers. The first examines bidder’s choice auctions using a field experiment and a laboratory experiment. Auctions are frequently used to allocate resources efficiently. However, we show that the bidder’s choice mechanism, which is sensitive to the effects of risk aversion, does not command as high of a premium as previously predicted by the literature under certain circumstances. Thus, sellers should evaluate their bidders and situation to determine whether the bidder’s choice format is likely to be helpful.

In the second chapter, theoretical modelling is used to show that managers in public organizations can use task incentives to motivate agents when pecuniary methods are unavailable. Managers in schools and government agencies are frequently faced with this constraint. Tournaments have long been used to model competition in the labor market. This paper illustrates how agents can be motivated to put forth optimal effort in a distinct setting characterized by task assignment, though inefficiency is created.

Finally, inefficiency in the market for “natural” food is highlighted in Chapter 3. Although the phrase “natural” is unregulated in the United States, most consumers do believe that the label indicates at least one characteristic such as “no artificial colors” or “no genetically modified ingredients”. This study is comprised of both a survey where beliefs are elicited and an incentive-compatible purchase procedure where consumers are motivated to reveal their willingness to pay for food items with “natural” and standard labels. The results indicate that a policy which regulated or defined the “natural” phrase would be welfare-enhancing.

VITA

Julianna Butler is from La Plata, Maryland. She attended Salisbury University in Maryland for her undergraduate education and earned a Bachelor's degree in economics. After working for a year at a transportation logistics company in Delaware, she accepted a graduate assistantship at the University of Tennessee, Knoxville. Julianna has taught a number of classes and has been involved in a variety of research projects as part of her assistantship. In 2010, she completed her Master of Arts degree in economics. She is planning to finish her Ph.D. in economics during the summer of 2014 and will begin working at the University of Delaware in September 2014.