Risky Choices over Goods

Patrick DeJarnette*

This Draft: September 1st, 2024

Abstract

This paper examines how risk preferences differ over goods and in-kind monetary rewards. I study an experiment in which control subjects allocate self-selected Amazon.com goods over uncertain states, whereas treated subjects allocate temporally-restricted Amazon.com credit instead. Under perfect information, allocations would be identical between these groups. In practice, subjects demonstrate considerable differences, with credit allocations 4 times more likely to be riskless. Using an information treatment, I find no evidence that price or product uncertainty explains these differences. An additional experiment demonstrates these differences do not exist in a risk-free environment requiring real effort in exchange for credit or goods.

Keywords: risk preferences, goods, uncertainty.

*Faculty of Political Science & Economics, Waseda University. pdejar@waseda.jp. I thank Jeremy Tobacman, Iwan Barankay, Daniel Gottlieb, Rob Jensen, and Judd Kessler for detailed feedback. This paper has also greatly benefited from the comments of Charles Sprenger, Muriel Niederle, and participants of the Asia Pacific Economic Science Association Conference, the World Meeting of the Economic Science Association, and the Stanford Institute of Theoretical Economics. This research was made possible thanks to research funding from the Mack Institute for Innovation Management and the Ministry of Science and Technology. All errors are my own. The most recent version of this paper may be found at www.patdejar.net.
JEL Codes: D80, D12, D83, C91, D90
Keywords: Multidimensional risk, good lotteries, product search

1 Introduction

Many important decisions involve risk, including insurance, portfolio choice, health behaviors, and labor relations. One standard way of encapsulating these risk preferences is to model a utility function over a single wealth variable. These measures of risk aversion have been shown to be predictive for a number of other risky behaviors. For example, Anderson and Mellor [2008] report monetary risk aversion correlates with cigarette smoking, heavy drinking, obesity and seat belt non-use. In Dohmen et al. [2011], measures of risk preferences from self-reported surveys predict stock investment, self-employment, smoking, and sports participation.[1]

However, in other settings, risk preferences seemed less stable across contexts. Sutter et al. [2013] found experimental measures of risk and ambiguity preferences only weakly predicted other behaviors among adolescents. In Barseghyan et al. [2011], demand for insurance (as measured by deductibles) was substantially different over two different goods, houses and cars. In Einav et al. [2012], demand for different types of insurance appears to be correlated, but does not correlate well with the riskiness of 401(k) investments.[2] Perhaps as a result, recent literature advocates to expand measurement methodology and test the stability of risk preferences across domains [O'Donoghue and Somerville, 2018, Schildberg-Hörisch, 2018].

As many important risks involve non-monetary outcomes, such as smoking and seat belt use, it might be beneficial to assess risk preferences over non-monetary lotteries. However, when quantifying these preferences, it may be unclear how to create "comparable" risks between monetary and non-monetary lotteries. For example, lotteries over varying number of peanuts would constitute risk over non-monetary outcomes, but may not be very generalizable to other settings, especially if the subject was allergic to peanuts. That is, the utility curvature over one dimension (one good) may not be very predictive of utility curvature over other dimensions.

One solution for how to create 'equivalent' monetary and non-monetary lotteries is to elicit monetary equivalents for each non-monetary outcome, and then investigate the implied

[1] Surprisingly, they also report that a single question about risk taking was the best all-around predictor of risky behavior. Also worth noting is the module recently designed by Falk et al. [2023], which identifies a hypothetical survey approach to best predict incentivized answers at lower cost.

[2] It may be worth noting that Einav et al. [2012] focuses on how individuals rank relative to their peers. As a result, they are less focused on testing "absolute" differences between risk preferences over different goods and instead at the reliability of how individual will rank relative to others.

risk aversion using standard techniques. When doing so in a between-subject study, Gneezy et al. [2006] find that subjects value risky non-monetary lotteries at less than the certain equivalent of the worst of the two outcomes. This *uncertainty effect* occured both in a laboratory setting using goods such as bookstore gift cards, but also in the field using baseball cards at sportscard shows.[3] This initial finding has largely[4] been replicated, though research on what causes this uncertainty effect is ongoing [Simonsohn, 2009, Andreoni and Sprenger, 2011, Newman and Mochon, 2012, Wang et al., 2013, Mislavsky and Simonsohn, 2018].[5]

This paper presents a novel approach to testing for equivalence of risk preferences over non-monetary goods and monetary lotteries that does not require eliciting monetary equivalents. Rather than elicit monetary equivalents over predetermined goods, this methodology instead allows subjects to choose optimal non-monetary bundles for a given 'budget' across uncertain states. To map these lotteries to economic decisions out of the laboratory, I employ a familiar marketplace, Amazon.com, which features millions of goods.

In this experiment, the control group allocates self-chosen Amazon.com goods (books, clothing, etc.) across uncertain states. One state is then selected randomly with equal probability.

In the primary treatment group, subjects instead allocates temporally-restricted Amazon.com credit amounts ($1, $2, $2.57, etc.) across states. As before, one state is then selected randomly with equal probability. After this resolution, subjects must spend credit, as informed to them in advance. To remove savings concerns, any unspent credit is of no value and discarded.[6]

Thus, one could view the difference between control and treatment as a 'framing' device to look at how individuals perceive monetary and non-monetary risk. Because at the end of the realization, both control and treatment subjects ultimately receive a single Amazon.com good of their own choosing. The only difference is the timing of the selection of goods –

[3]However, when exploring the concept using within-subject elicitations, 29 out of 30 subjects satisfied the internality axiom.

[4]The only study known to the author that fails to replicate appears to be Rydval et al. [2009] who proposes that subjects did not understand the instructions for the lotteries.

[5]Note that both Gneezy et al. [2006] and Newman and Mochon [2012] find the uncertainty effect disappears for within-subject studies. This might indicate a different decision process, perhaps a demand for consistency or arguably experimenter demand, as the 'norm' for the internality axiom may be strong. However, whether between-subject or within-subject risk attitudes for non-monetary outcomes is a better predictor for risky behaviors remains an unanswered question.

[6]Subjects were quizzed on this and several other topics to ensure understanding. See design section for more details.

whether subjects select goods before the final state is randomly chosen (control) or only after the final state is randomly chosen with the allocated credit (treatment).

To explore different stakes and test secondary hypotheses, in both control and treatment, the total value of the allocations is predetermined to be either \$20 or \$100. As a simple example, suppose an individual is given \$20 of credit to allocate across two potential states. If the individual allocates \$10 to each of the states, she would receive \$10 of credit for sure, which could be used to purchase any good on Amazon.com priced at \$10 or less. When allocating goods rather than credit, she can select up to 2 goods totaling \$20 of value. Intuitively, if an individual allocates a \$15 good to state 1 and a \$5 good to state 2 when selecting goods, it might be surprising that she allocates \$10 to both states when selecting credit. Specifically, since the credit must be spent immediately, why would the individual not allocate \$15 and \$5 of credit, and purchase the corresponding good after the uncertainty is resolved? Or put another way, if the individual prefered a risk-less allocation of \$10 credit for sure, why did they not choose the risk-less allocation of a \$10 (or less) good for sure?

I expand on this intuition theoretically by proving that if individuals treat self-selected goods and time-constrained credit identically *without* uncertainty, there should be no difference between good and credit allocations *with* uncertainty. Thus if a subject allocates a \$8 item and a \$12 item in the control, with perfect information he would theoretically allocate a \$8 credit and a \$12 credit if in the credit treatment. Although this theoretical result relies on transitivity and a modified independence assumption, it does not assume anything else regarding a utility function or risk preferences of the individual.

Contrary to this prediction, subjects exhibited considerably more aversion to risky allocations when selecting credit. Subjects were four times as likely to place "riskless" allocations with credit than they were with goods. Furthermore, the average standard deviation of credit allocations was about two-thirds that of the average standard deviation of good prices. To analyze whether these differences could be driven by price uncertainty, subjects are randomly forced to spend more time on Amazon.com, but this does not seem to influence the allocations (with a rather precise zero effect).

Although this is the first research (known to the author) to explicitly test risk preferences between monetary and *self-selected* non-monetary lotteries, an earlier theoretical literature uncovered several implicit assumptions about uncertainty over goods and money. Grant et al. [1992] addresses requirements for when preferences over goods lotteries can be reconstructed from preferences over monetary lotteries. The authors then go on to establish implications for what risk aversion over monetary lotteries implies about risk aversion over good lotteries. However, rather than assuming preferences over monetary lotteries are induced by underlying preferences for good lotteries, I outline precisely what assumptions will generate indifference between a monetary lottery and self-selected good lotteries.[7]

An additional approach to testing for risk (or ambiguity) aversion is to provide additional uncertainty to existing non-monetary lotteries. Employing a field experiment with attendees at a coin conference, Harrison et al. [2007] investigated risk aversion across monetary and coin lotteries. Each coin had either a known or unknown professional coin grading, and subjects displayed considerably higher risk aversion parameters in unknown grade coin lotteries. Therefore this study, like Gneezy et al. [2006], would suggest that the non-monetary allocations should have displayed more aversion to risk (more riskless allocations) compared to in-kind credit allocations, and are unable to explain the differences.

One possibility for the observed difference in risk preferences over money and goods is that individuals only have noisy estimations of their indirect utility function, akin to the model of "cognitive uncertainty" in Enke and Graeber [2023]. In other words, only when they observe a $20 Star Wars Blu-ray do they realize the true 'utility' of $20 of credit. This might also be motivated by the extensive literature on the endowment effect, which suggests that subjects who receive a good value that good more than subjects who do not [c.f. Knetsch, 1989, Kahneman et al., 1991, Bordalo et al., 2012].[8] This 'noisy estimation' of value is also supported by a growing literature on salience over monetary risk, as in Bordalo et al. [2010], and on salience for goods without uncertainty, as in Bordalo et al. [2012], Kőszegi and Szeidl [2013], Gabaix [2014].

While the aforementioned information treatment should reduce these errors, I employ two additional tests of this hypothesis. First, I analyze subjects in the risky choice experiment who initially received the "goods" control in period 1 and were later given the "money" treatment in period 2. If the cognitive costs of selecting goods is the primary concern, then having recently chosen specific goods should reduce the propensity for an equal distribution of credit. I find that, perhaps surprisingly, this does not seem to be the case.

[7] To the credit of Grant et al. [1992], they acknowledge this alternate approach in footnote 7, even though it was not the main focus of that study.

[8] On a related note, Bushong et al. [2010] also finds differences between in-person valuations of goods compared to text or image displays. However, in the experiments presented here, all goods were selected through image-and-text displays via Amazon.com or PCHome.com.tw.

Lastly, an additional real-effort laboratory experiment was conducted in a risk-free setting. In this experiment, subjects were given up to 10 pages of text-reversal tasks and a schedule of potential earnings. In the control group, subjects were asked to select one good for each possible earnings level from a large online retailer.[9] In the treatment, subjects first earned credit and then selected a good on the website after working. Thus, as in the risky choice experiment, the primary difference is the timing of selecting the good (with all subjects knowing in advance they would ultimately receive a self-selected good).

If subjects have an incorrect approximation of their own indirect utility function, one would also expect differences in effort between control and treatment in this risk-free experiment. There were no such significant differences despite a comparable sample size. As in the risky choice experiment, an information treatment was also implemented and had no significant impact on effort allocation.

Consumers also face decisions daily about whether to purchase products running promotional contests [Dhar and Simonson, 1992]. These contests pose somewhat of a mystery, given that they often feature pre-selected prizes rather than equivalent cash values. This study also indicates another possibility – individuals may wish to engage in risk over goods but prefer to avoid risk with equivalent cash prizes. This finding is broadly consistent with findings from the marketing literature on these promotions [Goldsmith and Amir, 2010, Laran and Tsiros, 2013]. This may have important implications for government run lotteries, which often serve to fund public programs [Landry and Price, 2007]. By adding physical items to these lotteries, it may be possible to encourage risk seeking behavior from participants and generate additional revenue for publicly funded programs.

The remainder of the paper is organized as follows. Section 2 demonstrates theoretical predictions. Section 3 outlines the primary experiment design. Section 4 presents the main results. Section 5 provides the additional evidence from the real-effort experiment without uncertainty, and Section 6 concludes.

[9] As Amazon did not deliver to Taiwan, where the risk-free experiment was conducted, the closest equivalent (PCHome.com.tw) was used instead. At the time, PCHome.com.tw was the largest online marketplace in Taiwan with 5 million SKUs, 2 million of which had 24 hour-or-less delivery.

2 Theory

In this section, I demonstrate that under perfect information of prices, risk preferences across money and goods should be the same in a static setting.

2.1 Theoretical Equivalence of Allocations

To start, I define a concept of an extended lottery that includes both a monetary "good" and non-monetary goods.

Let the set $\mathbb{S}$ refer to the set of states $s = 1, 2, ..., S$ which occur with associated probabilities γ_1, γ_2, ..., γ_S. Assume there are goods $n = 1, 2, ...N$ whose consumption is state contigent, $g_{n,s}$ an element of the good-specific set $G_n \subset \mathbb{R}_+$, as well as a monetary good for each state, $m_s \in \mathbb{R}_+$. Thus, any particular lottery L is defined by the matrix $(\gamma_1, m_1, g_{1,1}, g_{2,1}, ..., g_{N,1}; \gamma_2, m_2, g_{1,2}, g_{2,2}, ..., g_{N,2}; ...; \gamma_S, m_S, g_{1,S}, g_{2,S}, ..., g_{N,S})$, an element of $([0, 1] \times \mathbb{R}_+ \times G_1 \times G_2 \times ... \times G_N)^S$ For simplicity, I will denote this product space $\mathbb{L}^S$.[10] Thus, both the probabilities and the outcomes of the states are included in the lottery structure.

Any such lottery can also be written as a combination of degenerate lotteries L_s, where each $L_s \equiv (1, m_s, g_{1,s}, g_{2,s}, ..., g_{N,s}) \in \mathbb{L}^1$. Thus any lottery L may be written[11] as $L = (\gamma_1 L_1,\ \gamma_2 L_2,\ ...,\ \gamma_S L_S)$ where the product $\gamma_s L_s$ refers to $(\gamma_s, m_s, g_{1,s}, g_{2,s}, ..., g_{N,s}) \in \mathbb{L}^1$. I also assume the states are mutually exclusive and at most one state will occur, that is $\sum \gamma_s \leq 1$ and $Prob(s \,\text{occurs} \,|\, s' \neq s \,\text{occurs}) = 0$, which may make later assumptions of state independence more plausible.[12]

In addition, define the market for state s as a vector of prices $P_s = (p_{1,s}, p_{2,s}, ..., p_{N,s}) \in \mathbb{R}^N_+$. The market consists of a vector that consists of the individual market states $P =$

[10] A very similar structure could be constructed in which the goods have state-specific sets $G_{n,s}$. For example, to accomodate a 50% chance to order on Amazon.com and a 50% chance to order on Walmart.com, which feature different good sets available. This extension would primarily just require additional notation.

[11] That is, the operation $a \cdot L_s : [0, 1] \times \mathbb{L}^1 \to \mathbb{L}^1$ is defined as $(a\gamma_s, m_s, g_{1,s}, g_{2,s}, ..., g_{N,s})$ or equivalently $L_S \times \begin{bmatrix} a & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}$ but is omitted for clarity.

[12] If states were not mutually exclusive (with exactly one state occuring), it would be hard to believe that properties of other states would not cause preference reversals. For example, one might prefer a 100% chance of a chocolate to a 100% chance of a marshmallow. But if we had a 100% chance of marshmallow with an additional 50% chance of chocolate, this may now be preferred to a 100% chance of a chocolate with an additional 50% chance of additional chocolate.

$(P_1, P_2, ..., P_S)$.[13] The decision maker has a preference relation $\succsim_P$ over extended lotteries $\mathbb{L}^S$ for a given market P, which I further assume is a weak order relation.[14] For notational simplicity, if there are lotteries $A, B \in \mathbb{L}^R$ with $R < S$, I write $A \succsim_P B$ as a shorthand for $(A, \vec{0}) \succsim_P (B, \vec{0})$ where $\vec{0} \in \mathbb{L}^{S-R}$. In words, even though preferences are over lotteries for the entire S states, I insert the remaining states with zeroes for notational convenience of using the same preference relation.

In addition to the weak order assumption, I assume the preferences have two additional properties: (a) Monetary Equivalence Under Certainty and (b) Independence.

Monetary Equivalence Under Certainty. For a given market P,

(i) For any degenerate lottery $L_s = (1, m_s, g_{1,s}, g_{2,s}, ..., g_{N,s})$, the decision maker weakly prefers the degenerate monetary-only bundle $L_s^{'} = (1,\ m_s + \sum_n p_{n,s} g_{n,s}, 0, 0, ..., 0)$, that is $L_s^{'} \succsim_P L_s$.

(ii) For any degenerate lottery $L_s = (1, m_s, g_{1,s}, g_{2,s}, ..., g_{N,s})$, there exists a weakly preferred degenerate lottery $L_s^{''} = (1, 0, g_{1,s}^{''}, g_{2,s}^{''}, ..., g_{N,s}^{''})$ such that $\sum_n p_{n,s} g_{n,s}^{''} \leq m_s + \sum_n p_{n,s} g_{n,s}$ and $L_s^{''} \succsim_P L_s$.

In words, Monetary Equivalence (i) states that in a case with no uncertainty, the decision maker is at least as well off by converting any particular bundle into the money it would cost to purchase that bundle. Since this is true for all degenerate lotteries, including the optimal bundle(s) of goods, it also implies that there are no transaction costs to converting money into goods. For example, individuals would prefer a \$20 to socks that cost \$20, because they can always use the \$20 to purchase the socks.

Monetary Equivalence (ii) states that in a case with no uncertainty, the decision maker has no particular preference for holding onto money. In other words, money is only as useful as the things it can buy. It is worth noting that this assumption does not require that every dollar must get spent in an optimal bundle or that consumers cannot be satiated.[15] For

[13] This assumption of linear pricing is for ease of notational simplicity and could be instead considered as a vector function.

[14] In this case, because the monetary good is allowed to enter directly into the bundle, market prices may influence preferences over bundles.

[15] Although this model is being presented as a static one, the same item at different periods could be thought of as different goods – as long as the uncertainty is resolved in one period with discrete and finite time periods, the same results hold true intertemporally.

example, if goods are discrete rather than continuous, it may not be optimal to spend every last dollar.[16] However, what this assumption indicates is that any money left over after purchasing the optimal goods bundle would have no value (as they would be indifferent between that and the same goods bundle with no money). In the experiments discussed in later sections, subjects were informed that unused credit would be lost to ensure this condition should be true experimentally.

Independence Property. For a given market P and any lotteries L and $L^{'}$ in $\mathbb{L}^R$ with $R < S$, preferences are independent if $L \succsim_P L^{'}$ implies $\forall \alpha \in (0,1)$ and for all degenerate lotteries $L^{''} \in \mathbb{L}^1$, $(\alpha L,\, (1-\alpha)L^{''}) \succsim_P (\alpha L^{'},\, (1-\alpha)L^{''})$.

With monetary equivalence under uncertainty and the independence property, we can establish the following result with a bit more notation: Let $G_L = (g_{1,1}, g_{2,1}, ..., g_{N,1}, g_{1,2}, ..., g_{N,2}, ..., g_{1,S}, g_{2,S}, ...g_{N,S}) \in G_1 \times G_2 \times ... \times G_N \times G_1 \times ... \times G_N$ denote the vector of goods for a given lottery L in which all monetary values are 0. Let $G(P,\, I) = \{G_L\, s.t.\, \sum_s \sum_n p_s g_{n,s} \leq I\}$, with $sup\, G(P,\, I)$ defined using the weak ordering $\succsim_P$ in which all monetary values are set to 0. By a similar notation, let $M_L = (m_1, m_2, ..., m_S) \in \mathbb{R}_S$ denote the vector of monetary values for a lottery L in which all non-monetary goods are 0. And $M(P,\, I) = \{M_L\, s.t.\, \sum_s m_s \leq I\}$, with $sup\, M(P,\, I)$ defined using the partial ordering $\succsim_P$ in which all non-monetary goods are set to 0.

Theorem 2.1 (Monetary Equivalence Over Uncertainty)**.** *Under the assumptions of Monetary Equivalence Under Certainty and Independence, if a lottery of goods is optimal, then the monetary lottery (with equivalent value in each state) will also be optimal. In notation established above, if* $G^* = (g_{1,1}, g_{2,1}, ..., g_{N,1}, g_{1,2}, ..., g_{N,2}, ..., g_{1,S}, g_{2,S}, ...g_{N,S}) \in sup\, G(P,\, I)$*, then* $M^* = (p_{11}g_{11} + p_{21}g_{21} + ... + p_{N1}g_{N1}, p_{12}g_{12} + p_{22}g_{22} + ... + p_{N2}g_{N2}, ..., p_{1S}g_{1S} + p_{2S}g_{2S} + ... + p_{NS}g_{NS}) \in sup\, M(P,\, I)$.

Proof. For proof by contradiction, assume that the condition is true, that $G^* \in sup\, G(P)$

[16] For example, if I am buying discrete apples at \$2 and bananas at \$3 with a total of \$7 to spend, I may indeed prefer 2 bananas even though I have \$1 left over. But according to Monetary Equivalence (ii), I am indifferent between \$0 and 2 bananas and \$1 and 2 bananas in this case – because the extra \$1 cannot be used to purchase anything that would make me better off.

but that, as defined above, $M^* \notin sup\, M(P,\, I)$. Note that G^* can be rewritten as the concatenation of degenerate lotteries $G^* = \gamma_1 L_1^* + \gamma_2 L_2^* + ... + (1 - \sum \gamma_s)\, L_S^*$ where $L_s^* = (1, 0, g_{1,s}^*, g_{2,s}^*, ..., g_{N,s}^*)$ and $+$ represents the concatenation operator. Individually, each of these degenerate lotteries is weakly dominated by the degenerate lottery $L_s^{'} = (1, p_{1,s} g_{1,s}^* + p_{2,s} g_{2,s}^* + ... + p_{N,s} g_{N,s}^*, 0, ..., 0)$ via property (i) of Monetary Equivalence under Certainty. By multiple applications of the Independence assumption, this means that $(\gamma_1 L_1^*,\, \gamma_2 L_2^*, ..., \gamma_S L_S^*) \precsim_P (\gamma_1 L_1^{'},\, \gamma_2 L_2^{'}, ..., \gamma_S L_S^{'})$. However, this newly constructed compound lottery corresponds precisely to M^*.

However as $M^* \notin sup\, M(P,\, I)$ but $M^* \in M(P,\, I)$, that implies there is some $M^{**} \in M(P,\, I)$ with $M^{**} \succ_P M^*$.[17] We can rewrite this lottery as a combination of degenerate lotteries $(\gamma_1, m_1^{**}, 0, ..., 0; \gamma_2, m_2^{**}, 0, .., 0; ...; \gamma_S, m_S^{**}, 0, ..., 0) = (\gamma_1 L_1^{**},\, \gamma_2 L_2^{**},\, ...,\, \gamma_S L_S^{**})$. Yet for each of these degenerate lotteries L_s^{**}, property (ii) of Monetary Equivalence Under Certainty states that there exists a degenerate lottery $L_s^{''} = (1, 0, g_{1,s}^{**}, g_{2,s}^{**}, ..., g_{N,s}^{**})$ such that $L_s^{''} \succsim_P L_s^{**}$. Repeated application of the Independence property gives us $G^{''} \equiv (\gamma_1 L_1^{''},\, \gamma_2 L_2^{''},\, ...,\, \gamma_S L_S^{''}) \succsim_P (\gamma_1 L_1^{**},\, \gamma_2 L_2^{**},\, ...,\, \gamma_S L_S^{**})$. Thus by transitivity of weak orders, $G^{''} \succsim_P M^{**} \succ_P M^* \succsim_P G^*$. This is a contradiction, however, as G^* was in the supremum of $G(P,\, I)$ and now there is a new lottery $G^{''}$in $G(P,\, I)$ which strictly dominates it. □

2.2 Discussion regarding Equivalence

The assumptions that drive the theory in this case warrant additional discussion. First, if the preference relation is a weak order, that implies that the decision maker has both transitive and complete preferences. Transitivity of preferences over risk has been discussed as early as Tversky [1969] but more recent empirical evidence suggests that preferences can largely be summarized as transitive [c.f. Birnbaum and Gutierrez, 2007, Birnbaum and Schmidt, 2010, Regenwetter et al., 2011].[18] Completeness of preferences is harder to test, as indecision between two lotteries might be interpreted as indifference rather than an inability to prefer one to the other. This is especially difficult to test given the great number of goods available on Amazon.com.

[17]Note this M^{**} can be established even without establishing that sup M(P,I) exists. M^{**} is not necessarly in the supremum set, but the fact that M^* is not ensures that there must be something better off.

[18]However, this is an ongoing field of research. It is also possible that past research may not apply to the lotteries employed in this study, as they are arguably more intricate than some lotteries previously studied. Yet intransitive preferences would also make choosing a bundle more difficult in this setting given the subject has considerable freedom in how to allocate the credit or goods.

Regarding Monetary Equivalence under Certainty, part (i) states that the decision maker would be at least as well off with an amount of money equal to the total cost of a goods bundle. However, if decision makers are somewhat unaware of the goods available or the prices of the goods, this may not be the case. This lack of awareness could influence perceptions of what can be purchased with a given amount of money. For example, such a decision maker might strictly prefer a \$10 book to \$11, if they believed the book cost \$15 or was out of stock. This possibility will be discussed in more detail in the following subsection.

Monetary Equivalence under Certainty part (ii) states that money holds no inherent value above and beyond what can be purchased with it. In other words, with a given amount of money, the decision maker can always find a bundle that makes them at least as happy. Yet this assumption makes no mention of the psychic costs that may be associated with finding the bundle in question. In addition, this assumption may be true in our static model and the (static) experiment, but intertemporally decision makers may want to hold on to some of their money as future prices are not perfectly known.[19]

The Independence property is similar to the Independence property assumed for von Neumann-Morgenstern utility functions. In that set up, lotteries are probability distributions over fixed outcomes. If all goods are discrete, then as possible lotteries are bounded by the endowment income, then the lottery structure in section 2 could be rewritten under that framework,[20] and the Independence property would be identical.

However the Independence property has been criticized as potentially too strong an assumption. In particular, the famous Allais 'paradox' in which the chance of another lottery may cause preference reversals. Yet the relative importance and frequency of these non-independent lotteries for decision making is an ongoing debate [c.f. Rubinstein, 1988, Allais and Hagen, 2013].

[19]If future prices were perfectly known, then a good in different time periods could enter the model as different goods. However in addition to price uncertainty, there may be quantity uncertainty, e.g. car stolen, that might make Monetary Equivalence under Certainty (ii) unlikely to be true over time.

[20]If all goods are discrete and bounded by the endowment, then with a finite number of goods and states, there would only be a finite number of possible bundles. As a result, one could rewrite every possible goods bundle in every different state as a different (fixed) outcome. Though the assumption of discrete goods might be strong, this is intended only meant to highlight the relationship between the Indepedence assumptions – one does not need discrete goods in the proof above.

2.3 Theoretical Explanation

Based on the above theoretical result, substantial differences between good and monetary allocations indicate that one or more of the assumptions must not hold true. I have identified three categories of theories in the broader literature are capable of causing Monetary Equivalence Under Certainty (i) or (ii) to fail.

Although these theories can all result in disparities between allocations in goods and monetary units, additional tests can distinguish among them. These tests are discussed in more detail below and in subsequent sections detailing the results.

Noisy Ex-Ante Utility Maximization

Enke and Graeber [2023] introduces a model of "cognitive uncertainty" wherein decision makers encounter cognitive noise at the utility maximization decision. Within this framework, the default decision d may be an even allocation of money across uncertain states, and only subjects who allocate goods recieve a signal s (or a more precise signal as captured by $\lambda(N)$). This would result in allocations over goods to differ from allocations over money, as the decision maker effectively has more information about their optimal action.

While an information treatment should presumably result in an increased precision of cognitive signals, the effectiveness of this treatment might be limited. In particular, one unexplored question in Enke and Graeber [2023] is the optimal timing to resolve cognitive uncertainty. One theory of "thinking aversion", presented in Ortoleva [2013], suggests that individuals may wish to delay this cognitive process. In this case, a non-binding information treatment might have no impact on the action chosen by the individual.

Such psychic costs would threaten both Monetary Equivalence under Certainty (i) or (ii). For example, a decision maker may suffer psychic costs of translating money into optimal good bundles; this cost may be large enough that the individual strictly prefers the money to the optimal good bundle, violating (ii). Alternatively, the individual may find the decision making process distasteful enough that they would be prefer a "non-optimal" good bundle to the equivalent amount of money, violating (i).

To address this theory further, both subsample analysis and a risk-free real effort experiment were conducted. For the subsample analysis, individuals who have already completed the cognitive process of selecting goods in period 1 should no longer be impacted by the credit treatment in period 2, but this is not observed. Likewise, in a risk-free setting, if the main cause is cognitive costs, we should see a similar treatment effect for allocating money instead of goods. Instead, there is a rather precise null effect.

Ex Ante Endowment Effect

In other work, there's been mixed evidence of an endowment effect of receiving goods, possibly driven by expectation-based reference-dependent preferences [c.f. Marzilli Ericson and Fuster, 2014]. In these models, 'owning' the good can result in increased willingness-to-accept (WTA) for the good. One possibility is that decision makers that select goods may be somewhat 'endowed' with possibility of receiving the selected good, or adjust their expectations to a greater degree.

In the experiment setting, as subjects browse goods on Amazon.com, they don't want to lose the possibility of an observed good, and through their elevated WTA, are more willing to engage in risk. Evidence in Bushong et al. [2010] suggests this effect should be minimal, as the goods are only available through image and text descriptions. However, the additional risk-free experiment also serves to test this theory, as individuals should be willing to work harder to earn goods that are mentally endowed. As mentioned above, there is no evidence of this.

Regret Theory

Also in the realm of subjective uncertainty, Sarver [2008] explores preferences-over-menus with potential regret. In this framework, agents may wish to limit their options to reduce the chance of regret if their selection was ex-post inferior to alternatives.

In the context of this paper, choosing a risk-free monetary allocation reduces options across different states. In other words, choosing \$5 and \$15 allocation might result in more ex post regret than a sure-fire \$10, because there are many \$13 goods the individual could have potentially purchased if they were awarded \$5. In this framework, having chosen a menu of \$10 for sure, the \$13 goods are no longer alternatives at the time of resolution.

In comparison, when choosing goods, the decision maker is already limiting the number of goods that could be potentially compared. The \$5 good selected might still be compared to the \$15 good selected, but no others, as they are no longer alternatives. This reduction of alternatives can in turn result in less (expected) ex-post guilt, allowing for more risky behaviors.

This is similar in concept to the differences observed between "Feedback" and "No Feedback" conditions in recent experimental studies of regret aversion. In particular, Fioretti et al. [2022] explore the dynamics of regret and find evidence that individuals who receive feedback about future prices hold onto assets for longer. Thus, regret of future purchasing options may play a role when allocating money, but less salient when chosing goods.

Of particular note, versions of this theory could allow for subjects to allocate a risky distribution when allocating goods in period 1, and also risk-less distribution when allocating money in period 2. Especially if the regret is triggered by the salience of the comparison (e.g. additional attributes lower the salience of the price).

3 Experiment Design Overview

In order to test the above theories, I conducted an incentivized experiment with 124 undergraduate students at the Wharton Behavioral Lab in March 2016. Subjects sit at a computer with internet access located in a cubicle. Upon entry, they are informed and quizzed about the upcoming experiment (see Appendix for details).

During this experiment, the subjects ultimately selected goods on Amazon.com. Depending on their treatment, they select either temporally-restricted Amazon.com credit (monetary allocations) or Amazon.com goods over several mutually-exclusive uncertain states. A wide variety of goods are available on Amazon.com and most students report being previously familiar with the site,[21] making this an ideal environment for measuring risk preferences over goods and money.

The static decision is a 2x2x2 design, with subjects allocating either {credit or goods} worth a total of {$20 or $100} and is {required or not required} to spend an extra 3 minutes browsing Amazon.com. Subjects perform this procedure over two rounds (randomly ordered).[22] In one round, subjects are given an allocation total of $20 but there are only 2 potential mutually-exclusive states, each occuring with 50% probability. In the other round, subjects are given an allocation total of $100, but to keep the average payout the same, there are 10 potential mutually-exclusive states.

For example, subjects are given $20 of Amazon.com credit to allocate over two states, each of which occurs with 50% probability. In this case, a typical "risk averse" decision

[21] As of 2015, it is estimated that there are between 300 and 400 million unique items sold on Amazon.com.
[22] Every subject receives both $20 or $100 treatments, but it is randomized which occurs first.

would be allocate \$10 of credit to State 1 and \$10 of credit to State 2, thus ensuring that regardless of which state occurs, \$10 of Amazon.com credit will be selected. Subjects are then allowed to spend the awarded credit on a single Amazon.com good, with any unspent credit being lost.

Alternatively, the subject might be given \$20 of Amazon.com credit, but rather than asked to allocate the credit, the subject selects Amazon.com goods whose prices add up to \$20 or less. In other words, the subject determines what goods to "spend" the credit on before the uncertainty is resolved.

One important consideration is the intertemporal fungibility of the Amazon.com credit. To properly test the theory outlined above, any "unspent" credit must have no value. If this were not the case, the Monetary Equivalence Under Uncertainty assumption is unlikely to hold and optimal credit allocations may indeed be different from goods allocations (due to motive to save). As a result, subjects were informed and quizzed that no matter what state is selected for payout, all unspent credit would be lost.[23] That is at most a single item would be selected at the end of the session.[24]

To remove concerns about "shrouded attributes" [c.f. Gabaix et al., 2006, Chetty et al., 2009, Brown et al., 2010], only the list price of the good is considered. Subjects are informed and quizzed that only the list price will count toward their total, not shipping or tax. In addition, for any URL entered, the browser instantaneously used the Amazon Affiliate API to calculate the price of the item. At the same time, a "total value counter" at the bottom of the page informed subjects about the remaining credit available. By employing these measures, subjects can be fully informed of the price of any searched good. Thus, the only remaining "price uncertainty" should be over the unsearched goods. For example, I do not know the prices of all books, but once I search for The Wealth of Nations, price uncertainty for that particular book should be resolved. Without tax or shipping concerns, there are no further mental calculations required.[25]

[23] 2 subjects selected physical Amazon.com gift cards using either used their Amazon.com credit or Amazon.com items. This was not explicitly discouraged and could be viewed as a subject attempting to make both the Amazon.com goods and Amazon.com credit fungible. However, the in-laboratory credit still held the property that unspent credit was lost – that is, purchasing a \$15 Amazon.com gift card with \$20 of credit would result in \$5 lost. However, dropping these individuals makes no difference to the qualitative results or significance.

[24] To be precise, a single "good" means a single Amazon.com url. If there is a url that has a 6-pack of soda, this still qualifies as a single "good".

[25] To further simplify things, subjects are informed that only the "default" seller price matters. This was primarily Amazon.com itself.

To further address concerns about price uncertainty, a random sample of subjects were selected to spend an extra 5 minutes on Amazon.com to help understand the potential for price or product uncertainty driving potential results. This will be discussed in more detail in Section 4.

Prior to being allowed to start each period, the subjects had to correctly answer questions about the upcoming period, as seen in Appendix Figures 1 and 2. These procedures were implemented to ensure subjects fully understood the incentives they faced. To remove any subject overlap, the computer cookies and browsing history were also cleared in between sessions.

3.1 Risk Measurement

The equivalence result in Section 2.2 relies on minimal assumptions to clarify the role of preferences in determining optimal allocations across money (credit) and goods. However, as the theory above does not generate a parameterization of risk preferences, it bears further discussion how to measure "risk aversion" in this setting.[26]

In the case where Monetary Equivalence Under Certainty holds, if an optimal (given budget) degenerate lottery with goods exists, the individual would be indifferent between that lottery and the degenerate lottery with money equal to the total cost of the goods. Thus, holding prices fixed, one could use the preferences across the subset of optimal good-only degenerate lotteries to construct an indirect utility function over money, $V(w)$, which would be weakly increasing in w.

As a sketch of a proof that such a $V(w)$ would be weakly increasing in w, Monetary Equivalence Under Certainty claims that for a given monetary budget w, there is a good allocation that is affordable and is weakly preferred, which I denote $\vec{g}$. A higher $w' > w$ could always be used to purchase the previous goods-only bundle $\vec{g}$. Thus, any optimal goods bundle for w' must be weakly preferred to that $\vec{g}$ which is in turn weakly preferred to w.

This $V(w)$ approach bears similarity to previous treatments of indirect utility functions without uncertainty, which could then allow for definitions and parameterizations of risk

[26]Note that given that market prices would influence preferences between monetary and non-monetary goods, it would in general not support a utility function that is independent of prices. In other words, the choice of $3 versus a sock depends critically on what goods can be purchased with $3, which is determined by prices.

for monetary lotteries holding prices fixed. Yet note that for non-monetary lotteries, this approach would not tell us how to compare a lottery of L_1 ={\$3 pair of shoelaces with 50% probability and \$3 notebook with 50% probability} to a lottery L_2 = {\$10 CD with 50% probability and a \$10 DVD with 50% probability}. Even though the latter goods are more expensive, if the individual does not like CDs or DVDs, they may prefer L_1 to L_2. In other words, the theoretical approach above would only allow for comparison of self-selected bundles, that is, bundles that are optimal given the budget constraint, not any bundle of goods.[27]

Among the subset of optimal good bundles, one method to compare risk aversion across subjects would be the an elicitation of the certainty equivalence (CE). If a \$15 copy of *A Game of Thrones* is the optimal \$15 good for subject A, and a \$5 pen is the optimal \$5 good for subject A, then researchers could elicit the certainty equivilance for a lottery with a 50% chance of receiving either. This CE could be elicited either in money or credit, or potentialy in the form of a good. Another subject, subject B, may not like *A Game of Thrones*, so we would not be able to compare the CE of that precise goods lottery. However, one could still elicit the CE for the lottery that has a 50% chance of receiving the optimal \$15 good (for subject B) and a 50% chance of receiving the optimal \$5 good (for subject B). Then, the subject with the lower CE could be classified as relatively more risk-averse than the other subject (for this lottery, given prices and goods available).

However, in this paper's experiment examining preferences over money vs goods, eliciting the CE may be problematic. For example, if the CE is measured in Amazon.com credit, then some treatments might influence the subjects valuation of Amazon.com credit (e.g. in the presence of imperfect information or salience). In other words, it is possible that the value of Amazon.com credit itself is determined partly by the subject during the experiment, and this may vary differentially by treatment.

If the CE is measured in actual fungible \$, as opposed to Amazon.com credit, differences in time preferences may complicate the CE across subjects. In other words, an individual who choses a fungible \$8 for sure over a {50% chance of \$5 of time-limited Amazon.com credit, 50% chance of \$15 of time-limited Amazon.com credit} may not actually be risk averse but instead greatly values the ability to save. And given that fungible \$ could be spent on Amazon.com at a later date, there's still the potential for information treatments

[27] Empirically this is likely all we can hope for as well, given the data restrictions of testing pairwise preferences for all 300+ million goods is infeasible.

to influence the valuation of fungible \$. For example, learning that there is a \$150 guitar on Amazon.com might increase my valuation of fungible \$ relative to time-limited Amazon.com credit, as none of the budgets in the experiment were large enough to select the guitar. Some of these issues may be possible to overcome, though there's still debate about how to best experimentally elicit CE even in traditional settings [c.f. Harrison and Rutström, 2008, Charness et al., 2013].

Thus rather than certainty equivalent or a direct risk aversion parameterization, I instead focus on three measures of risk when comparing allocations:

1. Defining a risk-free allocation to be one in which the allocation is uniform (i.e. converting the lottery into a "sure thing").[28]

2. Non-parametric methods to test differences in distribution of credit amounts and good prices, both raw distributions and those correcting for the discreteness of good prices.

3. Regressions of the standard deviation of the allocation, both raw distributions and those correcting for the discreteness of good prices.

4 Experiment Results

The first question is whether the primary treatment of allocating Amazon.com credit (rather than goods) impacted the distribution of the allocation. Recall that under the assumptions of section 2, there should be no difference between the monetary distributions and the good distributions. For example, if the decision maker preferred a 10% chance of a \$100 item to a 100% chance of a \$10 item, then when selecting monetary distributions, they should have also preferred a 10% chance of \$100 of credit to a 100% chance of \$10 of credit. As the credit needed to be spent immediately after awarded, there are no intertemporal savings, so any difference in distribution over the uncertain states would indicate one of the assumptions was not satisfied.

Result 4.1. *Contrary to the equivalence theorem presented, subjects exhibited greater aversion to high variance allocations when selecting credit amounts than they did when selecting goods. When selecting goods, subjects were also significantly less likely to select goods of the same value.*

[28] Note this can be done by putting the same good in each uncertain state, which subjects were informed and quizzed on. See discussion in following section and Appendix Figure 2.

There are several approaches to analyze these differences in allocations. I provide results using multiple methods, including regressions of the standard deviation, tests of allocating risk equally across all states, and nonparametric methods to test differences in distribution. All of these methods support the conclusion that subjects selecting credit allocations were more likely to spread out the total amount over multiple states, while subjects selecting goods chose more "risky" allocations (measured with the price of the items).

Specifically, when analyzing each subject's allocated distribution across the possible states, subjects distributing credit reduced the standard deviation of the distribution by a third compared to subjects in the goods treatment. As seen in Table 3A, OLS estimates suggest the standard deviation of prices were significantly reduced by -2.66 to -2.43 down from a mean of 7.61. However, further analysis of the allocations suggests that not only the standard deviation, but also the mean differs between credit and goods treatments. This may be because goods on Amazon.com are discrete – it is likely difficult or even suboptimal to spend precisely the \$20 or \$100 budget.[29] As a result of this discreteness, I also investigate the standard deviation after normalizing values by the total amount allocated. However, the results in Table 3B are nearly identical, with a one-third reduction of the standard deviation when selecting credit.[30]

In the experiment, 16% of subjects removed all risk by allocating a uniform distribution (equal values across all of the uncertain states). As seen in Table 4, this risk-less distribution were over 4 times as likely to occur when the subjects were selecting credit than when they were selecting goods ($p < 0.01$). Note that subjects were instructed and quizzed that they could place the same good in multiple slots, thus increasing the chance of having it selected (see Appendix Figure 2). Despite being quizzed about the possibility, there were only 8 cases where a subject selected a uniform (riskless) distribution of good prices, indicating a greater tolerance for risk.

However, it remains theoretically possible that even though a subject allocates goods non-uniformly, e.g. chooses the lottery {50% chance of \$5 good, 50% chance of \$15 good}, that they are not engaging in risk in "utility terms". In other words, if the subject is indifferent between the \$5 good and the \$15 good, it's actually a riskless allocation. Note

[29] Indeed, the average allocation is \$9.32 instead of the full \$10. Selecting credit instead of the goods increases this average by about \$0.50 (more details in Appendix Table 1).

[30] It is worth noting that the \$100 treatment effect changes sign. This is likely because with more money to spend the potential standard deviation of allocations can increase; but once normalized to fractions, this effect goes away.

for this example that this would imply that there is no good between \$5 and \$10 that is strictly preferred, or else they could have obtained that good with 100% probability. If the majority of allocations are on the order {50% chance of \$9.75 good, 50% chance of \$10.25 good} this would indeed be concerning for the interpretation of using price variation as as measure of risk. However, as shown in Appendix Figure 4, in order for all goods allocations to be riskless, it would require that 50% of subjects are indifferent toward earning an extra \$15 or more on Amazon.com. Limiting to the \$20 Round subsample, 50% of subjects would have to be indifferent toward earning an extra \$7.50 or more on Amazon.com. While additional tests could be conducted about subjects preferences for Amazon.com credit, it seems evident that at least some subjects are engaging in risky allocations when allocating goods.

In addition to the above regression results, one can also non-parametrically analyze the distribution of values. As these tests assume independence among observations, it is not possible to simply use every individual allocation datapoint. Instead, each allocation is transformed into a single variable that can then be non-parametrically tested across the two primary treatments (goods and credit). The \$20 treatment is a good starting point for this, as most of the information of an allocation can be summarized in a single number, specifically "What is the price of the lower-priced good?" These distributions across individuals are plotted in Figure 2A, and the associated Kolmogorov-Smirnov test rejects equality of the distributions ($p < 0.01$). Figure 2B plots distributions of a similar nature, that is the normalized price of the lower-priced good (in other words, what fraction of the total spent is on the lower-priced good). Kolmogorov-Smirnov suggests borderline significant rejection for equality of the distributions of this transformation ($p < 0.07$).

However, though widely used, the Kolmogorov-Smirnov test uses the largest difference between the distributions. As a result, it tends to underweight differences in the tails of the cumulative distributions – Mason and Schuenemeyer [1983], Kim and Whitt [2015]. Given the large share of subjects who place \$10 and \$10 when using credit, the Kolmogorov-Smirnov test may not be the most efficient. Alternatively, we can also use more information from the distributions, such as a Kolmogorov-Smirnov test of the within-allocation standard deviation. This approach also allows for the \$100 treatment observations to be included. In these cases, the Kolmogorov-Smirnov rejects equality of distribution both when using the distributions of standard deviations ($p < 0.01$) or the distributions of normalized standard deviations ($p < 0.01$).

Result 4.2. *When randomly selected to spend more time searching Amazon.com, subjects did not significiantly alter the distribution allocations of goods or credit.*

To test the possibility that the difference in risk for the good domain is being driven by product or price uncertainty, some subjects were randomly submitted to an information treatment. In this treatment, subjects were made to wait an extra 5 minutes before they could submit their allocations. During this time, subjects were only allowed to visit Amazon.com or sit quietly at the desk.[31] The intent was to lower the marginal cost of searching. It appears this treatment was indeed successful in inducing subjects to spend more time browsing – the average treatment effect was to spend an extra 8 minutes searching Amazon (5 minutes beyond the 3 minutes imposed). This extra 5 minutes spent searching could be the result of product search being unexpectedly interesting or that the 3 minute timer was not visible while browsing Amazon, causing subjects to spend more time browsing before realizing the 3 minutes was up.

As we can see in the OLS regressions in Table 5A, the treatment information had no significant direct impact on allocation distributions (as measured by the standard deviation). If the information treatment's effect on allocation would be through the time spent searching, we can also use the information treatment as an instrumental variable for time spent in a section. This allows a causal impact of time spent searching on the allocated distributions. Table 5B presents results of this instrumental variable regression, but once again, spending more time searching has no significant impact on the standard deviation of the allocation.

Two further tests are conducted to further clarify the role of product or price uncertainty. First, I use "switcher" subjects who initially allocate goods in Round 1 and then later allocate credit in Round 2 and compare them to subjects who allocate goods in both rounds. The basic intuition is that subjects who have already allocated goods in Round 1 have undergone a search process that might be more active than the "browse Amazon.com" information treatment, and thus, we might expect the subjects selecting credit in Round 2 to behave more like the subjects selecting goods in Round 2. However, as seen in Table 6, this does not seem to be the case – results are qualitatively and quantitatively similar to the full sample. In other words, it seems like subjects tend to allocate risky good allocations in

[31] This website restriction, as well as a no cellphone rule, was enforced by lab assistants monitoring the study.

Round 1, then revert to more riskless behavior in Round 2 despite having recently learned about relevant products and prices.

Lastly, an additional real-effort experiment was run in a setting without uncertainty, as described in more detail in the following section. To summarize, this final experiment also showed that it is unlikely to be inherent product or price uncertainty that drives the differences between credit and good allocations in the case of lotteries.

5 Real Effort Experiment without Uncertainty

In addition to the aforementioned experiment over risk, a second experiment was conducted to address alternate theories in Section 2. In particular, if subjects only a noisy approximation of their utility or if they face an ex-ante endowment effect, encouraging them to take on additional risks in the risky choice experiment. Although the information treatment did not seem to alter the results in the primary experiment, this may be seen as a secondary test of whether the differences between credit and goods persist even without risk.

5.1 Risk-free Experiment Design

To disentangle whether salience might play a role in the decisions over risk, I offered subjects the option to complete tedious text-reversal tasks and an earning schedule, e.g. first page is \$2, second page is \$1.50, and so on. Ultimately, each subject would earn at most a single physical good, similar to the previous experiment.[32] Subjects were presented with this information in advance and quizzed, as well as provided with practice tasks to ensure understanding of the task.

The primary treatment in this experiment is whether (i) the subject first completes tasks and then allocates the earned credit for a physical good, or (ii) the subject first determines what good they would want to earn for each possible level of earnings.

As in the previous experiment, subjects are informed and quizzed that they are allowed to browse the shopping website at any time. In addition, a similar information treatment is employed to investigate concerns about price and product uncertainty. In particular, subjects were allowed to browse only PCHome for 5 minutes. This is larger than the 3

[32]Since Amazon.com does not directly operate in Taiwan, the closest equivalent, PCHome.com.tw with 24 hour shipping was selected. As before, subjects were instructed to ignore shipping costs, only the list price of the good was used.

minutes of the risky experiment because the risk-free experiment had more time allocated to subjects, and in part because the information treatment had no impact in the risky experiment.

Lastly, a third treatment arm focuses on the role of payment. The control payment scheme follows most real effort experiments, with each additional task earning a smaller marginal amount of reward. This helps us find the marginal cost of effort, as we would expect the individual to keep working until the marginal utility from the reward is equal to the disutility from the marginal effort [c.f. Abeler et al., 2011]. In this payment, each additional page of tasks provides approximately 80% of the marginal earnings of the last page, starting with 50 $NTD ($\approx$ $1.79 USD) down to 40 $NTD, then 32 $NTD and so on. The tenth page offered only 7 $NTD (about $0.25 USD) for completion.

However, if subjects are partially unaware of products or prices, it's not ex ante clear what impact price information would have under a decreasing payment scheme. For example, if there's persistent price uncertainty (despite being allowed to browse the website), a subject may do an extra page for very low marginal earnings out of fear that their 'optimal' good is priced higher than they expect. Thus, informing subjects of precise prices may actually cause a reduction in effort. On the other hand, informing subjects of the products available may cause an increase in effort as they find new (more expensive) optimal items to add to their consideration set.

Thus, an additional "increasing marginal earnings" scheme was randomly assigned to some subjects.[33] This scheme featured the same payment scheme as "decreasing marginal earnings" described above, but in reverse. The first page only paid 7 $NTD, increasing about 25% for each page, up to 32 $NTD, 40 $NTD, and 50 $NTD for the final pages. Under this payment scheme with a large diversity of goods, a non-satiated perfect-information consumer with constant marginal cost of effort would either do 0 tasks or all 10 tasks. Indeed, subjects in this treatment were about 35 percentage points more likely to end up in one of the 'corner solutions' (OLS, $p < 0.001$). This is depicted visually in Figure 3, where the size of each circle corresponds to the fraction of subjects (in each payment scheme) who complete the corresponding number of page tasks.

If limited information plays a large role, then finding a product closer to the maximum

[33]Specifically subjects had a 34% chance to be assigned to this treatment as experimental pilots identified a similar 'corner solution' phenomena which may limit data variation, and as the "decreasing marginal earnings" approach makes the study easier to compare to previous literature.

amount of earnings possible may be motivation to do the task beyond the initial low-return payments. However even in the “increasing marginal earnings” subset, investigating either the information treatment or interaction effects are not significantly more likely to increase number of task pages or reaching the maximum number of pages.

5.2 Risk-free Experiment Results

In this section I will provide additional results from the design above. For this experiment, 107 undergraduate students from National Taiwan University were randomly treated in a cubicle enviornment at the TASSEL laboratory.[34] This experiment was pre-registered with the AEA RCT Registry as AEARCTR-0002637 with an intended 100 subjects.[35]

As can be seen in Table 7, earning credit (to convert into a good) vs allocating goods before employing the task did not significantly alter the number of pages completed. The information treatment also did not have a significant impact, nor does subsample analysis on either of the two different payment schemes suggest that allocating goods resulted in more or less effort.

Although the experiment had a roughly similar number of subjects, one limitation of the experimental results is that roughly one out of four subjects (26 out of 107) completed the maximum number of pages of tasks. Although the data is essentially “right censored” for all 26 subjects, this is primarily of concern for the 7 subjects in the “decreasing marginal earnings” treatment as it obscures their true reservation wage. Put another way, 19 of the maximum-effort subjects were in the “increasing payment” scheme described above, which might be expected given the intuition of achieving a corner solution. Unfortunately, the limited time available for the laboratory study (105 minutes for the task section) and already low marginal earnings of the final page (7 $NTD in decreasing treatment) limited further pages of tasks.

One other potential issue in comparing the “increasing” and “decreasing” marginal earning treatments is that the “decreasing” marginal returns had a very focal point for stopping. If the subjected finished exactly 7 pages of tasks, the subject would receive 200 NTD of

[34] Excluded from this count are two students who left as soon as they completed the practice problems, before seeing the treatment. These students were still able to earn the basic show up reward, in line with TASSEL policies.

[35] However, at the time of registration, I had intended to employ “slider” tasks, but an early pilot experiment demonstrated that the subject pool was overly engaged in the task, and a “text reversal” task was employed instead. At the time of the original experiment with risk in 2016, the author was unaware of the AEA RCT Registry, and therefore that experiment was not registered.

time-limited credit to spend on the retail website. This focal point was not intended, but rather an extension of the relatively simple 80% rule used for determining the payment schedule. However, the proportion of subjects who stop at this focal point is only borderline significantly higher (13 percentage points, $p < 0.068$) for the "decreasing marginal earnings" group. Importantly, this focal point should not impact the primary results focusing on goods vs credit, which was an orthogonal treatment arm.

6 Conclusion

Contrary to the equivalence theory for money and goods under uncertainty, subjects exhibited reduced risk taking when selecting credit amounts than they did when selecting goods. When selecting goods, subjects were also significantly less likely to select goods of the same value across the uncertain states. These findings alone might indicate a general uncertainty of Amazon.com goods or prices, but forcing subjects to spend more time investigating Amazon.com does not change these differences.

As a result, one of the remaining assumptions of the equivalence theory must be false to result in this behavior. In the theoretical section, I outlined 3 categories of existing literature that might explain this, particularly (i) cognitive uncertainty as noisy ex ante utility, (ii) endowment effects, and (iii) regret theory.

Of these three categories, subsample analysis and the additional risk-free experiment provide the strongest support for a form of regret theory. However, while Sarver [2008] provides a framework for discussing choices "as if" decision makers anticipate ex-post regret, the specific reasons for why this regret occurs in the first place may be warranted. One reason for this is that the 'utility' or 'enjoyment' for goods is arguably more shrouded.

In other words, a \$15 book is not necessarily 3 times 'better' than a \$5 pair of scissors. Thus, a subject who receives the \$5 pair of scissors may be 'unlucky' and have been 'better off' with a risk free allocation of \$10, but how much better? As this comparison is obscured, subjects may feel less ex-post regret at the situation. Knowing this in advance, they are willing to engage in more ex-ante risk. However, whether this 'as if' regret stems from internal factors (salience or psychological burden of making a 'mistake') or external social factors (embarrassment of making a 'mistake') is currently unknown in this setting.[36] I hope

[36] While previous studies have found differences in risk behavior between feedback and non-feedback settings, the author is unaware of any that specifically test for subject vs. experimenter feedback in a risk

this will be furtile ground for future research.

It may also be the case that by reducing the choice set, the differences between credit and goods could decrease. For example, if subjects were restricted to choosing from among ten goods for each uncertain state, one might expect a convergence of money and good risk taking. But in many markets outside of the laboratory, individuals face many possible uses for their money, which might impact external validity of those results.

While these are interesting possibilities and warrant further study, this does not change the primary empirical finding of this paper – that individuals react differentially to risk over goods and risk over (time-limited) money.

This finding has important implications for public policy. In 2014 U.S. government sponsored lotteries raised \$70 billion in revenues, helping fund state governments and programs. This paper suggests that individuals may be more willing to engage in lotteries that have goods, not just money. Indeed, U.S. companies often run sweepstakes with prizes (cars, cruises, etc.) rather than a pure lottery [Kalra and Shi, 2010].[37] For example, a prize of \$1,000,000 with a car worth \$50,000 may cause more engagement in risk than a lottery with \$1,050,000. Although the total ramifications of government sponsored lotteries are debatable, this greater willingness in risk could be used to reduce advertisement and overhead budgets without changing revenues.

In addition, this research on risk taking over goods may shed light on important public policy regarding virtual goods. In recent years, several countries have written or enacted legislation regarding risk taking of digital goods, an estimated \$30 billion industry in 2017. These digital goods are often a bit different from traditional gambling where money is offered for a random chance to acquire more money. Instead, the most common form is exchanging money[38] for a random selection of virtual goods, which cannot be resold within the game.[39] As a brief recap of recent international policy discussion: The Netherlands has recently analyzed ten video games offering random virtual goods for sale and determined that four of them contravened its Betting and Gaming Act. The Belgium Gaming Commission looked at four video games' random virtual goods and determined that three of them were

setting.

[37]Some of this is likely driven by the prize being offered to the sweepstake company at a reduced price, perhaps due to the marketing opportunity with a partnering company. However, state sponsored lotteries or sweepstakes may also be able to reduce costs through similar marketing agreements.

[38]In some cases, the company offers a virtual currency as an intermediary, but this currency can be purchased with money.

[39]However, some concerns regarding the games involve the potential to sell accounts.

actually games of chance and subject to Belgian gambling law.[40] State lawmakers in Haiwaii introduced four pieces of legislation that would limit the sale of such lootboxes to consumers under 21 as well as requiring labeling of such games, however the legislation failed to meet final deadlines in March 2018. The Japanese Consumer Affairs Agency declared in 2012 that certain random virtual goods called "Complete Gacha" fall under consumer protection law and are thus prohibited.[41] Following this legislation, the gaming industry enacted self-regulation requiring transparency of random good probabilities and strict measures against real world trading. China has also passed legislation regarding online games in general and random digital goods.[42] This legislation limits the ability to obtain loot boxes with real money or virtual currency; and importantly, requires any virtual items obtainable by loot boxes to be able to be purchased by other means (real money or virtual currency).

It is worth noting that these games are often social or multiplayer in nature and may feature an element of an "rat race" in which one often can purchase better equipment to perform better than others. However, these games could easily sell such improved equipment piecemeal rather than in random bundles. Thus, it may be not only the "rat race" effect that is driving these industries, but also the underlying preferences for risk over (virtual) goods. Yet whether risk preferences for physical and virtual goods are similar remains a topic of further empirical study.

[40]The initial four video games under question were FIFA 18, Overwatch, Counter Strike: Global Offensive, and Star Wars Battlefront 2. In all but Star Wars Battlefront 2, the games offered "loot boxes" that granted random rewards to increase strength or appearance. The Belgian Minister of Justice, Koen Geens, stated that: "It is often children who come into contact with such systems and we cannot allow that."

[41]In these "Complete Gacha" systems, the player might be awarded several virtual goods that must be combined to complete a more rare virtual good.

[42]This regulation was enacted in May 2017 by the Ministry of Culture and the State Administration of Publication, Press, Radio, Film and Television.

CONFLICT OF INTEREST STATEMENT: On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

Johannes Abeler, Armin Falk, Lorenz Goette, and David Huffman. Reference points and effort provision. *The American Economic Review*, 101(2):pp. 470–492, 2011. ISSN 00028282. URL `http://www.jstor.org/stable/29783680`.

Maurice Allais and GM Hagen. *Expected Utility Hypotheses and the Allais Paradox: Contemporary Discussions of the Decisions Under Uncertainty with Allais' Rejoinder*, volume 21. Springer Science & Business Media, 2013.

Lisa R Anderson and Jennifer M Mellor. Predicting health behaviors with an experimental measure of risk preference. *Journal of health economics*, 27(5):1260–1274, 2008.

James Andreoni and Charles Sprenger. Uncertainty equivalents: Testing the limits of the independence axiom. Technical report, National Bureau of Economic Research, 2011.

Levon Barseghyan, Jeffrey Prince, and Joshua C Teitelbaum. Are risk preferences stable across contexts? evidence from insurance data. *American Economic Review*, 101(2): 591–631, 2011.

Michael H Birnbaum and Roman J Gutierrez. Testing for intransitivity of preferences predicted by a lexicographic semi-order. *Organizational Behavior and Human Decision Processes*, 104(1):96–112, 2007.

Michael H Birnbaum and Ulrich Schmidt. Testing transitivity in choice under risk. *Theory and Decision*, 69(4):599–614, 2010.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience theory of choice under risk. Technical report, National Bureau of Economic Research, 2010.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience in experimental tests of the endowment effect. *American Economic Review*, 102(3):47–52, 2012.

Jennifer Brown, Tanjim Hossain, and John Morgan. Shrouded attributes and information suppression: Evidence from the field. *Quarterly Journal of Economics*, 125(2), 2010.

Benjamin Bushong, Lindsay M King, Colin F Camerer, and Antonio Rangel. Pavlovian processes in consumer choice: The physical presence of a good increases willingness-to-pay. *American Economic Review*, 100(4):1556–1571, 2010.

Gary Charness, Uri Gneezy, and Alex Imas. Experimental methods: Eliciting risk preferences. *Journal of economic behavior & organization*, 87:43–51, 2013.

Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009.

Ravi Dhar and Itamar Simonson. The effect of the focus of comparison on consumer preferences. *Journal of Marketing Research*, pages 430–440, 1992.

Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the european economic association*, 9(3):522–550, 2011.

Liran Einav, Amy Finkelstein, Iuliana Pascu, and Mark R Cullen. How general are risk preferences? choices under uncertainty in different domains. *The American Economic Review*, 102(6):2606–2638, 2012.

Benjamin Enke and Thomas Graeber. Cognitive uncertainty. *The Quarterly Journal of Economics*, 138(4):2021–2067, 2023.

Armin Falk, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science*, 69(4):1935–1950, 2023.

Michele Fioretti, Alexander Vostroknutov, and Giorgio Coricelli. Dynamic regret avoidance. *American Economic Journal: Microeconomics*, 14(1):70–93, 2022.

Xavier Gabaix. A sparsity-based model of bounded rationality*. *Quarterly Journal of Economics*, 129(4), 2014.

Xavier Gabaix, David Laibson, et al. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2): 505–540, 2006.

Uri Gneezy, John A List, and George Wu. The uncertainty effect: When a risky prospect is valued less than its worst possible outcome. *The Quarterly Journal of Economics*, 121 (4):1283–1309, 2006.

Kelly Goldsmith and On Amir. Can uncertainty improve promotions? *Journal of Marketing Research*, 47(6):1070–1077, 2010.

Simon Grant, Atsushi Kajii, and Ben Polak. Many good choice axioms: when can many-good lotteries be treated as money lotteries? *Journal of Economic Theory*, 56(2):313–337, 1992.

Glenn W Harrison and E Elisabet Rutström. Risk aversion in the laboratory. In *Risk aversion in experiments*. Emerald Group Publishing Limited, 2008.

Glenn W Harrison, John A List, and Charles Towe. Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica*, 75(2): 433–458, 2007.

Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives*, 5(1): 193–206, 1991.

Ajay Kalra and Mengze Shi. Consumer value-maximizing sweepstakes and contests. *Journal of Marketing Research*, 47(2):287–300, 2010.

Song-Hee Kim and Ward Whitt. The power of alternative kolmogorov-smirnov tests based on transformations of the data. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 25(4):24, 2015.

Jack L Knetsch. The endowment effect and evidence of nonreversible indifference curves. *The American Economic Review*, 79(5):1277–1284, 1989.

Botond Kőszegi and Adam Szeidl. A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1):53–104, 2013.

Craig E Landry and Michael K Price. Earmarking lottery proceeds for public goods: Empirical evidence from us lotto expenditures. *Economics Letters*, 95(3):451–455, 2007.

Juliano Laran and Michael Tsiros. An investigation of the effectiveness of uncertainty in marketing promotions involving free gifts. *Journal of Marketing*, 77(2):112–123, 2013.

Keith M Marzilli Ericson and Andreas Fuster. The endowment effect. *Annu. Rev. Econ.*, 6 (1):555–579, 2014.

David M Mason and John H Schuenemeyer. A modified kolmogorov-smirnov test sensitive to tail alternatives. *The Annals of Statistics*, pages 933–946, 1983.

Robert Mislavsky and Uri Simonsohn. When risk is weird: Unexplained transaction features lower valuations. *Management Science*, 64(11):5395–5404, 2018.

George E Newman and Daniel Mochon. Why are lotteries valued less? multiple tests of a direct risk-aversion mechanism. *Judgment and Decision Making*, 7(1):19–24, 2012.

Ted O'Donoghue and Jason Somerville. Modeling risk aversion in economics. *Journal of Economic Perspectives*, 32(2):91–114, 2018.

Pietro Ortoleva. The price of flexibility: Towards a theory of thinking aversion. *Journal of Economic Theory*, 148(3):903–934, 2013.

Michel Regenwetter, Jason Dana, and Clintin P Davis-Stober. Transitivity of preferences. *Psychological Review*, 118(1):42, 2011.

Ariel Rubinstein. Similarity and decision-making under risk (is there a utility theory resolution to the allais paradox?). *Journal of Economic Theory*, 46(1):145–153, 1988.

Ondřej Rydval, Andreas Ortmann, Sasha Prokosheva, and Ralph Hertwig. How certain is the uncertainty effect? *Experimental Economics*, 12:473–487, 2009.

Todd Sarver. Anticipating regret: Why fewer options may be better. *Econometrica*, 76(2): 263–305, 2008.

Hannah Schildberg-Hörisch. Are risk preferences stable? *Journal of Economic Perspectives*, 32(2):135–154, 2018.

Uri Simonsohn. Direct risk aversion: Evidence from risky prospects valued below their worst outcome. *Psychological science*, 20(6):686–692, 2009.

Matthias Sutter, Martin G Kocher, Daniela Glätzle-Rützler, and Stefan T Trautmann. Impatience and uncertainty: Experimental decisions predict adolescents' field behavior. *American Economic Review*, 103(1):510–531, 2013.

Amos Tversky. Intransitivity of preferences. *Psychological review*, 76(1):31, 1969.

Yitong Wang, Tianjun Feng, and L Robin Keller. A further exploration of the uncertainty effect. *Journal of Risk and Uncertainty*, 47:291–310, 2013.

7 Figures

Figure 1: Example of Good Selection

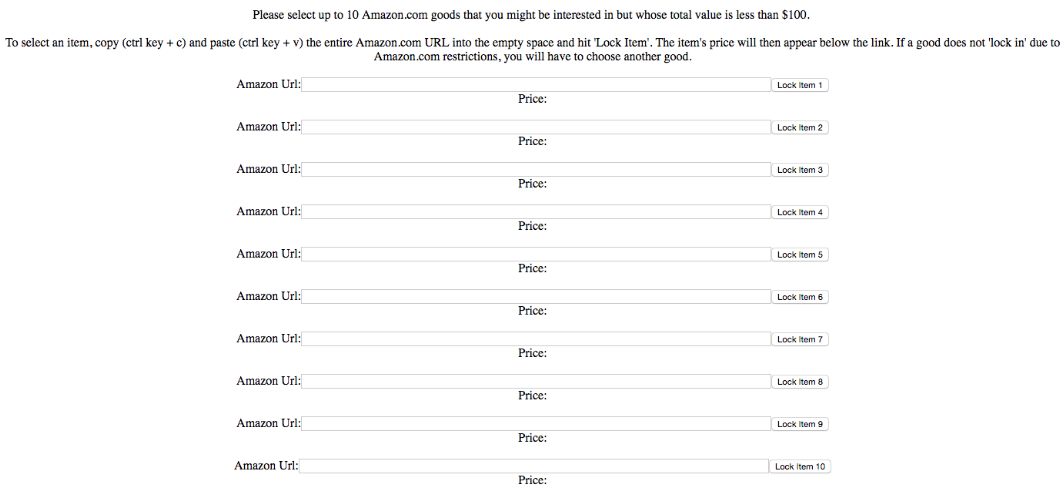

Please select up to 10 Amazon.com goods that you might be interested in but whose total value is less than $100.

To select an item, copy (ctrl key + c) and paste (ctrl key + v) the entire Amazon.com URL into the empty space and hit 'Lock Item'. The item's price will then appear below the link. If a good does not 'lock in' due to Amazon.com restrictions, you will have to choose another good.

Amazon Url: Lock Item 1
Price:

Amazon Url: Lock Item 2
Price:

Amazon Url: Lock Item 3
Price:

Amazon Url: Lock Item 4
Price:

Amazon Url: Lock Item 5
Price:

Amazon Url: Lock Item 6
Price:

Amazon Url: Lock Item 7
Price:

Amazon Url: Lock Item 8
Price:

Amazon Url: Lock Item 9
Price:

Amazon Url: Lock Item 10
Price:

Notes: Figure demonstrates a typical good selection screen faced by subject. Whether subject was asked to select Amazon.com goods (via URLs) or Amazon.com credit amounts was randomized. Whether subject was asked to find up to 10 items that totaled at most $100 or up to 2 items that totaled at most $20 was also randomized. See Experiment Design for more details.

Figure 2A: Distribution of the Smaller Value When Total is $20

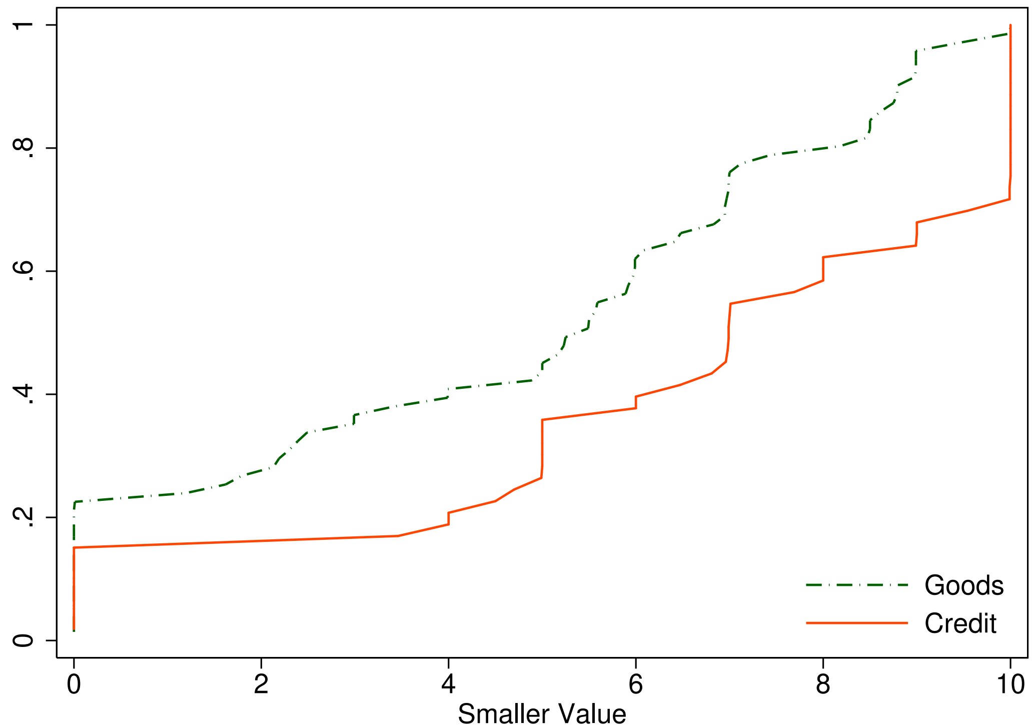

Notes: Plot shows two cumulative distributions of the value of the smaller good when total is $20. When given $20 to allocate, the subject chooses to allocate across 2 uncertain states, this represents the smaller of these two allocations. Vertical axis represents the frequency of that value occurring across the two different treatments (selecting goods or selecting credit allocations).

Figure 2B: Distribution of the Smaller Value (Normalized) When Total is $20

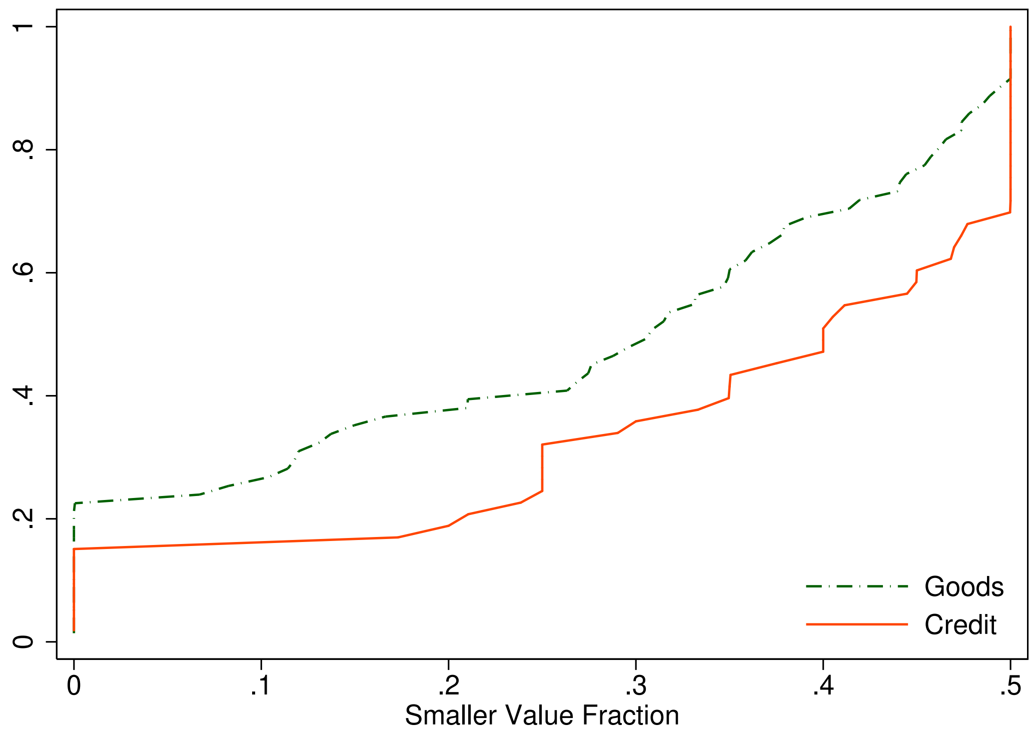

Notes: Plot shows two cumulative distributions of the (normalized) value of the smaller good when total is $20. Value is normalized by dividing by the total value allocated. When given $20 to allocate, the subject chooses to allocate across 2 uncertain states, this represents the smaller of these two allocations. Vertical axis represents the frequency of that (normalized) value occurring across the two different treatments (selecting goods or selecting credit allocations).

Figure 3: Experiment without Risk Earnings

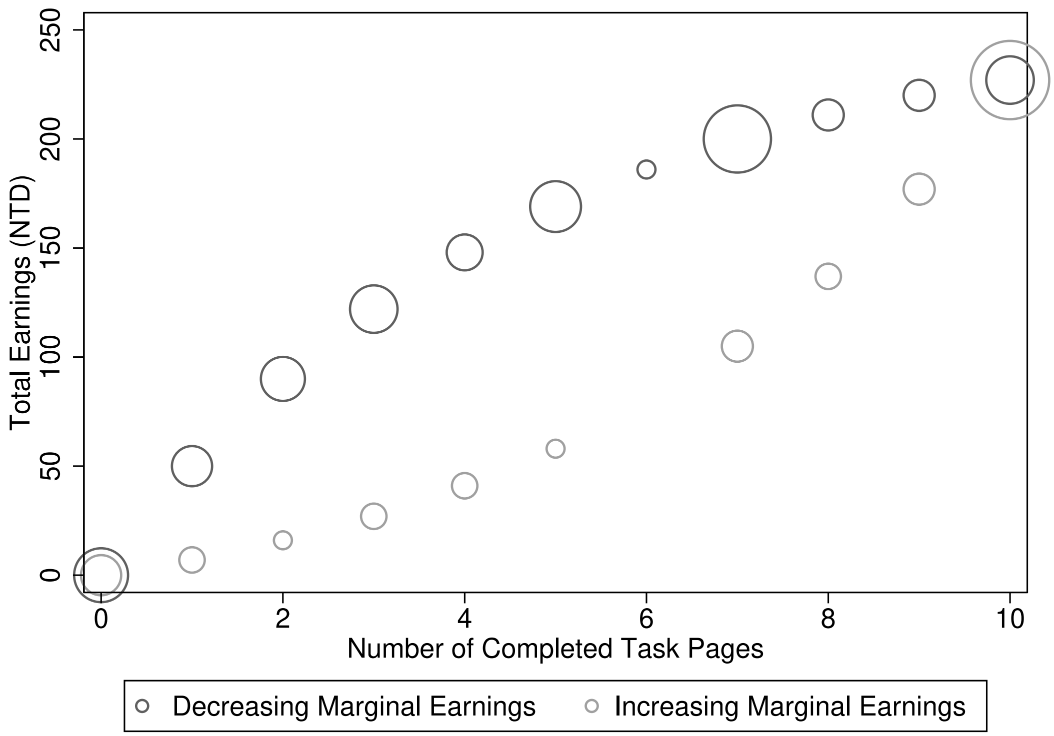

Notes: Plot shows the Total Earnings in New Taiwan Dollars (NTD) for 107 subjects in the real effort experiment. 67 subjects were paid according to a traditional "decreasing marginal earnings" scheme, while the lighter circles represent the 40 subjects paid under the "Increasing Marginal Earnings" payment scheme. The area of each circle corresponds to the fraction of subjects who, in that payment scheme, end up completing the corresponding number of task pages (n.b. doubling the radius would quadruple the area).

8 Tables

Table 1. Summary Statistics

	Mean	Standard dev	Min	Max	
Individual Level Variables					
Female	0.71	0.45	0	1	
Age	20.2	1.3	18	24	
SAT Math Score	733	62	540	800	(22 missing)
Computer Skill Test	2	0	2	2	(1 missing)
Number of Previous Lab Studies	26.6	24.8	1	133	
Period Level Variables					
Average Value of Entry	$9.33	1.07	4	10	
Standard Dev of Entry (within)	$7.06	6.92	0	31.6	
$100 Treatment Indicator	0.50	0.50	0	1	
Credit Treatment Indicator	0.46	0.50	0	1	
Time Spent Searching (seconds)	487	373	45	1699	
Number of Individuals	124				
Number of Treatment Periods	248				

Notes: Computer Skill Test was a demographic variable collected by the Wharton Behavioral Lab prior to the experiment, however among subjects above it had no variation. SAT Math score is missing for individuals who either took the ACT or otherwise did not wish to share that information with researchers.

Table 2. Randomization Check

Dependent Variable	Credit Treatment		Period # for $100 Treatment	
Female	−0.11	−0.01	−0.02	−0.01
	(0.07)	(0.11)	(0.10)	(0.11)
SAT Math Score ('00s of points)		0.01		0.01
		(0.08)		(0.09)
Previous WBL Studies		−0.001		−0.001
		(0.002)		(0.002)
F-test	2.59	0.24	0.05	0.24
p value	0.11	0.87	0.83	0.87
Dependent Variable Mean	0.46	0.48	1.54	1.54
Number of Observations	248	204	248	204
Number of Individuals	124	102	124	102

Notes: Standard Errors (clustered at individual level) presented in parentheses above. As every subject in experiment 1 receives both the $20 and $100 treatments, the dependent variable for $100 treatment is the period in which they received the treatment in question. If randomization was done properly, the pre-treatment variables should not predict the period they received this treatment. Indeed, the F-stats are all large enough that I fail to reject the hypothesis that all coefficients are zero under $\alpha = 0.05$. Thus, I conclude the randomization was adequately done. SAT Math score is missing for 22 individuals who either took the ACT or otherwise did not wish to share that information with researchers.

Table 3A. Credit and \$100: Impact on Standard Deviation of Selection Value

$$Std.Dev_{i,t} \quad = \quad \alpha \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \gamma X_i + \epsilon_{i,t}$$

Dependent Variable:	Specification			
Value Standard Deviation	(1)	(2)	(3)	(4)
Subject Allocates Credit	−2.66***	−2.61***	−2.56***	−2.43***
(Binary Treatment Var.)	(0.77)	(0.77)	(0.81)	(0.82)
\$100 Total Allocation	4.96***	5.00***	5.00***	4.99***
(Binary Treatment Var.)	(0.64)	(0.63)	(0.64)	(0.64)
First Period		0.59	0.59	0.61
		(0.64)	(0.66)	(0.67)
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	7.61	7.61	7.61	7.61
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.16	0.16	0.19	0.22

Notes: The dependent variable is the standard deviation of value of the entries in a single period. All specifications report results from OLS regressions and also include a constant term. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

Table 3B. Credit and \$100: Impact on Standard Deviation of Selection Value (Normalized)

$$Norm.Std.Dev_{i,t} \quad = \quad \alpha \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \gamma X_i + \epsilon_{i,t}$$

Dependent Variable:		Specification		
Normalized Value Std Dev	(1)	(2)	(3)	(4)
Subject Allocates Credit	−0.08***	−0.07***	−0.07***	−0.07***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
\$100 Total Allocation	−0.17***	−0.17***	−0.17***	−0.17***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
First Period		0.01	0.01	0.01
		(0.02)	(0.02)	(0.02)
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	0.20	0.20	0.20	0.20
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.20	0.20	0.25	0.27

Notes: The dependent variable is the standard deviation of normalized value of the entries in a single period. Values were normalized by dividing by the total value allocated. All specifications report results from OLS regressions and also include a constant term. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

Table 4. Credit and \$100: Impact on Riskless Allocation Selection

$$SameValue_{i,t} \quad = \quad \alpha \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \gamma X_i + \epsilon_{i,t}$$

Dependent Variable:	Specification			
Allocation is Riskless	(1)	(2)	(3)	(4)
Subject Allocates Credit	0.22***	0.21***	0.23***	0.21***
(Binary Treatment Var.)	(0.05)	(0.05)	(0.04)	(0.04)
\$100 Total Allocation	−0.06	−0.07*	−0.08*	−0.07*
(Binary Treatment Var.)	(0.04)	(0.04)	(0.04)	(0.04)
First Period		−0.13***	−0.13***	−0.13***
		(0.04)	(0.04)	(0.04)
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	0.16	0.16	0.16	0.16
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.10	0.13	0.21	0.22

Notes: The dependent variable is the whether the allocation is uniform, that is has an equal allocation for each state (e.g. \$10 and \$10). Values were normalized by dividing by the total value allocated. Thus, allocating one good priced \$9.50 good and the same good (or a different good priced \$9.50) across two states is also considered 'riskless'. All specifications report results from OLS regressions and also include a constant term. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

Table 5A. Timing: Impact on Standard Deviation of Selection Value

$$Norm.Std.Dev_{i,t} \quad = \quad \alpha \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \delta \cdot Info_{i,t} + \gamma X_i + \epsilon_{i,t}$$

Dependent Variable:	Specification			
Normalized Value Std Dev	(1)	(2)	(3)	(4)
Subject Allocates Credit	−0.08***	−0.08***	−0.07***	−0.07***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
$100 Total Allocation	−0.17***	−0.17***	−0.17***	−0.17***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
Information Treatment	0.01	0.01	0.03	0.04
	(0.02)	(0.02)	(0.03)	(0.03)
Round Fixed Effects		X	X	X
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	0.20	0.20	0.20	0.20
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.21	0.21	0.25	0.27

Notes: The dependent variable is the standard deviation of normalized value of the entries in a single period. Values were normalized by dividing by the total value allocated. All specifications report results from OLS regressions and also include a constant term. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

Table 5B. Timing: Impact on Standard Deviation of Selection Value

$$\begin{aligned} Clock_{i,t} &= \alpha_1 \cdot Info_{i,t} + \gamma_1 X_{i,t} + \nu_{i,t} \\ Norm.Std.Dev_{i,t} &= \alpha_2 \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \delta \cdot Clock_{i,t} + \gamma_2 X_i + \epsilon_{i,t} \end{aligned}$$

Dependent Variable:	Specification			
Normalized Value Std Dev	(1)	(2)	(3)	(4)
Subject Allocates Credit	−0.07***	−0.07***	−0.06***	−0.07***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
$100 Total Allocation	−0.17***	−0.17***	−0.17***	−0.17***
(Binary Treatment Var.)	(0.02)	(0.02)	(0.02)	(0.02)
Time Searching (Minutes)	0.001	0.001	0.003	0.004
	(0.003)	(0.003)	(0.003)	(0.003)
First Stage F Stat (IV)	176.9	168.8	159.2	144.2
Round Fixed Effects		X	X	X
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	0.20	0.20	0.20	0.20
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.21	0.21	0.25	0.27

Notes: The dependent variable is the standard deviation of normalized value of the entries in a single period. Values were normalized by dividing by the total value allocated. All specifications report results from GMM Instrumental variable regressions and also include a constant term. Time spent searching was instrumented by the information treatment, with F values from the first stage reported. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

Table 6. Allocations in Round 2 after Allocating Goods in Round 1

$$Tasks \quad = \quad \alpha + \beta \cdot T_i + \gamma X_i + \epsilon_{i,t}$$

	Dependent Variable (Round 2 Data Only):				
	Standard Dev	Norm. Stand. Dev	Norm. Stand. Dev	Norm. Stand Dev	Allocation is Riskless
Subject Allocates Credit	-3.47^{**}	-0.07	-0.08	-0.07^{**}	0.33^{***}
(Binary Treatment Var)	(1.55)	(0.04)	(0.11)	(0.02)	(0.10)
$100 Total Allocation	4.88***	-0.18***	(only $20 obs)	(only $100 obs)	0.16
(Binary Treatment Var.)	(1.50)	(0.05)			(0.11)
Round 1 Allocated Goods	Yes	Yes	Yes	Yes	Yes
$20 or $100 Total	Both	Both	$20 Total Only	$100 Total Only	Both
Dep. Var. Mean (Control)	10.34	0.22	0.34	0.15	0.17
Number of Observations	73	73	28	45	73
Number of Individuals	73	73	28	45	73
Adj-R^2	0.17	0.21	0.21	0.47	0.15

Notes: The dependent variables correspond to previous definitions (standard deviation of allocation, normalized standard deviation and binary indicator for riskless allocation). Where indicated, values were normalized by dividing by the total value allocated. All specifications report results from OLS with robust standard errors. Time spent searching was instrumented by the information treatment, with F values from the first stage reported. Individual Controls include sex, age, ethnicity bins, and number of previous WBL studies completed. Standard errors are given in parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p < 0.05$, $* * * = p < 0.01$.

Table 7. Risk-free Experiment: Task Pages Completed

$$Tasks_i \quad = \quad \alpha + \beta \cdot T_i + \gamma X_i + \epsilon_i$$

Dependent Variable:			Specification		
Pages of Tasks Completed	(1)	(2)	(3)	(4)	(5)
Subject Receives Credit	−0.81	−0.80	−0.55	−0.74	−0.27
(Binary Treatment Var.)	(0.70)	(0.70)	(0.68)	(0.78)	(1.25)
Information Treatment		0.17	−0.07	−0.65	0.92
(Binary Treatment Var.)		(0.70)	(0.69)	(0.80)	(1.31)
Payment Scheme			2.10***	Decreasing	Increasing
			(0.74)		
Dependent Variable Mean	5.57	5.57	5.57	4.76	6.93
Number of Observations	107	107	107	67	40
Number of Individuals	107	107	107	67	40
Adj-R^2	0.01	0.01	0.09	0.02	0.02

Notes: The dependent variable is the number of pages of text-reversal tasks completed. Credit treatment is whether the subject first earns credit that then must be spent on the online retail website (baseline is first allocating which good to earn for each potential earnings levels). Payment Scheme indicates the coefficient of having an increasing payment scheme in specification 3 and whether restricting the analysis to decreasing or increasing in specifications 4 and 5. Standard errors are given in parentheses and are robust (one observation per individual precludes clustering). $* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

9 Appendix: FOR ONLINE PUBLICATION

9.1 Appendix Tables

Appendix Table 1. Credit and $100: Impact on Mean of Selection Value

$$AverageValue_{i,t} \quad = \quad \alpha \cdot Credit_{i,t} + \beta \cdot 100Treatment_{i,t} + \gamma X_i + \epsilon_{i,t}$$

Dependent Variable	Specification			
Average Value	(1)	(2)	(3)	(4)
Subject Selects Credit	0.56***	0.57***	0.52***	0.52***
(Binary Treatment Var.)	(0.12)	(0.12)	(0.12)	(0.13)
$100 Total Allocation	0.29***	0.30***	0.30***	0.30***
(Binary Treatment Var.)	(0.11)	(0.11)	(0.11)	(0.11)
First Period		0.12	0.12	0.12
		(0.11)	(0.11)	(0.11)
Session Fixed Effects			X	X
Individual Controls				X
Dependent Variable Mean	9.32	9.32	9.32	9.32
Number of Observations	248	248	248	248
Number of Individuals	124	124	124	124
Adj-R^2	0.09	0.10	0.12	0.17

Notes: The dependent variable is the average value of the entries in a single per
also include a constant term. Individual Controls include sex, age, ethnicity bin
parentheses and clustered at the subject (individual) level. $* = p < 0.1$, $** = p$

9.2 Appendix Figures

Appendix Figure 1. Quiz for Introduction Instructions

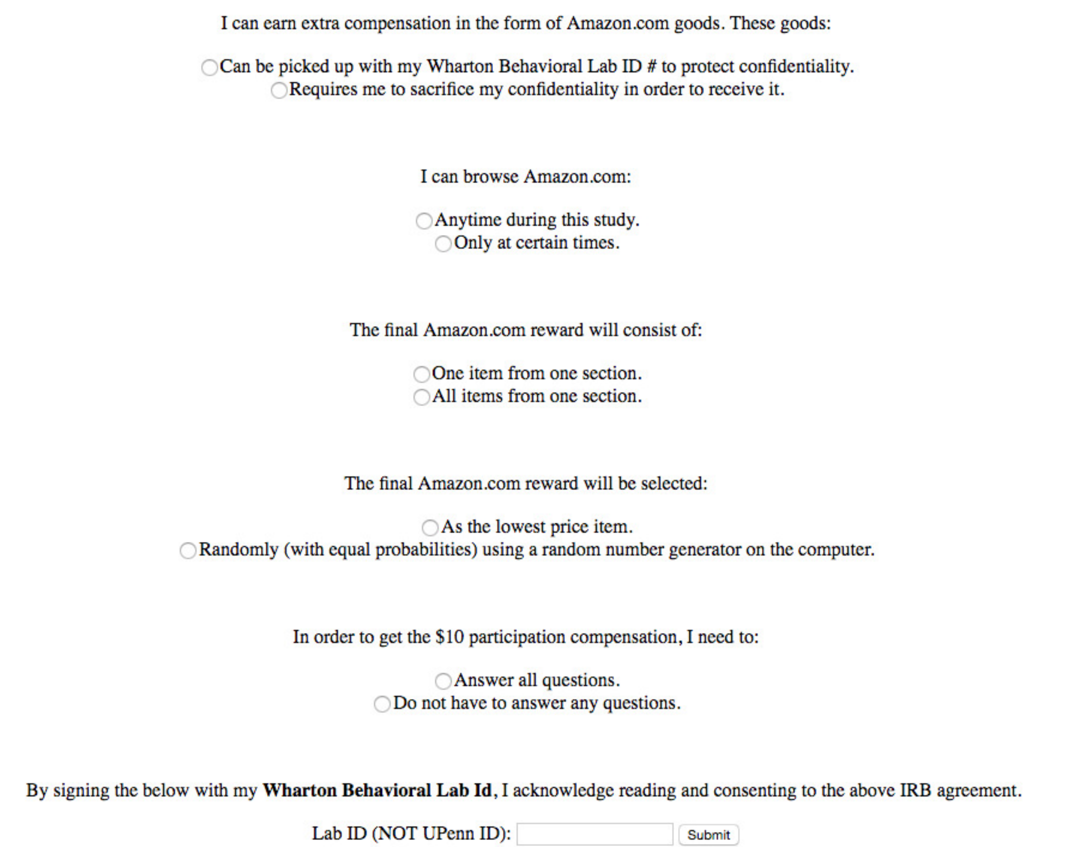

I can earn extra compensation in the form of Amazon.com goods. These goods:

○ Can be picked up with my Wharton Behavioral Lab ID # to protect confidentiality.
○ Requires me to sacrifice my confidentiality in order to receive it.

I can browse Amazon.com:

○ Anytime during this study.
○ Only at certain times.

The final Amazon.com reward will consist of:

○ One item from one section.
○ All items from one section.

The final Amazon.com reward will be selected:

○ As the lowest price item.
○ Randomly (with equal probabilities) using a random number generator on the computer.

In order to get the $10 participation compensation, I need to:

○ Answer all questions.
○ Do not have to answer any questions.

By signing the below with my **Wharton Behavioral Lab Id**, I acknowledge reading and consenting to the above IRB agreement.

Lab ID (NOT UPenn ID): Submit

Notes: Every participant had to answer questions after reading experiment instructions. Subjects had to answer all questions correctly to proceed. If the subject entered the wrong answers, the browser would alert them to this and ask for them to review the instructions again.

Appendix Figure 2. Quiz for Instructions Prior to Each Period

Please answer the questions below to continue

For this section:

The goods must be equal to or less than $ ________

◯ If I leave more slots empty, I am more likely to receive a reward.
◯ Each slot has an equal chance of being chosen regardless of whether it is empty or not.

◯ I do not add shipping to the prices.
◯ I need to add shipping to the prices.

When selecting goods

◯ I can place the same item in multiple slots to increase the chance of it being selected.
◯ I must put different items in every slot.

13 seconds until you can move on

Notes: Every participant had to answer questions prior to every period. If the subject entered the wrong answers, the browser would alert them to this and ask for them to review the instructions again.

Appendix Figure 3. Cubicle Environment

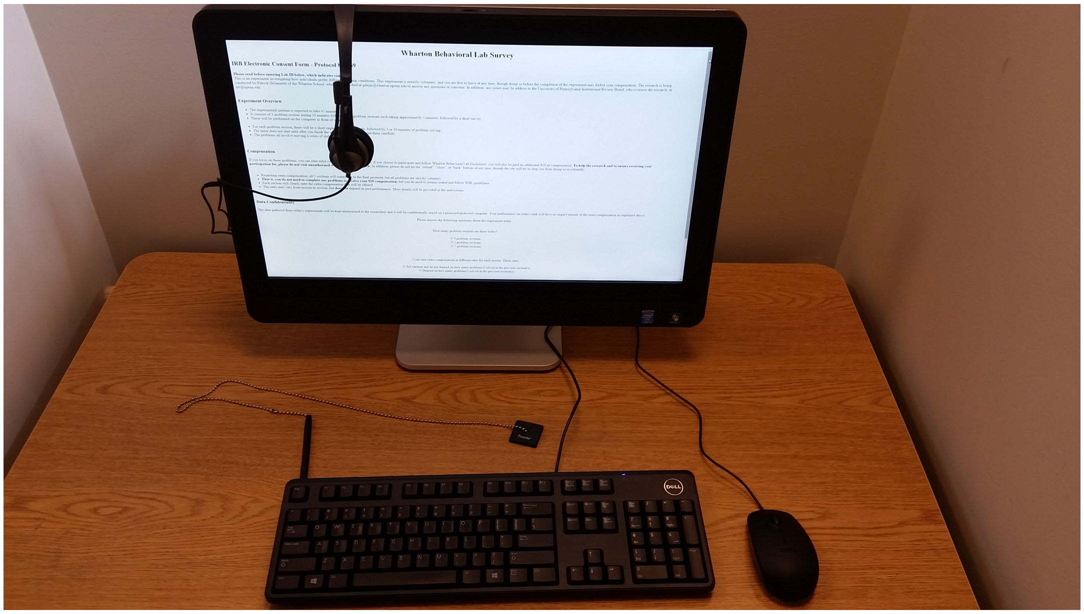

Notes: Every participant had access to an identical computer with headphones as pictured above. Cookies and browser history were cleared after every session to limit any subject overlap. It was not possible to see other subjects from within the cubicle. Google Chrome was employed as the browser. All instructions were written, but lab assistants were on site to answer any additional questions.

Figure 4: Cumulative Distribution of Maximum Price Difference in Good Allocation

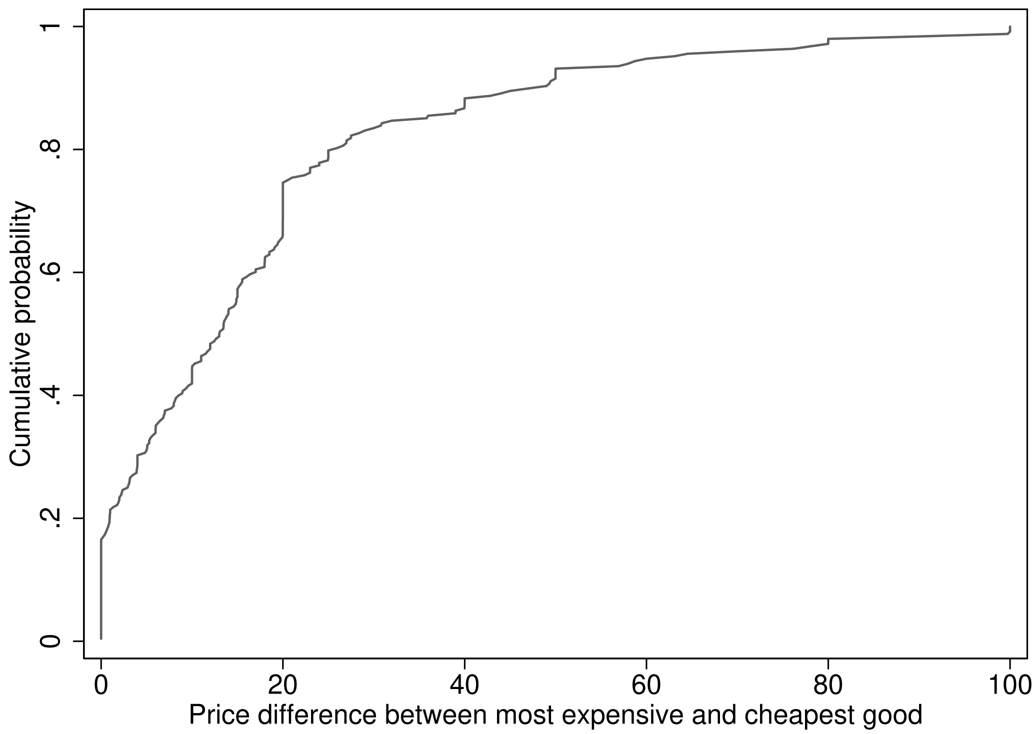

Notes: Plot shows cumulative distributions of the difference between the highest price and lowest price among allocated goods.