Taking a Load Off: Experimental Evidence of Preferences for Control with an Application to Residential Electricity Demand

Christina Korting [*]
Cornell University

This version: October 27, 2019

LINK TO LATEST VERSION

Abstract

The rising share of renewable electricity generation has led to an increased focus on demand-side mechanisms to balance the grid. For example, Direct Load Control (DLC) contracts allow utilities to curtail the electricity use of participating households at times of system stress. I use a novel experimental design to show that intrinsic preferences for control can significantly impact the rewards required to encourage consumers to participate in such contracts. In particular, I test for the existence, magnitude, and attributes of *control premia* in a lab environment which mimics basic features of the DLC context. I find that participants, on average, exhibit a control premium of 9–32% above the instrumental value of the decision. This premium responds to both the probability and stakes of ceding control. There is limited evidence for the existence of an endowment effect with respect to control. Participants' stated motivations underlying their decisions are consistent with an inflation of (perceived) option value that cannot be explained by probability weighting.

JEL Classification: C91, D12, Q40
Keywords: control premium, payoff autonomy, value of decision rights, direct load control

[*]Charles H. Dyson School of Applied Economics & Management, Ithaca, NY 14853. e: cmk272@cornell.edu

1 Introduction

The share of renewable electricity generation in the US is rising steadily, growing to 18% in 2018 compared to 9% in 2008.[1] While lowering average emissions, this creates new tensions for the grid. In addition to making supply more variable and harder to predict, the timing of solar generation relative to household demand for electricity can lead to more pronounced demand peaks in the afternoon and early evening. Utilities therefore increasingly rely on voluntary demand-response mechanisms such as peak pricing or Direct Load Control (DLC) in order to balance the grid. DLC contracts aim to reduce (peak) electricity demand by compensating consumers for the right to centrally switch off appliances such as air conditioners or water heaters at times of system stress. However, residential participation rates in these programs vary widely. In a 2016 report, E Source[2] found participation rates between 0.3% to 30.4% among the 32 programs in their study, with an average participation rate of 9.4%.[3] Studies exploring attitudes towards different demand response programs through surveys and focus groups frequently point to an aversion to being controlled coupled with consumers' lack of trust in energy companies (Fell et al., 2015, 2014; Mert et al., 2008). This raises the question whether DLC contracts trigger a willingess-to-pay (WTP) to retain control above and beyond the instrumental value of the electricity uses being curtailed.

Standard expected utility theory (EUT) does not account for intrinsic preferences for control since it is based on the principle of consequentialism: it assumes that individuals maximize utility over the available choice set independent of context. However, a growing body of work documents that individuals place an economically significant premium on retaining control over payoff-relevant decisions (Starr, 1969; Bartling et al., 2014; Owens et al., 2014; Bobadilla-Suarez et al., 2016).[4] An intrinsic preference for control can lead to under-delegation relative to the control-neutral benchmark (e.g. Bartling et al. (2014)); reduced motivation and performance by workers in the face of constraints on their effort choices (e.g. Falk and Kosfeld (2006)); an aversion to unmodifiable algorithms even if they outperform human forecasters (e.g. Dietvorst et al. (2016)); and increased non-compliance with treatments if the implementation or communication of the treatment triggers reactance (e.g. Quick and Considine (2008)). While such behaviors may seem costly, the cumulative evidence suggests that the control premium represents a genuine, stable preference. However, the contexts in which control premia arise are not yet well understood.

In this article, I show that control premia exist even in situations that do not involve conflicts of interest, distrust, ambiguity over outcomes under delegation, or the delegation of decision rights to a third party.[5] My aim is to establish how the qualitative patterns of control premia vary with

[1] `https://www.eia.gov/todayinenergy/detail.php?id=38752` (last accessed October 15th, 2019)

[2] E Source is a utility research and consulting firm working with utilities across the US and Canada.

[3] `https://www.esource.com/dlc`. The reported average participation rate is in line with the

[4] Starr was one of the first authors to discuss the existence of intrinsic preferences for control, pointing out a "difference by several orders of magnitude in society's willingness to accept voluntary and in-voluntary risk. As one would expect, we are loathe to let others do unto us what we happily do to ourselves" (p. 1235).

[5] For a discussion of the role of conflicts of interest in generating control premia, see Bartling et al. (2014). They

attributes of the decision environment with minimal procedural concerns, allowing me to clearly map the consequences of this preference into questions of DLC design.

My findings contribute to the literature in four ways. First, I provide evidence for the existence of a control premium in a novel experimental setting that speaks directly to the energy context. More broadly, my findings apply to instances of interruptible service or non-price rationing in which the reliability of service differs between consumers depending on their contract choices, such as the quality of alternative WIFI options in a hotel. Unlike existing research on the acceptability of DLC contracts, this result is based on incentive-compatible decisions in a controlled laboratory environment. Second, I replicate the earlier finding by Bartling et al. (2014) regarding the sensitivity of control premia to stake size. Third, I extend the literature by testing whether control premia respond to probabilities: while existing research focuses on one-shot delegation settings, I allow the probability of losing control to vary within subject. Lastly, I explore whether individuals exhibit an endowment effect with respect to control, i.e. whether increasing the probability of losing control triggers a stronger emotional response than regaining a commensurate amount of control.

On average, I find a control premium of 9–32% which is similar to prior findings in the literature.[6] At the individual level, 76% of participants exhibit a positive average control premium across the three main control decisions. I find that participants pay close attention to both the stakes and probability of losing control: in absolute terms, the observed control premium scales almost linearly in both of these aspects of the choice environment, while being broadly consistent with a fixed relative premium above instrumental value. There is limited evidence that a difference in losing versus gaining a commensurate amount of control amplifies standard endowment effects observed for purely financial lottery decisions.

This experiment leverages two commonly employed approaches for establishing intrinsic preferences for control in the literature. The first is to present participants with identical choices, once in the context of a decision involving control and once in a context-free setting (e.g. Bartling et al. (2014)). The second is to compare individuals' valuations to their theoretical control-neutral predictions (e.g. Owens et al. (2014)). I use an estimated control-neutral benchmark to gauge the magnitude of the control premium, and a comparison between equivalent decisions contexts with and without a control component to validate the interpretation of the estimates and to rule out alternative explanations.

The decision context participants face in this experiment borrows directly from electricity contract choice. Consider a household with a willingness-to-pay of \$3 to run their air conditioner

find that an increasing degree of conflict of interest between a principal and an agent reduces the control premium in their principal-agent setting. This result is likely in part due to fairness concerns: the principal can either choose a project which imposes a more negative outcome on the agent, or delegate this decision to the agent who may then select the project that is less favorable to the principal. Principals may therefore experience responsibility aversion triggered by not wanting to impose a negative outcome on the agent (Fleming and Bang, 2018).

[6]Bartling et al. (2014) report a control premium of 18-26% for the two relevant delegation decisions without conflict of interest; Owens et al. (2014) find a relative control premium between 8-15%.

on a hot afternoon. The utility relies partly on wind generation which is subject to some forecast error. On one of ten days, an exogenous shock to the level of wind generation requires the utility to temporarily reduce air conditioner use to ensure that the system remains balanced. The utility offers two alternative demand response programs: DLC and *Critical Peak Pricing* (CPP). CPP encourages conservation using price signals. On a normal day, the utility charges rates per kWh of electricity that translate into \$0.75 for one hour of air conditioner use. In case of a supply shock, the utility instead charges its CPP customers \$3.35 per hour of air conditioner use, and directly shuts off air conditioning for DLC households. These two contract structures translate into the lotteries illustrated in Figure 1.

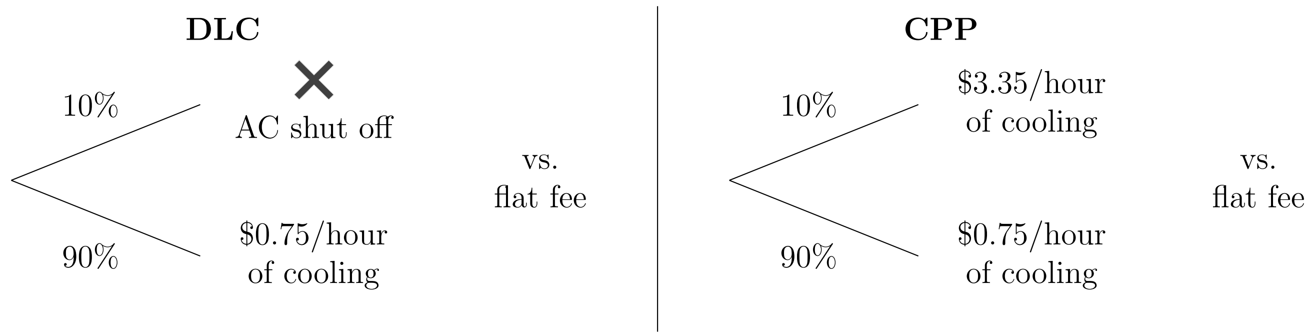

Figure 1: Stylized illustration of Direct Load Control (DLC) versus Critical Peak Pricing (CPP) contract choice relative to a flat fee for electricity use.

In this simple case with non-stochastic benefits to air conditioner use of \$3, which is below the peak price of \$3.35, these two contract types would result in identical actions by the household: running their air conditioner on normal days and having no air conditioning during control or peak-price events respectively. EUT therefore implies that households should value these two contracts equally. However, under CPP the latter outcome results from an active choice, whereas it was imposed on the consumer under DLC. An intrinsic value of control would therefore create a wedge between the valuations of these two contract offers, with the CPP contract being valued more highly. Which form this wedge would take is a priori unclear. It could be a fixed monetary value independent of the instrumental value of the decision, or respond to the stakes and probability of the control event. Such a wedge could also lead to decreased acceptance and hence low adoption rates of DLC contracts.[7]

I mimic this contract choice in a stylized lab environment in which participants complete a series of real-effort tasks. Participants have 30 seconds to complete each task, with the option of upgrading to 60 seconds in exchange for some cost. Under the implemented payment structure, increased time available translates into increased expected task earnings. Since air conditioning can

[7]In case the peak price is below the participant's true WTP, or in case the the true WTP is stochastic, the CPP contract has instrumental (option) value. In this case, the valuation difference between the CPP and DLC contract could arise in the form of a fixed add-on independent of the instrumental value, or as a relative add-on effectively inflating the instrumental value of the CPP contract. The existing literature does not touch on this potential link between control premia and (mis-)perceived option value.

also be viewed as a productivity-enhancing tool, participants' WTP for the additional 30 seconds serves as an experimental proxy for this instrumental value. After familiarizing participants with the task and the upgrade decision, I present participants with a series of *control-lottery decision*: each control lottery features some probability of being able to upgrade at some low cost, but forces the participant to complete the task in 30 seconds otherwise. I then elicit participants *equivalent flat fee* (EFF) for these lotteries, i.e. the maximum certain fee for upgrading they would prefer to face rather than playing the control lottery.

Participants provide EFFs for a series of control lotteries that differ with respect to the probability of losing control as well as the stakes of the upgrade decision. I then estimate the predicted EFF for a control-neutral participant. If stated EFFs exceed the predicted control-neutral level, I attribute this difference to a participant's intrinsic preference for control. In order to assess the validity of this interpretation, I also present participants with *peak-price lotteries* that potentially charge a higher upgrade cost instead of eliminating the opportunity to upgrade altogether. The peak-price to WTP-to-upgrade ratio varies between subjects. By comparing stated reservation costs for control and peak-price lotteries, I can rule out potentially confounding explanations of the control premium such as probability weighting or certainty effects.

To formalize my definition of the control premium, I begin by describing the relevant decision context. In Section 2, I derive predicted valuations for a control-neutral individual assuming preferences that either follow standard EUT or are based on cumulative prospect theory (CPT). I then propose absolute and relative measures of the control premium which are based on comparing participants' stated valuations to these predicted control-neutral levels. Section 3 describes the experimental design which elicits the valuations required to derive control premium estimates as well as to rule out competing explanations such as heterogeneity in the individual degree of risk aversion or probability weighting. In Section 4, I test for the significance of the control premium estimates, provide evidence supporting a control-based interpretation of these estimates, and discuss how the control premium varies with the stakes and probabilities of the control decision. Section 5 proposes takeaways for DLC and Section 6 concludes.

2 Theoretical Framework

Participants in this experiment face a series of decisions over *control lotteries*: lotteries that have some probability $(1-p)$ of being able to upgrade to a completion time of 60 instead of 30 seconds at a cost of c_L (which is the same across all participants), and some probability p of losing control over the upgrade decision and not being able to upgrade at all. Note that there is no ambiguity regarding outcomes under the no-control scenario under this setup: participants know that they will only have 30 seconds to complete the task in this case. Figure 2 illustrates the basic structure of control lotteries in this experiment.

For each control lottery decision, participants are asked to report the *equivalent flat fee* (EFF)

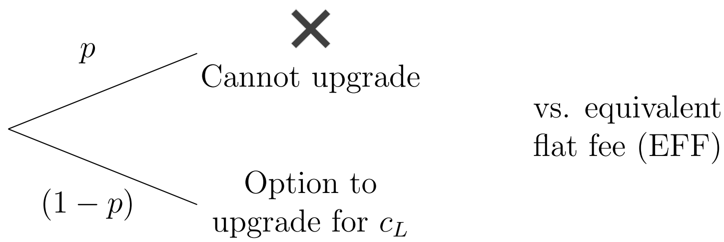

Figure 2: Control Lottery Illustration

they would prefer to face for certain instead of playing the control lottery. The proposed definition of the control premium relies on the difference between stated EFFs and their predicted control-neutral levels. In this section, I derive predicted control-neutral EFFs under standard expected utility theory (EUT) and cumulative prospect theory (CPT) which allows for probability weighting.

In order to study whether the control premium represents a fixed add-on or a relative multiplier depending on the intrinsic value of the decision I consider an absolute and a relative measure of the control premium. The absolute control premium is defined as the difference between the stated EFF of a control lottery, EFF_i, and its predicted control-neutral equivalent, $\widehat{EFF}_i$:

$$CP_{i,abs} = EFF_i - \widehat{EFF}_i \tag{1}$$

The relative control premium compares this difference to the level of the predicted control-neutral EFF:

$$CP_{i,rel} = \frac{EFF_i - \widehat{EFF}_i}{\widehat{EFF}_i} \tag{2}$$

The next two sections outline how $\widehat{EFF}_i$ is calculated under EUT and CPT respectively.

2.1 EFF Estimates under Standard Expected Utility Theory

The real-effort task in this experiment consists of a word search grid with ten search words. Participants thus have the chance to find between zero and ten words per task with a fixed reward of $1 per word found. The probability of finding $k \in \{0, ..., 10\}$ words depends on the amount of time available to complete the task ($s \in \{30, 60\}$). Denoting participant i's utility function by $u_i(\cdot)$ and the probability of finding k words in s seconds by $\mathbb{P}_i(k|s)$, the expected earnings of a s-second task are given by

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s) u_i(k). \tag{3}$$

Participant i's willingness-to-pay to upgrade from 30 to 60 seconds, WTP_i, equals the difference in expected earnings based on the amount of time available:

$$u_i(WTP_i) = \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k) - \sum_{k=0}^{10} \mathbb{P}_i(k|s=30)u_i(k). \tag{4}$$

Under control lottery decisions, participants are unable to upgrade with probability p, and can upgrade at a cost of c_L otherwise. Assuming that $c_L \leq WTP_i$, participant i's control-neutral equivalent flat fee $\widehat{EFF}_i$ solves

$$\begin{aligned} \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k - \widehat{EFF}_i) = (1-p) \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k - c_L) \\ + p \sum_{k=0}^{10} \mathbb{P}_i(k|s=30)u_i(k) \end{aligned} \tag{5}$$

As shown in appendix B.1, under risk neutrality this is the case when

$$\widehat{EFF}_i = (1-p)c_L + pWTP_i. \tag{6}$$

A participant's EFF is thus a linear combination of the low cost realized with probability $(1-p)$ and their maximum willingness to pay to upgrade WTP_i.[8] The same is true under the peak-price lottery whenever the peak price, $c_{i,H}$, exceeds the participant's willingness to pay to upgrade. When $c_{i,H} < WTP_i$, the participant always chooses to upgrade and the EFF satisfies

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k - \widehat{EFF}_i^{CPP}) = \sum_{k=0}^{10} \mathbb{P}_i(k|s=60) \left[(1-p)u_i(k - c_L) + pu_i(k - c_{i,H})\right]. \tag{7}$$

In this case, the EFF under risk neutrality is a linear combination of the two cost levels, and hence bounded by $c_{i,H}$:

$$\widehat{EFF}_i^{CPP} = (1-p)c_L + pc_{i,H} \tag{8}$$

Assuming risk neutrality, all components of a participant's control-neutral predicted EFF except WTP_i are thus exogenously given. Section 3 outlines how I elicit WTP_i. In addition, I describe how I estimate population-average preferences over money in order to relax the assumption of risk neutrality.

[8] In the exceptional case in which $c_L < WTP_i$ and the participant never chooses to upgrade under either the control or peak-price lottery, the EFF reduces to $\widehat{EFF}_i = WTP_i$.

2.2 EFF Estimates under Cumulative Prospect Theory

Under CPT, I will assume that participants evaluate their prospects in a given task relative to the anchor of finding zero words and therefore earning zero dollars. Denoting by $\mathbb{P}_i(\geq k|s)$ the probability of participant i finding at least k words in s seconds, the expected earnings from a given task can therefore be represented as

$$\begin{aligned}
&\pi_i(1)u(0) + \pi_i(\mathbb{P}_i(\geq 1|s))\left[u_i(1) - u_i(0)\right] + ... + \pi_i(\mathbb{P}_i(\geq 10|s))\left[u_i(10) - u_i(9)\right] \\
&= \left[\pi_i(\mathbb{P}_i(\geq 0|s) - \pi_i(\mathbb{P}_i(\geq 1|s)\right] u_i(0) + ... + \left[\pi_i(\mathbb{P}_i(\geq 11|s) - \pi_i(\mathbb{P}_i(\geq 10|s)\right] u_i(10) \\
&= \sum_{k=0}^{10} \left[\pi_i(\mathbb{P}_i(\geq k|s) - \pi_i(\mathbb{P}_i(\geq k+1|s)\right] u_i(k) \qquad (9) \\
&\equiv \sum_{k=0}^{10} V_i(s)u_i(k)
\end{aligned}$$

The shorthand notation $\sum_{k=0}^{10} V_i(s)u_i(k)$ represents the probability-weighted word-score utility under CPT. The cost that makes participants indifferent between upgrading to 60 seconds or completing the word search in 30 seconds is given by the WTP_i which satisfies

$$u(WTP_i) \equiv \sum_{k=0}^{10} V_i(60)u_i(k) - \sum_{k=0}^{10} V_i(30)u_i(k). \qquad (10)$$

Since participants are explicitly asked to form expectations around their performance in 30-second versus 60-second tasks in Part 2, I assume that they begin Part 3 with a fixed idea of $\sum_{k=0}^{10} V_i(30)u_i(k)$ and $\sum_{k=0}^{10} V_i(60)u_i(k)$. In this case, rather than evaluating the task performance and the lottery jointly, and hence applying probability weighting jointly across the two, I will instead assume that participants anchor on the control event and process performance and lottery outcomes sequentially. The relevant control-neutral EFF of a control lottery is then given by the $\widehat{EFF}_i$ which solves

$$\sum_{k=0}^{10} V_i(60)u_i(k - \widehat{EFF}_i) = \pi_i(1)\sum_{k=0}^{10} V_i(30)u_i(k) + \pi_i(1-p)\left[\sum_{k=0}^{10} V_i(60)u_i(k - c_L) - \sum_{k=0}^{10} V_i(30)u_i(k)\right] \qquad (11)$$

By rearranging and employing the same transformation as under EUT, we find that for a risk-neutral participant,

$$\widehat{EFF}_i = \pi_i(1-p)c_L + (1 - \pi_i(1-p))WTP_i. \qquad (12)$$

The corresponding EFFs for the peak-price lottery under CPT are similarly given by probability-weighted equivalents of their EUT counterparts. In general, CPT will lead to *higher* estimates of $\widehat{EFF}_i$ than EUT due to the relatively larger probability weight on the rare control (or peak-price) event. Since the control premium is estimated as the difference between stated and control-neutral predicted certainty equivalents, allowing for CPT preferences therefore leads to more conservative control premium estimates. Given the significant degree of probability weighting observed in my sample, estimates based on this sequential CPT approach represent my preferred measure of the control premium.

In the case of CPT, the components of $\widehat{EFF}_i$ are no longer fully observable since I do not collect individual-specific estimates of $\pi_i(\cdot)$. Instead, I operationalize the predicted level $\widehat{EFF}_i$ using the population-average probability weighting function $\pi(\cdot)$.

More severe probability weighting will tend to *inflate* the true level of $\widehat{EFF}_i$ (again due to the relatively larger weight on WTP_i). As a result, approximating $\widehat{EFF}_i$ using population-average probability weighting may bias the estimates of $\widehat{EFF}_i$ downward for some individuals, which would result in a larger estimated control premium. Individual heterogeneity in the degree of probability weighting therefore represents an important potential confound to my control premium estimates. To account for this fact, I run a series of robustness checks to rule out that my control premium estimates are driven by participants who over-weight low probabilities more than the average participant.

3 Experimental Design

The experiment is designed to generate individual control premium estimates at the participant level. The control premium is defined as the difference between the reported equivalent flat fee (EFF) a participant would prefer to face for certain instead of playing a control lottery and its predicted control-neutral level. In addition to the EFF itself, I need to elicit participants' risk preferences and WTP to upgrade in order to derive the control premium estimates.

3.1 Sequence of Events

The experiment consists of four parts. In Part 1, I elicit participants' preferences over monetary outcomes (including risk aversion and probability weighting). Part 2 consists of a series of real-effort tasks preceded by the opportunity to upgrade from 30 to 60 seconds of task completion time at a certain cost. Part 2 serves to calibrate participants on their WTP to upgrade to 60 seconds and familiarizes them with the upgrade decision. In Part 3, participants face a series of decisions over upgrade costs, control lotteries, and peak-price lotteries. These decisions determine the possibility and cost of upgrading from 30 to 60 seconds during one final word-search task at the end of Part 3. Part 4 is an exit survey which collects participant demographics, responses to the

Burger Desirability of Control Scale, and presents participants with some open-ended questions about the decisions earlier in the survey.

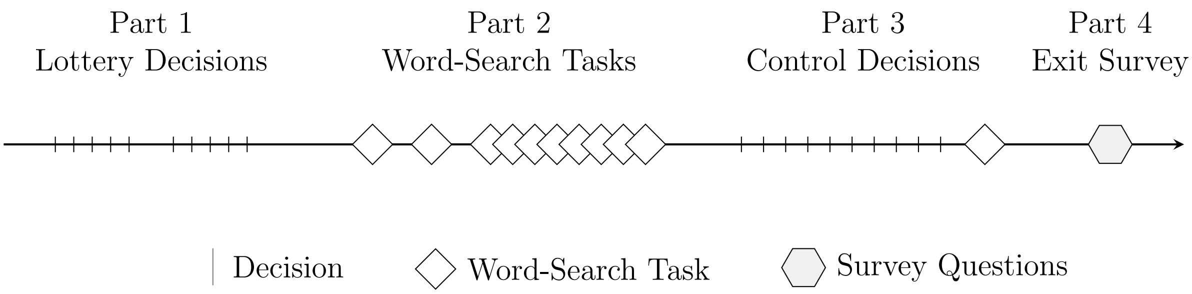

Figure 3: Sequence of Events

All decisions in Parts 1–3 are incentivized using a two-stage iterative Multiple Price List (MPL) approach, a modified version of the Becker-de Groot-Marshak (BDM) approach (Becker et al., 1964). MPL designs present participants with a table of paired options and ask them to make a decision for every row. The results are determined by drawing one row at random and implementing the participant's preferred option for that row. This elicitation method has the advantage of being easy to explain, and making the optimality of truthful revelation arguably more intuitive than under standard BDM. The two-stage iterative MPL (iMPL) allows participants to refine their initial stated valuation using a second table which offers values between the first value which participants initially accepted, and the last value participants initially rejected. In order to be able to present participants with a more granular price list based on their initial switching point, a single switching point is enforced in both the initial, more granular and the refined table.[9] The main advantage of the iMPL approach is that it provides relatively granular valuations without subjecting participants to an overwhelming number of choices.

3.1.1 Part 1: Eliciting Preferences over Money

Part 1 consists of ten lottery questions grouped into two sets of five decisions that differ with respect to their framing, but have equal potential payouts and probabilities. The order in which these two sets are shown to participants is randomized between subjects. Under the *standard frame*, a random draw determines whether the lottery pays out $5 or nothing in that round. Participants are asked to provide the minimum sure amount they would be willing to accept instead of playing this lottery.

However, this decision frame is somewhat different from the type of decisions participants face in Part 3. In Part 3, the lottery determines the cost (or possibility) of achieving a desirable outcome

[9] Andersen et al. (2006) compare estimates of risk aversion and discount rate parameters across different MPL elicitation approaches, such as standard MPLs, MPLs with an enforced unique switching point, and iterative MPLs. They find that the results are qualitatively similar across approaches, while the magnitude of risk aversion estimates appears a little more sensitive to the elicitation approach than estimates of discount factors.

(namely upgrading from 30 to 60 seconds). Participants in this case report the maximum certain cost they would prefer to playing the lottery. I will refer to this amount as the participant's *equivalent flat fee* (EFF). Assuming that this decision frame triggers a different degree of risk aversion or probability weighting, estimates using the standard frame might not adequately capture the preferences relevant to the decisions of interest.

To mirror this decision frame more closely, I therefore introduce a *token frame* in which the random draw determines whether participants purchase a \$5 token for free, or in exchange for \$5 (effectively making zero money in this round). Participants in this case report the maximum certain payment they would be willing to make for the token instead of playing the lottery. The token frame also serves to familiarize participants with this decision type.

The five probability levels considered in Part 1 are $p \in \{0.1, 0.2, 0.5, 0.8, 0.9\}$. Based on these lottery decisions, I construct (i) parametric estimates of the population-average degree of probability weighting using standard functional forms assumed in the literature; and (ii) non-parametric proxies for a participant's individual degree of risk aversion and probability weighting. I use these additional measures to rule out heterogeneity in preferences as a potential confound of my control premium estimates. In particular, a participant's individual degree of risk aversion can be approximated by their reported certainty equivalent for a standard lottery question with a 50% chance of winning. Since probability distortions are likely to be small in this probability region, lower reported certainty equivalents can be interpreted as evidence of heightened risk aversion. A negative and statistically significant rank correlation between a participant's certainty equivalent at this point and their estimated control premium would therefore suggest that at least part of the control premium should be attributed to risk aversion instead.

To capture the individual-specific degree of probability weighting, I leverage the ratio of reported certainty equivalents at two complementary probability levels: 90% and 10%. More over-weighting of low probabilities, and under-weighting of large probabilities, will bring the corresponding certainty equivalents closer together, leading to a lower ratio. In this case, a negative and statistically significant rank correlation between this measure and a participant's estimated control premium would therefore suggest that part of the estimated control premium is actually a result of heterogeneity in probability weighting.

3.1.2 Part 2: Word-Search Task

The word-search task consists of an eight-by-eight grid of letters with ten hidden words. Participants can find these words in any order.[10] To incentivize performance in Part 2, I pay participants a fixed reward of \$1 per word found in one randomly selected round. Figure 4 shows a task screen example.

[10] Search grids were created using Discovery Education's Puzzlemaker: `http://puzzlemaker.discoveryeducation.com/WordSearchSetupForm.asp`.

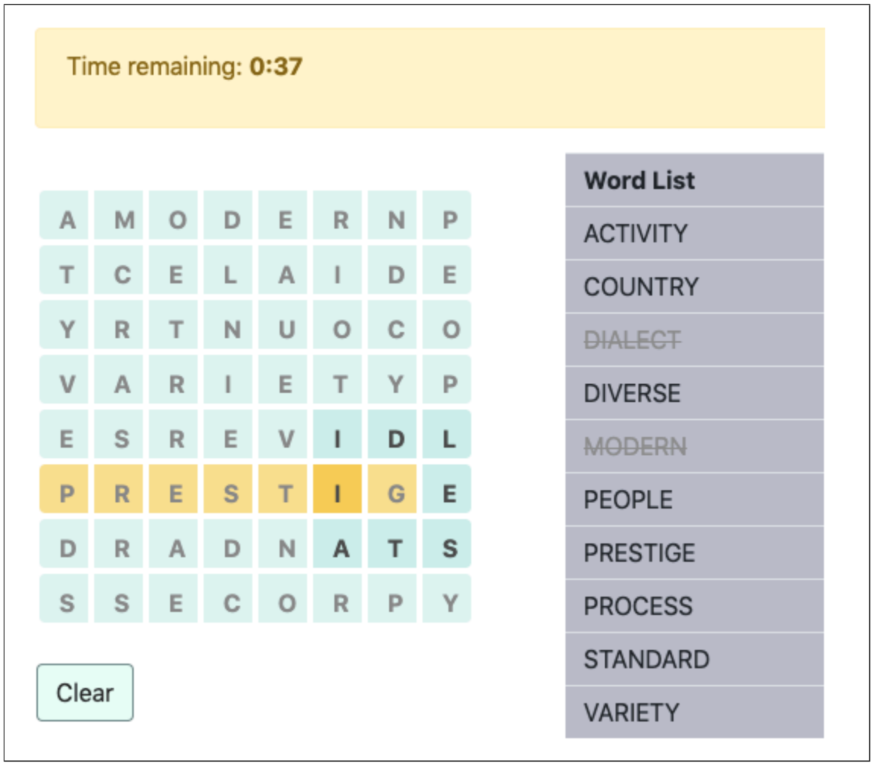

Figure 4: Word-Search Task Example

The first two rounds in Part 2 serve to familiarize participants with their relative performance depending on the time available: 60 seconds in round 1 versus 30 seconds in round 2. Before round 2, a prompt explicitly asks participants to pay attention to the impact of the reduced time on their performance. Before the start of each remaining round, participants have the opportunity to decide whether or not to upgrade from 30 to 60 seconds at a given cost. Figure 5 shows a screenshot of the upgrade screen. In this example, participants face the option of upgrading to 60 seconds in exchange for \$1.90. Round 3 features a relatively low cost of \$0.75 for every participant. Round 4 always features a relatively high cost of \$6.15. In each subsequent round, the cost a participant sees is randomly drawn from the list {\$1.35, \$1.90, \$2.45, \$3.00, \$3.55, \$4.10} without replacement. The cost of upgrading is only applied to earnings in the same round. Participants do not incur upgrade costs based on decisions for rounds that are not selected for payment at the end of the survey.

If a participant chooses to upgrade, they directly proceed to the word-search task having 60 seconds available to complete the round. If they choose not to upgrade, participants remain on the wait screen for 30 seconds demarcated by the large countdown clock on the left. They then proceed to the word-search task with only 30 seconds available to complete the round.

Part 2 serves to (i) allow participants to familiarize themselves with the task environment and to gauge the impact of time available on their performance; (ii) help participants translate the perceived performance impact of upgrading into a willingness to pay to upgrade; and (iii) align

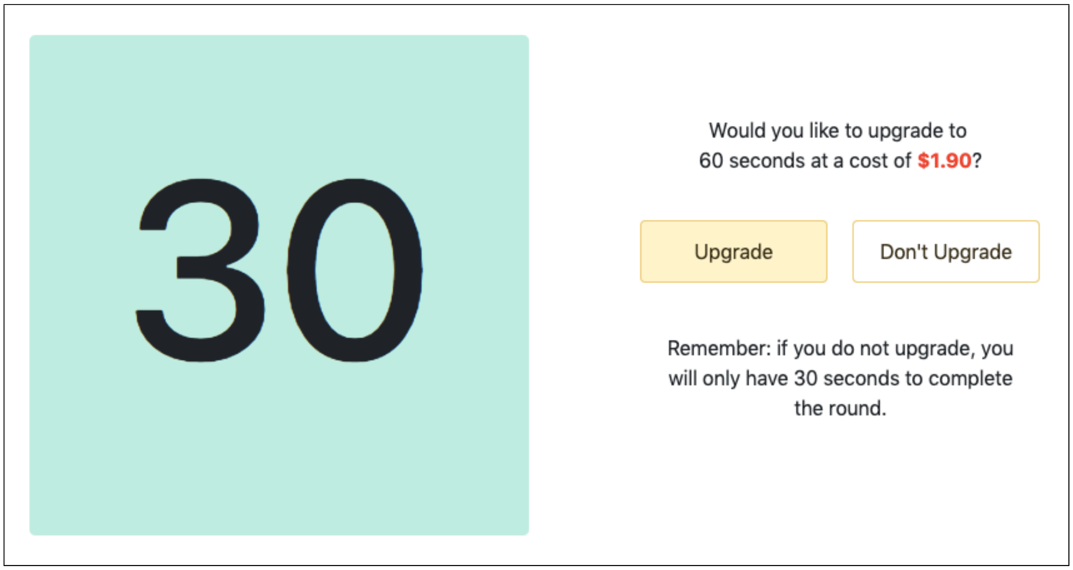

Figure 5: Upgrade Decision Screen

participants beliefs about the degree of control they have in the upgrade decision with reality. In particular, participants learn that the upgrade decision involves the opportunity, but not the obligation, to upgrade at a given cost. Minimizing the discrepancy between objective and perceived control mitigates the possibility of illusion of control as a potential confound.

3.1.3 Part 3: Control Decisions

In Part 3, participants face 10 lottery decisions followed by one final word-search task. The ten decisions can be classified into four groups: elicitation of the willingness to pay to upgrade (2/10); control lottery decisions (5/10); peak-price lottery decisions (1/10); and attention checks (2/10).

Elicitation of the Willingness to Pay to Upgrade

Decisions 1 and 7 elicit a participant's willingness to pay to upgrade, denoted by WTP_i, in two ways. One decision is framed directly as a choice between having 30 vs 60 seconds available to complete the final word-search task. The other decision is framed as the maximum flat fee below which participants would prefer to 'auto-upgrade' to 60 seconds. The order in which these question frames appear is randomized between participants. I included two separate elicitations of this value in order to test for its stability across decisions. If participants systematically update their indifference point based on the implicit value learning exercise of the intervening (control) questions, this might lead to a bias in my control-premium measure. In addition, re-eliciting this value allows me to test for decision fatigue which might affect the degree of measurement error in the later decisions of Part 3.

Control Decisions

Five decisions in Part 3 consist of control lotteries. Control lotteries in Part 3 have some probability p of losing control over the upgrade decision and being forced to complete the task in 30 seconds instead. With the remaining probability, participants have the opportunity to upgrade at a cost of $c_L = \$1.25$ which is the same across all participants.

The first three control decisions are designed to measure the individual-specific control premium as a function of the probability p of a control event occurring, as well as the stakes of the control event. The first decision provides a baseline measure at the usual stakes of \$1 per word found, and a control event probability of 20%. I will refer to this as the *base* control decision $\mathcal{B}$. Figure 6 shows a screenshot of the base control decision in the survey. Under control decision $\mathcal{P}$, the control event probability is doubled to 40%. Under control decision $\mathcal{S}$, the stakes are doubled from \$1 to \$2 per word found. The order of decisions $\mathcal{P}$ and $\mathcal{S}$ is randomized between participants.

For each of these three decisions, participants are asked to report their EFFs, i.e. the maximum certain flat fee for upgrading they would prefer to face rather than playing the control lottery. Note that participants are neither forced to upgrade at their EFF, nor at the cost realized under the control or peak-price lotteries at the end of Part 3. Rather, they have the opportunity to upgrade at this cost if desired.

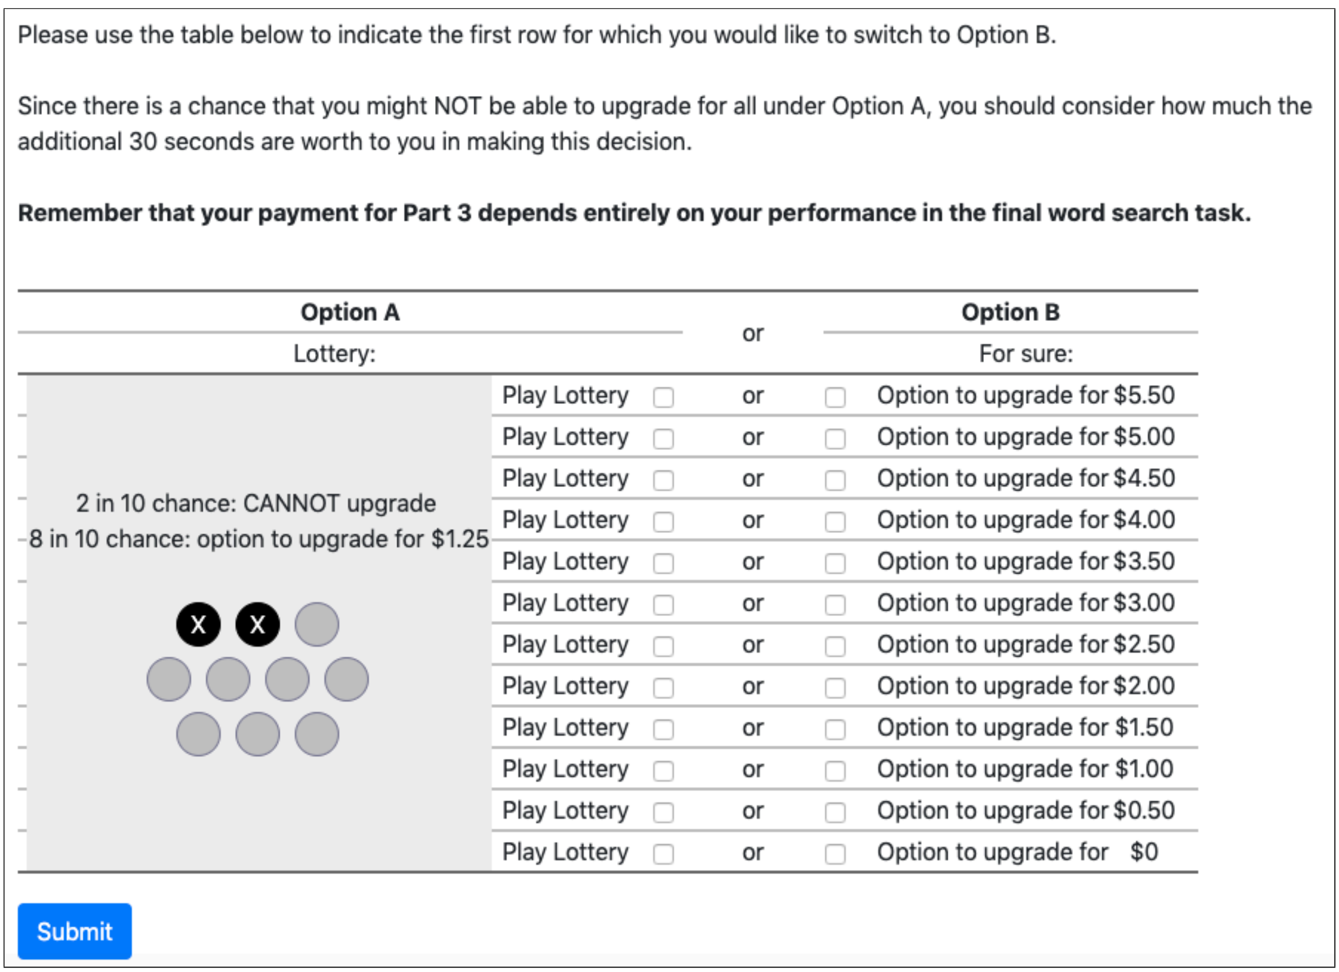
Please use the table below to indicate the first row for which you would like to switch to Option B.

Since there is a chance that you might NOT be able to upgrade for all under Option A, you should consider how much the additional 30 seconds are worth to you in making this decision.

Remember that your payment for Part 3 depends entirely on your performance in the final word search task.

Option A			or	Option B
Lottery:				For sure:
2 in 10 chance: CANNOT upgrade 8 in 10 chance: option to upgrade for \$1.25	Play Lottery	☐	or	☐ Option to upgrade for \$5.50
	Play Lottery	☐	or	☐ Option to upgrade for \$5.00
	Play Lottery	☐	or	☐ Option to upgrade for \$4.50
	Play Lottery	☐	or	☐ Option to upgrade for \$4.00
	Play Lottery	☐	or	☐ Option to upgrade for \$3.50
	Play Lottery	☐	or	☐ Option to upgrade for \$3.00
	Play Lottery	☐	or	☐ Option to upgrade for \$2.50
	Play Lottery	☐	or	☐ Option to upgrade for \$2.00
	Play Lottery	☐	or	☐ Option to upgrade for \$1.50
	Play Lottery	☐	or	☐ Option to upgrade for \$1.00
	Play Lottery	☐	or	☐ Option to upgrade for \$0.50
	Play Lottery	☐	or	☐ Option to upgrade for \$0

Submit

Figure 6: Control Lottery Screen (Baseline)

The last two control decisions test for the presence of endowment effects with respect to control. In this case, participants are randomly endowed with one control lottery and are offered the opportunity to switch to an alternative control lottery with a control event probability p that is twice (half) as large in exchange for money. In particular, participants are offered the opportunity to switch between a fixed cost of upgrading of c_L (i.e. $p = 0\%$) versus base control lottery $\mathcal{B}$ ($p = 20\%$) in question $\mathcal{E}_1$, and to switch between base control lottery $\mathcal{B}$ and lottery $\mathcal{P}$ with doubled control-event probabilities ($p = 40\%$) in decision $\mathcal{E}_2$. Once again, the order of decisions $\mathcal{E}_1$ and $\mathcal{E}_2$ is randomized between participants.

Peak-Price Lottery Decision

The validity of the proposed control-premium measure depends in large part on the correct specification of the predicted control-neutral EFF. The preferences elicited in Part 1 serve to construct a prediction which takes population-average risk preferences into account. The peak-price lottery $\mathcal{CPP}$ serves as an additional test to confirm that the observed valuation difference can indeed be attributed to an aversion to ceding control and cannot be fully explained by potentially confounding factors such as probability weighting.

Like lottery $\mathcal{B}$, the peak-price lottery leads to a 20% chance of being able to upgrade at $c_L = \$1.25$ (which lies below most participants' WTP to upgrade). Unlike lottery $\mathcal{B}$ however, participants still have the option to upgrade at some cost $c_{i,H}$ under the alternative outcome (see Figure 7 for illustration). If $c_{i,H} > WTP_i$, the participant should choose not to upgrade since the cost exceeds their WTP. In this case, the peak-price lottery therefore induces the same actions as the control lottery: a 20% chance of upgrading at \$1.25, and an 80% chance of completing the task in 30 seconds. Participants' EFFs to avoid playing these two lotteries should therefore be equal whenever $c_{i,H} > WTP_i$ irrespective of the participant's preferences over monetary outcomes. If participants report a higher EFF for lottery $\mathcal{B}$ in this case, the difference can therefore be directly attributed to an aversion to ceding control.[11]

For $c_{i,H} \leq WTP_i$, the picture is more nuanced. In this case, the peak-price lottery carries intrinsic value since participants can increase their expected earnings by upgrading. Participants should be more willing to play the peak-price lottery and should therefore report a lower EFF than under the control lottery. In addition, participants may value the peak-price lottery $\mathcal{CPP}$ more highly in this case due to a certainty effect over outcomes: if $c_{i,H} \leq WTP_i$, participants expect to upgrade regardless of the outcome of the lottery. Under the control lottery on the other hand, the realized completion time depends on the outcome of the lottery. If a significant certainty effect with respect to outcomes exists, we should therefore see the intrinsic value jump up (and hence the stated EFF jump down) to the left of the $c_{i,H} = WTP_i$ cutoff. Because the excess intrinsic value

[11]Note that in the extreme case of $c_{i,H} = \infty$, one would expect the participant to realize that they would never exercise their upgrade option at that cost. In this case, it is therefore likely that even a control-loving individual would report equal EFFs for the two lotteries. The region of interest therefore consists mainly of observations for $c_{i,H}$ approaching WTP_i from above.

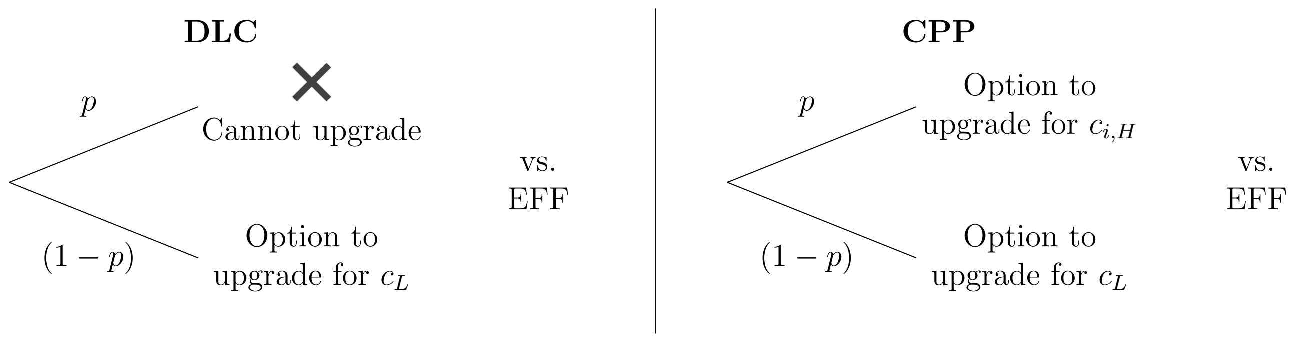

Figure 7: Control Lottery vs. Peak-Price Lottery

under $\mathcal{CPP}$ is only realized with probability p, this relationship is again subject to a potential probability-weighting confound. However, differences in EFFs that cannot be explained by any reasonable degree of probability weighting would again provide additional evidence supporting the existence of a control premium.

To avoid creating incentives to game the survey by basing $c_{i,H}$ on a participant's reported WTP to upgrade, I instead set it based on participants' performance in Part 2. In particular, I construct a performance measure $WTP_{i,est}$ as the average number of words found in the last 30 seconds across all 60-second rounds in Part 2. I use data from two pilots to estimate the linear relationship between participants' stated WTP_i and the observed $WTP_{i,est}$, and set $c_{i,H}$ based on this predicted relationship.[12]

Attention Checks

Two decisions in Part 3 served as attention checks. Participants were made aware of the presence of attention checks in the introduction to Part 3. Attention checks were deliberately easy to spot with the words "Attention Check" highlighted in bold, red font at the start of the prompt above the MPL table. The decisions looked otherwise similar to the other decisions in Part 3. To pass the attention check, participants were asked to ignore the information presented in the MPL and to select a pre-specified row instead. Attention checks were used to drop participants who did not pay attention to the provided instructions.

Before the final word-search task, one of the ten lottery decisions in Part 3 is randomly selected to be implemented for real. The outcome of the randomly selected decision determines a participant's cost and opportunity of upgrading to 60 seconds for this task. All decisions are implemented using the two-stage iterative MPL approach outlined in Part 3. Participants' EFF is estimated by the midpoint between the highest accepted cost and the rejected cost above in the second, more granular price list. Figure 8 summarizes the order of decisions in Part 3. Arrows indicate a randomization of decision order between subject.

[12]The predicted linear relationship consists of a constant of \$0.86 and a coefficient of 0.51 on $WTP_{i,est}$ (p-value=0.0021). The statistical significance of the coefficient on $WTP_{i,est}$ suggests that $WTP_{i,est}$ is a reasonable proxy for performance, and that participants take their performance into account in setting WTP_i.

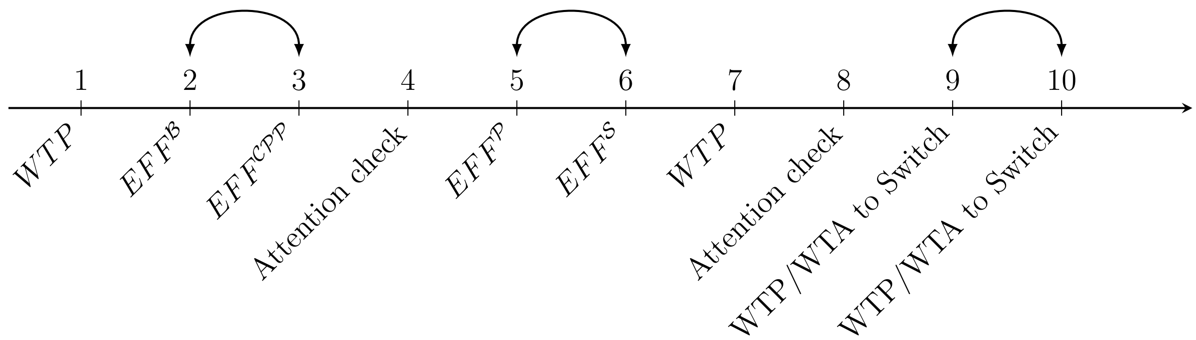

Figure 8: Decision Order in Part 3

3.1.4 Part 4: Exit Survey

Part 4 consists of basic demographics questions; the 20 questions of the Desirability of Control Scale (Burger and Cooper, 1979); and some open-ended questions relating to participants' relative preferences over a control versus a peak-price lottery for the case in which $c_{i,H}$ barely exceeds WTP_i (i.e. the case in which the peak-price lottery has no intrinsic value). Participants do not earn any additional money in Part 4.

3.2 Procedures

I use data from 180 participants recruited through Cornell's Business Simulation Lab (BSL). The subject pool consisted of currently enrolled Cornell graduate and undergraduate students. The experiment was programmed in oTree (Chen et al., 2016) and took about 45 minutes to complete.[13] This experiment was significantly longer and more computationally challenging than the average experiment run on the BSL online platform in 2019. Pilot studies revealed a significant share of participants who struggled with the lottery questions and/or the use of price lists. Based on these pilots, I designed and pre-registered four exclusion criteria to drop participants who (i) spent too little time on the decisions in Part 3; (ii) failed the attention checks in Part 3; (iii) reported a zero WTP to upgrade; and (iv) provided inconsistent answers during the standard lottery questions in Part 1. In total, 137 participants were dropped from the final sample based on these exclusion criteria. Table A1 in the Appendix shows how many participants were removed by each filter.

A pre-analysis plan including detailed sample exclusion criteria was pre-registered on the AEA RCT registry (Korting, 2019). Stringent exclusion criteria were implemented to ensure that participants (i) have a firm understanding of the MPL mechanism; (ii) are comfortable with lottery problems; and (iii) are paying enough attention to the control lottery decisions. The goal is to minimize the effect of responses driven by potential 'mistakes' rather than genuine preferences.

[13] The experiments proposed in BFH14 and OGF14 took 2-3 hours and 80 minutes on average to complete respectively, which suggests that the proposed setup in this paper provides a more time-effective way to elicit control premia.

Participants received a $2 participation fee and were paid a bonus based on one randomly drawn decision from each of Parts 1, 2 and 3. To avoid wealth effects, the outcomes of all randomly selected decisions and the resulting bonus payments were only announced after the survey was complete. On average, participants earned $14.27.

4 Results

Of the 180 participants included in this sample, 68% were between 18–22 years old, and 75% were female[14]. All subjects were currently registered Cornell graduate or undergraduate students. The following sections outline the population-level estimates of monetary preferences based on the lottery choices in Part 1; participants' performance during the 10 word-search tasks in Part 2; and the control premium estimates based on participants' responses to the control and peak-price lotteries in Part 3.

4.1 Preferences over Money

As outlined in section 3.1.1, I obtain parametric estimates of the population-average degree of probability weighting and risk aversion using the lottery decisions in Part 1. In a first step, I test whether the estimated preferences over money are statistically different between the standard and token frame. Assuming that participants evaluate the $5 token and the lottery jointly, a p% chance to win $5 or nothing is identical to a p% chance of winning a $5 token for free or in exchange for $5. The reported EFF of the latter should thus equal $5 minus the certainty equivalent of the simple lottery under both expected utility theory (EUT) and cumulative prospect theory (CPT). Using a paired t-test, I cannot reject that this relationship holds (p-value = 0.148).[15]

Parametric estimates of monetary preferences are obtained by jointly estimating the two non-linear least square regressions for the standard and token frame with standard errors clustered at the participant level:

$$\begin{aligned}\text{Standard Frame:} \quad & CE = u^{-1}(\pi(p)u(\$5)) \\ \text{Token Frame:} \quad & CE = 5 - u^{-1}(\pi(p)u(\$5))\end{aligned}$$

I either assume risk neutrality at small stakes[16], or allow for a curved utility function of the form $u(x) = x^{\alpha}$ as in Tversky and Kahneman (1992) and Bernheim and Sprenger (2019) to get a

[14]This is in line with the average ratio of females in the participant pool of 78%.

[15]If instead participants consider the $5 token as an initial endowment and do not compound the token and the payment lottery in their assessment, then $CE_{standard} = u^{-1}(pu(5))$ and $CE_{token} = u^{-1}((1-p)u(5))$ under EUT, with an equivalent probability-weighted relationship under CPT.

[16]This assumption is in line with the intuition presented in Rabin (2000).

sense of the population average degree of risk aversion. Note however that there is no variation in lottery payoffs. Estimates of α should therefore be interpreted as suggestive evidence for risk aversion. In addition, I test for probability weighting under both the functional form assumed in Tversky and Kahneman (1992) and the Prelec form (Prelec et al., 1998):

$$\text{TK (1992):} \quad \pi(p) = \frac{p^\gamma}{[p^\gamma + (1-p)^\gamma]^{\frac{1}{\gamma}}}$$
$$\text{Prelec (1998):} \quad \pi(p) = exp\left\{-(-ln(p))^\gamma\right\}$$

Table 1 summarizes the corresponding estimation results. A Wald test to assess whether the probability weighting parameter $\hat{\gamma}$ is equal across frames is rejected at the 10% significance level under assumed risk neutrality (p-value = 0.0955), but cannot be rejected when the risk aversion parameter is also allowed to vary (p-value = 0.228).

Table 1: Parametric Estimates of Monetary Preferences

	$\hat{\alpha}$	$\hat{\gamma}$
	Assuming Risk Neutrality	
Standard Frame		0.86
		(0.02)
Token Frame		0.83
		(0.02)
	Allowing for Risk Aversion	
Standard Frame	1.11	0.86
	(0.03)	(0.02)
Token Frame	1.00	0.83
	(0.03)	(0.02)

Estimates based on joint non-linear least squares estimation of the equations for the standard and token frame clustered at the participant level. Numbers in parentheses indicate standard errors. The probability weighting function is assumed to be of the form presented in Tversky and Kahneman (1992).

The estimates for $\hat{\alpha}$ suggest that the assumption of risk neutrality at small stakes is reasonable at the population level. $\hat{\alpha}$ is exactly one under the token frame, and larger than one under the standard frame which would suggest risk loving participants on average. Ignoring risk loving would lead to a downward bias in my control premium estimates. In this case my estimates would therefore serve as a lower bound on the actual control premium.

Results under the token frame do not change much depending on the assumed functional forms for $\pi(p)$ and assuming compounding of the value of the token and the payment determined by the lottery or not. In order to generate conservative estimates of the control premium, I therefore use the estimates from the token frame assuming zero risk aversion at small stakes, compounding of token and lottery, and using the functional form suggested in Tversky and Kahneman (1992) as my preferred estimate. Figure 9 highlights the estimated shape of the population-average probability weighting function.

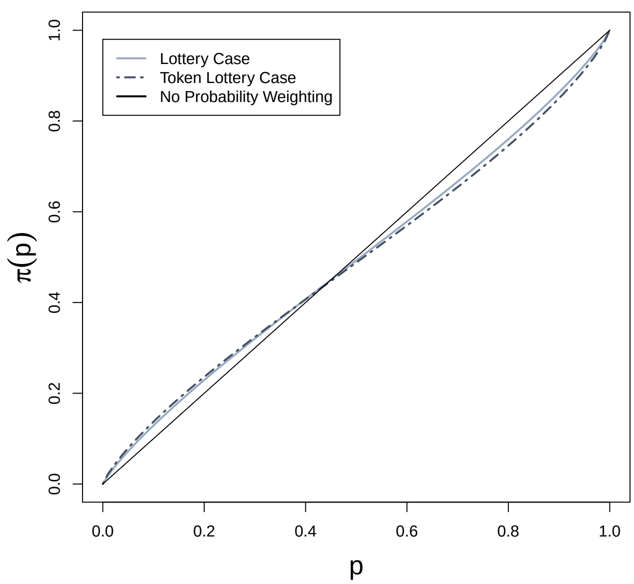

Figure 9: Estimated probability weighting functions under the standard and token frame assuming the parametric form suggested in Tversky and Kahneman (1992)

4.2 Task Performance and WTP to Upgrade

On average, participants found 3.6 words in 30 seconds compared to 7.3 words in 60 seconds. This suggest that the ability to upgrade to 60 seconds had a a significant effect on participants' performance. Figure A2 in the appendix highlights the distribution of task performance across rounds. Task performance does not seem to increase over time, which suggests that learning effects were either limited or offset by increasing task fatigue.

One important input for my predicted control-premium measure is participants' maximum WTP to upgrade from 30 to 60 seconds, which I collect in Part 3 using two separate elicitations. On average, participants were willing to pay \$2.66 to upgrade to 60 seconds (with a standard deviation of \$0.93 and a median of \$2.58). I find that participants' reported WTP at the participant level is positively correlated with a simple performance proxy for Part 2 which measures the difference in average earnings between 30- and 60-second tasks.[17] This suggests that participants take their

[17]Figure A3 in the appendix depicts this relationship at the participant level. Note that this simple measure does

Part 2 performance into account in forming their reported WTP to upgrade.

To assess the reliability of this WTP measure, I can compare the reported values of WTP_i across the two elicitation frames which are presented to participants six decisions apart. The Spearman rank correlation between the two measures is 78% (p-value < 0.001), which suggests that participants have stable and well defined preferences over their ability to upgrade. Figure A4 in the appendix highlights the relationship between these two values. Note that the second elicited value is slightly higher at \$2.89 (p-value < 0.001 based on a paired t-test). Interestingly, participants with an average control premium above the median value for the population appear to revise their answer upward more frequently. To account for this fact, I calculate control premium estimates for control-lottery decisions $\mathcal{P}$ and $\mathcal{S}$ using both elicitations as a robustness check.

4.3 The Control Premium

Based on participants' reported EFFs for the control lotteries in Part 3 and their reported WTP to upgrade to 60 seconds, I can calculate implied control premia at the participant level. Figure 10 compares the resulting estimates in absolute and relative terms across three alternative assumptions for the baseline lottery $\mathcal{B}$. The bars on the left represent control premium estimates assuming expected-value maximizers. These estimates take neither risk aversion nor probability weighting into account. The estimate is based on the midpoints of all relevant MPL ranges. The bar in the middle relies on the same assumptions but uses the most conservative bound of each price list (i.e. the lower bound of the chosen range for EFFs and the upper bound of the chosen range for stated WTP). The estimate is comparable in magnitude to the first bar and remains statistically significant. This rules out that my control premium estimates are driven by imprecise elicitation due to the price list granularity. The final column reflects estimates accounting for population-average probability weighting (again using the midpoint of each MPL price range).[18] Estimates are statistically significant in absolute and relative terms across all specifications.

Accounting for the population-average degree of probability weighting, participants exhibit an absolute control premium of \$0.29 corresponding to a relative control premium of 17.2%. A simple t-test to establish whether these means are different from zero has a p-value of less than 0.001. This suggests that participants place a substantial premium on retaining control over the upgrade decision. For comparison, Bartling et al. (2014) find an average relative control premium of 16.7%[19] in a principal-agent setting. Owens et al. (2014) find a relative control premium between 8-15%[20].

not account for how often a participant actually chose to upgrade in Part 2.

[18]Since the estimated risk aversion parameter under the token frame was exactly one, estimates do not change if taking the population-average degree of probability weighting and risk aversion into account.

[19]Bartling et al. (2014) calculate the relative control premium as the percentage difference between the context-free elicitation of certainty equivalents of a *delegation* and *control lottery* which principals reported being indifferent between in a control-driven context

[20]Owens et al. (2014) calculate this control-premium measure as the difference between the expected value of observed choices and predicted optimal choices based on stated beliefs about success probabilities of an expected-value maximizer

In total, 136 participants (76% of the sample) exhibited a positive value of control, on average, across the three main control lotteries $\mathcal{B}, \mathcal{P}$ and $\mathcal{S}$.

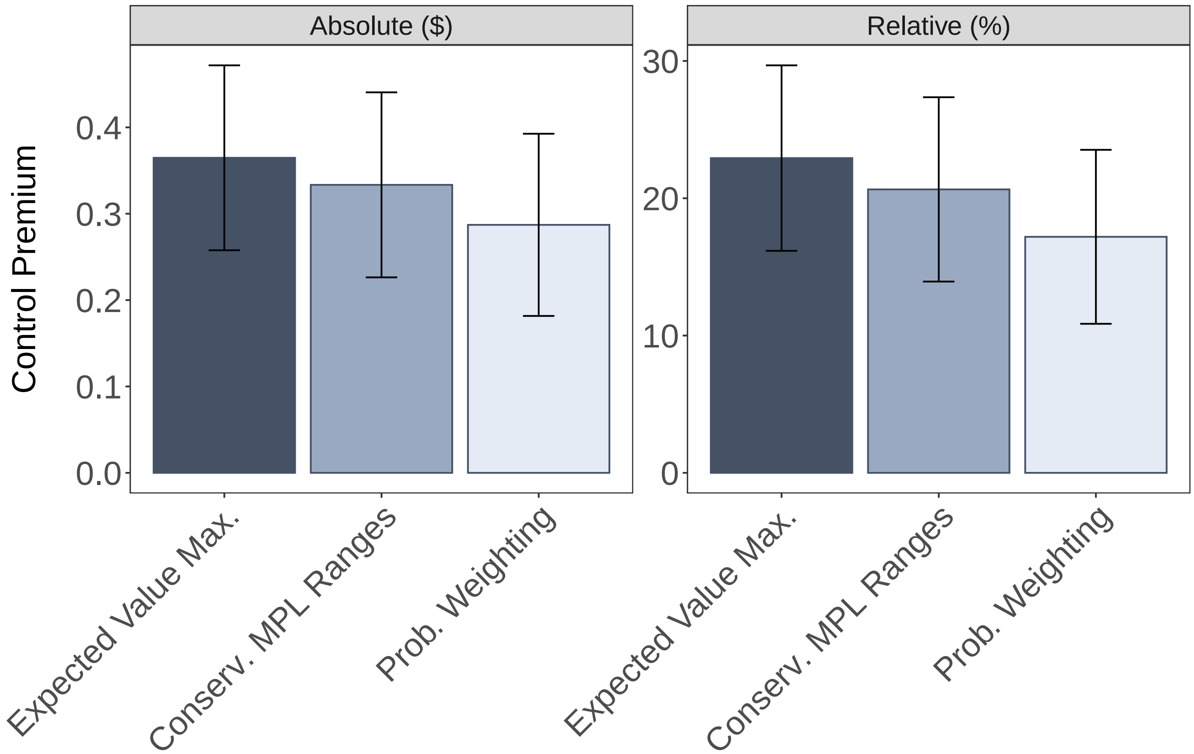

Figure 10: Model comparison under the base control lottery (error bars represent 95% confidence intervals).

Figure 11 highlights the individual heterogeneity in control premium estimates. Each dot in this picture represents one participant's average reported EFF across the three main control lotteries in Part 3 (vertical axis) relative to the average implied control-neutral certainty equivalent (horizontal axis). Points above the 45° line indicate that the reported EFFs exceed the estimated control-neutral level, which implies that the participant places a premium on retaining control. This is true for the majority of participants.[21]

To assess the validity of my interpretation of this measure as a control premium, I can leverage the comparison between control and peak-price lotteries in two ways. First, I can compare the stated EFF for the base control lottery $\mathcal{B}$ and the corresponding peak-price lottery $\mathcal{CPP}$ in Part 3. Recall that the peak-price of lottery $\mathcal{CPP}$ is calibrated based on participants' Part 2 performance,

[21] There is some evidence that points below the 45° are in part driven by noise rather than a genuine preference to cede control: Of the 44 participants who exhibit a negative control premium on average, 14 reported EFFs strictly below their logical lower bound for at least two of the three control lottery decisions considered here. The logical lower bound of participants' EFF is given by the minimum of c_L and their own reported WTP to upgrade. Participants should always prefer being able to upgrade at this amount for sure to playing the control lottery. Figure A5 in the appendix breaks down Figure 11 by how many times a participant reported an EFF below its logical lower bound.

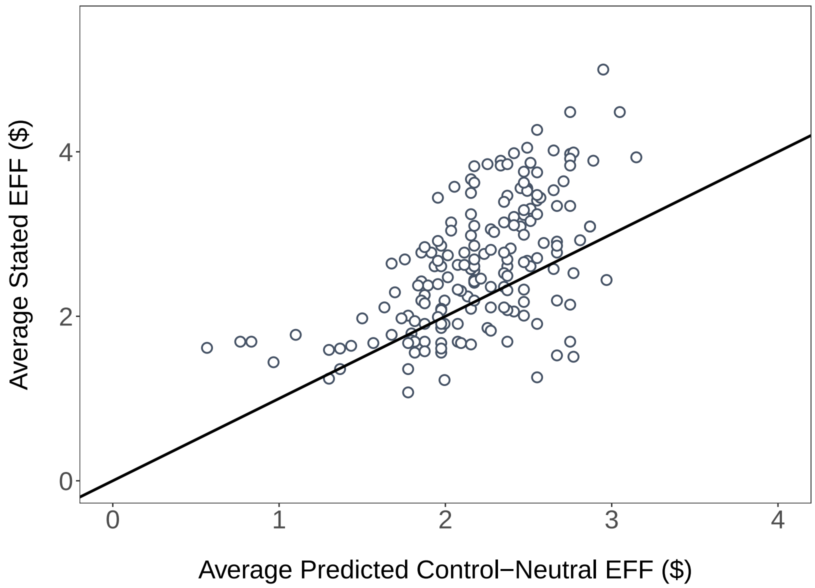

Figure 11: Individual participants' average reported equivalent flat fees (EFFs) across the three main control decisions versus the implied control-neutral EFFs (accounting for population-average probability weighting). Each dot represents one participant. Points above the 45 degree line imply a positive average control premium across decisions.

and does not take their stated WTP to upgrade into account. For some participants, the calibrated peak-price $c_{i,H}$ lies above their WTP to upgrade WTP_i, suggesting that participants would not choose to upgrade if $c_{i,H}$ was realized. For others $c_{i,H}$ lies below WTP_i, which means that the peak-price lottery has instrumental value above and beyond the control lottery.

Figure 12 plots the difference in EFFs as a function of the realized peak-price. The vertical axis reflects the EFF of the control lottery minus the EFF of the peak-price lottery. Recall that a *higher* EFFs suggests an increased desire to avoid playing the lottery. Positive values along this dimension therefore suggest a preference for the peak-price lottery. How strong this preference is may depend in part on the peak price: It is likely that participants would treat an exceedingly punitive peak price (e.g. a peak price of $100) as an effective restriction of their use, which might lead to both contracts being valued equally. The control premium add-on relative to the peak-price lottery may therefore be decreasing in the peak price. I find evidence that this is indeed the case.

The horizontal axis in Figure 12 represents the difference between the peak price and a participant's WTP to upgrade. The dark blue line shows a local-constant kernel fit of the observed EFF difference. The two lighter solid lines highlight the predicted relationship for a control-neutral individual (the upper line accounts for the population-average degree of probability weighting while the lower line does not). Below zero, the difference in EFFs exceeds the predicted control-neutral relationship, i.e. participants are reporting a higher EFF for control lottery $\mathcal{B}$ than peak-price lottery $\mathcal{CPP}$. This suggests that participants are more averse to playing the control lottery. This

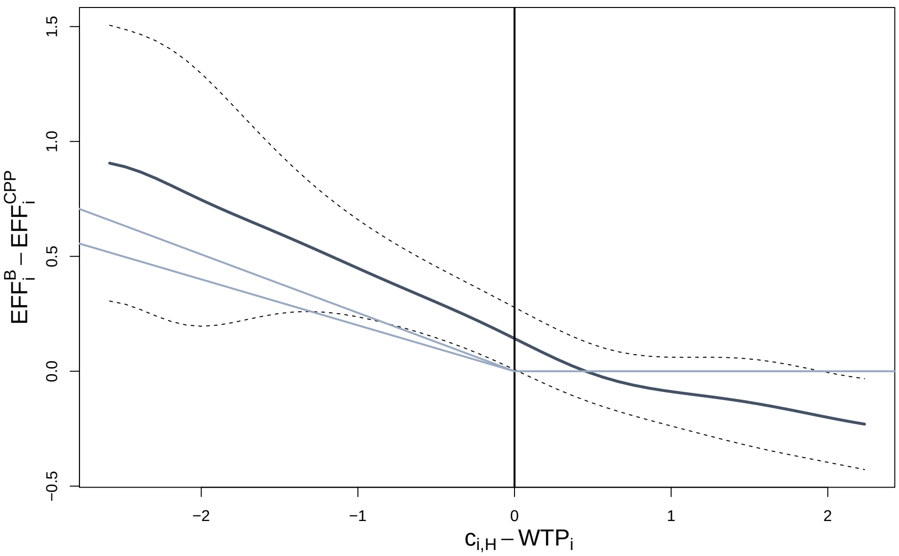

Figure 12: Comparing stated equivalent flat fees (EFFs) for base control lottery $\mathcal{B}$ and the peak-price lottery $\mathcal{CPP}$. Local-constant kernel fit of the difference in EFFs between the base control and peak-price lottery. The kernel bandwidth was selected using cross-validation. Estimates were generated using the *np* package (Hayfield et al., 2008). Dashed lines indicate 95% bootstrapped confidence intervals. Solid light blue lines indicate predicted control-neutral relationship (upper line: accounting for probability weighting; lower line: maximize expected-value).

difference remains positive near the $c_{i,H} = WTP_i$ cutoff at which control-neutral participants should become indifferent between the two lotteries, and at which exactly the same individual level risk preferences should apply. This positive difference therefore lends additional credence to a control explanation.

For values above the cutoff, Figure 12 suggests that this valuation difference goes to zero relatively quickly. In part, this indifference between the two lotteries for values above the cutoff may be explained by the fact that I elicit WTP_i just before asking participants to value lotteries $\mathcal{B}$ and $\mathcal{CPP}$, thereby making this level salient in participants' mind. While this approach ensures that I use the most relevant value of WTP_i in estimating the control premium, it may also mean that participants who face a peak price which clearly exceeds their previously stated WTP may feel a similar level of reactance to this lottery as to a control lottery.[22]

This comparison in EFFs across lotteries $\mathcal{B}$ and $\mathcal{CPP}$ also allows me to propose a lower bound

[22]Note that Owens et al. (2014), who use a quiz task in which participants can choose to get paid either on their own performance or the performance of some randomly paired match, also find a control premium only for those participants who think that their success probabilities are similar to their match's. Those who believe their match to be substantially more likely to correctly answer the question bid on the match at the predicted rate for an expected value maximizer. They conclude that "the control premium mainly affects choices for which the difference in expected returns for the two assets is small, but it plays a major role in these choices. In other words, participants are sensitive to the cost of control (...)" (p. 148).

for my baseline control premium estimate above. In particular, I find that participants provide an EFF for CPP lotteries with $c_{i,H} < WTP_i$ that is 8% higher than its predicted level (see equation 8). Subtracting this difference, which cannot be driven by preferences for control or reactance, from the difference in stated versus predicted EFFs for control lottery $\mathcal{B}$ of 17%, I arrive at a lower bound for my control premium estimate of 9%.

Finally, I can leverage the control vs. peak-price lottery comparison to understand the motivations underlying participants' reported EFFs by looking at participants' responses to an open-ended question in the exit survey. In this question, participants are asked to directly compare the base control lottery and a peak-price lottery with a peak price calibrated to just exceed participants' WTP to upgrade.[23] These two lotteries should lead to identical outcomes: playing the task in 60 seconds with probability $(1-p)$, and in 30 seconds with probability p. However, in case of the peak-price lottery, participants would play the round in 30 seconds by choice, while this outcome would be imposed on them under the control lottery. While this decision was purely hypothetical, the open-ended question asking participants to provide their reasoning for preferring one lottery over the other provides valuable insights into the potential psychological underpinnings of the control premium.

The corresponding responses suggest that the experimental design was successful in triggering an emotional response to the restriction of participants' choices under the control lottery. 70% of participants report preferring the peak-price lottery to the control lottery in this hypothetical scenario. This number is similar to the number of participants exhibiting a positive control premium on average across the three main lottery decisions in Part 3. Participants arguments for why they perceive the peak-price lottery more favorably suggest that this mainly has to do with valuing having the option to upgrade, even at an unattractive cost[24]. In some instances, participants expressed this preference for keeping the upgrade option open in terms of perceived riskiness of the situation[25]. The remaining 30% of participants reported either being indifferent between the two lotteries as predicted by standard expected utility theory; preferring the control lottery as a commitment device to avoid overpaying for upgrading[26]; due to it being simpler to think through or potentially involving one fewer upgrade decisions[27]; or in some instances due to the misconception that upgrading at the realized cost would be mandatory under the peak-price lottery in this hypothetical scenario[28].

Note that the experimental setting differs from real-world DLC contracts in two important

[23] $c_{i,H}$ is set at the maximum of the participant's two elicited WTPs to upgrade plus \$0.02

[24] e.g. "I would like the ability to upgrade even if i have to risk paying too much for it"; "I get the option to upgrade or not, even if I chose not to"; "I prefer knowing that I can have 60 seconds regardless of the outcome of the lottery"; "It gives more options"; "theres no fear of not being able to upgrade"

[25] e.g. "it seems less risky"; "It has a better 'worst case' scenario."; "It seems less risky to have the option to upgrade (even though the price might be higher) than to have no options at all"

[26] "not being aple to upgrade seems safer than paying \$3.62"

[27] e.g. "chances of upgrading for \$1.25 are the same for both, and I would not upgrade for \$2.52 in either case, so (the control lottery) minimizes an excessive choice."; "It gives me a simpler decision later."

[28] e.g. "Because I don't need to upgrade for \$3.02"

ways: First, the electricity context may be subject to perceived conflicts of interest, particularly if consumers have low trust in their utility company. Bartling et al. (2014) show that control premia respond to the degree of conflict of interest inherent in the decision environment. Second, there may be important uncertainty about future WTP to run the air conditioner at the time a control-event is called. This uncertainty in turn creates option value in the DLC decision which this paper abstracts away from in order to avoid confounding perceived option value and intrinsic preferences for control. In particular, all control decisions in this experiment relate to one single, final task which rules out the potential for additional (value) learning before the final upgrade decision.

4.4 Determinants of the Control Premium

In order to understand what drives control premia, I can compare estimated levels across the three main control lotteries in Part 3: the base lottery $\mathcal{B}$, lottery $\mathcal{P}$ which doubles the probability of a control event from 20 to 40%, and lottery $\mathcal{S}$ which doubles the stakes of the decision by doubling both the reward per word found and the lower cost level c_L. Using similar logic as outlined in equation 6, an expected value maximizer should report EFFs of

$$EFF_i^{\mathcal{P}} = (1-2p)c_L + pWTP_i = 2EFF_i^{\mathcal{B}} - c_L \tag{13}$$

and

$$EFF_i^{\mathcal{S}} = (1-p)2c_L + p2WTP_i = 2EFF_i^{\mathcal{B}}. \tag{14}$$

As shown in Figure A6 in the appendix, these implied theoretical relationships map well onto the observed changes in stated EFFs: very few participants double their stated EFF in response to the doubling in probability, with most participants reporting values 1-1.5 times higher than under the base lottery. For lottery $\mathcal{S}$, the majority of participants reports an EFF that is close to two times as high as under the base lottery.

As before, the control premia for lotteries $\mathcal{P}$ and $\mathcal{S}$ are calculated as the difference between these predicted, control-neutral EFFs and participants' stated values. Figure 13 shows the corresponding control premium estimates across the three lotteries accounting for the population-average degree of probability weighting. The left panel shows results in absolute terms, while the right panel provides relative control premium estimates. Under lottery $\mathcal{S}$, some stated EFFs were censored by the upper bound of the price list. The dashed version of the third bar therefore accounts for fitted values from a tobit regression of the stated EFF for lottery $\mathcal{S}$ on the EFFs of the other two lotteries.

I test for significance of the base control premium estimate, as well as for differences in mean between lotteries $\mathcal{B}$ and $\mathcal{P}$, and between lotteries $\mathcal{B}$ and $\mathcal{S}$, in absolute and relative terms. A

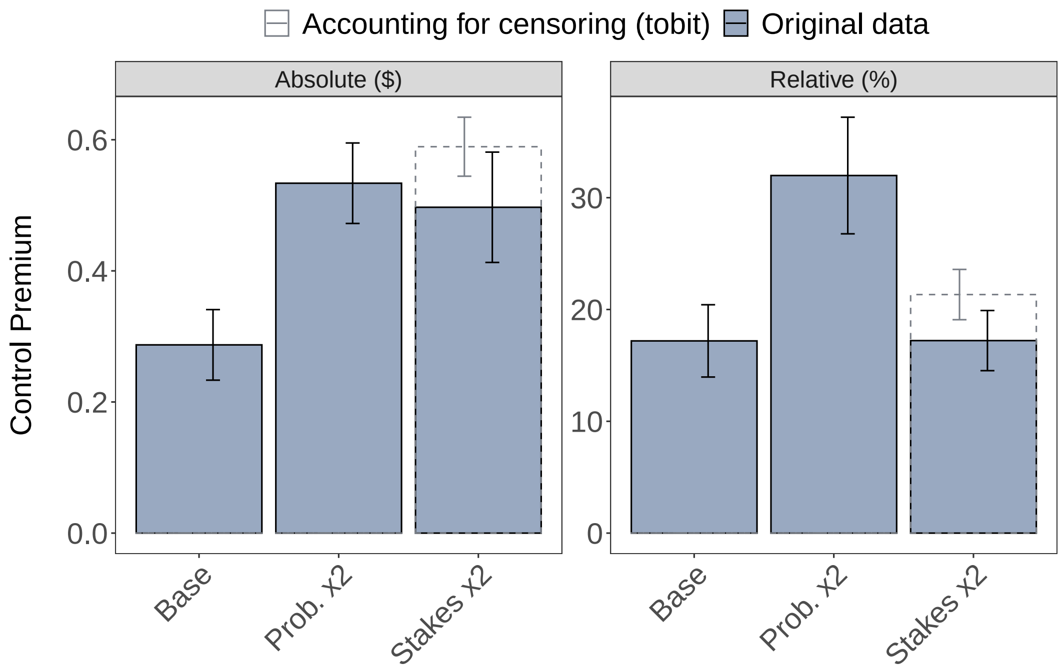

Figure 13: Control premium estimates across decision type allowing for probability weighting. The censoring-corrected estimate (dashed) for the decision with doubled stakes represents the mean of the fitted values from a Tobit regression of the stated lottery $\mathcal{S}$ EFF on the stated EFFs for lotteries $\mathcal{B}$ and $\mathcal{P}$. Error bars represent one standard deviation of the control premium estimate.

conservative approach to accounting for this multiple hypothesis testing is a Bonferroni correction which effectively divides the relative significance thresholds by the number of completed tests (six in my case). For example, significance at the 1% level would therefore be assessed by comparing the p-value to a threshold of 0.0017 instead.

4.4.1 Changes in Absolute Control Premium

Doubling the probability or stakes (accounting for the tobit-based censoring correction) significantly increases the control premium in absolute terms: paired, two-sided t-tests for a difference in means yield a p-value of 0.0007 and $<$0.0001 respectively. These effects remain significant even after Bonferroni correction.

At \$0.53 and \$0.59, the mean control premia under lotteries $\mathcal{P}$ and $\mathcal{S}$ are almost twice as high as the \$0.29 control premium in the base case. This suggest that the absolute control premium scales almost linearly with both probabilities and stakes, rather than consisting of some fixed value. While the finding on the importance of stakes mirrors the finding in Bartling et al. (2014), the importance of control event probabilities had not been previously analyzed in the literature.

4.4.2 Changes in Relative Control Premium

While the direction of changes remains the same, the difference in control premia is less significant in relative terms, especially when the Bonferroni-corrected criteria is used (p-value of 0.0159 for $\mathcal{P}$ vs $\mathcal{B}$, and 0.1885 for $\mathcal{S}$ vs $\mathcal{B}$ after censoring-correction). Interestingly, the effect of stakes is less pronounced in relative terms than the effect of control-event probability: at 21%, the average relative control premium after doubling stakes is only marginally higher than the base premium of 17%. For a doubling in probabilities on the other hand, the average relative control premium doubles to 32%.

To be conservative, I can re-calculate the control premium estimates for lotteries $\mathcal{P}$ and $\mathcal{S}$ using the second elicited value of WTP_i. Recall that I elicit participants' WTP to upgrade from 30 to 60 seconds twice, 6 decisions apart. As noted earlier, the second elicitation tends to be marginally higher for participants who exhibit a control premium based on the base control lottery. Using this second measure, the relative control premia under decisions $\mathcal{P}$ and $\mathcal{S}$ are 22% and 12% respectively, meaning that the same qualitative pattern between the two estimates still holds. The estimates also remain statistically significant (p-value < 0.0001). Neither estimate is statistically significantly different from the estimate for the base control lottery $\mathcal{B}$. Jointly, these finding support the idea that the control premium scales almost linearly with the stakes and probability of losing control in absolute terms, while representing a roughly constant add-on of around 10-30% above instrumental value.

A final caveat regarding my control-premium estimate concerns the use of multiple price lists. One drawback of using iMPLs is that multiple-price-list approaches are sensitive to framing effects: participants may be drawn to the middle of the list irrespective of their true valuations. In my case, this would lead to an overstatement of certainty equivalents for the baseline control-lottery decision in Part 3 for most participants (given their reported WTP to upgrade), which might in turn be incorrectly interpreted as a value of control. In order to mitigate this effect, I (i) provide participants with ample multiple price list training in Part 1 of the survey, with risk-neutral indifference points spread across a wide range of the table; and (ii) keep the bounds of all price lists identical within each part of the survey to avoid participants inferring perceived value signals based on the price list bounds. In addition, the average predicted certainty equivalent under the control decision with high stakes lies above the value in the center of the price list, suggesting that control premium estimates for this decision would be biased downward by this effect. Finally, I can once again leverage the comparison between the baseline control lottery and the preak-price lottery which should be affected equally by the iMPL procedure to show that the difference between stated and predicted certainty equivalents is at least partly due to control.

4.5 Endowment Effect

Does losing control trigger a more or less severe change in utility than a commensurate gain of control? Contingent valuation studies in environmental settings were among the first papers to note differences in stated valuations from the perspective of buying versus selling a good (Hammack and Brown, 1974; Heberlein and Bishop, 1986). Knetsch (1989) and Kahneman et al. (1990) interpreted these results as evidence for loss aversion and established the use of experiments to formally test for the existence of endowment effects in the lab.[29] Here, I will test for the presence of such valuation asymmetries without discussing their potential drivers. The presence of valuation asymmetries is of particular relevance in the DLC context: many utilities allow participants to *override* a given number of control events per year, i.e. to reverse the utilities decision to centrally switch off their appliance. How much participants value such an override option depends in part on whether re-gaining control matters as much (or more) as losing control to begin with[30]

I explore potential valuation differences between gaining and losing control using the final two lottery decision of Part 3: participants are randomly endowed with one of two control lotteries, and are offered to switch to an alternative lottery with a higher or lower control event probability in exchange for some payment.[31] Figure 14 summarizes the resulting average reported valuations. Note that the decisions to switch between lotteries is arguably more cognitively challenging than the standard lottery decisions in Part 3: First, participants have not been trained on this decision type. Second, reasonable EFFs for the standard lottery decision are naturally bounded between c_L and WTP_i, allowing participants to use the price list to guide their answers. No such natural upper bound exists for this lottery comparison. As a result, participants report very large valuations on average, which would suggest control premia which exceed the actual instrumental value of the lottery choice. I therefore focus exclusively on the difference in valuations between the WTA and WTP condition, and between probability steps.

Strikingly, the valuations going from 0 to 20% are almost identical to the valuations going from 20 to 40%. This finding is not driven by participants simply reporting the same value across both decisions: as shown in Figure A7 in the appendix, less than 25% of subjects report exactly the

[29] Note that loss aversion over monetary outcomes need not be related to loss aversion over control, nor to the presence of a control premium itself. While Bartling et al. (2014) only consider instances in which principals cede control to agents, they find no association between loss aversion over monetary outcomes and their control-premium measure. Bobadilla-Suarez et al. (2016), who equally only consider the delegation of control to an external advisor, find that control premia are comparable if choices affect outcomes in the gain versus loss domain.

[30] Note that overrides in practice also carry real option value since participants can choose during which control events to implement this option. Even if control premia are identical in the gain and loss domain, overrides might therefore be valued higher than a corresponding reduction in the number of control events.

[31] Note that unlike in other endowment effect studies, subjects in the WTP condition are not given an additional upfront payment to protect against potential house money effects: First, since no earnings for any part of the experiment have been announced by the point participants reach these questions, the effect of providing any upfront payment now might be perceived as a strong signal of value, thereby influencing participants reported EFFs. Second, the expected control-neutral valuation difference between lotteries is $\Delta p(WTP_i - c_L)$ which is around \$0.28 based on the mean WTP to upgrade. House money effects would therefore be expected to be negligible.

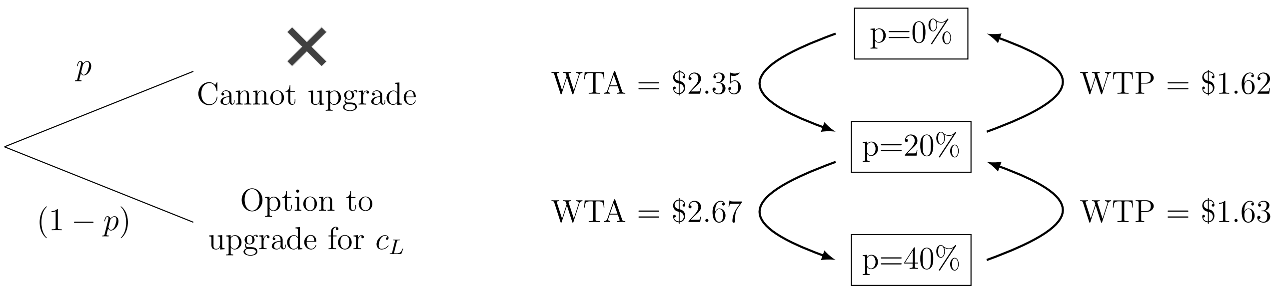

Figure 14: Test for Endowment Effect

same valuation for both lotteries. This result therefore provides further evidence for a near linear relationship of control premia and the underlying control event probability.

With respect to the endowment effect, we find a WTA/WTP ratio for the two lotteries of only 1.45 and 1.63 respectively. This is lower than the mean ratio for monetary lotteries of 2.1 reported in Horowitz and McConnell (2002) (which do not feature a control component). This result suggests that losing or gaining a commensurate degree of control are valued similarly, which suggests that participants should attach significant value to the control gain inherent in overrides.

4.6 Ruling out Risk Aversion

A potential confound in attributing the difference between stated EFFs and predicted control-neutral levels to an intrinsic preference for control is individual heterogeneity in the degree of risk aversion at small stakes. Consider a participant who always finds $\bar{k}_{i,60}$ words in 60 seconds, and $\bar{k}_{i,30} < \bar{k}_{i,60}$ words in 30 seconds. Under EUT, a control-neutral participant should choose the EFF $\widehat{EFF}_i$ which satisfies

$$
\begin{aligned}
u_i(\bar{k}_{i,60} - \widehat{EFF}_i) &= (1-p)u_i(\bar{k}_{i,60} - c_L) + p u_i(\bar{k}_{i,30}) && (15)\\
&= (1-p)u_i(\bar{k}_{i,60} - c_L) + p\left[u_i(\bar{k}_{i,60}) - u_i(WTP_i)\right] && (16)
\end{aligned}
$$

Under risk-neutrality, this expression simplifies to the familiar $\widehat{EFF}_i = (1-p)c_L + pWTP_i$. However, this relation no longer holds under risk aversion. By concavity of $u_i(\cdot)$ and assuming that WTP_i is reasonably capped at $\bar{k}_{i,60}$,

$$u_i(\bar{k}_{i,60}) - u_i(WTP_i) \leq u_i(\bar{k}_{i,60} - WTP_i)^{32}$$

and hence

[32] since

$$u_i(\bar{k}_{i,60} - WTP_i) + u_i(WTP_i) = u_i((1-\alpha)\bar{k}_{i,60}) + u_i(\alpha\bar{k}_{i,60}) \geq (1-\alpha)u_i(\bar{k}_{i,60}) + \alpha u_i(\bar{k}_{i,60}) = u(\bar{k}_{i,60})$$

for $\alpha = \frac{WTP_i}{\bar{k}_{i,60}} \in [0,1]$

$$
\begin{aligned}
u_i(\bar{k}_{i,60} - \widehat{EFF}_i) &\leq (1-p)u_i(\bar{k}_{i,60} - c_L) + p(u_i(\bar{k}_{i,60}) - WTP_i) \\
&\leq u_i((1-p)(\bar{k}_{i,60} - c_L) - p(\bar{k}_{i,60} - WTP_i)) \\
&= u_i(\bar{k}_{i,60} - ((1-p)c_L + pWTP_i)))
\end{aligned}
\tag{17}
$$

which implies $\widehat{EFF}_i \geq (1-p)c_L + pWTP_i$. The control-neutral $\widehat{EFF}_i$ in this case can therefore be higher than under risk-neutrality, a difference which should not be ascribed to a preference for control. To rule out heterogeneity in risk aversion as the main driver of my control premium estimates, I once again turn to participants' responses to the lottery questions in Part 1 of the survey.

Consider participants' stated certainty equivalents for the lottery that pays \$5 with 50% probability and zero otherwise. At this probability level, the impact of probability weighting on the certainty equivalent should be negligible. Lower reported certainty equivalents therefore indicate a larger degree of risk aversion. As a result, if my control premium estimates are confounded by individual-level heterogeneity in risk aversion, the Spearman rank correlation between the estimated control premium and the reported certainty equivalent for this standard lottery should be negative.

However, the estimated rank correlation between the two is only marginally negative at -0.0167 (p-value = 0.8238). This suggests that control premium estimates are not conflated by individual heterogeneity in risk aversion.

4.7 Ruling out Probability Weighting

Another possible confound of the control premium is individual-level heterogeneity in the degree of probability weighting. Recall that the predicted control-neutral EFF is given by

$$
\widehat{EFF}_i = (1-p)c_L + pWTP_i
$$

Probability weighting would lead to a relatively larger weight on the second, higher term, and hence a higher EFF estimate overall, which might bring predicted EFFs closer to stated levels. Once again, this difference should not be incorrectly interpreted as evidence of a control premium.

There are three simple tests which show that probability weighting does not appear to be a dominant driver of the earlier control premium estimates:

Control premium estimates based on lottery $\mathcal{P}$

For lottery $\mathcal{P}$, the probability of control events is 40%. In this probability range, the effect of probability weighting should be relative small. However, we find the largest control premium in relative terms for this lottery. It is therefore unlikely that the control premium I estimate for this lottery is explained by individual heterogeneity in probability weighting.

Correlation between control premium estimates and non-parametric probability-weighting proxy

To capture the individual degree of probability weighting which participants exhibit in the lottery questions in Part 1 (under both the standard and token frame), I propose a simple non-parametric probability weighting proxy: based on the s-shaped relationship between objective and subjective probabilities which arises under probability weighting, the reported certainty equivalent of a zero or \$5 lottery with a 90% objective winning probability should be closer to the corresponding certainty equivalent at a 10% objective winning probability the more a participant weights probabilities.[33] If heterogeneity in probability weighting was driving the control premium estimates, we would therefore expect a strong *negative* correlation between the ratio of certainty equivalents at these probability levels (with a lower ratio suggesting certainty equivalents that lie closer together) and the control premium.

The estimated Spearman rank correlation between this probability weighting proxy and the estimated control premium under the base decisions is -0.0631 (p-value = 0.3998). This correlation is relatively small and not significantly different from zero, which provides further evidence against a dominant role of heterogeneity in probability weighting in driving the observed valuation differences.

Comparing EFFs for the control and peak-price lottery

Finally, I can return to the evidence presented in Figure 12. Recall that this graph shows the difference in stated EFFs for base control lottery $\mathcal{B}$ and the peak-price lottery $\mathcal{CPP}$. The lower light blue line indicates the theoretically predicted relationship for an expected value maximizer using objective control event probabilities. In order to explain the difference between this predicted line and the mean observed EFF difference, participants would have to overweight the 20% probability of a control event occurring as being anywhere between 35-55% on average, which is much larger than the average degree of probability weighting I observe in the lottery decisions in Part 1.

4.8 Ruling out Alternative Explanations

I measure the control premium as the difference between stated and predicted, control-neutral EFFs for control lotteries. Individual-level heterogeneity in monetary preferences does not seem to contribute significantly to this gap as shown in sections 4.6 and 4.7. In this section, I consider additional alternative explanations.

[33]Note that risk aversion would also contribute to compressing this spread.

4.8.1 A Preference for Flexibility

In the presence of uncertainty about future tastes or mood, individuals may exhibit a 'preference for flexibility': striving to preserve larger choice sets for a choice to be made at a later time. Kreps (1979) formally modeled such a preference by considering choices over menu options, followed by a second stage of choices among the options contained in the chosen menu. The canonical example of an individual exhibiting a preference for flexibility is someone who prefers a menu that contains steak but no chicken to a menu that contains chicken but no steak. However, a menu containing both chicken and steak is *strictly* preferred to either of those menus.

In the current framework, control lotteries may lead to a contraction of the future choice set, with participants no longer having the option to upgrade to 60 seconds if a control event is realized. However, by eliciting participants' WTP to upgrade twice, several decisions apart, I am able to show that participants have relatively stable beliefs about the value of upgrading to 60 seconds. In addition, all lottery choices concern a single final task played at the end of Part 3. Participants therefore know that there will be no additional learning about their own skill in the word-search task or their relative performance under different amounts of time before facing the final decision. Participants should therefore not anticipate any changes, and hence uncertainty, in their relevant preferences for upgrading before the start of this final task.

4.8.2 Illusion of Control

In some instances, subjects' perceived control may differ from the objective degree of control available. Illusions of control arise if subjects believe themselves to be more in control of a decision or situation than is practically the case (Langer, 1975), which may lead to overconfidence.[34]

In this experiment, I carefully avoid any misalignment between perceived and actual contingency: participants experience ten tasks in Part 2 which serve to both gauge their task performance and ground them in terms of the execution of the upgrade decision. Participants in this setting should therefore be well calibrated with respect to the amount of control implied under the final upgrade decision at the end of Part 3.

Note that the same logic holds with respect to the idea of locus of control (i.e. the extent to which an individual beliefs they can control events that affect them as opposed to outcomes being driven by forces beyond the individual's control): participants in Part 2 experience that the upgrade decisions is entirely at their own discretion.

Finally, overconfidence regarding task performance is controlled for in this setting since I use participants' own reported WTP to upgrade to calculate predicted control-neutral EFFs. Overconfident beliefs should affect both participants' reported WTP to upgrade and their lottery decisions. The difference between stated and predicted control-neutral EFFs can therefore not be attributed to this effect.

[34]The opposite case may give rise to psychological constructs such as learned helplessness.

4.8.3 Loss Aversion

The lottery and peak-price decisions do not represent decisions over monetary losses. Instead, the realized upgrade costs and time available influence the potential earning during the final task at the end of Part 3. It is therefore unlikely that participants treat the lottery decisions as a monetary loss frame. Bartling et al. (2014) explicitly test for a relationship between loss aversion (over money) and their measure of the control premium and find no association.

4.8.4 Preference Reversals

Preference reversals arise when subjects prefer one lottery over another in a valuation task, but reverse their preference ordering over the same lottery in a direct choice comparison. Settings in which preference reversals arise generally involve lotteries with similar expected value, but very different probabilities and maximum potential earnings. Do participants in my study exhibit a preference for control in terms of their valuations that would not reflect in a choice setting?

In one of the open-ended questions in Part 4, I ask participants to directly choose between a control lottery and a peak-price lottery with a peak-price calibrated to just exceed their WTP to upgrade ($c_{i,H} = WTP_i + \$0.02$). While this choice is purely hypothetical, it allows me to observe how many participants would strictly prefer the peak-price lottery in this case. I find that 70% of subjects report preferring the peak-price lottery, which is in line with the 76% of subjects exhibiting a positive control premium on average across the three main control lottery decisions. This suggests that there is no significant discrepancy between actual choices and choices implied by stated valuations in this setting.

4.8.5 Failure to Elicit Correct Indifference Points

The validity of my control-premium measure relies critically on participants' reported WTP to upgrade. In order to ensure that participants are well calibrated with respect to this value, I provide an extended value learning exercise in the form of the eight upgrade decisions in Part 2 of the experiment. At the start of Part 2, I explicitly ask participants to pay attention to how much upgrading matters to them. Before reporting their WTP to upgrade in Part 3, I ask them to think back to how much the 30 versus 60 second time difference mattered to them in Part 2. In addition, I elicit the indifference point in two ways several decisions apart and find that the values are overall stable with a small tendency to report a higher indifference point during the second elicitation. Finally, in the final decision at the end of Part 3, all participants who encounter a cost of blocking at or below their reported WTP choose to block at this cost. Only four participants encounter a cost that exceeds their WTP. Of those four, only one participant decided to block at this excessive cost.

4.8.6 Certainty Effects

One final consideration concerns the validity of a direct comparison between control and peak-price lotteries, and hence the evidence in favor of interpreting the estimated difference as a control premium. As long as the peak-price, $c_{i,H}$ exceeds a participant's WTP to upgrade, the two lotteries result in equal outcomes, and should therefore be valued identically in the absence of an intrinsic preference for control. However, when comparing stated EFFs in the region for which $c_{i,H} > WTP_i$, there is an alternative explanation that could lead to lower EFFs for control lotteries. Since participants would choose to upgrade irrespective of the realized cost, these peak price lotteries may be subject to a *certainty effect*: subjects may be less averse to playing these lotteries since the effect of the lottery on time available is certain, and may therefore report a lower EFF.

To rule this out, I can leverage the fact that the peak-price, $c_{i,H}$, is assigned to participants as a function of a performance indicator for Part 2 rather than any stated preferences or valuations. In addition, $c_{i,H}$ is revealed in the peak-price lottery which follows the initial elicitation of the WTP to upgrade. As a result, participants are not able to manipulate the difference $c_{i,H} - WTP_i$ which can therefore be treated as a quasi-randomly assigned running variable in a regression discontinuity design (RD) with a cut-off of interest at $c_{i,H} - WTP_i = 0$.

The point estimate of the discontinuity assuming a local linear regression with one common mean-square-error-optimal bandwidth is small (at $-0.05) and not statistically significant (p-value = 0.866), which suggests that participants do not attach a significant certainty effect to this decision on average. Further corroborating evidence in this respect comes from Figure 14: the WTP to go from a 20% chance of losing control to no control events (and hence the certain opportunity to upgrade) is virtually identical to the WTP to go from a 20 to a 40% probability of control events.

5 Implications for Direct Load Control

As the US electricity mix shifts towards renewable sources which are intermittent in nature, it is becoming more challenging to continuously balance the grid. With the timing and magnitude of electricity supply becoming more variable and less predictable, utilities are seeking ways to manage electricity demand instead. Demand-response programs aim to control usage through price signals or a direct curtailment of the consumption of certain appliances (e.g. DLC).[35]

Studies exploring attitudes towards different demand-response programs through surveys and focus groups frequently point to an aversion to DLC programs due to autonomy concerns coupled with consumers' lack of trust in energy companies (Fell et al., 2015, 2014; Mert et al., 2008). Encouragingly, (hypothetical) acceptance increases significantly if consumers are given an option to override changes made by the utility company (Mert et al., 2008). However, as pointed out by

[35] According to EPRI, Detroit Edison was the first utility to implement a load control program in 1968 (Electric Power Research Institute, 1985).

Joskow and Wolfram (2012), mandatory changes to electricity contracts to include DR components may face significant political opposition. Many utilities are therefore exploring ways to increase the appeal of DLC to consumers in order to spur voluntary adoption.[36]

Utilities aim to maximize controllable load in order to increase their responsiveness to potential supply shortfalls. Historically, this was achieved by cycling air conditioners or other appliances for given periods of time using an external switch. For example, 50% cycling implies that a household's air conditioner alternates between being turned on for 30 minutes and off for 30 minutes in the course of one hour. To address variations in customer comfort during control events, the focus is now shifting towards integrating DLC contracts with existing smart thermostats at the household level. This "Bring-Your-Own-Thermostat" approach has the added benefit of potentially lower up-front costs since the utility no longer needs to apply a physical external switch on the appliance in question. XcelEnergy's *Saver's Stat Smart Thermostat* pilot program is currently exploring customer satisfaction under cycling versus fixed temperature offsets as part of a field experiment.

Table 3 provides an overview of key features of the contract design employed by the eight largest existing DLC programs by 2016 enrollment according to E Source (Ryder et al., 2017). Each of these programs had more than 100,000 participants in 2016. This overview highlights substantial differences between programs in terms of the cost and availability of overrides, the incentive structure, as well as the maximum number, duration and intensity of control events. For example, some programs limit the maximum number of control events per year, while others either do not specify any limit or only provide the historic average of control events without committing to a fixed cap.

In terms of incentives, all eight programs rely on either a flat incentive structure in the form of monthly or annual bill credits, or provide a special rate for all or part of participating households' electricity bills. However, other incentive structures also exist among smaller programs not listed here: The ConEd *CoolPoints* program for example offers a per-event incentive which pays $2.50 for every event which consumers participate in. In addition, participants receive $5 at the end of the cooling season for participating in all events.

Program implementations also diverge with respect to the cost and availability of overrides. While the largest program, XcelEnergy's *Saver's Switch*, offers no overrides at all, Rocky Mountain Power instead allows for an unlimited number of overrides, and Southern California Edison offers up to 5 overrides in exchange for a reduced annual bill credit. This suggests that utilities are still exploring the relative costs and benefits of providing overrides in boosting adoption. In offering overrides, utilities effectively trade-off potentially increased adoption rates against reduced controllability of the enrolled load.

My experimental findings suggest four key takeaways with respect to DLC contract design:

[36]For example, the Great River Energy Report "Shaping Our Future: The Future Grid Initiative" notes that "Great River Energy worked with Minnesota Valley Electric Cooperative and LREC on a study related to its cycled air conditioning program. The study revealed an opportunity for members to target their marketing efforts more effectively" (p. 3).

Takeaway 1: *DLC-style contract settings trigger substantial control premia which amplify the incentives required to spur adoption.*

This result holds even without capturing the issues around trust in utility companies and conflict of interest often raised in the related literature on DLC acceptance.

Takeaway 2: *Consumers' control premia respond strongly to the probability of control events.*

Assuming that behavior is similar with respect to the maximum number of events per season, this suggests that consumers do not exhibit a diminishing marginal sensitivity to control events. Such a diminishing marginal sensitivity would have bolstered the case for imposing a larger number of control events on a relatively smaller number of consumers in exchange for high one-time incentives. Under traditional DLC contracts, utilities may still seek to impose a relatively large number of events per customer due to the large upfront installation costs under traditional DLC contracts. However, future bring-your-own-thermostat DLC contracts will no longer be subjcet to this constraint.

Takeaway 3: *Consumers' control premia respond strongly to the stakes of losing control.*

In the DLC setting, *stakes* correspond to the temperature impact of control events on participants' homes. While traditional DLC contracts can only approximate the impact intensity by controlling the degree of air conditioner cycling, future DLC contracts will be able to rely on smart thermostats in combination with fixed thermostat set points. Assuming that households' productivity function is non-linear in temperature (as found by e.g. Seppanen et al. (2006)) which suggests that stakes themselves increase non-linearly as temperature moves away from a household's preferred setting, this provides support for an approach that exposes a larger number of households to minor deviations in temperature.

Takeaway 4: *Participants do not seem to exhibit a significant endowment effect with respect to control.*

This finding has potentially important implications for the role of overrides. If a substantial endowment effect existed, this would suggest that returning partial control to consumers in the form of overrides would have a limited impact on acceptability while significantly reducing the controllability of enrolled loads. Instead, I find suggestive evidence that consumers value losing and gaining control similarly in a probabilistic context. In addition, overrides in practice carry additional real option value since participants can choose during which control event to exercise them. Since my evidence on the control premium is consistent with an inflation of (perceived) option value not due to probability weighting, this supports the conclusion that overrides in practice may be a valuable tool in lowering the required enrollment incentive and reducing the risk of reactance in situations with conflict of interest.

One important factor in practice which this experiment does not speak to is the potential for targeting of DLC contracts to specific consumer groups based on their relative control preferences. Figure 11 highlights significant heterogeneity in the magnitude of participants' control premia. However, a simple regression of relative control premia on sample demographics as well as the elicited Burger Desirability of Control Scale (Burger and Cooper, 1979) only detects significant differences based on gender: female participants' control premia are 14 percentage points higher on average (p-value = 0.0059, see Table 2).[37] In part, this result may be due to the relatively narrow sample selection using only currently-registered Cornell students. More work therefore remains to be done on which observable factors tend to predict stronger intrinsic preferences for control.

Table 2: The Role of Demographics and Personality Traits

	Estimate	Standard Error	p-value
Constant	-0.0512	0.1893	0.7873
Female	0.1379	0.0494	0.0059
Age Group 18-22	-0.0630	0.0549	0.2529
DOC Scale	0.0021	0.0018	0.2375

OLS regression of the average relative control premium across decisions $\mathcal{B}$, $\mathcal{P}$ and $\mathcal{S}$ on sample demographics and the Burger Desirability of Control scale.

6 Discussion

This article provides the first experimental evidence for the existence of control premia in a contract setting which mimics interruptible service or non-price rationing contracts such as DLC. I use within-subject variation in the stakes and probability of losing control to show that participants' control premia scale almost linearly in both of these aspects of the choice environment, while representing an approximately constant add-on in relative terms. This finding suggests that the marginal effect of additional control events does not diminish as the number of control events increases, which has important implications for DLC contract design.

More broadly, this article contributes to the literature on the existence and properties of intrinsic preferences for control. I replicate Bartling et al.'s (2014) finding regarding the importance of

[37] This finding is in line with Owens et al. (2014) who also found that the Desirability of Control Scale did not correlate with their individual-specific control premium estimates.

stakes, and extend the current state of knowledge by considering the importance of control event probabilities as well as by testing for the existence of endowment effects with respect to control. My findings therefore provide additional evidence regarding relevant stylized facts about control premia which can inform formal models of preferences for control.

I also collect evidence regarding participants' motivations in wanting to retain control. Experimental evidence in combination with participants' open-ended responses at the end of the experiment suggest that the control premium may be in part explained by an inflation of (perceived) option value that is not driven by probability weighting. This interpretation is in line with the conceptual framework of control outlined by Skinner (1996) who views autonomy as an antecedent to control: She defines control as "the extent to which an agent can intentionally produce desired outcomes and prevent undesired ones" (p. 554). She argues that for an individual to have control, the individual must have access to some pathway (means) which can produce a desired outcome (end). Control therefore requires the joint existence of effective agent-means and means-end relations. Neither autonomy, i.e. the freedom to choose one's own actions, nor the presence of choice itself satisfy this definition in her view since neither implies the existence of a meaningful agent-means-end relation. For example, a student's autonomous effort choice does not imply full control over his or her grade if the grade is also determined by innate ability (implying a weakened agent-means relation). Similarly, adding a decoy good to an existing choice set does not increase objective control over monetary outcomes since it does not provide a meaningful new means-end relation.

In spite of this theoretical delineation, the definitions of 'choice', 'autonomy' and 'control' exhibit significant overlap in terms of their practical applications. As a result, autonomy and control are often used interchangeably. In other cases, autonomy is treated as one of three collectively exhaustive motivations for control, the other two being power, i.e. the ability to affect the outcomes of others, and non-interference, i.e. the ability to avoid having one's own outcomes be determined by another individual (see e.g. Neri and Rommeswinkel (2015)). Rather than adding to the literature that attempts to experimentally disentangle these motivations, I isolate the effects of autonomy in a setting without delegation or conflict of interest. Specifically, in my setting participants' outcomes are either dependent on their own choices or the result of a random draw but are never determined or affected by the actions of a third party. This shows that control premia not only arise in settings involving a third party, but potentially in any setting with decision rights that carry inherent option value.

This experiment was carefully designed to disentangle control premia from option value or anticipated learning effects. In particular, all control decisions in this experiment relate to one single, final task which rules out the potential for additional (value) learning before the final upgrade decision. I therefore cannot directly test for the role of option value in this setting. However, I plan to explicitly explore this relationship in future work by exogenously varying the degree of real option value inherent in participants' decisions. Future work will also include a more direct test for loss aversion with respect to control.

Finally, I find that some participants report wanting to retain the option to upgrade even at a cost at which they explicitly acknowledge they would not choose to exercise this option. This suggests that participants are partly sophisticated about the effect which intrinsic preferences for control have on their decisions. Testing for the degree of naivetë versus sophistication, and studying whether some participants would be willing to pay for commitment devices to more closely align their decisions with the control-neutral optimum, represents another interesting avenue for future research.

Table 3: Overview of Eight Largest DLC Programs by 2016 Enrollment According to E Source (Ryder et al., 2017)

Program Name	State	Event Type	Max. Number of Events	Relevant Months & Max. Duration	Override Availability	Incentive Structure
XcelEnergy *Saver's Switch*	CO, MN, NM, WI, ND, SD, TX	50% cycling	Unspecified (10 days avg.)	Unspecified 4 hours	None	$40 bill credit
Rocky Mountain Power *Cool Keeper*	UT	50% cycling	100 hours per year (2 hours avg.)	May - Sep 4 hours	2 per year	$30 bill credit
PG&E *Smart AC*	CA	50% cycling	Unspecified	May - Oct 6 hours	Unlimited (except emergency cycling events)	$50 enrollment incentive
BGE *PeakRewards*	MD	Choice between 50%, 75% or 100% cycling	Unspecified	Jun - Sep 7 hours	2 per year (except emergency cycling events)	$50 bill credit for 50% cycling, $75 for 75%, $100 for 100%
FPL *On Call*	FL	Choice between 50% or 100% cycling	Unspecified	Apr - Oct 6 hours	None	$21 bill credit for 50% cycling, $63 for 100%
Southern California Edison *Summer Discount Plan*	CA	Choice between 50% or 100% cycling	Unspecified	Jun - Oct 6 hours	Override plan: 5 per year Standard plan: none	$80 bill credit for 50% cycling, $160 for 100% (half the credit with overrides)
DTE Energy *CoolCurrents*	MI	50% cycling	Unspecified	Jun - Oct 8 hours	None	Special rate (for AC only)
Great River Energy *Cycled Air Conditioning*	MI	50% cycling	200 hours per year	May - Sep 6 hours	None	Special rate (total use)

References

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006): "Elicitation using multiple price list formats," *Experimental Economics*, 9, 383–405.

Bartling, B., E. Fehr, and H. Herz (2014): "The Intrinsic Value of Decision Rights," *Econometrica*, 82, 2005–2039.

Becker, G. M., M. H. DeGroot, and J. Marschak (1964): "Measuring utility by a single-response sequential method," *Behavioral science*, 9, 226–232.

Bernheim, B. D. and C. Sprenger (2019): "Direct Tests of Cumulative Prospect Theory," *Available at SSRN 3350196*.

Bobadilla-Suarez, S., C. R. Sunstein, and T. Sharot (2016): "Are Choosers Losers? The Propensity to Under-Delegate in the Face of Potential Gains and Losses," *The Propensity to Under-Delegate in the Face of Potential Gains and Losses (February 15, 2016)*.

Burger, J. M. and H. M. Cooper (1979): "The desirability of control," *Motivation and emotion*, 3, 381–393.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTreeAn open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Dietvorst, B. J., J. P. Simmons, and C. Massey (2016): "Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them," *Management Science*, 64, 1155–1170.

Electric Power Research Institute (1985): "The Demand-side Management Information Directory: Interim Report," 3426.

Falk, A. and M. Kosfeld (2006): "The hidden costs of control," *American Economic Review*, 96, 1611–1630.

Fell, M. J., D. Shipworth, G. M. Huebner, and C. A. Elwell (2014): "Exploring perceived control in domestic electricity demand-side response," *Technology Analysis & Strategic Management*, 26, 1118–1130.

——— (2015): "Knowing me, knowing you: the role of trust, locus of control and privacy concern in acceptance of domestic electricity demand-side response," ECEEE.

Fleming, S. M. and D. Bang (2018): "Shouldering responsibility," *Science*, 361, 449–450.

Hammack, J. and G. M. Brown (1974): *Waterfowl and wetlands: toward bio-economic analysis*, [Washington]: Resources for the Future; distributed by the Johns Hopkins University Press, Baltimore, http://newcatalog.library.cornell.edu/catalog/461924.

Hayfield, T., J. S. Racine, et al. (2008): "Nonparametric econometrics: The np package," *Journal of statistical software*, 27, 1–32.

Heberlein, T. A. and R. C. Bishop (1986): "Assessing the validity of contingent valuation: Three field experiments," *Science of the Total Environment*, 56, 99–107.

Horowitz, J. K. and K. E. McConnell (2002): "A review of WTA/WTP studies," *Journal of environmental economics and Management*, 44, 426–447.

Joskow, P. L. and C. D. Wolfram (2012): "Dynamic Pricing of Electricity," *The American Economic Review*, 102, 381–385.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1990): "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 98, 1325–1348.

Knetsch, J. L. (1989): "The Endowment Effect and Evidence of Nonreversible Indifference Curves," *The American Economic Review*, 79, 1277–1284.

Korting, C. (2019): "Establishing Properties of the Control Premium," *AEA RCT Registry.*

Kreps, D. M. (1979): "A representation theorem for" preference for flexibility"," *Econometrica: Journal of the Econometric Society*, 565–577.

Langer, E. J. (1975): "The illusion of control." *Journal of personality and social psychology*, 32, 311.

Mert, W., J. Suschek-Berger, and W. Tritthart (2008): "Consumer acceptance of smart appliances (report SMART-A D5.5)," *Smart domestic appliances in sustainable energy systems (Smart-A).*

Neri, C. and H. Rommeswinkel (2015): "Freedom, Power and Interference," *Available at SSRN 2485107.*

Owens, D., Z. Grossman, and R. Fackler (2014): "The Control Premium: A Preference for Payoff Autonomy," *American Economic Journal. Microeconomics; Pittsburgh*, 6, 138–161.

Prelec, D. et al. (1998): "The probability weighting function," *ECONOMETRICA*, 66, 497–528.

Quick, B. L. and J. R. Considine (2008): "Examining the Use of Forceful Language When Designing Exercise Persuasive Messages for Adults: A Test of Conceptualizing Reactance Arousal as a Two-Step Process," *Health Communication*, 23, 483–491, pMID: 18850395.

Rabin, M. (2000): "Risk-aversion for small stakes: A calibration theorem," *Econometrica*, 68, 1281–1292.

Ryder, K., R. R. Buckley, and C. Framularo (2017): "DLC Program Benchmarking: Results from the 2016 E Source Direct Load Control Program Study," Tech. rep., E Source.

Seppanen, O., W. J. Fisk, and Q. Lei (2006): "Room temperature and productivity in office work," .

Skinner, E. A. (1996): "A guide to constructs of control." *Journal of personality and social psychology*, 71, 549.

Starr, C. (1969): "Social benefit versus technological risk," *Science*, 1232–1238.

Tversky, A. and D. Kahneman (1992): "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.

A Exclusion Criteria

Table A1: Overview of Participants Dropped by Exclusion Criterion

Exclusion Criterion	# Affected
Criteria based on Part 1	
> 2 extreme valuations	29
≥ 2 inconsistent answers	71
Cumulative	90
Criteria based on Part 3	
Attention check	73
Time spent	32
Zero WTP to upgrade	5
Cumulative	89
Total Cumulative	137

68% of included participants were female, and 78% were between 18-22 years old. There is no significant difference in gender between included and excluded participants (p-value = 0.2344), but excluded participants were significantly more likely to be 23 years old or older based on a Pearson χ^2-test (p-value = 0.048).

Figure A1 shows the raw lottery valuations for the participants included in the sample, split out by the token vs. lottery frame.

B Derivation of Theoretical Certainty Equivalents

B.1 Expected Utility Theory

Case 1: Control lottery with $c_L < WTP_i$

Participants choose the equivalent flat fee EFF_i they would be willing to accept instead of playing the control lottery. Under control-neutrality, EFF_i therefore satisfies

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u(k-EFF_i) = p\sum_{k=0}^{10} \mathbb{P}_i(k|s=30)u_i(k) + (1-p)\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-c_L) \quad (18)$$

Rearranging yields

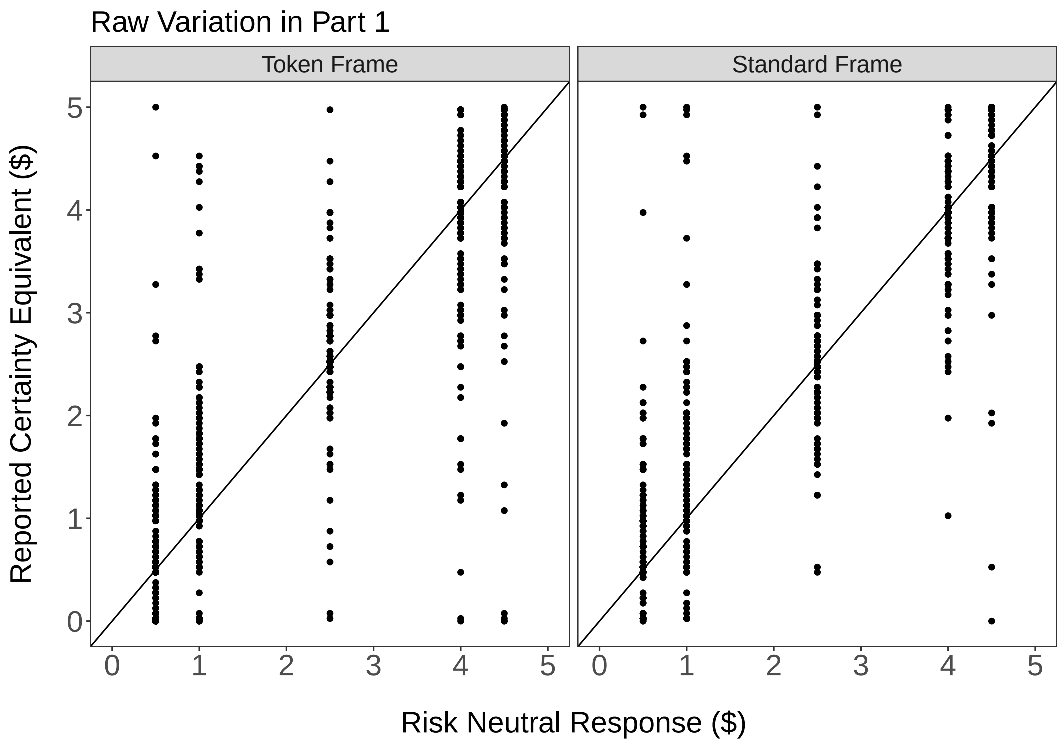

Figure A1: Raw variation in lottery valuations by lottery type.

$$
\begin{aligned}
&(1-p)\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k-c_L)\right] \\
+&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u(k-EFF_i)-\sum_{k=0}^{10}\mathbb{P}_i(k|30s)u_i(k)\right]=0
\end{aligned} \tag{19}
$$

By adding and subtracting and leveraging the definition of WTP_i, we can transform the term on the right as follows:

$$
\begin{aligned}
&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u_i(k-EFF_i)-\sum_{k=0}^{10}\mathbb{P}_i(k|s=30)u_i(k)\right] \\
=&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k)\right]\right. \\
&\left.+\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u_i(k)-\sum_{k=0}^{10}\mathbb{P}_i(k|s=30)u_i(k)\right] \\
=&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k)\right]+u_i(WTP_i)\right]
\end{aligned} \tag{20}
$$

Equation 19 can therefore be re-written as

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)\left[(1-p)(u_i(k-EFF_i)-u_i(k-c_L))+p(u_i(k-EFF_i)-u_i(k)+u_i(WTP_i))\right]=0$$

Assuming that participants are approximately risk neutral at small stakes, i.e. $u(x)=x$, this reduces to

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)\left[(1-p)(c_L-EFF_i)+p(WTP_i-EFF_i)\right]=0 \tag{21}$$

$$\Rightarrow (1-p)c_L+pWTP_i-EFF_i=0 \tag{22}$$

which yields

$$EFF_i=(1-p)c_L+pWTP_i \tag{23}$$

Case 2: Peak-price lottery with $c_{i,H}>WTP_i$

When $c_{i,H}>WTP_i$, participants choose not to upgrade whenever the high cost is realized. In this case, the EFF of this lottery, EFF_i, once again solves

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u(k-EFF_i)=p\sum_{k=0}^{10} \mathbb{P}_i(k|s=30)u_i(k)+(1-p)\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-c_L) \tag{24}$$

which again yields

$$EFF_i=(1-p)c_L+pWTP_i \tag{25}$$

under risk neutrality. The EFF of the control lottery and a peak-price lottery for which the peak price exceeds participant i's maximum WTP to upgrade are therefore identical for control-neutral individuals.

Case 3: Peak-price lottery with $c_{i,H}\leq WTP_i$

Whenever $c_{i,H}\leq WTP_i$, the peak price lottery has additional value to a participant since they would be able to upgrade at favorable costs under either outcome. In this case, EFF_i solves

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-EFF_i) = \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)\left[(1-p)u_i(k-c_L)+pu_i(k-c_{i,H})\right]. \tag{26}$$

Similar math as for the control lottery yields that in this case,

$$EFF_i = (1-p)c_L + pc_{i,H} \tag{27}$$

This EFF is naturally lower than the EFF of the corresponding control lottery since participants are more inclined to accept the lottery.

Case 4: Control and peak-price lottery when $c_L > WTP_i$

If the low cost, c_L, exceeds a participant's maximum WTP to upgrade, WTP_i, the lottery always results in playing the task in 30 seconds. The maximum EFF a participant should be willing to pay to avoid this lottery should therefore be their maximum WTP to upgrade, i.e. $EFF_i = WTP_i$.

B.2 Cumulative Prospect Theory

Under CPT, the EFF of a control lottery with $c_L < WTP_i$ solves

$$\begin{aligned}\sum_{k=0}^{10} V_i(60)u_i(k-EFF_i) &= \pi_i(1)\sum_{k=0}^{10} V_i(30)u_i(k) + \pi_i(1-p)\left[\sum_{k=0}^{10} V_i(60)u_i(k-c_L) - \sum_{k=0}^{10} V_i(30)u_i(k)\right] \\ &= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)u_i(k-c_L) + (1-\pi_i(1-p))\sum_{k=0}^{10} V_i(30)u_i(k).\end{aligned} \tag{28}$$

Once again adding and subtracting a term to make WTP_i appear, we can re-write this as

$$
\begin{aligned}
0 &= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)\left[u_i(k-EFF_i)-u_i(k-c_L)\right]] \\
&+ (1-\pi_i(1-p))\left[\sum_{k=0}^{10} V_i(60)(u_i(k-EFF_i)-u_i(k)) + \sum_{k=0}^{10} V_i(60)u_i(k) - \sum_{k=0}^{10} V_i(30)u_i(k)\right] \\
&= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)\left[u_i(k-EFF_i)-u_i(k-c_L)\right]] \\
&+ (1-\pi_i(1-p))\sum_{k=0}^{10} V_i(60)\left[(u_i(k-EFF_i)-u_i(k))+u_i(WTP_i))\right].
\end{aligned}
\tag{29}
$$

Assuming $u(x)=x$, this simplifies to

$$
0=\sum_{k=0}^{10} V_i(60)\left[\pi_i(1-p)(c_L-EFF_i)+(1-\pi_i(1-p))(WTP_i-EFF_i)\right] \tag{30}
$$

which yields

$$
EFF_i=\pi_i(1-p)c_L+(1-\pi_i(1-p))c_{i,H} \tag{31}
$$

since $\sum_{k=0}^{10} V_i(60)=1$ by construction.

The other three cases follow as shown in appendix B.1

C Supplementary Results

C.1 Task Performance Measures for Part 2

Figure A2 highlights the distribution of the number of words found per task split out by tasks of 30 versus 60 seconds in Part 2. Recall that every participant had 60 and 30 seconds available for rounds 1 and 2 respectively in order to familiarize themselves with the impact of time available on their performance. The left panel (performance distribution in 30 seconds) therefore has no information for round 1, while the second panel has no information for round 2.

This figure shows that there is no significant improvement in task performance over time. This suggests that learning was either limited, or that any learning effects were mitigated by task fatigue over the course of the ten rounds.

Figure A3 shows the relationship between the reported WTP to upgrade, WTP_i, and a simple performance measure based on the task performance in Part 2: the estimated benefit of upgrading is approximated as the difference in the average number of words found in 60- versus 30-second tasks.

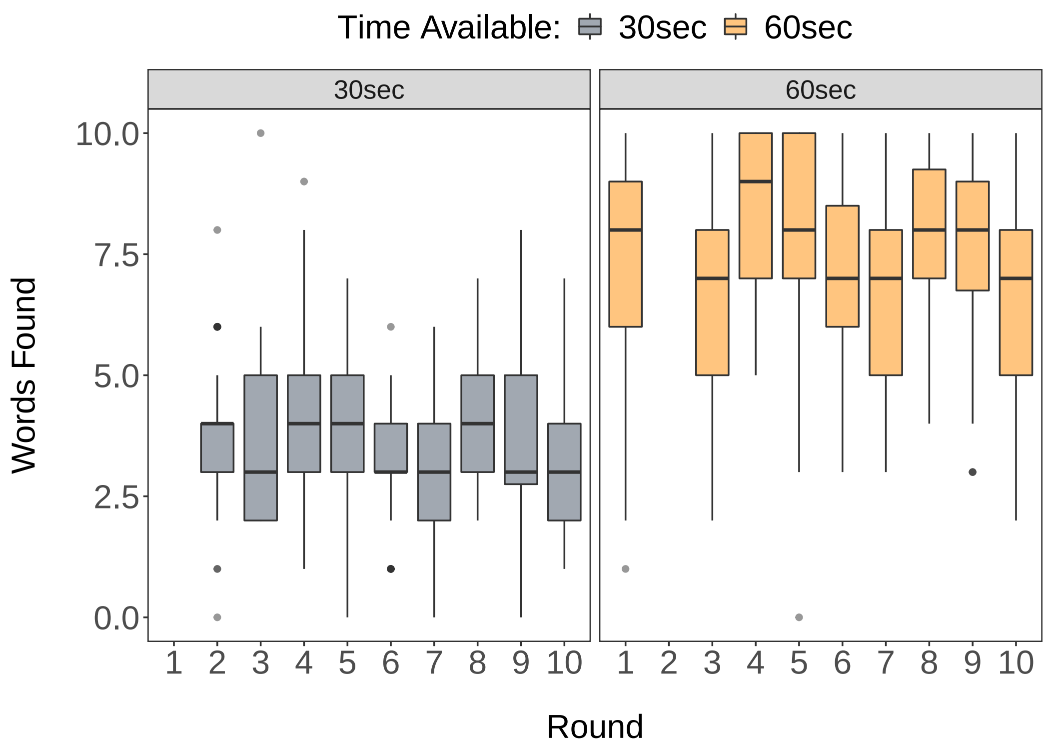

Figure A2: Box plot of the number of words found per task for tasks of 30 versus 60 seconds. The box reflects the first, second and third quartiles (i.e. the 25th percentile, median, and 75th percentile). The whiskers extend to the largest value no further than 1.5 times the inter-quartile range. Data beyond the end of the whiskers are plotted individually.

I find a positive and statistically significant relationship between these two measures, suggesting that participants take their Part 2 performance into account in forming their reported WTP to upgrade in Part 3.

Figure A4 highlights the relationship between the two elicitations of participants' WTP to upgrade. Most of the points fall close to the 45° line despite the two responses being six decisions apart. This suggests that participants have stable and well defined preferences over their ability to upgrade. Participants with an average control premium above the median value for the population appear to revise their answer upward more frequently. A higher indifference point leads to higher predicted EFFs, and thereby lower control premium estimates. To account for this fact, I calculate control premium estimates for control-lottery decisions $\mathcal{P}$ and $\mathcal{S}$ using both elicitations as a robustness check.

C.2 Stated EFFs

Figure A5 visually splits out Figure 11 by the number of times the participant reported an EFF below its logical lower bound across the main control lottery decisions $\mathcal{B}$, $\mathcal{P}$ and $\mathcal{S}$. The figure

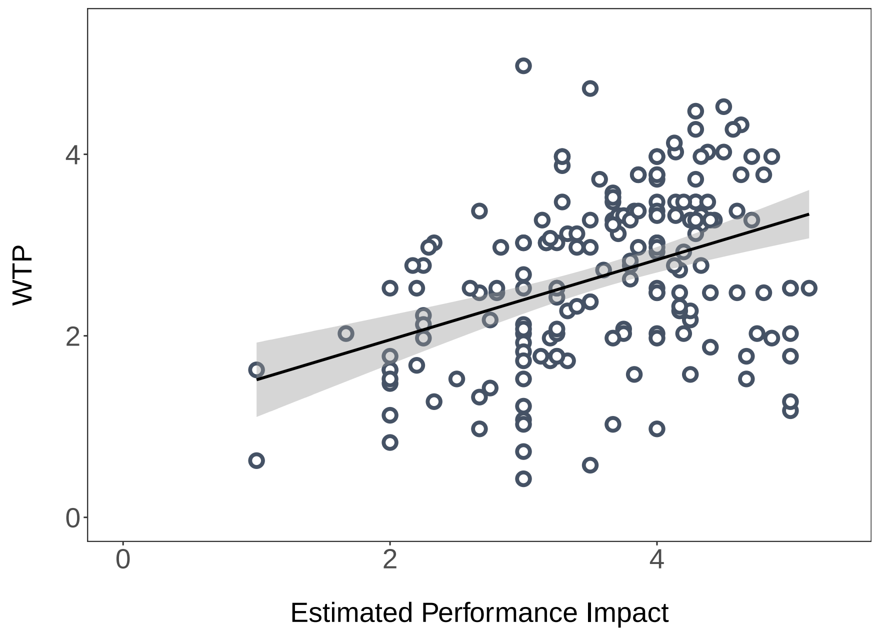

Figure A3: Stated WTP to upgrade versus simple performance measure (average number of words found in 60 versus 30 seconds in Part 2). The black linear represents the linear regression of WTP on the estimated performance impact which has a positieve and statistically significant slope coefficient (p-value < 0.001).

highlights that most of the participants whose valuations fall far below the 45 degree line frequently reported such valuations.

Figure A6 highlights the ratio of reported EFFs after increases in control-event probability and stakes relative to the stated EFF under the base control lottery. In line with the theoretically predicted levels, more participants double their stated EFFs in response to a doubling of stakes.

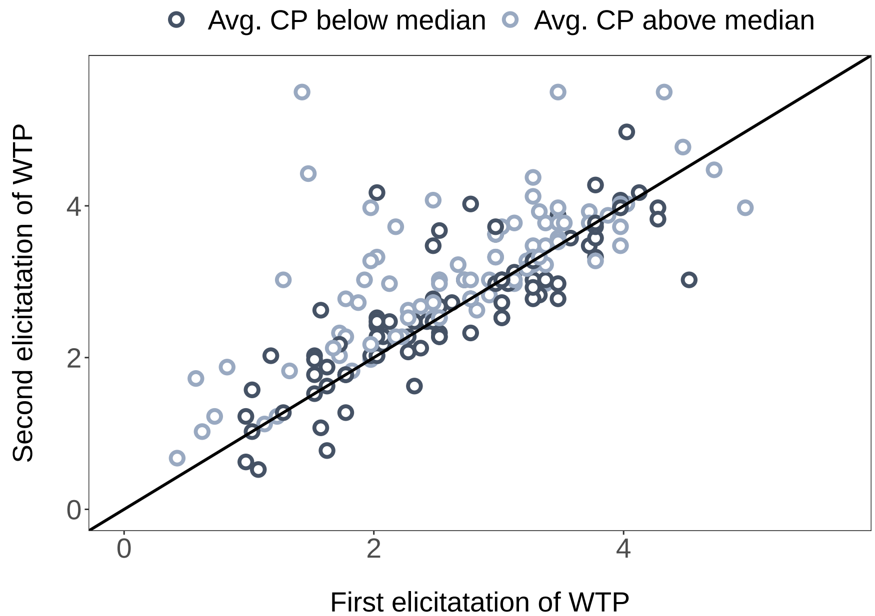

Figure A4: Difference in first and second reported WTP to upgrade to 60 seconds. The color of each circle indicates whether the average control premium of the individual across the three main control lottery decisions in Part 3 was above or belove the population median.

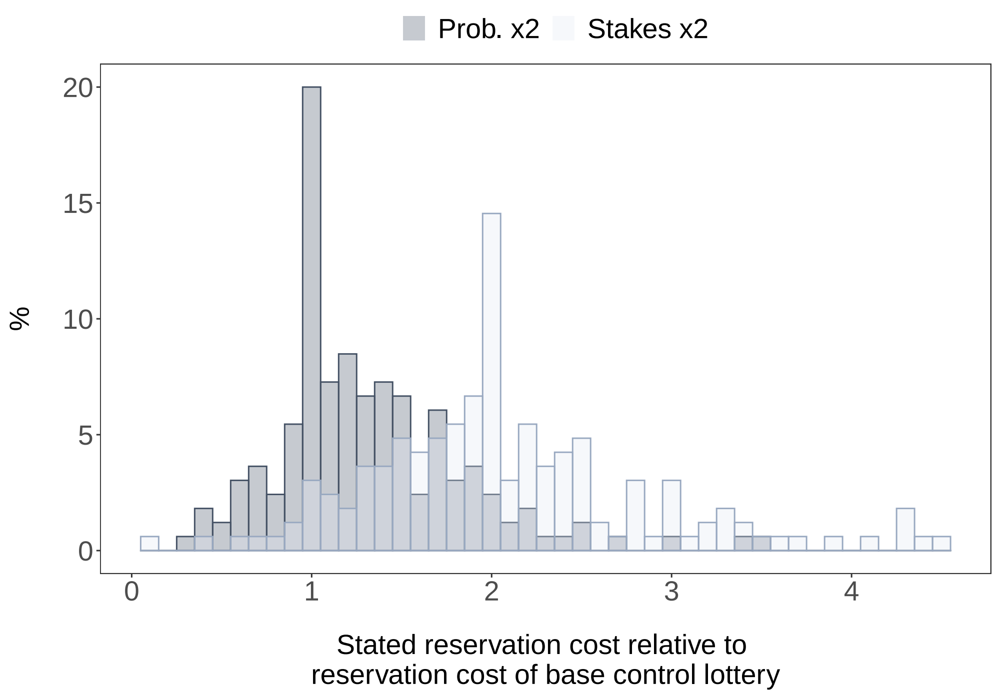

Figure A6: Histogram of the ratio of reported equivalent flat fees (EFFs) for control-lotteries $\mathcal{P}$ (doubled probability of control events) and $\mathcal{S}$ (doubled stakes) relative to the base control lottery $\mathcal{B}$.

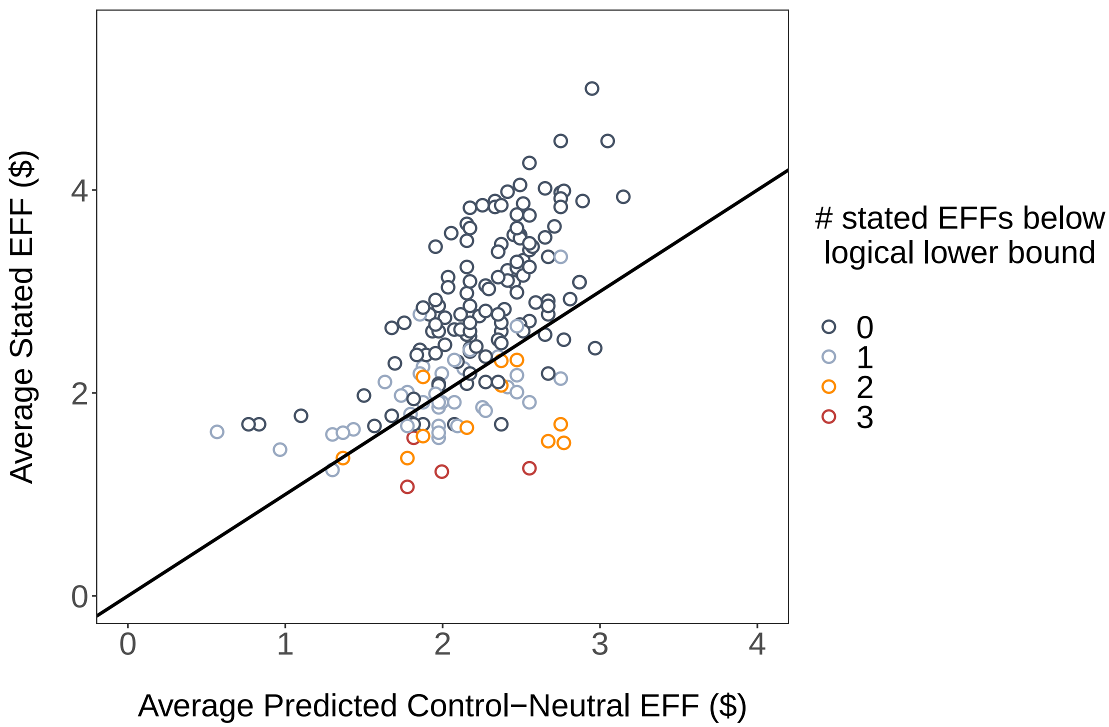

Figure A5: Individual participants' average stated equivalent flat fees (EFFs) across the three main control decisions versus the implied control-neutral EFFs (accounting for population-average probability weighting). Each dot represents one participant. Points above the 45 degree line imply a positive average control premium across decisions. Colors indicate how many times a participant reported an EFF below its logical lower bound.

C.3 Endowment Effect Questions

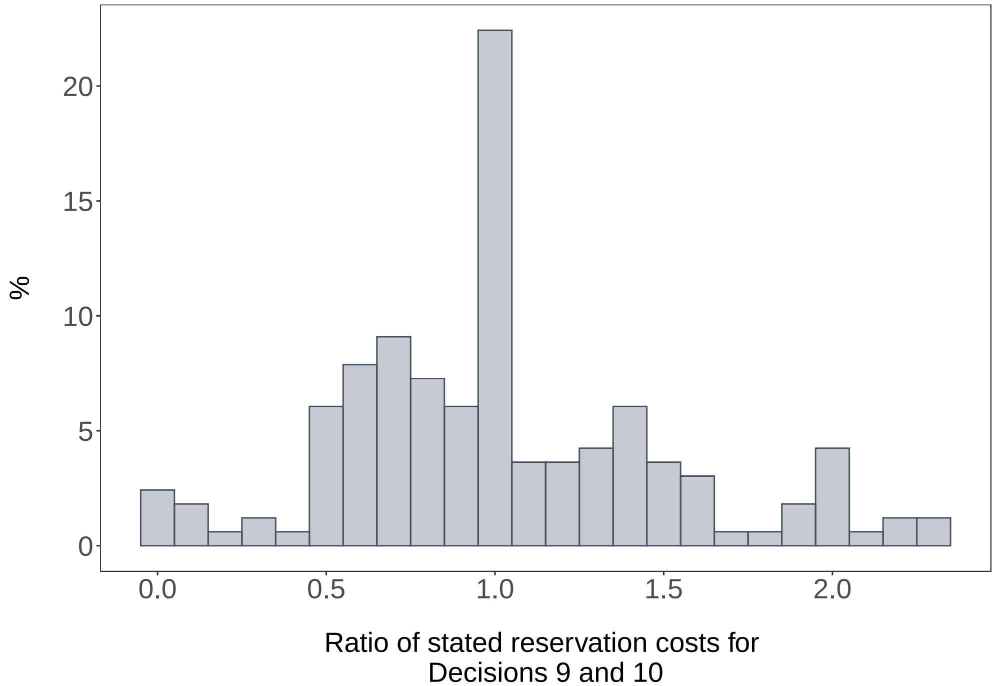

Figure A7: Histogram of the ratio of reported equivalent flat fees (EFFs) for the control lotteries in decisions 9 and 10 of Part 3.

D Burger Desirability of Control Scale

Participants performed similarly to usual student samples on the Burger Desirability of Control (DOC) scale, with a mean of 99.3 (standard deviation of 12.8) which is close to the usual average of around 100. In addition, Cronbach's alpha for the 20 sub-items of this scale was 0.81 in my sample, suggesting high internal consistency of the questions. However, students' score on this scale was not significantly correlated with their average control premium across the three main control decisions (Spearman rank correlation of 0.08, p-value = 0.309). This finding is in line with results by Owens et al. (2014) who equally find no relationship between the DOC scale and their individual-level control-premium estimates. One explanation could be that the questions on this scale are usually focused on conflicts involving third parties, rather than preferences for choice or autonomy. However, I also find no correlation between my control-premium estimates and the subset of questions (questions 1, 8 and 9) which only relate to individual autonomy or choice.

E Survey Screens

E.1 Part 1 Lottery Questions

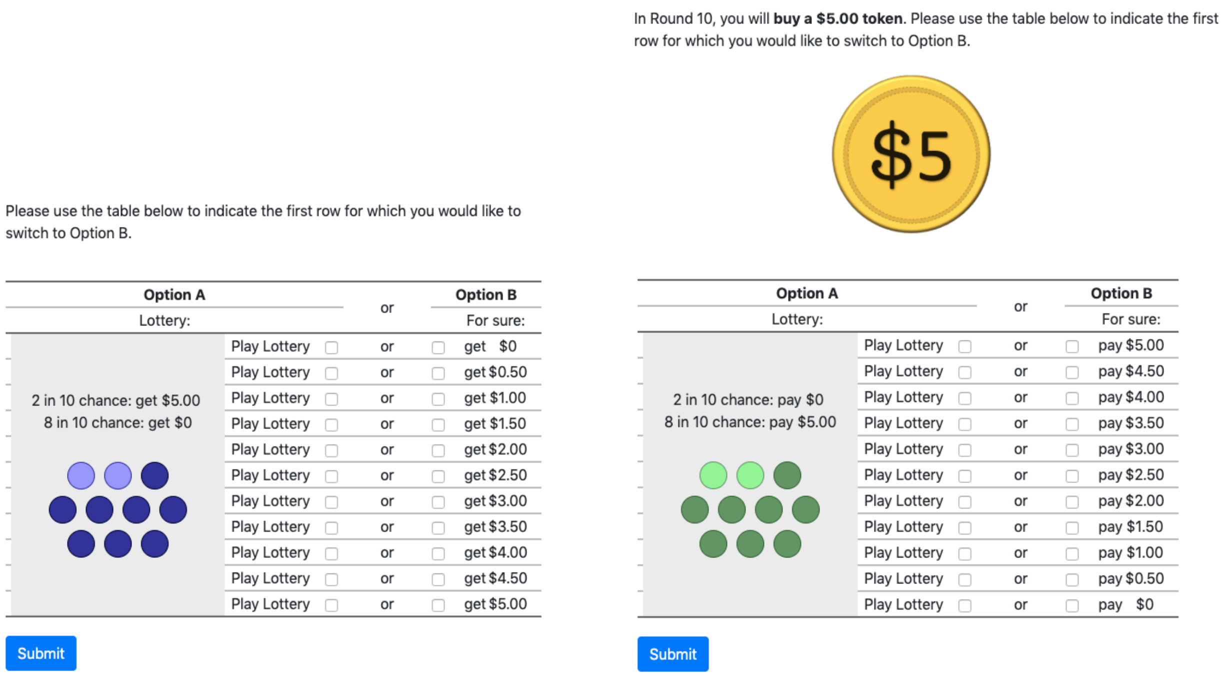

Please use the table below to indicate the first row for which you would like to switch to Option B.

Option A		or	Option B
Lottery:			For sure:
2 in 10 chance: get \$5.00 8 in 10 chance: get \$0	Play Lottery	or	get \$0
	Play Lottery	or	get \$0.50
	Play Lottery	or	get \$1.00
	Play Lottery	or	get \$1.50
	Play Lottery	or	get \$2.00
	Play Lottery	or	get \$2.50
	Play Lottery	or	get \$3.00
	Play Lottery	or	get \$3.50
	Play Lottery	or	get \$4.00
	Play Lottery	or	get \$4.50
	Play Lottery	or	get \$5.00

Submit

(a) Standard frame

In Round 10, you will **buy a \$5.00 token**. Please use the table below to indicate the first row for which you would like to switch to Option B.

Option A		or	Option B
Lottery:			For sure:
2 in 10 chance: pay \$0 8 in 10 chance: pay \$5.00	Play Lottery	or	pay \$5.00
	Play Lottery	or	pay \$4.50
	Play Lottery	or	pay \$4.00
	Play Lottery	or	pay \$3.50
	Play Lottery	or	pay \$3.00
	Play Lottery	or	pay \$2.50
	Play Lottery	or	pay \$2.00
	Play Lottery	or	pay \$1.50
	Play Lottery	or	pay \$1.00
	Play Lottery	or	pay \$0.50
	Play Lottery	or	pay \$0

Submit

(b) Token frame

Figure A8: Part 1: Lottery Questions

E.2 Part 2: Word-Search Task

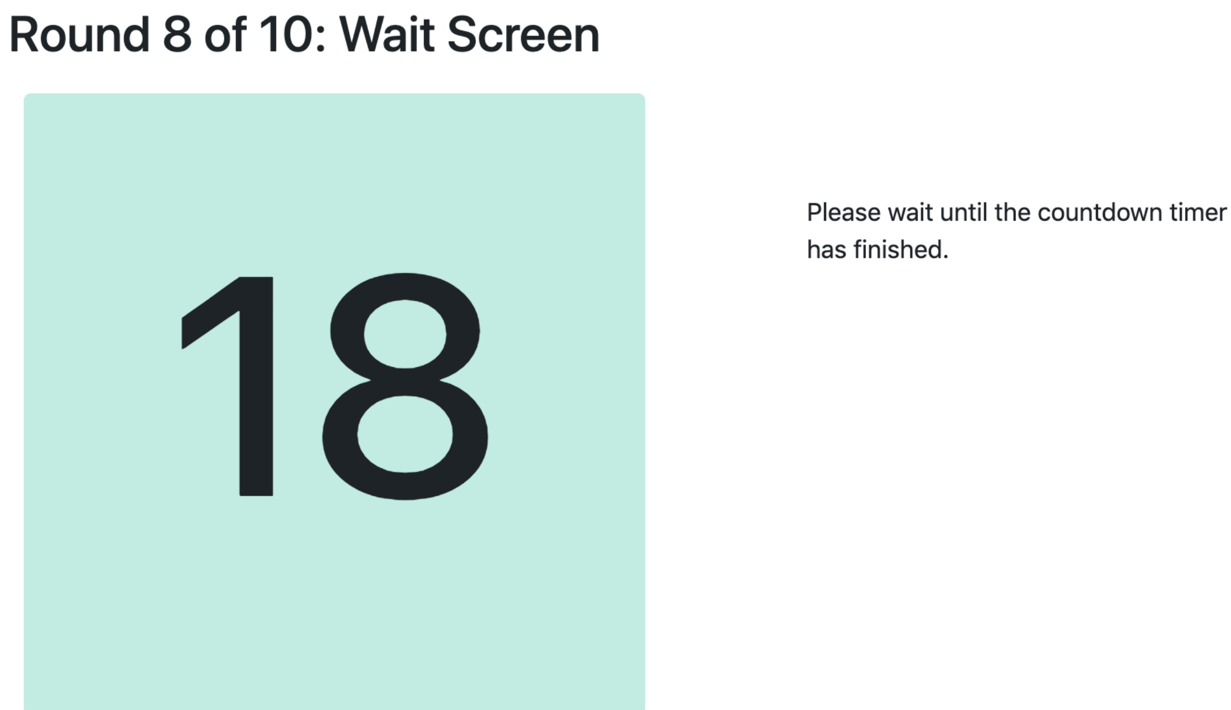

Figure A9: Screenshot of wait screen during 30 second countdown

E.3 Part 3: Control Lottery and Peak-Price Lottery Decisions

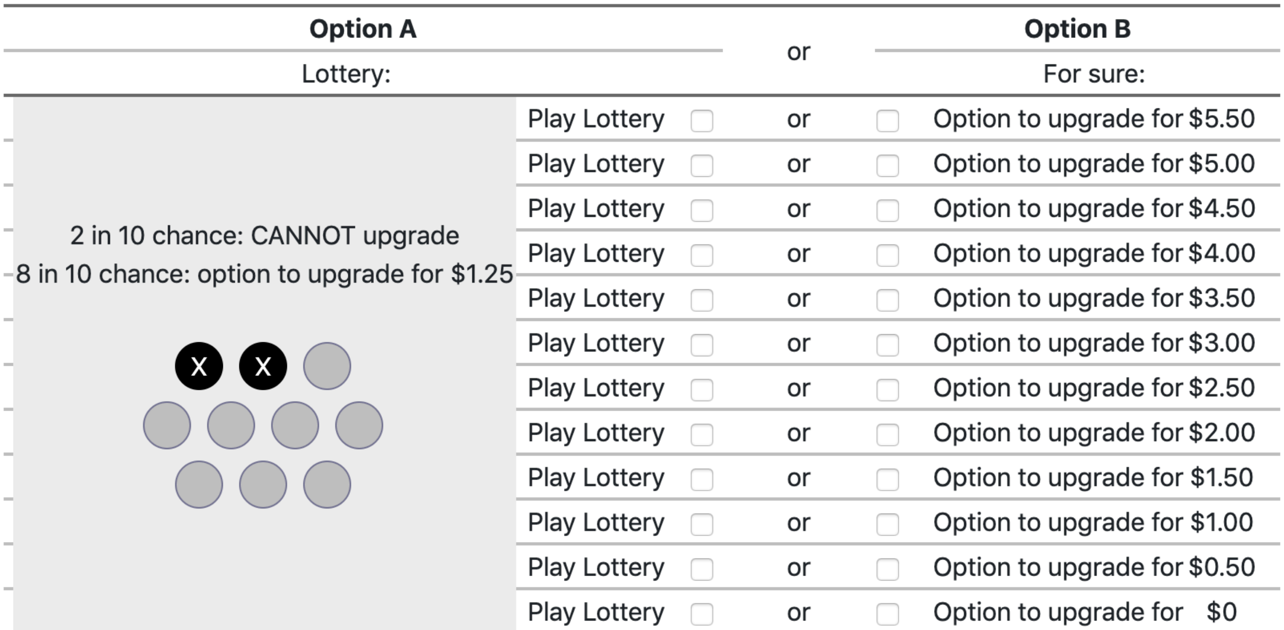

Option A Lottery:			or		Option B For sure:
2 in 10 chance: CANNOT upgrade 8 in 10 chance: option to upgrade for $1.25	Play Lottery	☐	or	☐	Option to upgrade for $5.50
	Play Lottery	☐	or	☐	Option to upgrade for $5.00
	Play Lottery	☐	or	☐	Option to upgrade for $4.50
	Play Lottery	☐	or	☐	Option to upgrade for $4.00
	Play Lottery	☐	or	☐	Option to upgrade for $3.50
	Play Lottery	☐	or	☐	Option to upgrade for $3.00
	Play Lottery	☐	or	☐	Option to upgrade for $2.50
	Play Lottery	☐	or	☐	Option to upgrade for $2.00
	Play Lottery	☐	or	☐	Option to upgrade for $1.50
	Play Lottery	☐	or	☐	Option to upgrade for $1.00
	Play Lottery	☐	or	☐	Option to upgrade for $0.50
	Play Lottery	☐	or	☐	Option to upgrade for $0

Figure A10: Screenshot of base control lottery ($\mathcal{B}$)

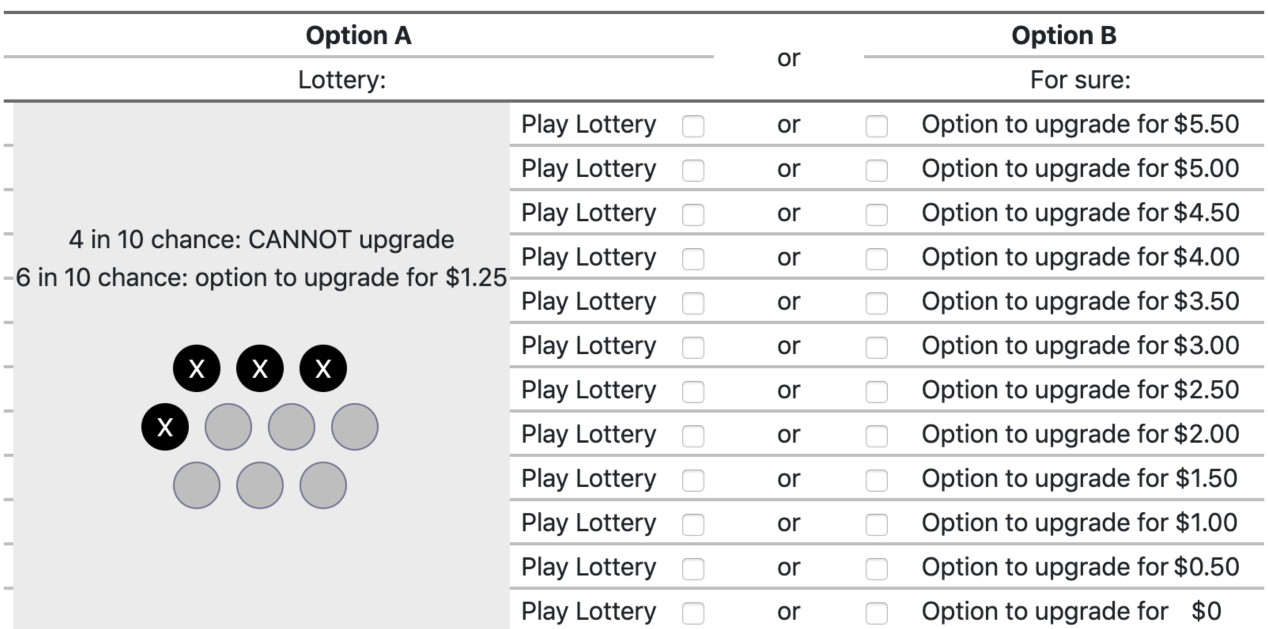

Option A Lottery:			or		Option B For sure:
4 in 10 chance: CANNOT upgrade 6 in 10 chance: option to upgrade for $1.25	Play Lottery	☐	or	☐	Option to upgrade for $5.50
	Play Lottery	☐	or	☐	Option to upgrade for $5.00
	Play Lottery	☐	or	☐	Option to upgrade for $4.50
	Play Lottery	☐	or	☐	Option to upgrade for $4.00
	Play Lottery	☐	or	☐	Option to upgrade for $3.50
	Play Lottery	☐	or	☐	Option to upgrade for $3.00
	Play Lottery	☐	or	☐	Option to upgrade for $2.50
	Play Lottery	☐	or	☐	Option to upgrade for $2.00
	Play Lottery	☐	or	☐	Option to upgrade for $1.50
	Play Lottery	☐	or	☐	Option to upgrade for $1.00
	Play Lottery	☐	or	☐	Option to upgrade for $0.50
	Play Lottery	☐	or	☐	Option to upgrade for $0

Figure A11: Screenshot of control lottery with doubled probability ($\mathcal{P}$)

Decision 6 of 10

Under Decision 6, and Decision 6 ONLY, your payment in the word search task at the end of Part 3 will be DOUBLED. You normally earn $1.00 per word you find. If Decision 6 is implemented, you will earn $2.00 per word you find instead. If you find all 10 words in the task at the end of Part 3, you thus have the chance to win up to $20.00 in this part.

Next

(a) Announcement that stakes will be doubled

Remember that under Decision 6, your payment in the word search task at the end of Part 3 will be DOUBLED.

Please use the table below to indicate the first row for which you would like to switch to Option B.

Option A		or	**Option B**
Lottery:			For sure:
2 in 10 chance: CANNOT upgrade 8 in 10 chance: option to upgrade for $2.50	Play Lottery	or	Option to upgrade for $5.50
	Play Lottery	or	Option to upgrade for $5.00
	Play Lottery	or	Option to upgrade for $4.50
	Play Lottery	or	Option to upgrade for $4.00
	Play Lottery	or	Option to upgrade for $3.50
	Play Lottery	or	Option to upgrade for $3.00
	Play Lottery	or	Option to upgrade for $2.50
	Play Lottery	or	Option to upgrade for $2.00
	Play Lottery	or	Option to upgrade for $1.50
	Play Lottery	or	Option to upgrade for $1.00
	Play Lottery	or	Option to upgrade for $0.50
	Play Lottery	or	Option to upgrade for $0

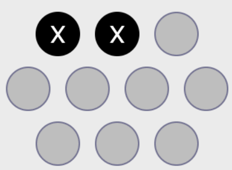

(b) Control lottery screen

Figure A12: Screenshot of control lottery with doubled stakes ($\mathcal{S}$)

Option A			or		Option B
Lottery:					For sure:
2 in 10 chance: Option to upgrade for $2.81 8 in 10 chance: Option to upgrade for $1.25 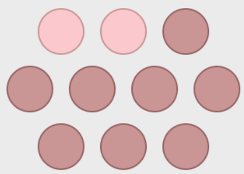	Play Lottery	☐	or	☐	Option to upgrade for $5.50
	Play Lottery	☐	or	☐	Option to upgrade for $5.00
	Play Lottery	☐	or	☐	Option to upgrade for $4.50
	Play Lottery	☐	or	☐	Option to upgrade for $4.00
	Play Lottery	☐	or	☐	Option to upgrade for $3.50
	Play Lottery	☐	or	☐	Option to upgrade for $3.00
	Play Lottery	☐	or	☐	Option to upgrade for $2.50
	Play Lottery	☐	or	☐	Option to upgrade for $2.00
	Play Lottery	☐	or	☐	Option to upgrade for $1.50
	Play Lottery	☐	or	☐	Option to upgrade for $1.00
	Play Lottery	☐	or	☐	Option to upgrade for $0.50
	Play Lottery	☐	or	☐	Option to upgrade for $0

Figure A13: Screenshot of peak-price lottery $\mathcal{CPP}$

E.4 Part 3: Attention Checks

Part 3

Please think about your decisions carefully and make sure to read all prompts in detail. Some prompts in this section may feature **attention checks**. If you fail an attention check, you will not earn any money in Part 3.

Next

(a) Attention check announcement at the start of Part 3

Decision 4 of 10

Decision 4 is an **attention check**. If you pass this attention check and Decision 4 is randomly selected to be implemented, you will be able to upgrade for free. If you fail the attention check, you will earn no money in Part 3.

In order to pass the attention check, please select to switch to Option B starting in the first row of the table below (i.e. indicate that you would prefer Option B in all cases).

Option A Lottery:		or	Option B For sure:
8 in 10 chance: Option to upgrade for $3.20 2 in 10 chance: Option to upgrade for $1.55 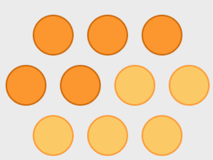	Play Lottery	or	Option to upgrade for $5.50
	Play Lottery	or	Option to upgrade for $5.00
	Play Lottery	or	Option to upgrade for $4.50
	Play Lottery	or	Option to upgrade for $4.00
	Play Lottery	or	Option to upgrade for $3.50
	Play Lottery	or	Option to upgrade for $3.00
	Play Lottery	or	Option to upgrade for $2.50
	Play Lottery	or	Option to upgrade for $2.00
	Play Lottery	or	Option to upgrade for $1.50
	Play Lottery	or	Option to upgrade for $1.00
	Play Lottery	or	Option to upgrade for $0.50
	Play Lottery	or	Option to upgrade for $0

Submit

(b) Attention check screen

Figure A14: Screenshot of attention check