# ESSAYS ON ECONOMIC CHALLENGES TO RENEWABLE ENERGY INTEGRATION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Christina Maria Ursula Korting

May 2020

ESSAYS ON ECONOMIC CHALLENGES TO RENEWABLE ENERGY INTEGRATION

Christina Maria Ursula Korting, Ph.D.
Cornell University 2020

As awareness regarding the adverse climate and health impacts of fossil-based energy sources grows around the world, so does the need for rigorous evaluation of possible interventions aimed at promoting the use of renewable energy alternatives. The introduction of renewable resources often creates unforeseen tensions because they differ from the fossil-based energy sources they replace either in physical composition (as in the case of biofuels) or due to the nature and timing of their production (as in the case of renewable electricity sources). In this dissertation, I use numerical simulation methods and experiments to study two such challenges in detail—blending constraints for biofuels in transportation and contracts aiming to address solar-induced peaks in electricity demand. The analyses developed here aim to inform regulatory decision making by quantifying the potential challenges and highlighting previously undocumented effects of renewable energy integration.

This dissertation consists of three essays. The first two chapters study the market effects and incidence of a physical constraint on biofuel blending; the so called "ethanol blend-wall". The existence of this constraint substantially affects the chosen compliance channels and hence welfare implications of the Renewable Fuel Standards (RFS) which mandate the use of biofuels in transportation in the US. The third chapter provides experimental evidence for the existence of a "control premium" (an intrinsic preference for retaining control over a de-

cision right above and beyond its instrumental value) in a context relevant to solar electricity integration: due to the timing of solar production, an increasing share of solar generation exacerbates demand peaks in the early evening relative to demand during the surrounding hours. Peak demand creates the need for costly short-term generation capacity which is often associated with higher marginal costs and increased emissions. One demand-side tool to address these peaks are Direct Load Control (DLC) contracts which strive to reduce or shift peak electricity demand by compensating consumers for granting utilities the right to turn off certain appliances remotely in times of tight supply. However, I experimentally show that the compensation which consumers require to adopt this type of contract exceeds the value of the usage benefits they forgo due to an intrinsic preference to retain control.

The first chapter, "Demystifying RINs: A Partial Equilibrium Model of U.S. Biofuels Markets," co-authored with David Just, examines the available compliance channels under the RFS and highlights how their relative use depends on the prevailing mandate requirements. As market pressures increase due to rising total renewable mandates in the presence of binding ethanol infrastructure constraints, the simulation results provide evidence for two important compliance channels not usually emphasized in the literature: overage from nested mandate categories and a contraction of the market for low-ethanol blend fuels such as E10 in order to reduce the overall compliance base. In fact, I show that the overall markets for motor gasoline and diesel fuel may contract in order to accommodate the mandates. In addition, the paper studies the price formation of the main mandate compliance instrument, so called Renwable Identification Numbers (RINs), and points out important inconsistencies in the usual practice of equating the price of RINs to the gap between ethanol supply and demand evaluated at the mandate level.

The second chapter, "Who Will Pay for Increasing Biofuel Mandates? Incidence of the Renewable Fuel Standard Given a Binding Blend Wall," coauthored with David Just and Harry de Gorter, extends this analysis to study the resulting welfare implications. This analysis fills an important gap in the literature by explicitly taking the nested mandate structure and joint compliance into account. We show that these two regulatory features effectively create a dual link between gasoline and diesel markets with the result that the cost of increasing biofuel mandates given a binding ethanol blend wall falls disproportionately on diesel fuel consumers. This result is likely to have significant general equilibrium ramifications through indirect channels such as inflation since the main consumers of diesel fuel in the U.S. are trucks and trains. Overall, these two chapters provide important insights into the market and welfare consequences of the ethanol blend wall which has important implications for the future implementation of the RFS.

The third chapter, "Taking a Load Off: Experimental Evidence of Preferences for Control with an Application to Residential Electricity Demand," uses a novel experimental design to show that intrinsic preferences for control can significantly impact the rewards required to encourage consumers to participate in DLC-style contracts. My findings relate to earlier work outside of the energy domain showing that individuals value retaining control over payoffs or delegation rights above instrumental value. This paper makes important contributions to both the behavioral economics literature and the literature on the cost and effectiveness of demand response programs. 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 earlier findings 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. I find that participants, on average, exhibit a *control premium* of 9-32%, and are sensitive to both the stakes and probability of losing control.

## BIOGRAPHICAL SKETCH

Christina Korting is a doctoral candidate at Cornell University. Christina is an applied microeconomist with a focus on energy and the environment. Her research uses experimental, empirical, and numerical simulation methods to analyze the effects of energy policies, to study consumer behavior in energy markets, and to understand stated valuations in environmental conservation settings. Before attending Cornell, Christina received a Masters degree in economics from Humboldt University Berlin and a Masters degree in statistics from the Ecole Nationale de la Statistique et de l'Administration Economique (ENSAE) Paris.

## ACKNOWLEDGEMENTS

I would like to thank my committee: David Just, William Schulze, and Harry de Gorter. It was your guidance and support over the last few years that have made this work possible. I would also like to thank Ted O'Donoghue for his generous mentorship; our stimulating discussions have contributed greatly to the last chapter of this dissertation. Working with Zhuan Pei has been a constant source of inspiration over the last two years and has profoundly affected my approach to research. Thank you for giving me this opportunity.

I am grateful to my cohort who have been the best friends and peers anyone could wish for–especially to Matt, whose advice, love and encouragement is woven into every page of this dissertation. Above all, I would not be where I am today without the steadfast support of my parents and sister. Thank you for everything!

The first two chapters of this dissertation are the result of collaborations with David Just and Harry de Gorter. Thank you for sharing your insights and institutional knowledge about biofuels, and for patiently guiding me through my first review and publication process. These chapters reflect only a fraction of the things I have learned from you.

I gratefully acknowledge the financial support received for the first two chapters research from (i) the National Research Initiative of the National Institute of Food and Agriculture, US Department of Agriculture (USDA; Washington, DC, USA; Grant No 2014-67023-21820 under the Agricultural and Food Research Initiative (AFRI)).; and (ii) USDA Cooperative Agreement 58-0111-15-018; and (iii) Cornell University Agricultural Experiment Station federal formula funds (project no. NYC-121-7444) received from the Cooperative State Research, Education and Extension Service, USDA. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the view of the U.S. Department of Agriculture. Research in the third chapter of this dissertation has been generously supported by the Cornell CALS Andrew W. Mellon Student Research Grant.

# TABLE OF CONTENTS

## LIST OF TABLES

## LIST OF FIGURES

CHAPTER 1

# DEMYSTIFYING RINS: A PARTIAL EQUILIBRIUM MODEL OF U.S. BIOFUELS MARKETS

## 1.1 Abstract

We explore four fundamental channels of mandate compliance available under current U.S. biofuels policy: increased ethanol blending through E10 or E85, increased biodiesel blending, and a reduction in the overall compliance base. Simulation results highlight the interplay and varying importance of these channels at increasing blend mandate levels. In addition, we establish how RIN prices are formed: The value of a RIN in equilibrium is shown to reflect the marginal cost of compensating the blender for employing one additional ethanol-equivalent unit of biofuel. This contrasts with existing research equating the price of RINs to the gap between ethanol supply and demand evaluated at the mandate level. We demonstrate the importance of this distinction in case of binding demand side infrastructure constraints such as the ethanol blend wall: as percentage blend mandates increase, the market for low-ethanol blends may contract in order to reduce the overall compliance base. This has important implications for implied ethanol demand in the economy.

## 1.2 Introduction

U.S. biofuels policy has reached a critical junction. By lowering the final 2014-2016 blend mandate requirements in December 2015, the EPA has acknowledged the difficulty of complying with the original 2010 renewable fuel targets due to important demand side bottlenecks. To gauge the intensity of the feasibility debate currently raging in the biofuels space, it suffices to look at the number of comments received on proposed rulemakings by the EPA: 23 in 2011, 488 in 2012, 94 in 2013, and 344,947 in 2014 [1].

Given the current state of the market, it is more important than ever to understand the available compliance mechanisms, their impacts and limitations. Important groundwork in this context was laid by [14] and [49] who analyze the general market effects and incidence of a blend mandate. [62] provide estimates of potential demand for high-ethanol blends given current infrastructure constraints, which [63] integrate into a short term partial equilibrium model accounting for existing market rigidities. Forthcoming work by [46] explores the dynamic nature of the mandate.

We contribute to this discussion by extending the existing literature in two important ways: First, we make explicit which compliance channels are available under the current biofuels policy and when they are employed. To do so, we propose a partial equilibrium model which takes into account the nested mandate structure of the RFS2. As market pressures increase due to higher total renewable mandates in the presence of ethanol infrastructure constraints, our simulation results provide evidence for two important compliance channels not usually emphasized in the literature: overage from nested mandate categories and a contraction of the market for low-ethanol blend fuels such as E10 in order to reduce the overall compliance base.

[1] see www.regulations.gov, dockets EPA-HQ-OAR-2010-0133-0001 for 2011, EPA-HQ-OAR-2010-0133-0102 for 2012, EPA-HQ-OAR-2012-0546-0001 for 2013, EPA-HQ-OAR-2013-0479-0037 for 2014

Second, we let the exchange of RINs enter the model endogenously as an additional decision variable for blending and refining operations that are distinct and not vertically integrated[2]. This allows us to conclusively establish how RIN prices are formed and what they represent. Contrary to most of the existing literature, we find that the *core value* of a RIN in equilibrium reflects the marginal cost of compensating the blender for employing one additional ethanol-equivalent unit of biofuel[3]. Previously, this core value was usually equated to the gap between the ethanol supply and demand curves evaluated at the mandate level (e.g. [53], [82] and [51]). Figure 1.1 illustrates this idea. While the two definitions of RIN prices are conceptually very similar, we emphasize the importance of this distinction by showing the strong dependence of the notion of implied ethanol demand on prevailing percentage mandate levels. Our RIN price formula therefore provides a more concise and reliable way to explain the price of RINs.

## 1.3 Renewable Fuel Standards

The Renewable Fuel Standards of 2005 (RFS1) and 2007 (RFS2), passed as part of the Energy Policy Act (EPAct) and the Energy Independence and Security Act (EISA) respectively, mandate the use of specific amounts of biofuels in the

[2]According to a study by [58], refiners and importers of fossil fuels have limited blending operations and blend only a small fraction of their own production for retail use. See A.1 for further details on the fuel market structure.

[3]RIN prices are usually broken down into three components: (i) the core value, or cost of complying with the mandate, (ii) the marginal value of transferring physical blending opportunities from high-cost to low-cost blenders, and (iii) the real option value of meeting binding mandates in the future thanks to the (limited) bankability of RINs. Our model focuses on the first component since it tends to dominate RIN prices when mandates are binding. However, our model could be extended to a dynamic setting including heterogeneous cost structures for blenders in order to capture the remaining RIN price components.

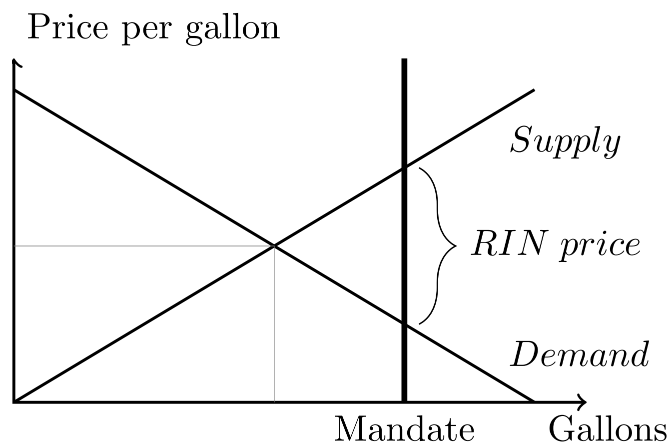


**Figure 1.1:** Proposed RIN Price Formation in Existing Literature; Market for Ethanol

transportation sector. The reasons for promoting the use of biofuels are manifold: (a) protecting against rising fossil fuel prices; (b) reducing emissions from the transportation sector; (c) promoting energy security by reducing the dependency on fossil fuel imports; and (d) increasing farm income ([69, 52]). The Renewable Fuel Standards share many of their environmental goals with other existing energy policies such as fuel taxes and fuel efficiency standards[4].

The RFS2 imposes a series of annual volumetric targets which are subsequently converted into blend mandates for the year ahead using forecast gasoline and diesel consumption levels. The obligated parties are refiners and importers of fossil fuels who often do not directly control the final blend of consumer motor fuels. Compliance is therefore monitored through financial instruments called Renewable Identification Numbers (RINs) which represent one ethanol-equivalent unit of biofuel blended. RINs are generated at production or import of a biofuel, and become detached and separately tradable at blending.

The RFS2 is designed to be ‘technology forcing’, governing both the pace and

[4]For more details about the RFS, please see A.1

the intensity of the shift to more environmentally friendly fuels using a nested mandate structure. For this reason, four nested categories were established under the RFS2: both *cellulosic biofuels* and *biomass-based diesel* (BBD) are nested under the *advanced biofuels* category, which requires a greenhouse gas emissions (GHG) reduction of at least 50% compared to the fossil fuel being replaced; the advanced biofuels mandate in turn is part of the *total renewable fuels* category (TR) which requires GHG savings of at least 20%. Figure 1.2 provides a graphical representation of this nested structure.

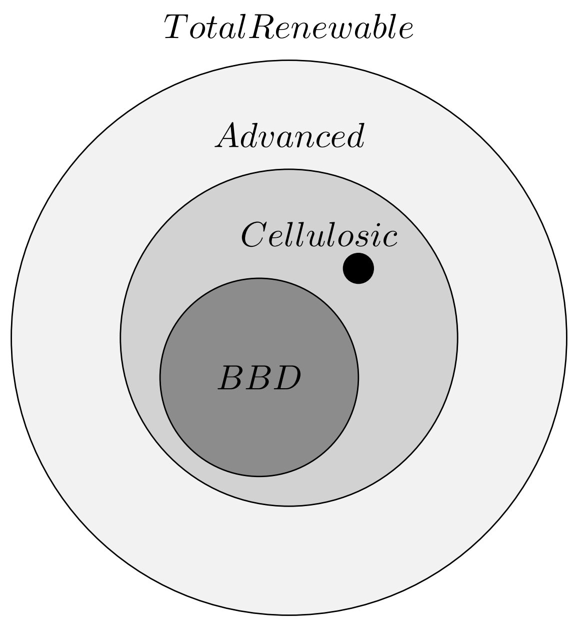


**Figure 1.2:** Nested Mandate Structure under the RFS2

The residual part of the total renewable fuels mandate not met through advanced biofuels is often referred to as conventional biofuels. It is usually filled using U.S. corn-based ethanol or biodiesel which did not meet the advanced biofuel GHG savings requirements[5]. However, overage from the advanced biofuels

[5]Based on data from the EPA Moderated Transaction System (EMTS), 252mn D6 RINs were generated from biodiesel or renewable diesel in 2013 (https://www.epa.gov/fuels-registration-reporting-and-compliance-help/2013-renewable-fuel-standard-data). We maintain a single biodiesel supply curve for simplicity in this model. It may be interesting to analyze the dichotomy between D4 and D6 biodiesel inputs and the resulting market dynamics in future work.

category could equally be used for compliance towards the overall mandate. It is therefore important to note that the RFS2 does not impose explicit ethanol volume requirements.

At the end of each year, every obligated party has to comply with all four sub-categories of the RFS2. For instance, a pure gasoline importer would still have to provide BBD RINs to the EPA in order to meet his renewable volume obligation (RVO). Table 1.1 highlights some volumetric and percentage blend targets by category and introduces the labeling convention for RINs. For example, cellulosic biofuels generate D3 RINs. The nested mandate structure leads to an implicit pricing relationship between the categories since any D3 or D4 RIN can also be used to comply with the advanced or total biofuels mandate: the prices must always satisfy $p_{D3}, p_{D4} \geq p_{D5} \geq p_{D6}$. These pricing relationships will hold with equality if the wider mandates become binding and there is a need to attract overage from some of the nested sub-categories.

2013 was the first year when BBD and conventional RIN prices approached parity. The fundamental driver behind this convergence was the attainment of the so called *ethanol blend wall*: ethanol is corrosive and therefore risks damaging engines and fuel tanks at concentrations of more than 10% in cars that are not specifically equipped to use it. In 2013, most motor gasoline was already sold with a 10% ethanol share. There are currently four distinct ethanol-gasoline blends sold in the U.S.: (i) E0 which contains no ethanol, (ii) E10 which contains up to 10% ethanol and can be used in all cars, (iii) E15 which contains up to 15% ethanol, is approved for use in models newer than 2001, but does not meet some car manufacturer warranties and is not currently widely available, and (iv) E85 which contains between 51-83% ethanol. E85 is designated for use in so

**Table 1.1:** RFS2 Mandates by Category

| | | 2015 | | 2022 |
|---|---|---|---|---|
| Mandate Category | Label | Volumetric Mandate (bn GAL) | Percentage Mandate | Volumetric Mandate (bn GAL) |
| Cellulosic biofuel | D3 | 0.123 | 0.069% | 16 |
| Biomass-based diesel | D4 | 1.73 | 1.49% | TBD |
| Advanced biofuel | D5 | 2.88 | 1.62% | 21 |
| Renewable fuel | D6 | 16.93 | 9.52% | 36 |

Volumetric mandates are shown in billion gallons of ethanol-equivalent except BBD which was originally introduced as a diesel standard and is therefore represented on a biodiesel-equivalent energy basis under the RFS. All percentage blend mandates, including D4, are shown in ethanol-equivalent terms.
†: Based on 2010 Final Rule [20]. The final 2022 mandates will likely be revised downward to reflect the slower than expected growth in cellulosic biofuel production and ethanol demand.
Source: [20] and [21]

called flexible-fuel vehicles (FFVs) that can run on higher blends but currently represent a small percentage of the U.S. fleet. We focus on E10 and E85 as the two dominant blends in the U.S. market[6].

## 1.4 Methodology

### 1.4.1 Model

Like [63], we propose a short term model of the U.S. biofuels market which explicitly captures the rigidities imposed by demand side infrastructure constraints. However, our framework extends their setting in two important ways: First, we model the creation of RIN prices more directly by allowing blenders

[6]The most recent EPA final rule document for 2014-2016 predicts E15 consumption of about 320mn GAL in 2016 assuming an optimistic rate of growth in retail availability. However, the relatively higher ethanol content in E15 is expected to be roughly offset by small amounts of E0 sales. Their most favorable infrastructure scenario for E85 would lead to an estimated 400mn GAL consumed in 2016. As a reference point, 2014 consumption is estimated at 150mn GAL ([21], pp. 77460-4)

and refiners to choose the quantity of RINs endogenously, and to then trade RINs between each other subject to a market clearing constraint. Second, we capture the nested structure of the U.S. biofuels mandate by explicitly modeling the biodiesel space and allowing for strategic overage of biodiesel RINs in order to meet the total renewable mandate.

Generally, existing models of RIN prices and the RFS2 can be differentiated along four dimensions: (i) short vs. long term approaches (e.g. considering the blend wall or abstracting away from current infrastructure constraints) (ii) link to agricultural markets and trade vs. closed economy, fuel-only models (iii) nesting vs. ethanol only and (iv) static vs dynamic settings. In order to obtain a parsimonious yet meaningful representation of the core value of RIN prices, and to study all available channels of mandate compliance, we have chosen a static, closed economy model considering only fuels and focusing exclusively on D4 and D6 RINs, but taking the nested mandate structure and short term infrastructure constraints into account[7]. An explicit distinction between conventional and flex-fuel vehicle drivers differentiates our approach from [12] who largely abstract away from the ethanol consumption bottleneck by assuming a perfectly inelastic motor gasoline demand which can be arbitrarily split between E10 and E85.

Based on our model choice, the refiner ($R$) solves the problem of maximizing revenue from refined product sales minus the cost of refining ($C^R$), subject to meeting the BBD mandate requirement as well as the residual total renewable requirement not met by BBD overage.

[7]D3 RIN generation from cellulosic biofuels has historically been very low. D5 RINs are in large part generated from imported sugarcane ethanol and would have a very similar effect to D6 RINs in our model.

$$
\begin{aligned}
& \max_{\{q_G,\, q_D,\, q^R_{D4},\, q^R_{D6}\}} \Pi^R = \\
& p_G q_G + p_D q_D - C^R(q_G, q_D) \\
& - p_{D4} q^R_{D4} - p_{D6} q^R_{D6} \qquad (1.1) \\
s.t. \quad & q^R_{D4} \geq \kappa_{BBD}(q_G + q_D) \\
and \quad & q^R_{D4} + q^R_{D6} \geq \kappa_{TR}(q_G + q_D)
\end{aligned}
$$

Throughout this paper, *motor gasoline* ($MG$) denotes finished gasoline including ethanol blending components, while *diesel fuel* ($DF$) refers to finished diesel including biodiesel for transportation. $G$ and $D$ on the other hand symbolize gasoline and diesel derived from crude oil. RFS2 percentage blend mandates are denoted by $\kappa$, which represents the ratio of required renewable to fossil fuels[8]. All quantities and prices are denoted as $q_\circ$ and $p_\circ$ respectively, where $\circ$ stands for a generic subscript. An exhaustive list of variable descriptions is provided in A.3.

The first constraint specifies that the quantity of D4 RINs retired has to be at least commensurate with the BBD percentage mandate requirement ($\kappa_{BBD}$) multiplied by the total amount of fossil fuels consumed in the economy. Given the nested mandate structure, D4 RINs also count towards the total renewable mandate. The second constraint therefore imposes that the sum of D4 and D6 retirements has to exceed the TR percentage mandate applied to the total compliance base.

[8]Note that the blend wall is a physical limitation which relates to the amount of ethanol relative to the overall fuel quantity instead. We will therefore sometimes convert mandate amounts into these terms for illustration

In stating the blender's problem ($B$), we will assume that E85 always has an ethanol content of 74%. This assumption is in line with the average E85 specifications reported by the EPA and used in the literature (e.g. [40]). The E10 ethanol content on the other hand, represented by $\theta_{E10}$, is allowed to vary between zero and 10% . This permits blenders to use less than 10% ethanol at low mandate levels and thereby implicitly captures the physical reality of E0 sales. The biodiesel blend ratio in diesel fuel ($\theta_{DF}$) is allowed to vary freely since there are no blend wall constraints in the biodiesel market. Biodiesel blends up to 5% require no separate labelling at the pump, and blends up to 20% do not require engine modifications and are commonly used in the U.S.

The blender generates revenue by selling E10, E85 and diesel fuel and incurs a cost of blending denoted by $C^B$. The ethanol content in E10 as well as the biodiesel content in diesel fuel are endogenous to his decision. The blender's constraints represent the process of RIN generation: the number of units generated has to be proportional to the amount of biofuels blended.

All percentage blend mandates, and therefore all RIN quantities, are set in ethanol-equivalent terms by the RFS. The first constraint for the blender, which represents $D4$ RIN generation, therefore applies an equivalence value of 1.5 in order to account for the higher energy content of biodiesel compared to ethanol. The second constraints reflects the generation of D6 RINs through motor gasoline sales. The last constraint effectively imposes the blend wall, constraining the maximum ethanol content in E10 to 10%.

$$
\begin{aligned}
\max_{\substack{\{q_{E10},\ q_{E85},\ \theta_{E10}\\ q_{DF},\ q^B_{D4},\ q^B_{D6},\ \theta_{DF}\}}} \quad & \Pi^B = \\
& q_{E10}(p_{E10} - t_{MG}) + q_{E85}(p_{E85} - t_{MG}) \\
& + q_{DF}(p_{DF} - t_{DF}) \\
& + p_{D6}q^B_{D6} + p_{D4}q^B_{D4} + \theta_{DF}q_{DF}tc_{BD} \\
& - ((1-\theta_{E10})q_{E10} + 0.26q_{E85})p_G \\
& - (\theta_{E10}q_{E10} + 0.74q_{E85})p_E \\
& - (1-\theta_{DF})q_{DF}p_D - \theta_{DF}q_{DF}p_{BD} \\
& - C^B_{MG}(q_{E10}, q_{E85}) - C^B_{DF}(q_{DF}) \\
s.t. \quad & q^B_{D4} \leq 1.5\theta_{DF}q_{DF} \\
and \quad & q^B_{D6} \leq \theta_{E10}q_{E10} + 0.74q_{E85} \\
and \quad & \theta_{E10} \leq 0.1
\end{aligned}
\tag{1.2}
$$

The blender has to deduct gasoline and diesel fuel taxes, $t_{MG}$ and $t_{DF}$, on every gallon sold, but receives a 1 USD/GAL biodiesel blender tax credit, $tc_{BD}$, on the amount of biodiesel he blends[9]. $C^{MG}_B$ and $C^{DF}_B$ represent the blending costs for motor gasoline and diesel fuel respectively.

All supply and demand functions are assumed to be of the constant elasticity form

$$q = Ap^{\epsilon}$$

[9]This tax credit has recently been extended to December 2016 through House of Representatives Bill 2029, Section 185

with an elasticity of $\epsilon$ and a scaling factor of $A$. Throughout this paper, subscripts are used to differentiate between cost, supply and demand functions $(C, S, D)$, and then indicate the relevant product. To model the consumer choice between E10 and E85, we adopt a neo-classical approach: first, consumers are split into flexible-fuel vehicle (FFV) owners and conventional vehicle (C) owners. Rather than allowing for heterogeneous preferences for environmental quality and hence gradual switching behavior, we assume that all FFV drivers will switch from E10 to E85 whenever E85 prices become equally or more attractive on an energy-equivalent basis. We denote the demand functions for E10 and E85 by $D_{\circ}(p_{E10}, p_{E85})$, but will drop the price arguments going forward for notational convenience. Denoting by $\lambda$ the energy-equivalence factor between E10 and E85, we therefore obtain the following piecewise demand functions:

- Case 1: $p_{E85} > \lambda p_{E10}$

  No E85 will be consumed and all FFV drivers will choose to consume E10 instead

$$D_{E85} = 0$$
$$D_{E10} = A_{D_{FFV}} p_{E10}^{\epsilon_{D_{MG}}} + A_{D_C} p_{E10}^{\epsilon_{D_{MG}}}$$

- Case 2: $p_{E85} = \lambda p_{E10}$

  FFV drivers are indifferent between E10 and E85 and will therefore consume any quantity of E85 between zero and their total fuel demand. Any residual demand will be consumed in the form of E10.

$$
\begin{aligned}
D_{E85} &\in [0, A_{D_{FFV}} p_{E10}^{\epsilon_{D_{MG}}}] \\
D_{E10} &= A_{D_{FFV}} p_{E10}^{\epsilon_{DMG}} + A_{DC} p_{E10}^{\epsilon_{D_{MG}}} - D_{E85}
\end{aligned}
$$

- Case 3: $p_{E85} < \lambda p_{E10}$

  FFV drivers exclusively use E85

$$
\begin{aligned}
D_{E85} &= A_{D_{FFV}} \left( \frac{1}{\lambda} p_{E85} \right)^{\epsilon_{D_{MG}}} \\
D_{E10} &= A_{DC} p_{E10}^{\epsilon_{DMG}}
\end{aligned}
$$

The equilibrium in our model is governed by the interplay of first order and complementary slackness conditions for blender and refiner as well as market clearing equations. The full list of equations is provided in A.2. Before presenting the simulation results for this model in section 1.5.3, we will first discuss the data we use in calibrating the model to market.

### 1.4.2 Data and Calibration

The scale parameters $A_{\circ}$ for the fuel demand functions are calibrated to annual price and quantity data from 2015. Table 1.2 below highlights the elasticity estimates from the literature that were chosen for our calibrations:

In order to calibrate the motor gasoline demand functions, we use E85 demand estimates from Figure 7 of [62]. Based on current infrastructure limitations, the E85 demand at a price of 1 USD/GAL in this figure is at around 1.2

**Table 1.2:** Elasticity Estimates from the Literature

| Variable | Description | Value | Source |
|---|---|---|---|
| **Supply Elasticities** | | | |
| $\epsilon_{S_E}$ | Ethanol | 2 | [50] |
| $\epsilon_{S_{BD}}$ | Biodiesel | 2 | (at 1.2b GAL) [4] |
| **Demand Elasticities** | | | |
| $\epsilon_{D_{MG}}$ | Motor Gasoline | -0.25 | [62] |
| $\epsilon_{D_{DF}}$ | Diesel Fuel | -0.07 | [13] |

bGAL. While [62] allow for gradual switching behavior from E10 to E85, the 1 USD/GAL price level is far enough below the assumed energy-equivalent E10 price to suppose that virtually all switching would have occurred. Any further consumption increases beyond this point can therefore be attributed to increased driving demand at low fuel prices. The motor gasoline demand of conventional drivers in our model is calibrated to the residual of total motor gasoline demand minus FFV demand.

In order to calibrate the cost functions, we rely on the perfect competition assumption and choose the scale parameter $A$ which equalizes marginal cost and marginal revenue (price) in 2015. In this case, this is equivalent to imposing a zero profit condition. Table A.1 highlights key data sources as well as the corresponding realizations for 2015.

## 1.5 Results

### 1.5.1 Compliance Channels

Using a sequence of simulation results at changing mandate levels, we are able to establish the existence of four distinct channels for mandate compliance in our model:

1. Increasing the blend ratio of ethanol in E10 (up to 10%)
2. Increasing E85 sales
3. Increasing the biodiesel share in diesel fuel beyond the BBD mandate
4. Decreasing the overall compliance base by selling less diesel fuel and/or motor gasoline

The first two compliance channels rely on increased ethanol blend ratios in motor gasoline (which could also be viewed as a decrease in E0 sales). The third channel makes use of the nested mandate structure, calling on RINs generated through biodiesel overage to comply with the total renewable mandate. To illustrate how the fourth compliance channel operates, consider an economy in which only motor gasoline is sold (i.e. there is no diesel fuel market), and the maximum E85 demand by FFV drivers is fixed at 1 bn GAL. A mandate level of $\kappa_{TR} = 11.\overline{11}\%$ in this economy implies an ethanol blend ratio of 10%, i.e. a blend ratio just at the blend wall. In this case, any amount of motor gasoline sales would be feasible under the mandate. If the mandate was raised to to $\kappa_{TR} = 12\%$ instead, the mandate would effectively impose a cap on total motor gasoline sales. To see this, consider the requirement

$$\frac{q_E}{q_G} = \frac{0.1q_{E10} + 0.74 * 1}{0.9q_{E10} + 0.26 * 1} \geq 12\% \tag{1.3}$$

Solving for $q_{E10}$, we find a maximum of 89.6 bn GAL of E10 sales in order to ensure mandate compliance. This fourth channel, which is rarely mentioned in the existing literature, plays a key role in practice since other channels grow more costly as mandates tighten. In a model without nesting in which surplus biodiesel RINs cannot be used to overcome the blend wall, this channel becomes the only option for compliance once the blend wall has been hit and E85 demand has been exhausted.

Our simulation results in section 1.5.3 provide evidence for the existence of all four channels. They also highlight the interplay of the different compliance channels as blend mandates increase given the 2015 market environment.

### 1.5.2 The Price of RINs

Based on the behavioral equations outlined in A.2, we derive the pricing formula for D4 RINs given in equation 1.5. The term

$$\frac{\partial C_B^{DF}}{\partial q_{DF}} = A^B_{C_{DF}} \epsilon^B_{C_{DF}} q_{DF}^{(\epsilon^B_{C_{DF}}-1)} \tag{1.4}$$

represents the marginal cost of blending diesel fuel. The 'core value' of a D4 RIN thus represents the marginal cost of compensating the blender for employing one additional ethanol-equivalent unit of biodiesel. The blender faces the input costs for the two blending components, incurs a marginal cost of blending, and is able to sell the final product at the diesel fuel price minus tax. If the costs of generating an additional unit of diesel fuel are higher than the price which can be achieved in the market, the blender demands a positive RIN price as compensation for blending since he is not himself obligated under the RFS.

Similarly, for D6 RINs we find the pricing relationship outlined in equation

$$p_{D4} = \underbrace{\frac{1}{1.5\theta_{DF}}}_{\substack{\text{Scaling Factor:}\\ \text{Diesel Fuel to D4 RINs}}} \left[ \underbrace{(1-\theta_{DF})p_D + \theta_{DF} p_{BD} + \frac{\partial C_B^{DF}}{\partial q_{DF}}}_{\substack{\text{Marginal Cost of Blending one}\\ \text{Additional Unit of Diesel Fuel}}} - \underbrace{(p_{DF} - t_D)}_{\substack{\text{Marginal Revenue from Selling one}\\ \text{Additional Unit of Diesel Fuel}}} \right] \tag{1.5}$$

1.6 assuming both fuel blends are sold in equilibrium. The interpretation of terms in these equations is equivalent to the diesel fuel case.

By establishing these concise pricing formulas, we provide an alternative to the widely established simplification equating the price of RINs to the gap between ethanol supply and demand at the mandated level. Beside the obvious abstraction away from the nested mandate structure, there are two key problems with this notion:

- *Implied ethanol demand is not well defined:* The ethanol demand schedule is usually defined as the implied demand for ethanol through E10 and E85 as ethanol prices vary. However, due to the existence of the four different compliance channels, and the potential reduction of low-ethanol blends at

$$\begin{aligned} p_{D6} &= \underbrace{\frac{1}{\theta_{E10}}}_{\substack{\text{Scaling Factor:}\\ \text{E10 to D6 RINs}}} \left[ \underbrace{(1-\theta_{E10})p_G + \theta_{E10} p_E + \frac{\partial C_B^{MG}}{\partial q_{E10}}}_{\substack{\text{Marginal Cost of Blending one}\\ \text{Additional Unit of E10}}} - \underbrace{(p_{E10} - t_{MG})}_{\substack{\text{Marginal Revenue from Selling one}\\ \text{Additional Unit of E10}}} \right] \\ &= \underbrace{\frac{1}{0.74}}_{\substack{\text{Scaling Factor:}\\ \text{E85 to D6 RINs}}} \left[ \underbrace{0.26 p_G + 0.74 p_E + \frac{\partial C_B^{MG}}{\partial q_{E85}}}_{\substack{\text{Marginal Cost of Blending one}\\ \text{Additional Unit of E85}}} - \underbrace{(p_{E85} - t_{MG})}_{\substack{\text{Marginal Revenue from Selling one}\\ \text{Additional Unit of E85}}} \right] \end{aligned} \tag{1.6}$$

high mandate levels in particular, the notion of implied ethanol demand is highly sensitive to the prevailing percentage mandate levels. Figure 1.3 illustrates this effect by showing simulated demand schedules for different total renewable blend mandates ($\kappa_{TR}$ =0%, 9% and 11%). Clearly, for any given ethanol volume, the free-market supply demand gap is substantially different from the supply-demand gap given a binding mandate[10].

- *Equilibrium ethanol quantities do not equal volumetric mandates:* Even assuming a well defined implied ethanol demand schedule and ignoring the fact that the RFS2 does not impose any direct mandates for ethanol, the implied volumetric ethanol mandate is not a meaningful quantity to consider to assess the price of RINs. Percentage mandate requirements are calculated using forecast motor gasoline consumption which will not usually be fulfilled exactly as predicted.

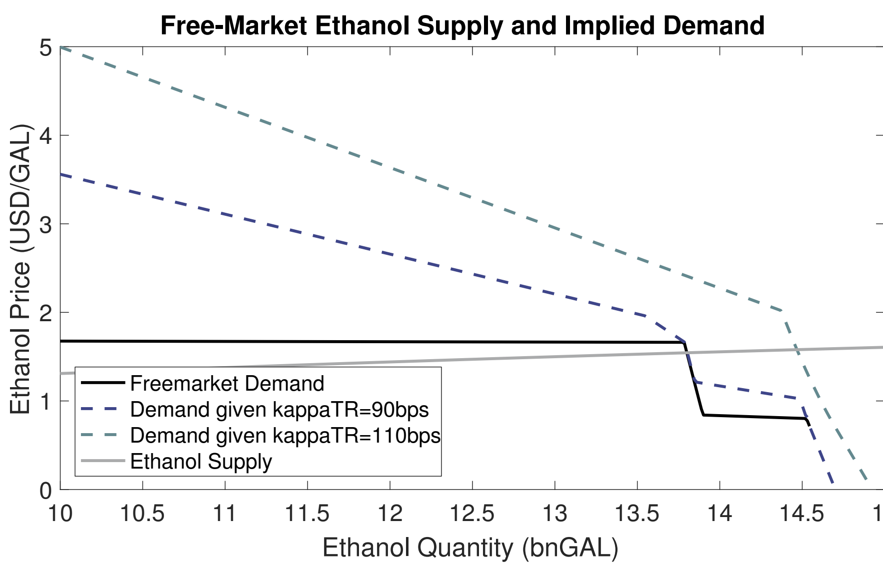


**Figure 1.3:** Implied Ethanol Demand Schedules at Different Mandate Levels

[10]At 0% and 9% mandate levels, we first see increased demand thanks to higher ethanol blend ratios in E10, and finally a jump in demand as ethanol becomes cheap enough to induce E85 sales. The 11% demand schedule only features one kink when channel four starts to dominate and the market contracts

This description of RIN prices therefore represents an inaccurate and highly impractical representation of the core value of RINs. However, the notion of the supply-demand gap is highly correlated to the more accurate pricing formula we provide: both are a function of the elasticity of ethanol supply as well as the potential ethanol demand given the blend wall. For example, the D6 equilibrium RIN price in equation 1.6 depends negatively on $p_{E85}$. This means that if the price of E85 has to adjust downwards faster due to demand side bottlenecks, the RIN price will increase faster as mandates rise.

In order to reinforce this idea and to study the evolution of equilibrium pricing relationships at varying mandate levels, we provide comparative static results for a simplified motor-gasoline-only model (see A.5 for a detailed model description). By calculating the comparative statics results based on this reduced framework, we can solve for upper bounds on $\kappa_{TR}$ which guarantee that $p_{E85}$ declines with mandate levels as would be expected. The corresponding bounds are provided in equation 1.7.

The bound on $p_{E85}$ is a function of the scaled distance to the blend wall ($\frac{0.1}{0.9}$ in an ethanol-only world) plus an additional term which depends on the quantity of low-ethanol blend sales ($q_{E10}$). The resulting expression will become small if the imposed blend mandate requires a heavy reliance on channel four with $q_{E10}$ sales declining in order to lower the compliance base and $p_{D6}$ exploding. In a very extreme scenario in which motor gasoline markets effectively shut down, we could thus encounter E85 prices that rise with the mandate level. Otherwise,

$$\frac{dp_{E85}}{d\kappa_{TR}} \leq 0 \quad \Leftrightarrow \quad \kappa_{TR} \leq \frac{0.1}{0.9}\left(\frac{p_E}{\epsilon_{S_E} p_{D6}} + 1\right) + \frac{q_{E10}}{0.9 p_{D6} B_{D_C}} \tag{1.7}$$

$p_{E85}$ is guaranteed to depend negatively on the prevailing mandate level.

A.5 provides a similar formula for an upper bound on mandate levels which guarantees increasing RIN prices as a function of percentage standards. Our comparative statics results therefore show that under ’normal’ market environments, RIN prices will increase with blend mandate levels while E85 prices decrease. This implies the negative correlation between RIN prices and the price of high-ethanol blends pointed out above. However, under very extreme circumstances, the fourth compliance channel may become so dominant that these natural relationships are no longer guaranteed to hold.

### 1.5.3 Simulation Results

Figure 1.4 shows the impact of an increasing $\kappa_{TR}$ mandate in the market for consumer fuels. Throughout these simulations, we hold the biomass-based diesel mandate fixed at its 2015 level of around 1.5%. The first row of graphs represents quantity changes, while the second row highlights changes in equilibrium prices. It is important to note that these graphs do not have a time component, but rather represent the market outcomes if a given mandate level was imposed in a market environment similar to the U.S. in 2015.

The quantity of E85 jumps up sharply in response to the blend wall[11]. This change of consumption is induced by a drastic reduction in E85 prices which incentivizes FFV drivers to switch to the cheaper fuel. As some consumers

[11] As noted previously, the blend wall occurs in the ethanol-only model whenever $\frac{E}{G}$ becomes larger than 10%, which corresponds to an ethanol-only mandate of $11.\overline{11}\%$. However, in a market with both motor gasoline and diesel, $\kappa_{TR} = \frac{E+BD}{D+G}$ which is why the blend wall now occurs before the salient level of $11.\overline{11}\%$

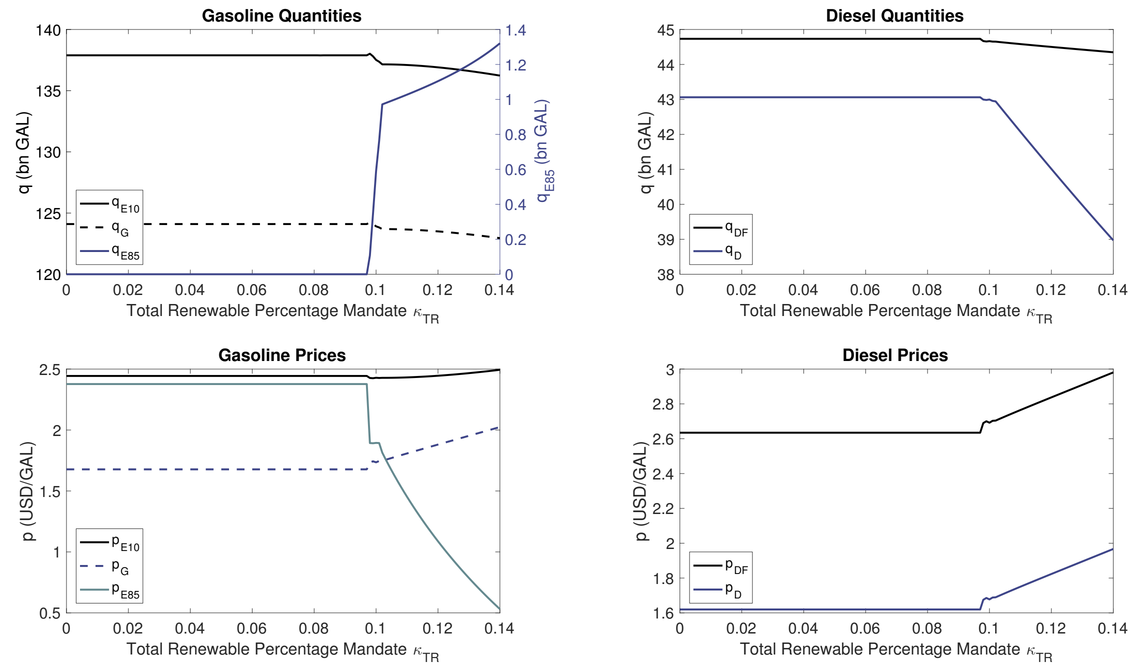


**Figure 1.4:** Price and Quantity Changes for Consumer Fuels. All Prices are Shown in USD/GAL. All Quantities are Shown in bn GAL.

abandon E10 for E85, the E10 price drops, thereby enabling an increase in E10 consumption by conventional vehicle owners. This phenomenon explains the small uptick in E10 consumption in the first panel of Figure 1.4.

The E10/E85 switching corresponds to compliance channel two in action: more ethanol is blended into motor gasoline in order to meet the mandate. Overall motor gasoline use increases as low prices induce more E85 consumption by FFV drivers. We also find clear evidence for the fourth compliance mechanism: as the mandate tightens and the inducement of additional E85 consumption becomes more costly, total motor gasoline consumption begins to decline again as shown in Figure 1.5. Because we are modelling E85 demand using a constant-elasticity function, the expansion in E85 here dominates the curtailment of E10. However, if we were to impose a cap on E85 consumption, we would expect overall motor gasoline use to decline as channel four would begin to dominate once channel two has been exhausted.

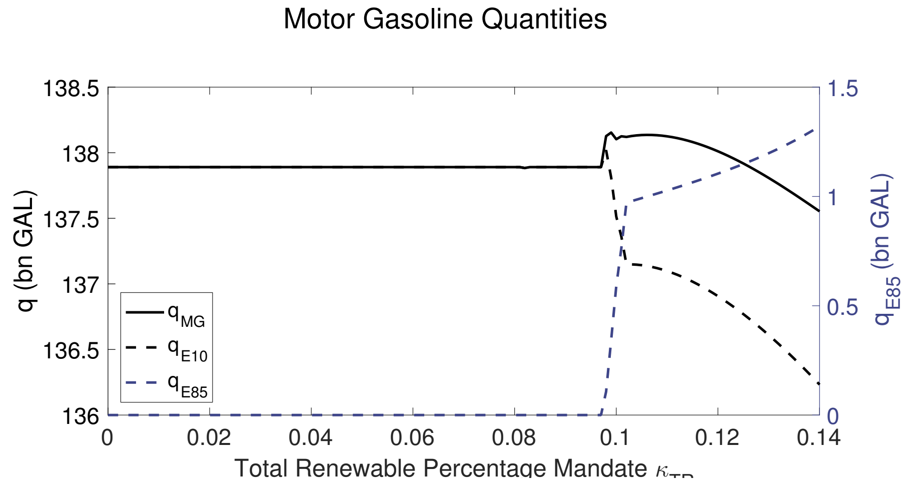


**Figure 1.5:** Quantity Changes in Motor Gasoline (bn GAL)

Looking at the diesel fuel market as shown in the right hand side panels of Figure 1.4, it becomes evident that the dominant compliance channel in our model is an increase in the biodiesel blend ratio, accompanied by increasing diesel fuel prices. The decrease in diesel fuel sales is thus driven by a supply shock in the form of increased marginal costs due to higher biodiesel use.

The important role which biodiesel plays in overcoming the blend wall is also evident in the choice of blend ratios. Our calibrated ethanol supply curve leads to ethanol prices cheap enough to encourage full use in E10, but not cheap enough to also encourage the production of E85 at low mandate levels. The E10 blend ratio is therefore stable at 10% regardless of the percentage blend mandate[12]. The diesel fuel blend on the other hand changes significantly beyond the blend wall, increasing from around 3.7% to 7.8% as shown in Figure 1.6. This change is purely driven by ethanol demand constraints as the BBD mandate level itself remains fixed at 1.5% throughout our simulations.

[12]The free-market blend ratio which refiners would choose in the absence of a blend wall in our model is 12.5%

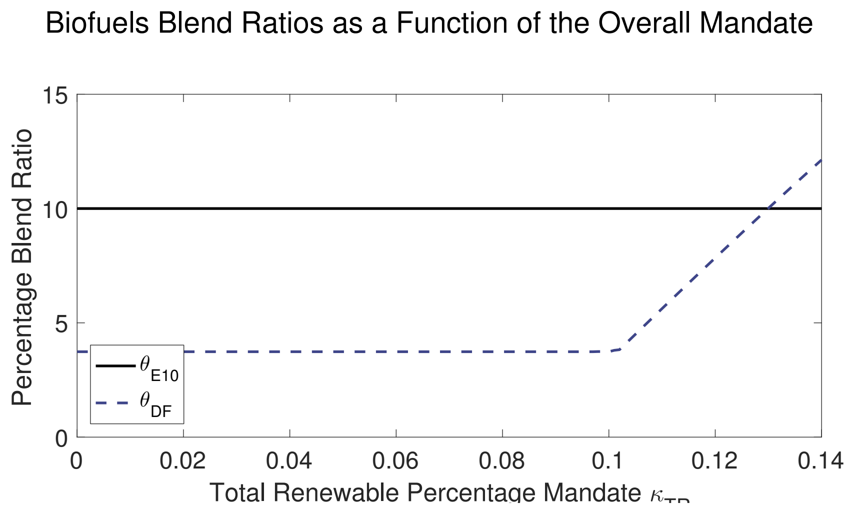


**Figure 1.6:** Changes in Fuel Blend Ratios (Percent)

To emphasize this point, Figure 1.7 highlights the order in which the four compliance channels are activated. Figure 1.8 shows the relative reliance on each of the channels.

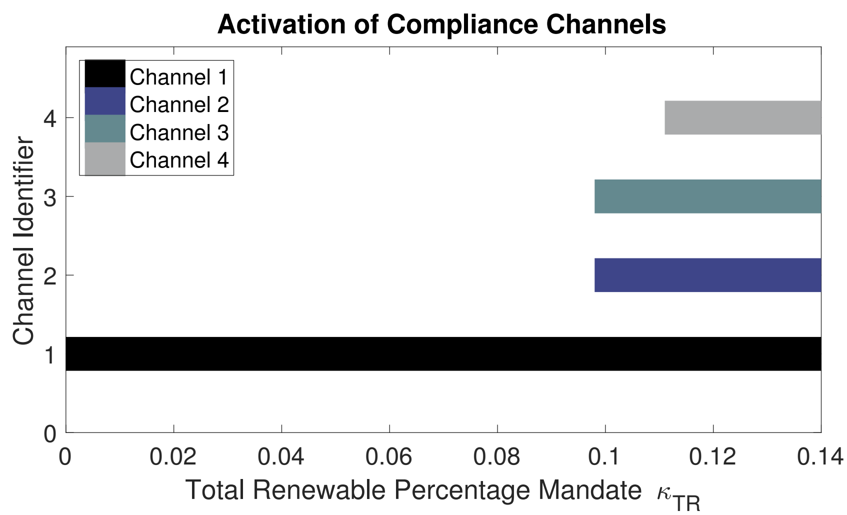


**Figure 1.7:** Order of Activation of the Four Compliance Channels

While E85 quantities jump up once the price discount incentivizes switching by FFV drivers and exhibit a slow rate of growth beyond this point, biodiesel overage ramps up slightly later but increases at an almost constant rate to accommodate the increasing mandates. The fourth compliance channel is used as a measure of last resort and only becomes active at significantly higher mandate levels. Once initiated however, the reduction in the total compliance base also proceeds at a near constant rate. As A.4 shows, the relative activation points of the different channels depend on the prevailing supply and demand elasticities. A.4 also explains how we estimate the impact of channel four for this Figure.

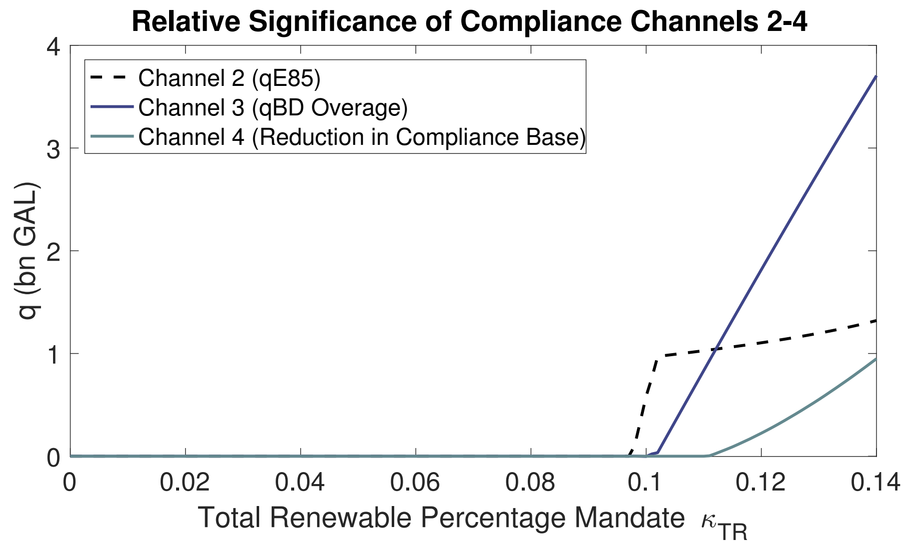


**Figure 1.8:** Relative Significance of Compliance Channels

Finally, Figure 1.9 depicts the changing equilibria in the biofuels market[13]. The increase in ethanol and D6 RIN quantities is effectively capped due to the blend wall. The increase in biodiesel and D4 RIN quantities on the other hand is unconstrained and progresses in line with the increasing biodiesel blend ratio we observed in the previous figure. The difference between biodiesel and D4 quantities is due to the 1.5 equivalence value applied in biodiesel RIN generation.

As noted in section 1.3, the nested mandate structure leads to an implicit

[13]The erratic behavior of D6 RIN quantities below the blend wall is due to over-blending of ethanol at low mandate levels. Given a zero RIN price and no bankability, the chosen quantity in our model oscillates freely as neither blender nor refiner have an incentive to exchange the superfluous RINs. Once the mandate starts binding and RIN prices become positive, the quantities of ethanol and generated D6 RINs in our model converge as expected

pricing relationship between D4 and D6 RINs: the prices must always satisfy $p_{D4} \geq p_{D6}$. This relationship holds with equality when the wider mandate becomes binding and needs to attract overage from the nested sub-category. This effect becomes evident in the second panel of Figure 1.9.

We also observe a sharp increase in biodiesel prices compared to ethanol prices due to the demand constraints in E85. This suggests that the diesel market, and hence the transportation sector, carries a large part of the economic burden imposed by the ethanol consumption bottleneck.

## 1.6 Conclusion

We propose a parsimonious partial equilibrium framework which demonstrates the available compliance channels under current U.S. biofuels policy. Our simulation results provide evidence for a strong reliance on both biodiesel overage and a reduction in low-ethanol motor gasoline blends in order to meet the total renewable mandate given a binding blend wall. Since heavy trucks and trains account for most of the diesel consumption in the U.S., this suggests impor-

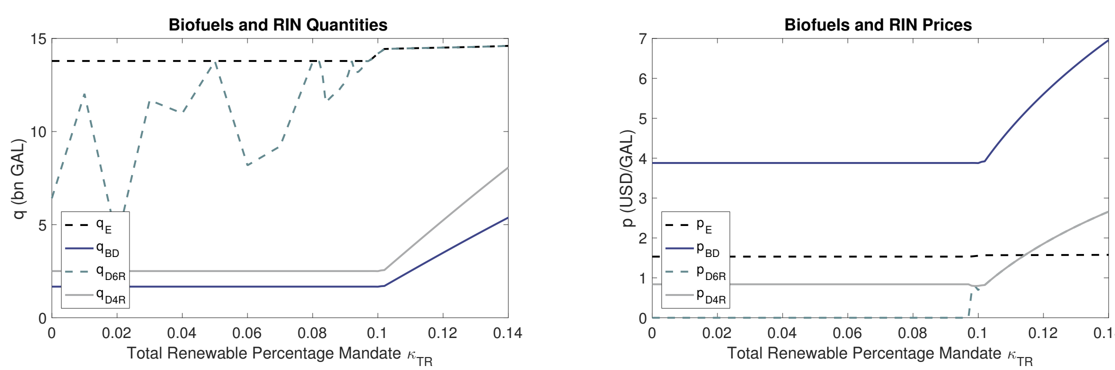


**Figure 1.9:** Price and Quantity Changes for Biofuels. All Prices are Shown in USD/GAL. All Quantities are Shown in bn GAL.

tant general equilibrium ramifications as the higher cost of transportation will likely be passed on in the form of higher consumer prices. In addition, assuming limited biodiesel production capacity in the short term, we expect the fourth compliance channel to play an increasingly important role going forward. This implies significant welfare losses for drivers of cars without flexible-fuel capacity.

We also present a concise formula for the core value of RINs. The value of a RIN in equilibrium is shown to reflect the marginal cost of compensating the blender for employing one additional ethanol-equivalent unit of biofuel. Previous research had often represented RIN prices as the gap between free-market ethanol supply and demand at the mandate level. Our simulation results highlight the important approximation error induced by this notion of RIN rices by showing the strong dependence of implied ethanol demand on the prevailing blend mandate requirements.

Future work will focus on establishing the feasibility of the proposed RFS mandate progression under different infrastructure scenarios. If the technology forcing potential of the RFS2 cannot be realized as expected, it is important to understand which additional policy tools could be employed in order to overcome the current bottleneck in ethanol consumption. Our model provides a convenient starting point to explore the effect of additional incentive structures such as increased biodiesel tax credits.

CHAPTER 2

# WHO WILL PAY FOR INCREASING BIOFUEL MANDATES? INCIDENCE OF THE RENEWABLE FUEL STANDARD GIVEN A BINDING BLEND WALL

## 2.1 Abstract

We show that the cost of increasing biofuel mandates given a binding ethanol blend wall falls disproportionately on diesel fuel consumers. The extent of the burden on diesel fuel consumers is explained neither by their relatively more inelastic demand nor by blenders seeking to capitalize on the biodiesel tax credit. Relaxing the blend wall constraint by increasing the potential demand for high-ethanol blends is the only effective lever to insulate diesel fuel drivers from the one-sided welfare impacts of rising mandate levels. The independent effects of the nested mandate structure and the joint compliance base under the Renewable Fuel Standard (RFS) generate the link between motor gasoline and diesel fuel markets. Our results highlight the importance of evaluating the incidence of the RFS in a holistic framework taking both ethanol and biodiesel into account.

## 2.2 Introduction

The U.S. Renewable Fuel Standard (RFS) was introduced with the joint objectives of (a) protecting against rising fossil fuel prices; (b) reducing emissions from the transportation sector; (c) promoting energy security by reducing the dependency on fossil fuel imports; and (d) increasing farm income and creating new jobs ([69]; [52]). The recent shale gas boom and the consequent drop in oil prices have somewhat attenuated the concerns around the reliance on foreign fuel and its price. Findings regarding the RFS' performance with respect to the other two goals as well as its overall welfare impact have been mixed.

Work in progress by [54] notes that the RFS constitutes an expensive mechanism to reduce carbon emissions in the short run, with an estimated cost of USD 300 per metric ton of CO2 avoided. [56] find a net positive long-term welfare impact of the RFS at current mandate levels thanks to beneficial terms of trade effects, but their analysis also suggests significant distributional impacts, with farm sector surplus increasing while consumers experience welfare losses.

We contribute to this literature by focusing on the effects of the RFS on fuel consumer welfare in the short term. Building on work by [64], we analyze the impact of ethanol blend mandates taking short-term rigidities into account. In particular, we expand on on the partial equilibrium model proposed in [43] which accounts for the effects of the ethanol blend wall given constrained demand for high-ethanol blends and extends the [64] framework to capture the nested mandate structure of the RFS.[1] We find that the welfare effects of rising mandate levels are most severely felt by diesel fuel consumers.[2]

[1]Due to the corrosive nature of ethanol, cars have to be equipped with special tanks and fuel pumps in order to sustain the use of motor gasoline blends with a high ethanol content. As a result, the ethanol share in E10, the most common motor gasoline blend in the U.S., is legally capped at 10% ethanol, a fact often referred to as the ethanol blend wall. While so called flexible fuel vehicles (FFVs) can operate on blends of up to 85% ethanol (referred to as E85), these cars currently represent only a small fraction of the total U.S. fleet. According to Table 40 of the 2016 Annual Energy Outlook there were 18.3mn FFVs in 2015 compared to a total vehicle fleet of 240mn [79]. The number of gas stations able to dispense high-ethanol blends is also limited. Additional blends such as E15 currently capture only a very small share of the consumer fuel market and are therefore omitted here for simplicity.

[2]Note that we restrict our model to a closed economy framework of U.S. biofuels markets in order to maintain the necessary flexibility to explore alternative mandate structures and demand scenarios. Unlike [54] and [56], our model abstracts away from agricultural inputs and

We provide numerical simulation results for a number of alternative market and policy scenarios to systematically investigate the drivers of this result. In particular, we show that most of the burden on diesel fuel consumers can be directly attributed to the ethanol blend wall. We support this result by highlighting that (i) in a model without a blend wall, the effect of rising total renewable mandates is largely borne by motor gasoline consumers; and (ii) neither the more inelastic diesel fuel demand nor the effect of the biodiesel tax credit can explain the extent of the burden placed on diesel fuel consumers.

In addition, we investigate how the particular structure of the RFS contributes to this result. While the nested mandate structure encourages efficient switching between biofuel sources to meet the total renewable fuel mandate, we show that this is not the only factor tying the fate of diesel fuel consumers to the ethanol blend wall. In fact, the loss of diesel fuel consumer welfare is exacerbated in a model without nesting. This can be explained by the additional link between consumer fuels generated by the joint compliance base of the RFS.[3]

Since the main consumers of diesel fuel in the U.S. are trucks and trains[4], the relative burden share between fuel consumers is likely to have significant general equilibrium ramifications. The findings of this analysis therefore have important policy implications concerning the increase in biofuel mandate requirements in the face of continued demand side infrastructure constraints for high-ethanol blends.

---

focuses exclusively on fuel markets. Readers interested in terms of trade effects, endogenous oil prices and food price impacts of the RFS are referred to their work. More details about the infrastructure constraints limiting the sale of high-ethanol blends can be found in [62].

[3]All percentage blend mandate requirements under the RFS apply to the total imports or production of gasoline *and* diesel by obligated parties, see next section for details.

[4]According to tables 40 and 50 of [18] less than 1% of conventional cars and light trucks for personal use in 2015 consumed diesel, while 77% of the U.S. truck fleet relied on diesel as their main fuel source.

This article is organized as follows. The next section provides an introduction to the RFS and discusses the particular policy aspects which shape the welfare results we uncover. The third section introduces our model and its adaptations to alternative market scenarios. Section four summarizes data sources and outlines our calibration approach. The penultimate section provides numerical simulation results and explores the drivers behind the strong reliance on diesel fuel consumers to overcome the blend wall. The last section concludes.

## 2.3 Renewable Fuel Standard

The Renewable Fuel Standards of 2005 (RFS1) and 2007 (RFS2), introduced as part of the Energy Policy Act (EPAct) and the Energy Independence and Security Act (EISA) respectively, contain annual volumetric targets for biofuel use in the transportation sector. In order to operationalize these targets, they are converted into percentage blend mandates by dividing the required amount of renewable fuel for the year ahead by the total forecast amount of gasoline and diesel consumption. The forecasts are obtained from the November issue of the Short-Term Energy Outlook[5] preceding the mandate year.

Each obligated party's renewable volume obligation (RVO) is calculated by applying the percentage blend mandate requirements to their total imports or production of gasoline and diesel. The sum of gasoline and diesel therefore represents the *joint compliance base* of the RFS2 mandates. As noted in the 2010 final rule, the EPA considered creating separate standards for gasoline and diesel, but deemed this alternative mandate structure unnecessarily more complex to

[5] `http://www.eia.gov/forecasts/steo/outlook.cfm`

implement ([20], page 14716). Blend mandates are thus given by the fraction of volumetric mandates divided by the joint compliance base.

The RFS2 encompasses four distinct categories of biofuel requirements which are defined as a function of both the type of biofuel and the relative amount of greenhouse gas (GHG) savings achieved compared to the petroleum fuel being replaced. Table 1.1 introduces the different categories and compares mandate levels for 2015 and 2022. Cellulosic biofuel - produced from cellulose, hemicellulose or lignin - has to achieve GHG reductions of at least 60%. Biomass-based diesel (BBD) and advanced biofuels share a lifecycle GHG emissions reduction target of 50%. BBD is often made from soybean or canola oil while advanced biofuels can be made from any type of renewable biomass except corn starch. The remainder of eligible renewable fuels are required to reduce GHG emissions by 20% or more.

The four compliance categories under the RFS2 are nested to enable the use of more efficient biofuels in GHG terms towards compliance with the larger mandate: The biomass-based diesel and cellulosic biofuel standards are nested under the advanced category which is itself nested within the total renewable mandate [20]. This structure allows for strategic overage from nested categories, if desirable, based on the relative cost of compliance. For example, additional units of BBD can be used to meet the advanced and total renewable mandate requirements.[6] Note that the RFS2 does not impose a specific ethanol mandate: in an extreme scenario, cellulosic biofuels and biomass-based diesel could be

[6]Biodiesel not meeting the D4 GHG reduction threshold, but providing sufficient savings compared to the total renewable level to earn D6 RINs instead, can also be used to comply with the total renewable mandate, but is omitted from our analysis for simplicity. According to the EPA Moderated Transaction System (EMTS), 252mn D6 RINs were generated from biodiesel or renewable diesel in 2013 (https://www.epa.gov/fuels-registration-reporting-and-compliance-help/2013-renewable-fuel-standard-data)

used to meet the entire total renewable mandate. As a result, both increased ethanol blending and increased biodiesel blending can help to overcome the ethanol blend wall.[7]

Obligated parties under the RFS2 are refiners and importers of petroleum-based fuels rather than the blenders and distributors who ultimately control the final blends at the pump. As noted in the 2010 final rule, this decision was based on a desire to minimize the number of obligated parties ([20] page 14722). To monitor compliance, financial instruments called Renewable Identification Numbers (RINs) were created: each RIN represents one ethanol-equivalent unit of biofuel blended for transportation. RINs are generated when biofuels are produced or imported and become detached and separately tradable once the biofuel has been blended for use in the transportation sector. Table 1.1 indicates the type of RIN earned per mandate category. At the end of each year, every obligated party has to present a number of RINs commensurate with each of the four mandate requirements. Based on the nested structure, the total renewable mandate can be met by a combination of D3, D4, D5 or D6 RINs.

The nature of the RFS2 and its implications for fuel and agricultural input markets have been a subject of active research: [52], [53] and [80] study the regulatory framework and provide insights on the nature of RINs. **(author?)** [14] and [49] explore the nature of blend mandates and derive their incidence and general market effects. **(author?)** [16] analyze the economic relationships created by biofuel blending and highlight the resulting link between crude oil and agricultural markets.

Recent work has focused particularly on the feasibility and incidence of the

[7]According to [34], the D4 / D6 price relationship indicates that biodiesel was the marginal gallon for compliance with the overall mandate in late 2015 and early 2016

RFS2 given the ethanol blend wall: [62] provide estimates of potential demand for high-ethanol blends (E85) given current infrastructure constraints, which [64] integrate into an ethanol-only model accounting for short-term market rigidities. Work by [41], [45], [65] and many others analyzes the pass-through of RIN prices to fuel consumers and fuel producers: While [41] find evidence of limited pass-through of RIN prices to end consumers in the form of an E10-E85 discount, results in [65] and [45] suggest near complete pass-through to rack and retail prices respectively. The difference in findings might be due to the more disaggregated data sources used in the latter papers, or the inclusion of a greater number of lagged wholesale and component prices. [43] study the available channels of mandate compliance under the RFS2 and discuss the difference between the core value of RINs and its often invoked interpretation as the gap between free-market ethanol supply and demand at the volumetric mandate level. A working paper by [47] explores the dynamic nature of the mandates: since RINs are bankable and borrowable across compliance years subject to constraints, blenders face an inter-temporal arbitrage opportunity, adding a speculative component to RIN prices.

## 2.4 Model

We build on the partial equilibrium model of U.S. biofuel markets introduced in [43]. The static, closed economy model is based on a representative blender and refiner who operate in a non-integrated setting[8] and trade D4 and D6 RINs for

[8]This assumption is consistent with the majority of active market participants as shown by [58]: 40% of US refiners in 2015 were neither integrated with blenders nor retailers. In addition, [64] show that the incentives facing integrated blenders and refiners are the same as the incentives of a non-integrated blender trading RINs with a separate refiner.

compliance. D3 and D5 RINs are omitted for simplicity.[9] Since we are interested in the short-term welfare dynamics associated with the ethanol blend wall, we assume that blenders and refiners respond to the annual blend mandate requirements under the RFS.[10] The model captures the nested mandate structure of the RFS2 and allows for strategic overage: the refiner can choose to retire more D4 RINs than required under the percentage BBD mandate ($\kappa_{BBD}$) to meet the overall total renewable mandate ($\kappa_{TR}$).

Note that our treatment of the compliance obligation as a direct constraint on the refiner's profit maximization problem differs from [56] and [54], who instead introduce the blend mandates as a market clearing constraint. Our model therefore most closely follows the actual incentive structure of the RFS. In addition, we allow for joint refining of gasoline and diesel, while [54] introduces a separate blender and refiner for each fuel type.

Our model has an annual time horizon and does not allow for uncertainty or inter-temporal considerations such as the banking and borrowing of RINs. Simulated RIN prices therefore only capture the core value of RINs and do not reflect any inherent option value. For the sake of parsimony, the cellulosic and advanced mandate categories are not explicitly modeled.[11] In addition, we assume a fixed, exogenous price of oil. Throughout this paper, quantities, prices

[9] D5 RINs are largely generated from sugarcane ethanol imports from Brazil, which are not captured in our closed economy setting. D3 RINs are derived from cellulosic biofuels. Since growth in this sector has lagged significantly behind expectations, the EPA has provided waiver credits in recent years which may be purchased by obligated parties instead of obtaining D3 RINs from blenders. Both categories are small relative to the number of D4 and D6 RINs produced: As shown in [61], estimated generation net of exports and retirements for non-compliance purposes of D4 and D6 RINs in 2015 was at 2.6bn and 14.3bn RINs respectively, compared to only around .1bn RINs each from the D3 and D5 category.

[10][56] study long term welfare impacts of the RFS assuming sufficient infrastructure investments to make increased ethanol blending feasible, and therefore directly impose the volumetric mandate schedule instead.

[11]See footnote 9.

and blend ratios are represented by the letters $q$, $p$ and $\theta$ respectively. The subscripts $S$, $D$ and $C$ distinguish between supply, demand and cost parameters. Product types are shown in (double) subscripts, while superscripts denote the refiner and blender ($R$, $B$).

### 2.4.1 The Refiner's Problem

The representative refiner maximizes his profit by selling gasoline and diesel net of production costs $C^R(q_G, q_D)$ which capture both input and operational costs. To reflect the reality of joint production in the refining process, we model the refiner cost function as jointly CES as shown in equation 2.1 below. For $\epsilon_{C_R} < 0$, this functional form captures economies of scope in refining while ensuring convexity of the cost function in gasoline and diesel. Note that this specification does not account for a link to endogenous oil prices. Interested readers are referred to the analysis in [56] which takes input markets into account.

$$\text{Refiner:} \qquad C^R(q_G, q_D) = \left( A^R_{C_G} q_G^{\frac{1}{\epsilon_{C_R}}} + A^R_{C_D} q_D^{\frac{1}{\epsilon_{C_R}}} \right)^{\epsilon_{C_R}} \tag{2.1}$$

The refiner purchases RINs from the blender for compliance. The quantity of RINs chosen by the refiner is denoted by $(q^R_{D4}, q^R_{D6})$. Throughout this paper, we will refer to petroleum-based diesel and gasoline simply as diesel ($D$) and gasoline ($G$), while the final blends at the pump including biofuels are designated as diesel fuel ($DF$) and motor gasoline ($MG$). Since refiners are the obligated party under the RFS2, the profit maximization is subject to the constraint of meeting the nested mandate requirement by providing a sufficient number of D4 and D6 RINs against the total compliance base of petroleum fuel sales ($q_G + q_D$).[12] Consistent with the RFS2, RIN quantities and percentage mandates are represented in ethanol equivalent terms.

$$
\begin{aligned}
\max_{\{q_G,\, q_D,\, q^R_{D4},\, q^R_{D6}\}} \quad & \Pi^R = p_G q_G + p_D q_D - C^R(q_G, q_D) - p_{D4} q^R_{D4} - p_{D6} q^R_{D6} \\
s.t. \quad & q^R_{D4} \geq \kappa_{BBD}(q_G + q_D) \\
and \quad & q^R_{D4} + q^R_{D6} \geq \kappa_{TR}(q_G + q_D)
\end{aligned}
\tag{2.2}
$$

We allow for strategic overage by adding the mandate requirements as inequality constraints to the refiner's problem rather than pre-specifying the chosen compliance strategy in the form of a RIN bundle (see for example [77]): refiners can choose to over-comply with the BBD mandate and use the resulting excess of D4 RINs towards compliance with the residual total renewable mandate. Note that it is easy to transform the refiner's problem to consider a non-nested mandate: The refiner's constraints in this case no longer allow for the retirement of D4 RINs towards the total renewable mandate requirement. Refiners therefore face an effective ethanol mandate requirement of ($\kappa_{TR} - \kappa_{BBD}$). The corresponding refiner constraints are shown in equation 2.3.

$$
\begin{aligned}
& q^R_{D4} \geq \kappa_{BBD}(q_G + q_D) \\
& q^R_{D6} \geq (\kappa_{TR} - \kappa_{BBD})(q_G + q_D)
\end{aligned}
\tag{2.3}
$$

[12] See the discussion about the nested mandate structure under the RFS in the previous section.

Since the EPA's reasons for introducing nested mandates and a joint compliance base are independent, we maintain the assumption of a joint compliance base of $(q_G + q_D)$ in our model without nesting.

### 2.4.2 The Blender's Problem

The blender purchases gasoline and diesel as well as ethanol ($E$) and biodiesel ($BD$) as inputs to the blending of motor gasoline and diesel fuel. For simplicity, we only consider two distinct types of motor gasoline: E10, with a blend ratio of up to 10%, and E85, which we assume to have a constant blend ratio of 74% ethanol in line with the average blend assumed by the EPA and used in the literature (e.g. [41]).[13] The blender can endogenously determine the blend ratios in E10 and diesel fuel $(\theta_{E10}, \theta_{DF})$, but $\theta_{E10}$ is capped at 10% in the short run by the ethanol blend wall. The blender incurs separate blending costs for each fuel. Given the short-term nature of our model, we assume constant marginal costs of blending for E10, E85 and diesel fuel.

The blender's revenue is based on his fuel sales net of taxes $(t_G, t_D)$, his sale of RINs to the refiner, as well as the one dollar biodiesel tax credit which he earns on the amount of biodiesel blended $(tc_{BD})$. The blender tax credit was extended through December 2016.[14] The two equality constraints of the blender reflect the process of RIN generation by detaching them from the biofuels used for blending. Recall that biodiesel has a higher energy value than ethanol and that RINs are measured in ethanol equivalent terms. We therefore apply an

[13] In practice, some gas stations also offer ethanol-free motor gasoline (E0), as well as E15 which contains up to 15% ethanol, is approved for use in models newer than 2001, but does not meet some car manufacturer warranties. Note that an increase in the E10 blend ratio could also be viewed as a reduction in E0 sales.

[14] House of Representatives Bill 2029, Section 185

equivalence value of 1.5 to transform the amount of biodiesel blended into the available amount of D4 RINs.

$$
\begin{aligned}
\max_{\substack{\{q_{E10},\ q_{E85},\ \theta_{E10} \\ q_{DF},\ q^B_{D4},\ q^B_{D6},\ \theta_{DF}\}}} \quad & \Pi^B = q_{E10}(p_{E10} - t_{MG}) + q_{E85}(p_{E85} - t_{MG}) + q_{DF}(p_{DF} - t_{DF}) \\
& \quad + p_{D6}q^B_{D6} + p_{D4}q^B_{D4} + \theta_{DF}q_{DF}tc_{BD} \\
& \quad - ((1 - \theta_{E10})q_{E10} + 0.26q_{E85})p_G - (\theta_{E10}q_{E10} + 0.74q_{E85})p_E \\
& \quad - (1 - \theta_{DF})q_{DF}p_D - \theta_{DF}q_{DF}p_{BD} \\
& \quad - C^B_{MG}(q_{E10}, q_{E85}) - C^B_{DF}(q_{DF}) \\
s.t. \quad & q^B_{D4} = 1.5\theta_{DF}q_{DF} \\
and \quad & q^B_{D6} = \theta_{E10}q_{E10} + 0.74q_{E85} \\
and \quad & \theta_{E10} \leq 0.1
\end{aligned}
\tag{2.4}
$$

If we want to simulate a long term model without a binding blend wall, it suffices to drop the blend wall constraint of our baseline model and consider a single type of motor gasoline ($q_{MG}$) with a freely determined blend ratio ($\theta_{MG}$). This scenario is represented in equations 2.5. The two equality constraints once again represent the process of RIN generation by detaching them from the biofuels used for blending.

$$
\begin{aligned}
\max_{\substack{\{q_{MG},\ \theta_{MG} \\ q_{DF},\ q^B_{D4},\ q^B_{D6},\ \theta_{DF}\}}} \quad & q_{MG}(p_{MG} - t_G) + q_{DF}(p_{DF} - t_D) \\
& + p_{D6}q^B_{D6} + p_{D4}q^B_{D4} + \theta_{DF}q_{DF}tc_{BD} \\
& - (1-\theta_{MG})q_{MG}p_G - \theta_{MG}q_{MG}p_E - C^B_{MG}(q_{MG}) \\
& - (1-\theta_{DF})q_{DF}p_D - \theta_{DF}q_{DF}p_{BD} - C^B_{DF}(q_{DF}) \\
s.t. \quad & q^B_{D4} \leq 1.5\theta_{DF}q_{DF} \\
and \quad & q^B_{D6} \leq \theta_{MG}q_{MG}
\end{aligned} \tag{2.5}
$$

All other agents in the model are represented by their respective supply and demand curves. For simplicity, ethanol and biodiesel supply as well as the demand for diesel fuel are assumed to be of the constant elasticity form $q = Ap^{\epsilon}$ where $\epsilon$ is the relevant supply or demand elasticity and $A$ represents a scaling parameter.

### 2.4.3 Motor Gasoline Demand

Our E10 and E85 demand functions are local quadratic regression approximations of the fuel demand relationships derived in [62] and employed in [64], who kindly shared their demand data with us. [62] use locational data of FFV ownership in combination with the exact location of fuel stations in order to account for supply side infrastructure constraints such as (i) limited E85 pump capacities at the retail level; as well as (ii) the effort cost of finding E85 stations given the relative location of FFV vehicles in the U.S. In addition to explicitly modeling search costs, they also allow for heterogeneous preferences for E10 and E85, based for example on environmental concerns or the relative range of the two fuels ([2], [71]). The demand for motor fuel in [62] is assumed to be of a constant elasticity form with a short-run demand elasticity of -0.25. To arrive at aggregate E85 demand estimates, Pouliot and Babcock integrate over consumer types and distances to the nearest pump. In our simulation results, we show that the constraints on E85 demand play an important role for the relative burden share between motor gasoline and diesel fuel consumers. Ignoring this effect leads to an understatement of the impact of the blend wall at high mandate levels.

The functional form employed in [54] also accounts for heterogeneous preferences for the two ethanol blends, but does not capture the demand constraints implied by the interplay of limited supply side infrastructure and effort costs of finding E85. However, we show in an online appendix that this demand specification performs similarly for total renewable mandate levels of up to 12%. [56], like [43], assume that consumers value fuel purely based on miles traveled [15]. In addition, they assume a linear functional form for fuel demand. We show that the first assumption does not significantly change simulation results at high mandate levels using the conditional demand specification employed in [43]: demand functions derived from heterogeneous preferences behave like a smoothed version of the demand functions implied by discrete price-based switching (see online appendix). The linearity assumption on E85 demand on the other hand likely underestimates the stringency of the blend wall in the short term. The scenarios presented in [56] at high mandate levels should therefore be viewed as long-term analyses, once the ethanol blend wall has shifted sufficiently to accommodate mandate requirements.

The online appendix provides a graphical comparison of the three functional forms of E85 demand employed in [43], [54] and [64].

In total, our baseline model consists of 25 equations representing first order, complementary slackness and market clearing conditions. A complete list of equations is provided in the online appendix. We will now discuss our data sources and provide details on the calibration assumptions we make in our simulations. We will then present the results of our partial equilibrium analysis at varying mandate levels.

## 2.5 Data

Our model is calibrated to annual data from 2015 by finding the parameter values that minimize the error on the first order conditions for the blender and refiner. Table 2.1 lists the data sources and provides an overview of the average prices and annual quantities we employ. We use national average federal and state tax rates provided by the EIA.[15] To calibrate the constant elasticity supply and demand functions, we rely on elasticity estimates from the literature: we assume ethanol and biodiesel supply elasticities of $\epsilon_{S_E} = 2$ [50] and $\epsilon_{S_{BD}} = 2$ ([33] at 1.7bn gal.) respectively, and a diesel fuel demand elasticity of $\epsilon_{D_{DF}} = -0.07$ [13]. The refiner cost function is assumed to have an exponent of $\epsilon_{C_R} = 0.75$ to ensure convexity. A sensitivity analysis with respect to this parameter revealed that simulation results do not change significantly for values between nonjoint production ($\epsilon_{C_R} = 1$) and our assumed value of 0.75.

Note that we maintain E85 demand estimates exactly as provided in [64],

[15] $t_G$ = 44.89 cents/gal. of which 18.4 cents/gal. are federal; $t_D$ = 51.64 cents/gal. of which 24.4 cents/gal. are federal (`http://www.eia.gov/tools/faqs/faq.cfm?id=10&t=10`)

but recalibrate E10 demand to match our observations for 2015 in table 2.1. To do so, we uniformly shift the whole E10 demand surface up by a constant.

In order to calculate the amount of carbon savings from biofuels, we use lifecycle analysis emission factors provided by the EPA.[16] These emission factors account for all 'well-to-wheel' and 'field-to-wheel' emissions including feedstock production and the resulting land use change, fuel production and distribution as well as emissions generated during the combustion of the finished fuel. The online appendix contains detailed information on the emission factors by fuel type. We use the 2015 central social cost of carbon estimate of USD 36 per metric ton of CO2 to convert emission reductions into the value of carbon savings.[17]

The online appendix contains calibration results for model parameters such as the supply and cost function multipliers. Table 2.1 provides a comparison of simulation results at 2015 mandate levels to the observed market data for the same year. Overall, our simulation results provide a good fit for the 2015 data. The main difference in simulated versus actual values is the higher E85 demand under our baseline model. We believe that this is due to inconsistencies in data reporting between the sources for E85 demand in [62] and our paper. Our E85 quantity inputs rely on the sum of (i) EIA data on Weekly Refiner & Blender Net Production of conventional gasoline with an ethanol content of more than 55% and (ii) EIA data on motor gasoline supply by renewable fuel and oxygenate plants from the Petroleum & Other Liquids report.[18] [62] derive the demand for

[16] https://www.epa.gov/fuels-registration-reporting-and-compliance-help/lifecycle-greenhouse-gas-results

[17] In 2007 dollars assuming a central discount rate of 3%. Impact analysis revised July 2015: https://www3.epa.gov/climatechange/Downloads/EPAactivities/social-cost-carbon.pdf

[18] This is in line with the measurement of E85 quantities in industry reports (e.g. [3]).

E85 using locational data on FFV ownership and retail fuel stations.

In addition, our model yields a lower D6 RIN price than observed in practice. This is in part due to the fact that our model does not capture any option value driven by the dynamic nature of the RFS, and in part due to the difference between the theoretical blend wall level of 10% imposed in our model and the "effective blend wall": due to factors such as small refinery exemptions, the effective blend wall is often assumed to be slightly lower than its theoretical level (see e.g. [16], p. 180; [57], Appendix N). In our model, we obtain a simulated 2015 D6 RIN price of USD 0.59 at an effective blend wall level of 9.8%, suggesting that this may be closer to the prevailing 2015 blend wall level in practice.

**Table 2.1:** Data Sources and Empirical Verification

| | Description | Simulated | 2015 | Units | Source |
|---|---|---|---|---|---|
| | Percentage Standards (%) | | | | |
| $\kappa_{TR}$ | Renewable Fuel | 9.5% | 9.52% | Percent | EPA |
| $\kappa_{BBD}$ | BBD | 1.5% | 1.49% | Percent | EPA |
| | Blend Ratios (%) | | | | |
| $\theta_{E10}$ | E10 Ethanol Content | 10% | 9.9% | Percent | Calculated |
| $\theta_{DF}$ | Biodiesel Content | 3.2% | 3.0% | Percent | Calculated |
| | Quantities (bn gal.) | | | | |
| $q_{E10}$ | E10 Consumption | 140.87 | 140.66 | bn gal. | Calculated |
| $q_{E85}$ | E85 Consumption | 0.51 | 0.07 | bn gal. | EIA |
| $q_{DF}$ | Diesel Fuel in Transport. | 55.67 | 55.65 | bn gal. | EIA |
| $q_G$ | Gasoline in Transport. | 126.91 | 126.78 | bn gal. | EIA |
| $q_D$ | Diesel in Transport. | 53.86 | 53.96 | bn gal. | EIA |
| $q_E$ | Ethanol in Transport. | 14.46 | 13.95 | bn gal. | EIA |
| $q_{BD}$ | Biodiesel in Transport. | 1.81 | 1.69 | bn gal. | EIA |
| | Prices (USD/gal.) | | | | |
| $p_{E10}$ | Regular Motor Gasoline, All Areas | 2.44 | 2.43 | USD/gal. | EIA |
| $p_{E85}$ | E85 Prices | 2.17 | 1.96 | USD/gal. | e85prices.com |
| $p_{DF}$ | On-Highway Diesel Fuel Price | 2.69 | 2.71 | USD/gal. | EIA |
| $p_G$ | Refiner Price of MG for Resale | 1.71 | 1.73 | USD/gal. | EIA |
| $p_D$ | Refiner Price of No. 2 DF for Resale | 1.64 | 1.67 | USD/gal. | EIA |
| $p_E$ | Ethanol Rack, FOB Omaha | 1.64 | 1.61 | USD/gal. | NEO |
| $p_{BD}$ | U.S. Retail Fuel Prices B99/B100 | 3.74 | 3.62 | USD/gal. | AFDC |
| $p_{D4}$ | BBD RINs (D4) | 0.73 | 0.75 | USD/RIN | OPIS |
| $p_{D6}$ | Ethanol RINs (D6) | 0.29 | 0.55 | USD/RIN | OPIS |

Notes: Numerical simulations were generated under the baseline model holding the total renewable and BBD blend mandates constant at their 2015 level of $\kappa_{TR} = 9.5\%$ and $\kappa_{BBD} = 1.5\%$ respectively.

We now turn to the results of our fuel market simulations.

## 2.6 Simulation Results

Figure 2.1 presents the evolution of welfare changes under our baseline model as a function of total renewable blend mandate levels. Welfare changes are calculated with respect to the "no-mandate case" ($\kappa_{TR} = \kappa_{BBD} = 0$) under the baseline model. Note that the figure does not have a time component. Instead, the graph reflects the simulated results given the U.S. market environment in 2015 under varying assumed levels of the total renewable mandate, $\kappa_{TR}$, while holding the BBD mandate, $\kappa_{BBD}$, fixed at its 2015 level of 1.5%. We thus provide market and welfare outcomes holding all exogenous components of the model except for the total renewable mandate level constant.

The second panel of this figure highlights the unequal effects of rising mandate levels on diesel fuel and motor gasoline consumers.[19] The x-axis in this graph reflects levels of $\kappa_{TR}$ varying from 8% to 12%. We have omitted simulation results at lower mandate levels for convenience since they are similar to results at the 8% mandate level. The blend wall generates sharp losses in diesel fuel consumer surplus as the biodiesel blend ratio increases, leading to a higher price at the pump and discouraging demand. Motor gasoline consumers on the other hand temporarily benefit from lower prices as blenders try to encourage E85 sales by decreasing its price, leading to higher motor gasoline sales overall.

It is important to note that the ethanol blend wall does not become binding

[19] Since E10 and E85 are substitutes whose demands each depend on the two prices, $(p_{E10}, p_{E85})$, the consumer surplus calculations based on Marshallian demand functions are subject to path dependence. To minimize the approximation error, we calculate consumer surplus as a sequence of line-integrals along linear paths between the successive equilibrium results for different mandate levels. [28] estimate a publication bias-corrected income elasticity of motor gasoline demand of 0.1. This low estimate implies that Willig's condition, which guarantees a small approximation error from using consumer surplus measures under certain conditions, is satisfied in the context of motor gasoline demand (see for example [37], p. 138)

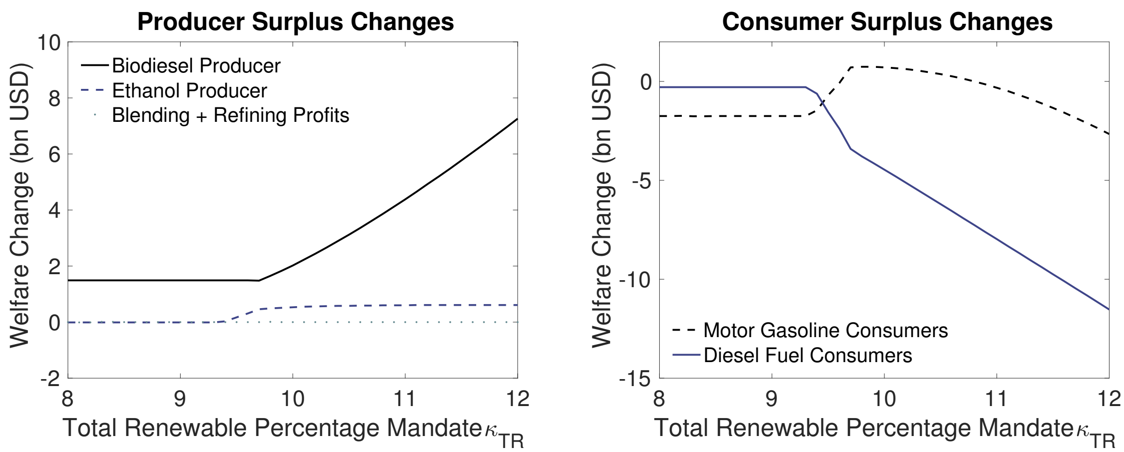


**Figure 2.1:** Welfare results under the baseline model for total renewable mandate levels up to 12% (bn USD)

Notes: Numerical simulations hold the BBD blend mandate constant at its 2015 level of $\kappa_{BBD} = 1.5\%$ while varying the total renewable blend mandate level $\kappa_{TR}$. Welfare results are calculated as welfare changes compared to the no-mandate case ($\kappa_{TR} = \kappa_{BBD} = 0$).

at a particular total renewable mandate level. As emphasized previously, the RFS2 does not mandate specific amounts of ethanol use. Rather, ethanol is used to fill the gap between the BBD and total renewable mandate not met through BBD overage. In addition, the joint compliance base implies that the amount of ethanol blended is measured against the sum of petroleum gasoline and diesel, rather than gasoline alone. The blend wall on the other hand is a function of the amount of ethanol relative to gasoline. This means that the total renewable mandate level at which the blend wall starts binding is endogenous. In our numerical simulation results, we therefore rely on the point at which D6 RIN prices become non-zero as a reliable indication of when the blend wall is reached.

Throughout this section, we will refer to table 2.2 which provides a comparison of simulated market and welfare outcomes under the baseline model to results under the alternative market and policy frameworks we explore. The table contains results for the no-mandate case ($\kappa_{TR} = \kappa_{BBD} = 0$) as well as snapshots of the numerical simulations for the other scenarios at the 2015 total renewable mandate level of $\kappa_{TR} = 9.5\%$ and a hypothetical, increased total renewable mandate level of $\kappa_{TR} = 11.5\%$.

As mentioned previously, we argue that the welfare loss of diesel fuel consumers is largely attributable to the effects of the ethanol blend wall. We show that (i) in a model without a blend wall, the effect of rising total renewable mandates is largely borne by motor gasoline consumers instead; and (ii) neither the more inelastic diesel fuel demand nor the effect of the biodiesel tax credit can explain the extent of the burden placed on diesel fuel consumers. Our simulation results indicate that only an expansion of E85 demand, and hence a relaxation of the ethanol blend wall, can alleviate the consumer surplus losses by diesel fuel consumers. Finally, the last part of this section highlights the dual link between motor gasoline and diesel fuel consumers generated by the particular nature of the RFS2. Both the nested mandate structure and the joint compliance base drive the burden share between consumer types.

### 2.6.1 Effects of the Ethanol Blend Wall

First, we will show that the welfare losses that diesel fuel consumers experience are indeed a consequence of the ethanol blend wall. For this purpose, we compare welfare results from our baseline model to a model without blend wall. The 'No Blend Wall' columns in table 2.2 highlight the corresponding simulation results. As mentioned in previous sections, this alternative model assumes that there is only one type of motor gasoline with freely varying ethanol content.

Under the baseline model, the biodiesel blend ratio increases from 3.2% at

$\kappa_{TR} = 9.5\%$ to 6.6% at $\kappa_{TR} = 11.5\%$. Recall that the required BBD blend mandate itself is held constant at 1.5%, meaning that we are seeing an increase in BBD overage. This indicates that ramping up biodiesel blending is the preferred compliance channel in the presence of a binding blend wall.[20]

Without the blend wall constraint, blenders continue to choose an optimal biodiesel blend ratio of 3.2% at $\kappa_{TR} = 11.5\%$, and instead choose to ramp up ethanol blending to 12.6%. As a result, diesel fuel consumers experience a net welfare loss of only USD -1.6bn at the higher mandate level, compared to a welfare loss of USD -9.8bn under the baseline scenario. The blend wall thus clearly represents the bulk of diesel fuel consumer surplus losses.

The first panel of figure 2.2 provides a graphical representation of this phenomenon. The black line represents diesel fuel consumer surplus changes relative to the no-mandate case under the baseline model. The light grey line represents the same in a world without a binding blend wall. As this figure shows, diesel fuel consumer surplus decreases only marginally in a world without blend wall.

The total deadweight cost of the blend wall at 11.5% mandate levels is USD -6.8bn.

[20]See [43] for an analysis of the four compliance channels available under the RFS2, and their relative importance at different mandate levels.

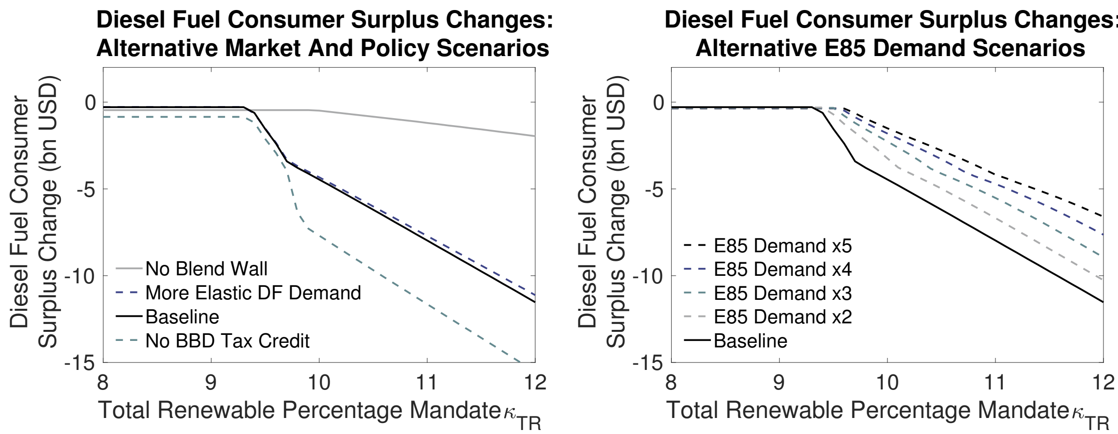


**Figure 2.2:** Diesel fuel consumer surplus changes under the baseline model compared to alternative market and policy scenarios (bn USD)

Notes: Numerical simulations hold the BBD blend mandate constant at its 2015 level of $\kappa_{BBD} = 1.5\%$ while varying the total renewable blend mandate level $\kappa_{TR}$. Welfare results are calculated as welfare changes compared to the no-mandate case ($\kappa_{TR} = \kappa_{BBD} = 0$) under the baseline model.

### 2.6.2 Ruling out Alternative Explanations for the Relative Burden Share by Diesel Fuel Consumers

There are two possible alternative explanations for why diesel fuel consumers shoulder most of the effect of the ethanol blend wall. First, the more inelastic demand for diesel fuel could make blenders more likely to pass through costs to these consumers.[21] Second, the biodiesel tax credit could add to the relative attractiveness of biodiesel blending compared to larger E85 price discounts. However, we show that neither of these two factors explains the disproportional incidence on diesel fuel consumers.

- *Effect of the Biodiesel Tax Credit*

  The biodiesel tax credit does not add to the imbalance of welfare effects

[21] [35] first proposed the principle of *physical neutrality*, noting that economic incidence is independent of the party physically charged with paying a new tax or cost. The pass-through relationship is instead determined by the relative elasticities, with the more inelastic side of the market carrying more of the economic burden of the tax (see e.g. [81]).

for diesel fuel consumers. Instead, it insulates diesel fuel consumers from even greater losses by subsidizing the relatively more expensive biodiesel being blended. In a world without the tax credit, E85 sales are slightly higher than in the baseline case as the biodiesel / ethanol trade-off shifts marginally towards ethanol. However, this effect is dominated by the increased price of diesel fuel as the biodiesel subsidy disappears. At the 11.5% mandate level, diesel fuel prices are 7 cents higher than in the baseline case despite a similar fuel composition (see table 2.2). The net welfare effect of eliminating the tax credit is roughly unchanged across mandate levels, at around USD -0.9bn. However, this net effect hides additional diesel fuel consumer surplus losses of USD -3.8bn, partly offset by the positive change in government tax revenues.

- *Effect of More Elastic Diesel Fuel Demand*

  We also consider a model in which diesel fuel demand elasticity is increased to be on par with the elasticity of motor gasoline demand. In this case, we choose $\epsilon_{D_{DF}} = -0.25$ as in [62] and obtain a corresponding cost function multiplier of $A_{D_{DF}} = 71.38$ based on calibrations to 2015 data. We find almost no welfare changes at the 9.5% total renewable mandate level relative to the baseline results. At 11.5%, we see a slight increase in the diesel fuel consumer surplus of USD 0.3bn, largely offset by losses by motor gasoline consumers. This suggests that an increased diesel fuel demand elasticity does not significantly alter the trade-off between ethanol and biodiesel use. As table 2.2 suggests, most quantities and prices are unchanged with the exception of a net reduction in diesel fuel. The composition of diesel fuel remains unchanged at 6.6% biodiesel as in the baseline case.

### 2.6.3 Relaxing the Blend Wall Constraint through Additional E85 Demand

Having ruled out the relative elasticity of diesel fuel demand as well as the effect of the biodiesel tax credit as dominant factors determining the diesel fuel consumer surplus loss, we now show that an increase in E85 demand can mitigate the welfare impacts of rising mandates. Note that such an increase effectively makes the ethanol blend wall less binding. We consider the effect of scaling up the level of E85 demand by fixed multipliers while adjusting down the demand for E10 to maintain the net motor gasoline demand levels observed in 2015.

The second panel of figure 2.2 depicts the change in diesel fuel consumer surplus losses as a function of increasing levels of E85 demand. If E85 demand increases fivefold, the diesel fuel consumer surplus loss at the 11.5% mandate level drops sharply to USD -5.3bn. This highlights the importance of reducing the demand-side bottleneck for high-ethanol blends in order to insulate diesel fuel consumers from the effects of the ethanol blend wall.

### 2.6.4 No Nesting

Do the differential welfare impacts disappear if biodiesel overage can no longer be used to meet the total renewable mandate? Our simulation results for a model without nesting show that the link between motor gasoline and diesel fuel consumers is not just an artifact of nesting, but also of the joint compliance base. Due to this dual link, the added flexibility provided by the nested mandate structure actually acts as a net welfare enhancement for diesel fuel consumers and the biofuel sector overall.

As discussed previously, the EPA's reasons for imposing a joint compliance base were distinct from the nested mandate structure choice. When designing our market framework without nesting, we therefore maintain the assumption of a joint gasoline and diesel compliance base. Simulation results for the two models are roughly equal between the baseline and non-nested scenario for mandate levels of up to 9.7%. At 9.8%, diesel fuel consumer surplus losses are 50% more severe in the case of no nesting (USD -5.9bn compared to USD - 3.8bn under the baseline model). For mandate levels of 9.9% or higher, the non-nested model no longer solves.

This difference in diesel fuel consumers surplus can be explained by a reliance on decreasing the overall compliance base to satisfy the mandate requirement[22] combined with the blender's first order condition with respect to the quantity of diesel fuel (equation 2.6).

$$p_{DF} - t_{DF} = \underbrace{\frac{\partial C_B^{DF}}{\partial q_{DF}}}_{\substack{\text{Marginal Cost}\\\text{of Blending}}} + \underbrace{(1-\theta_{DF})p_D + \theta_{DF}(p_{BD} - tc_{BD})}_{\text{Input Costs}} - \underbrace{1.5\theta_{DF}p_{D4}}_{\substack{\text{Value of}\\\text{RINs detached}}} \tag{2.6}$$

In the baseline model, diesel fuel consumers are hurt by rising diesel fuel prices as more of the more costly biodiesel is being blended at higher mandate levels. In the model without nesting, diesel fuel prices increase at an even faster rate, but this increase can no longer be explained by increased biodiesel blending. Instead, the price change is due to the D4 RIN price. Since additional D4 RINs can no longer be used to meet the total renewable mandate, D4 RIN prices

[22]See [43] for a detailed explanation of this compliance channel.

no longer move in line with D6 RIN prices.The overall reduction in the compliance base makes compliance with the BBD mandate easier, leading to a reduction in D4 prices as mandate levels rise. Accordingly, all else equal the blender now has to charge a higher diesel fuel price to maintain equality of marginal benefits (LHS) and marginal costs (RHS).

This highlights that the fate of motor gasoline and diesel fuel consumers are linked through two distinct policy aspects written into the RFS: the joint compliance base and the nested mandate structure.

## 2.7 Concluding Remarks

Using a short term model of current U.S. biofuel policy, we find that under 2015 market conditions, diesel fuel consumers are disproportionately affected by the cost of increasing blend mandates given a binding ethanol blend wall. If 2015 total renewable blend mandates had been 2% higher, diesel fuel consumers would have suffered significant welfare losses while motor gasoline consumers would have remained largely unaffected. Using a series of simulation results under alternative market and policy scenarios, we are able to show that diesel fuel consumers' welfare losses are entirely driven by the ethanol blend wall, and that only an increase in potential demand for high-ethanol blends can effectively alleviate the pressure on this consumer group.

It is important to note that our model is restricted to the perfectly competitive fuel sector of a closed economy in the absence of uncertainty. The EPA's decision to alleviate short-term pressure by cutting 2014-2016 mandate requirements underscores the considerable degree of policy uncertainty affecting the

**Table 2.2:** Market and Welfare Outcomes under Alternative Scenarios

| | Baseline | | | No Tax Credit | | DF Elasticity | | No Blend Wall | |
|---|---|---|---|---|---|---|---|---|---|
| $\kappa_{TR}$ | *No Mandate* | 9.5% | 11.5% | 9.5% | 11.5% | 9.5% | 11.5% | 9.5% | 11.5% |
| | | | | Blend Ratios (%) | | | | | |
| $\theta_{E10}$ | 10% | 10% | 10% | 10% | 10% | 10% | 10% | – | – |
| $\theta_{DF}$ | 1.6% | 3.2% | 6.6% | 3.2% | 6.6% | 3.2% | 6.6% | 3.2% | 3.2% |
| $\theta_{MG}$ | – | – | – | – | – | – | – | 10.8% | 12.6% |
| | | | | Quantities (bn gal.) | | | | | |
| $q_{E10}$ | 141.2 | 140.9 | 140.2 | 140.8 | 140.2 | 140.9 | 140.2 | – | – |
| $q_{E85}$ | 0.2 | 0.5 | 1.3 | 0.5 | 1.4 | 0.5 | 1.3 | – | – |
| $q_{MG}$ | 141.4 | 141.4 | 141.5 | 141.3 | 141.6 | 141.4 | 141.5 | 141.1 | 141.0 |
| $q_{DF}$ | 55.7 | 55.7 | 55.5 | 55.7 | 55.4 | 55.7 | 55.0 | 55.7 | 55.7 |
| $q_G$ | 127.1 | 126.9 | 126.5 | 126.8 | 126.6 | 126.9 | 126.5 | 125.9 | 123.3 |
| $q_D$ | 54.8 | 53.9 | 51.8 | 53.9 | 51.7 | 53.9 | 51.4 | 53.9 | 53.9 |
| $q_E$ | 14.3 | 14.5 | 15.0 | 14.5 | 15.0 | 14.5 | 15.0 | 15.2 | 17.7 |
| $q_{BD}$ | 0.9 | 1.8 | 3.7 | 1.8 | 3.6 | 1.8 | 3.6 | 1.8 | 1.8 |
| | | | | Prices (USD/gal.) | | | | | |
| $p_{E10}$ | 2.43 | 2.44 | 2.45 | 2.45 | 2.45 | 2.44 | 2.45 | – | – |
| $p_{E85}$ | 2.37 | 2.17 | 1.20 | 2.18 | 0.77 | 2.17 | 1.21 | – | – |
| $p_{MG}$ | – | – | – | – | – | – | – | 2.45 | 2.46 |
| $p_{DF}$ | 2.66 | 2.69 | 2.84 | 2.70 | 2.91 | 2.69 | 2.83 | 2.67 | 2.69 |
| $p_G$ | 1.67 | 1.71 | 1.87 | 1.72 | 1.94 | 1.71 | 1.87 | 1.68 | 1.69 |
| $p_D$ | 1.62 | 1.64 | 1.79 | 1.65 | 1.86 | 1.65 | 1.79 | 1.62 | 1.64 |
| $p_E$ | 1.63 | 1.64 | 1.67 | 1.64 | 1.67 | 1.64 | 1.67 | 1.68 | 1.81 |
| $p_{BD}$ | 2.62 | 3.74 | 5.34 | 3.74 | 5.32 | 3.75 | 5.32 | 3.74 | 3.71 |
| $p_{D4}$ | – | 0.73 | 1.70 | 1.39 | 2.30 | 0.74 | 1.69 | 0.74 | 0.71 |
| $p_{D6}$ | – | 0.29 | 1.70 | 0.29 | 2.30 | 0.29 | 1.69 | 0.00 | 0.13 |
| | | | | Producer Surplus (bn USD) | | | | | |
| Ethanol | 0.0 | 0.2 | 0.6 | 0.2 | 0.6 | 0.2 | 0.6 | 0.8 | 3.0 |
| Biodiesel | 0.0 | 1.5 | 5.8 | 1.5 | 5.7 | 1.5 | 5.7 | 1.5 | 1.4 |
| | | | | Consumer Surplus (bn USD) | | | | | |
| MG | 0.0 | -0.7 | -1.3 | -1.9 | -1.1 | -0.7 | -1.5 | -1.5 | -2.7 |
| DF | 0.0 | -1.6 | -9.8 | -2.1 | -13.6 | -1.6 | -9.5 | -0.5 | -1.6 |
| | | | | Government Revenue and Value of Carbon Savings (bn USD) | | | | | |
| Tax Rev. | 0.0 | -1.0 | -2.9 | -0.1 | -0.1 | -0.9 | -3.1 | -1.0 | -1.1 |
| CO2 | 0.0 | 0.3 | 1.0 | 0.4 | 1.0 | 0.3 | 1.2 | 0.6 | 1.1 |
| Total | 0.0 | -1.2 | -6.6 | -2.1 | -7.4 | -1.2 | -6.5 | -0.1 | 0.2 |

Notes: All numerical simulations hold the BBD blend mandate constant at its 2015 level of $\kappa_{BBD} = 1.5\%$ except for the no mandate column under the baseline model which reflects the case $\kappa_{TR} = \kappa_{BBD} = 0$. Welfare results are calculated as welfare changes compared to the no-mandate case under the baseline model.

blender and refiner decision problem. Our analysis should therefore be viewed as a complement to recent work on the dynamic effect of policy shocks on RIN prices, the pass-through of RIN to retail fuel prices, as well as the long-term effects of mandates on input markets and terms of trade.

While [56] find net positive long term effects of the RFS, mostly due to favorable changes in the terms of trade and benefits to the agricultural sector, our results caution that the short term effects of the ethanol blend wall on diesel fuel consumers are substantial. Our results therefore underscore the importance of information campaigns targeted at drivers as well as of E15 and E85 infrastructure projects at the pump and distribution level. Diesel fuel consumer surplus losses are likely to have important general equilibrium ramifications: since heavy trucks and trains account for most of the diesel fuel consumption in the U.S., the increased cost of transportation will likely be passed on in the form of consumer good price inflation. It is therefore becoming increasingly clear that industry and policy makers need to find a joint way forward to keep the mandates both physically and economically feasible. The USDA's commitment of USD 100mn towards industry projects investing in additional E15 and E85 infrastructure under its Biofuel Infrastructure Partnership (BIP), requiring matching contributions from industry partners, may prove to be an important first step in that direction[23].

Our results also highlight the importance of evaluating the incidence of the Renewable Fuel Standard in a holistic framework taking both ethanol and biodiesel into account. While ethanol-only models can add important intuition about the nature of blend mandates, they do not adequately capture the nu-

[23] https://www.fsa.usda.gov/programs-and-services/energy-programs/index

ances of the burden share between consumer groups implied by the dual link generated by the nested mandate structure and the joint compliance base under the RFS.

CHAPTER 3

# TAKING A LOAD OFF: EXPERIMENTAL EVIDENCE OF PREFERENCES FOR CONTROL WITH AN APPLICATION TO RESIDENTIAL ELECTRICITY DEMAND

## 3.1 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.

## 3.2 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 [24, 23, 55]. 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

[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

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 [76, 5, 60, 9].[4] An intrinsic preference for control can lead to under-delegation relative to the control-neutral benchmark (e.g. [5]); reduced motivation and performance by workers in the face of constraints on their effort choices (e.g. [22]); an aversion to unmodifiable algorithms even if they outperform human forecasters (e.g. [17]); and increased non-compliance with treatments if the implementation or communication of the treatment triggers reactance (e.g. [67]). 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 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 ev-

[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 [5]. They 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 [25].

idence 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 [5] 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.

[6][5] report a control premium of 18-26% for the two relevant delegation decisions without conflict of interest; [60] find a relative control premium between 8-15%.

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. [5]). The second is to compare individuals' valuations to their theoretical control-neutral predictions (e.g. [60]). 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 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 3.1.

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

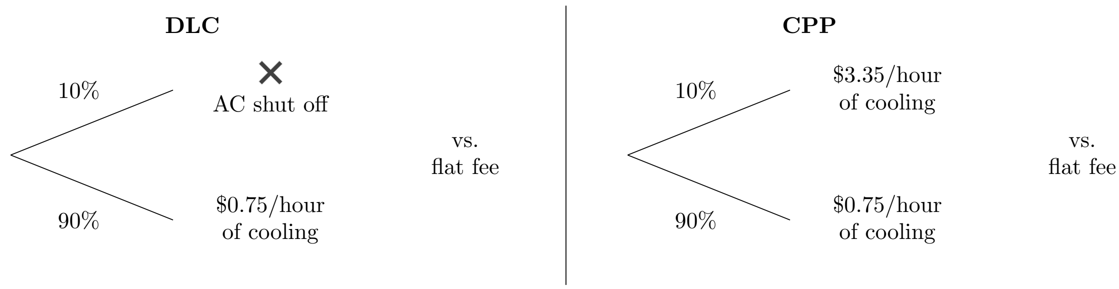


**Figure 3.1:** Stylized illustration of Direct Load Control (DLC) versus Critical Peak Pricing (CPP) contract choice relative to a flat fee for electricity use.

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 also be viewed as a productivity-enhancing tool, participants' WTP for the additional

[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.

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 3.3, 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 3.4 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 3.5, 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 3.6 proposes takeaways for DLC and Section 3.7 concludes.

## 3.3 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 3.2 illustrates the basic structure of control lotteries in this experiment.

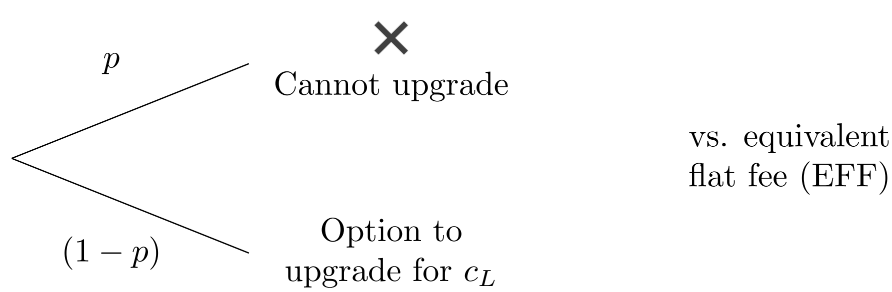


**Figure 3.2:** Control Lottery Illustration

For each control lottery decision, participants are asked to report the *equiv-*

*alent flat fee* (EFF) 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, $\widehat{EFF}_i$:

$$CP_{i,abs} = EFF_i - \widehat{EFF}_i \tag{3.1}$$

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

$$CP_{i,rel} = \frac{EFF_i - \widehat{EFF}_i}{\widehat{EFF}_i} \tag{3.2}$$

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

### 3.3.1 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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s)u_i(k). \tag{3.3}$$

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:

$$u_i(WTP_i) = \sum_{k=0}^{10} \mathbb{P}_i(k|s = 60)u_i(k) - \sum_{k=0}^{10} \mathbb{P}_i(k|s = 30)u_i(k). \tag{3.4}$$

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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-\widehat{EFF}_i) = (1-p)\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-c_L) + p\sum_{k=0}^{10} \mathbb{P}_i(k|s=30)u_i(k) \tag{3.5}$$

As shown in appendix C.2.1, under risk neutrality this is the case when

$$\widehat{EFF}_i = (1-p)c_L + pWTP_i. \tag{3.6}$$

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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-\widehat{EFF}_i^{CPP}) = \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)\left[(1-p)u_i(k-c_L) + pu_i(k-c_{i,H})\right]. \tag{3.7}$$

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

$$\widehat{EFF}_i^{CPP} = (1-p)c_L + pc_{i,H} \tag{3.8}$$

[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 $\widehat{EFF}_i = WTP_i$.

Assuming risk neutrality, all components of a participant's control-neutral predicted EFF except $WTP_i$ are thus exogenously given. Section 3.4 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.

### 3.3.2 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

$$
\begin{aligned}
&\pi_i(1)u(0) + \pi_i(\mathbb{P}_i(\geq 1|s))\,[u_i(1) - u_i(0)] + ... + \pi_i(\mathbb{P}_i(\geq 10|s))\,[u_i(10) - u_i(9)] \\
&= [\pi_i(\mathbb{P}_i(\geq 0|s) - \pi_i(\mathbb{P}_i(\geq 1|s)]\,u_i(0) + ... + [\pi_i(\mathbb{P}_i(\geq 11|s) - \pi_i(\mathbb{P}_i(\geq 10|s)]\,u_i(10) \\
&= \sum_{k=0}^{10} [\pi_i(\mathbb{P}_i(\geq k|s) - \pi_i(\mathbb{P}_i(\geq k+1|s)]\,u_i(k) \\
&\equiv \sum_{k=0}^{10} V_i(s)u_i(k)
\end{aligned}
\tag{3.9}
$$

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

$$u(WTP_i) \equiv \sum_{k=0}^{10} V_i(60)u_i(k) - \sum_{k=0}^{10} V_i(30)u_i(k). \tag{3.10}$$

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

$$\begin{aligned} \sum_{k=0}^{10} V_i(60)u_i(k - \widehat{EFF}_i) = \pi_i(1) \sum_{k=0}^{10} V_i(30)u_i(k) \\ + \pi_i(1-p)\left[\sum_{k=0}^{10} V_i(60)u_i(k - c_L) - \sum_{k=0}^{10} V_i(30)u_i(k)\right] \end{aligned} \tag{3.11}$$

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

$$\widehat{EFF}_i = \pi_i(1-p)c_L + (1 - \pi_i(1-p))WTP_i. \tag{3.12}$$

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.

## 3.4 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.

### 3.4.1 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.

All decisions in Parts 1–3 are incentivized using a two-stage iterative Mul-

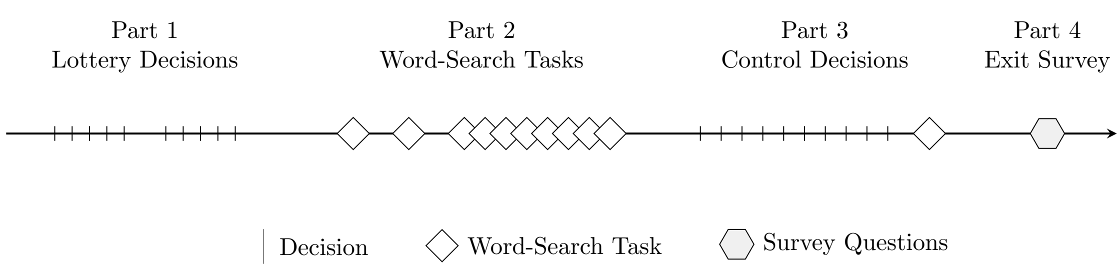


**Figure 3.3:** Sequence of Events

tiple Price List (MPL) approach, a modified version of the Becker-de Groot-Marshak (BDM) approach [6]. 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.

**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

[9] [1] 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.

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 (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.

**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 3.4 shows a task screen 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 partici-

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

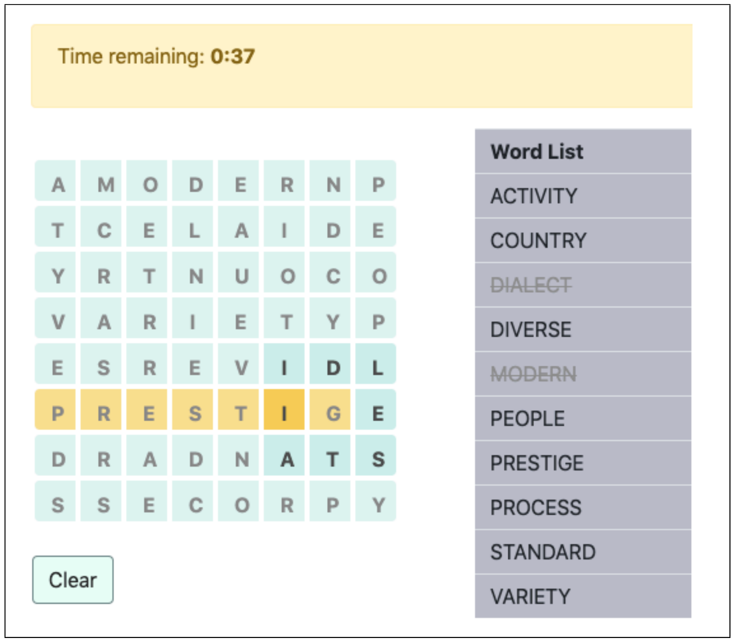

**Figure 3.4:** Word-Search Task Example

pants 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 3.5 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

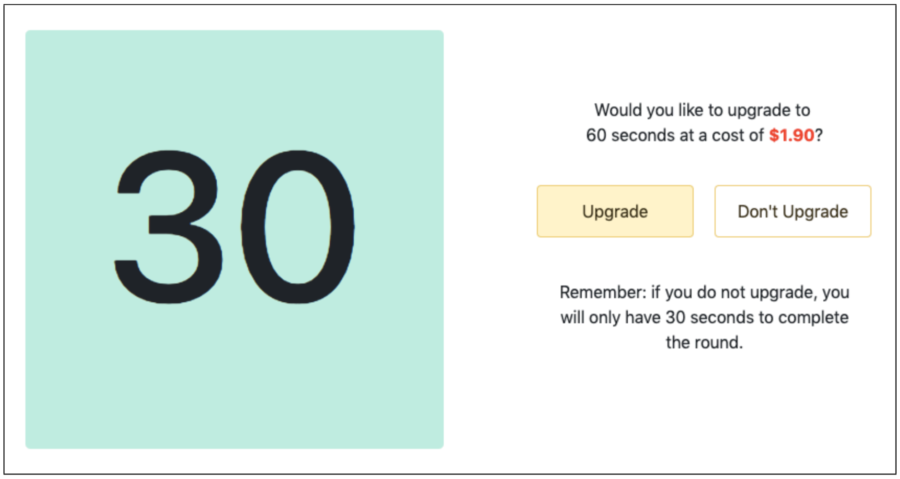


**Figure 3.5:** Upgrade Decision Screen

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 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.

#### 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 3.6 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**<br>Lottery: | | or | **Option B**<br>For sure: |
|---|---|---|---|
| 2 in 10 chance: CANNOT upgrade<br>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 |

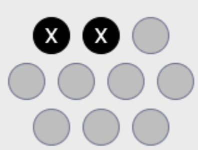


Submit

**Figure 3.6:** 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 3.7 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 under $\mathcal{CPP}$

[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.

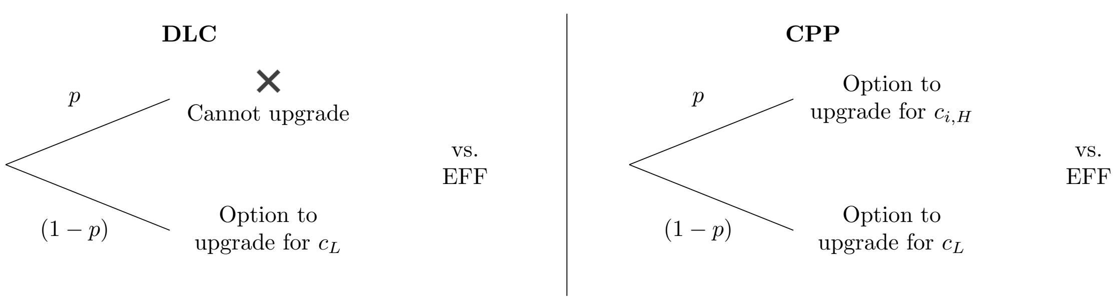


**Figure 3.7:** Control Lottery vs. Peak-Price Lottery

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

[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 $WTP_i$.

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 3.4. 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 3.8 summarizes the order of decisions in Part 3. Arrows indicate a randomization of decision order between subject.

### Part 4: Exit Survey

Part 4 consists of basic demographics questions; the 20 questions of the Desirability of Control Scale [10]; 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.

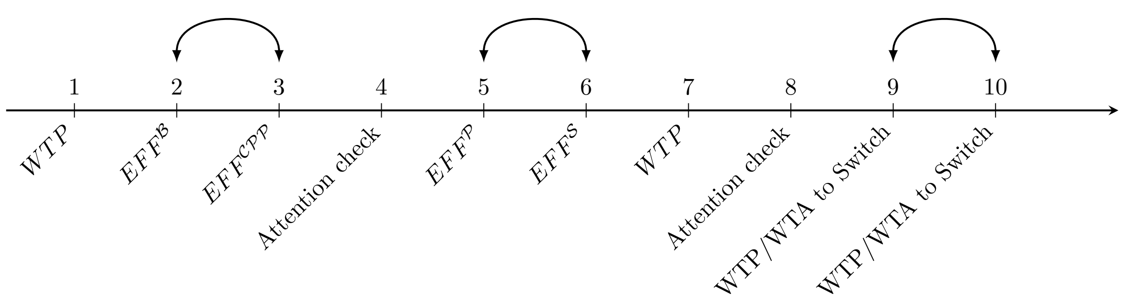


**Figure 3.8:** Decision Order in Part 3

### 3.4.2 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 [11] 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 C.1 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 [42]. 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.

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

[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.

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.

## 3.5 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.

### 3.5.1 Preferences over Money

As outlined in section 3.4.1, 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 the-

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

ory (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 [78] and [8] to get a sense of the population average degree of risk aversion. Note however that there is no variation in lottery payoffs. Estimates of $\alpha$ should therefore be interpreted as suggestive evidence for risk aversion. In addition, I test for probability weighting under both the functional form assumed in [78] and the Prelec form [66]:

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

[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 [68].

Table 3.1 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 3.1:** 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 [78].

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 [78] as my preferred estimate. Figure 3.9 highlights the estimated shape of the population-average probability weighting function.

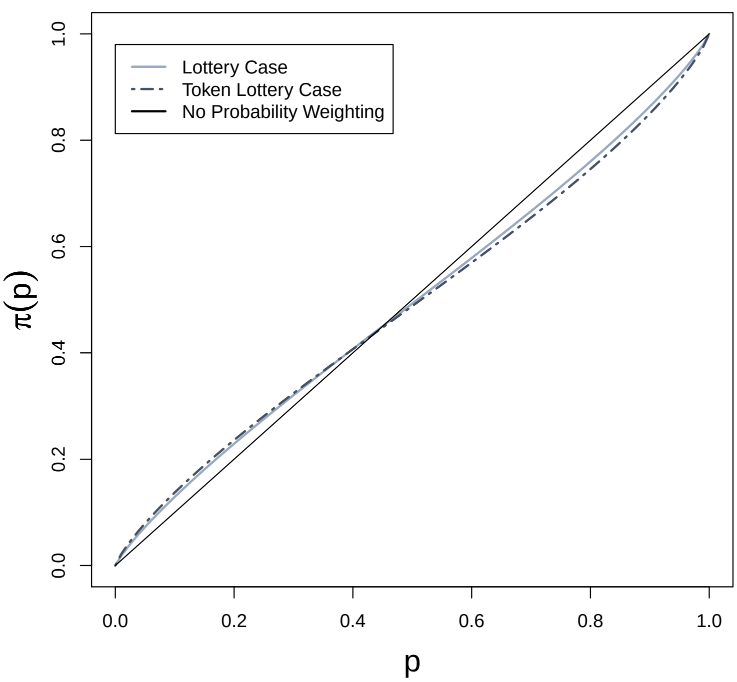


**Figure 3.9:** Estimated probability weighting functions under the standard and token frame assuming the parametric form suggested in [78]

### 3.5.2 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 C.2 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 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 C.4 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.

[17] Figure C.3 in the appendix depicts this relationship at the participant level. Note that this simple measure does not account for how often a participant actually chose to upgrade in Part 2.

### 3.5.3 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 3.10 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, [5] find an average relative control premium of 16.7%[19]

[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][5] calculate the relative control premium as the percentage difference between the context-

in a principal-agent setting. [60] find a relative control premium between 8-15%[20]. 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 $\mathcal{S}$.

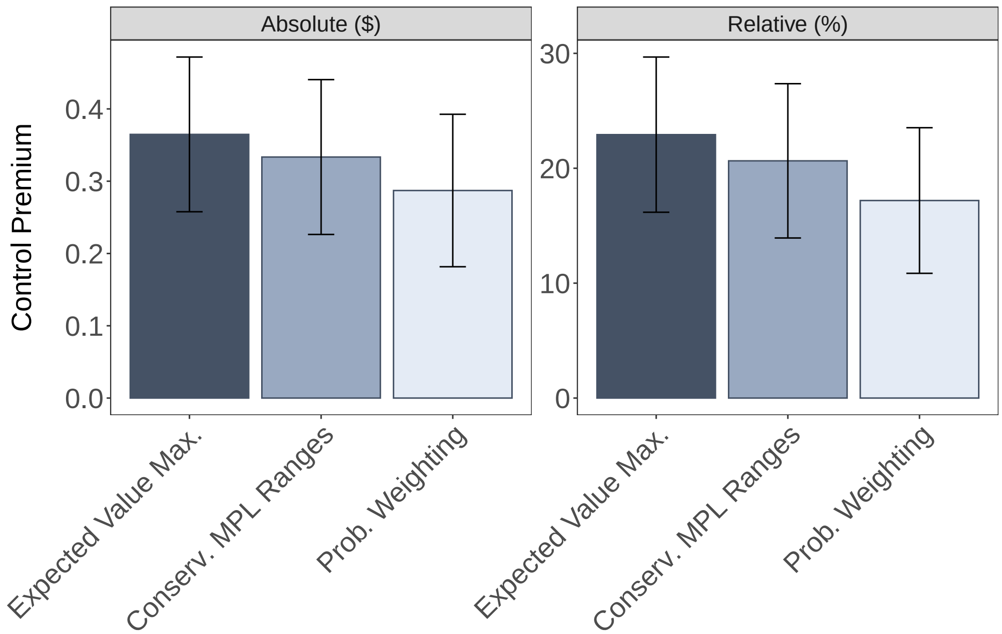


**Figure 3.10:** Model comparison under the base control lottery (error bars represent 95% confidence intervals).

Figure 3.11 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]

free elicitation of certainty equivalents of a *delegation* and *control lottery* which principals reported being indifferent between in a control-driven context

[20][60] 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 proba-

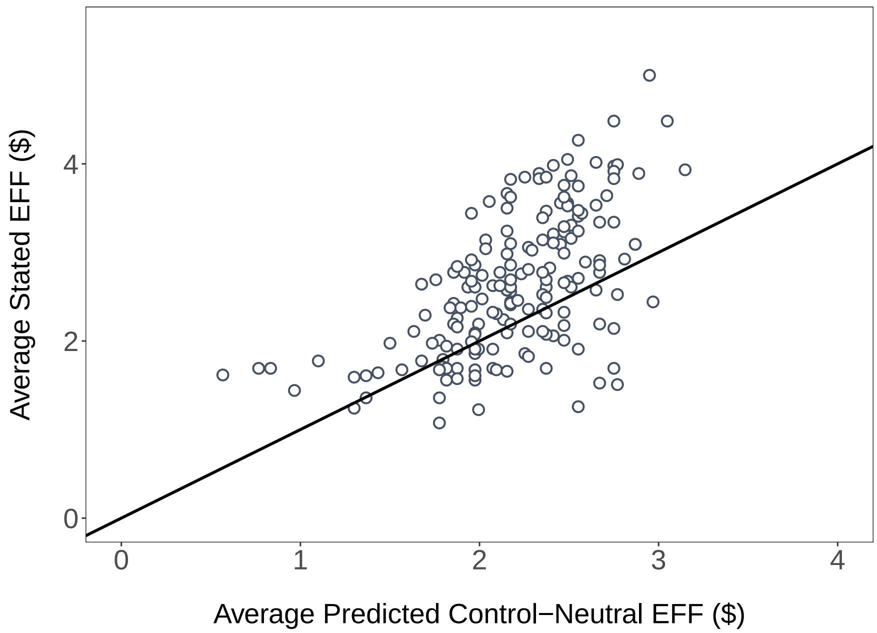


**Figure 3.11:** Individual participants' average reported equivalent flat fees (EFFs) across the three main control decisions versus the implied control-neutral EFFs (accounting for population-average probability weighting). Each dot represents one participant. Points above the 45 degree line imply a positive average control premium across decisions.

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, 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

bilities of an expected-value maximizer

[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 C.5 in the appendix breaks down Figure 3.11 by how many times a participant reported an EFF below its logical lower bound.

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 3.12 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 3.12 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 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

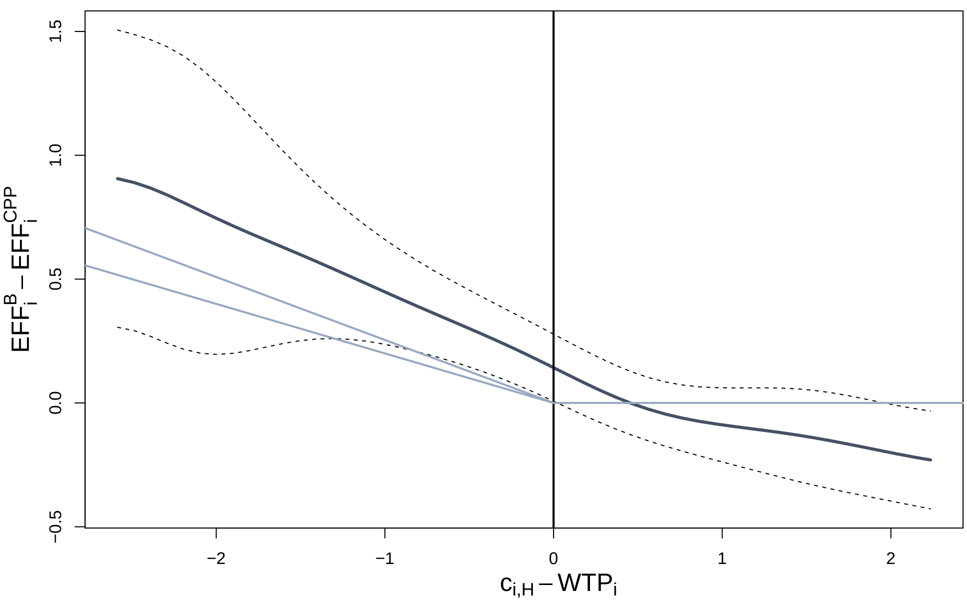


**Figure 3.12:** Comparing stated equivalent flat fees (EFFs) for base control lottery $\mathcal{B}$ and the peak-price lottery $\mathcal{CPP}$. Local-constant kernel fit of the difference in EFFs between the base control and peak-price lottery. The kernel bandwidth was selected using cross-validation. Estimates were generated using the *np* package [29]. 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).

additional credence to a control explanation.

For values above the cutoff, Figure 3.12 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 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 3.8). 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.

[22]Note that [60], 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).

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

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 ways: First, the electricity context may be subject to perceived conflicts of interest, particularly if consumers have low trust in their utility com-

[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 [sic] 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"

pany. [5] 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.

### 3.5.4 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 3.6, an expected value maximizer should report EFFs of

$$EFF_i^{\mathcal{P}} = (1-2p)c_L + pWTP_i = 2EFF_i^{\mathcal{B}} - c_L \tag{3.13}$$

and

$$EFF_i^{\mathcal{S}} = (1-p)2c_L + p2WTP_i = 2EFF_i^{\mathcal{B}}. \tag{3.14}$$

As shown in Figure C.6 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 3.13 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 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.

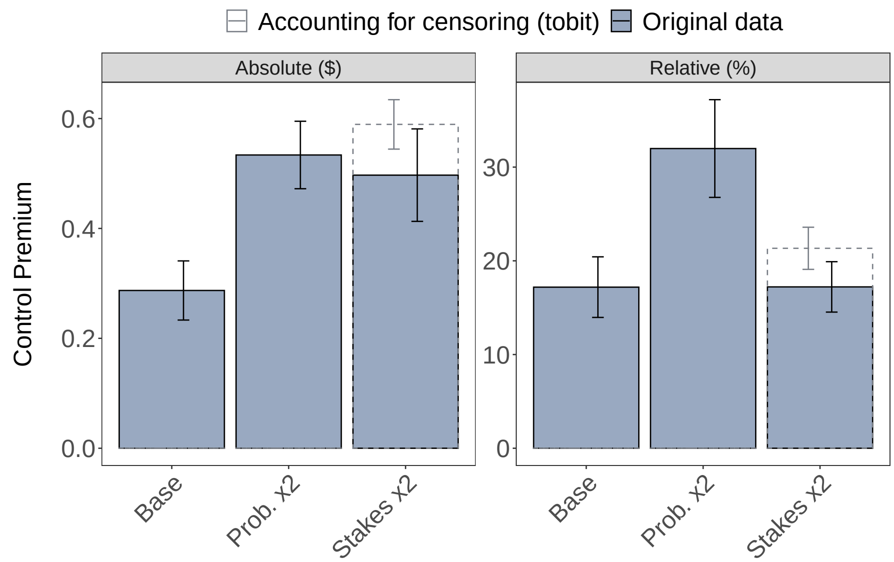


**Figure 3.13:** Control premium estimates across decision type allowing for probability weighting. The censoring-corrected estimate (dashed) for the decision with doubled stakes represents the mean of the fitted values from a Tobit regression of the stated lottery $\mathcal{S}$ EFF on the stated EFFs for lotteries $\mathcal{B}$ and $\mathcal{P}$. Error bars represent one standard deviation of the control premium estimate.

### 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 [5], the importance of control event probabilities had not been previously analyzed in the literature.

### 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.

### 3.5.5 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 [27, 30]. [39] and [38] 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,

[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 [5] only consider instances in

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 3.14 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

which principals cede control to agents, they find no association between loss aversion over monetary outcomes and their control-premium measure. [9], 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.

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 C.7 in the appendix, less than 25% of subjects report exactly the 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 [32] (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.


**Figure 3.14:** Test for Endowment Effect

### 3.5.6 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

$$u_i(\bar{k}_{i,60} - \widehat{EFF}_i) = (1-p)u_i(\bar{k}_{i,60} - c_L) + pu_i(\bar{k}_{i,30}) \tag{3.15}$$

$$= (1-p)u_i(\bar{k}_{i,60} - c_L) + p\left[u_i(\bar{k}_{i,60}) - u_i(WTP_i)\right] \tag{3.16}$$

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

$$
\begin{aligned}
u_i(\bar{k}_{i,60} - \widehat{EFF}_i) &\leq (1-p)u_i(\bar{k}_{i,60} - c_L) + p(u_i(\bar{k}_{i,60}) - WTP_i) \\
&\leq u_i((1-p)(\bar{k}_{i,60} - c_L) - p(\bar{k}_{i,60} - WTP_i)) \qquad (3.17) \\
&= u_i(\bar{k}_{i,60} - ((1-p)c_L + pWTP_i)))
\end{aligned}
$$

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.

### 3.5.7 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 3.12. Recall that this

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

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.

### 3.5.8 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 3.5.6 and 3.5.7. In this section, I consider additional alternative explanations.

#### 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. [44] 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.

**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 [48], 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

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

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 believes 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.

**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. [5] explicitly test for a relationship between loss aversion (over money) and their measure of the control premium and find no association.

### 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.

### 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.

**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 3.14: 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.

## 3.6 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 pro-

[35]According to EPRI, Detroit Edison was the first utility to implement a load control program in 1968 [19].

grams due to autonomy concerns coupled with consumers' lack of trust in energy companies [24, 23, 55]. Encouragingly, (hypothetical) acceptance increases significantly if consumers are given an option to override changes made by the utility company [55]. However, as pointed out by [36], 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 3.3 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 [70]. Each of these programs had more than 100,000 participants

[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).

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 *Cool-Points* 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:

**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. [74]) 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 3.11 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 [10] only detects significant differences based on gender: female participants' control premia are 14 percentage points higher on average (p-value = 0.0059, see Table 3.2).[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 3.2:** 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.

## 3.7 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 al-

[37]This finding is in line with [60] who also found that the Desirability of Control Scale did not correlate with their individual-specific control premium estimates.

most 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 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 [75] 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. [59]). 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 3.3:** Overview of Eight Largest DLC Programs by 2016 Enrollment According to E Source [70]

| 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) |

APPENDIX A

**APPENDIX FOR CHAPTER 1**

## A.1 Additional Information on the RFS

The RFS2 is implemented as a set of annual volumetric mandates for the 48 contiguous states and Hawaii that are subject to final review by the EPA[1]. In order to apportion these requirements to the obligated parties, the mandates are transformed into percentage blend obligations by dividing the required amount of renewable fuel for the year ahead by the total forecast amount of gasoline and diesel consumption. The forecasts are obtained from the November issue of the Short-Term Energy Outlook[2] preceding the mandate year. The RFS2 thus effectively operates as a blend mandate for biofuels.

Compliance is monitored through so called renewable identification numbers (RINs). RINs represent an accounting mechanism in which a unique 38 digit tracking number is assigned to every gallon of biofuel produced domestically or imported into the U.S. Once the underlying gallon has been physically blended for final consumption in the transportation sector, the RIN becomes detached from the renewable fuel it accompanies and turns into an independently tradable financial instrument. The advantage of this mechanism is that the obligated party can be a step removed from the physical blending process: the necessary amount of RINs for compliance can be procured through the blending of biofuels, purchases in the RIN market or a combination thereof. Similar to a cap-and-trade scheme, RINs therefore allow for efficiency gains through

[1]Alaska, Hawaii and non-contiguous U.S. territories are exempted from the program unless they opt in. Unlike Hawaii, Alaska has not yet chosen to exercise this option.

[2]*http://www.eia.gov/forecasts/steo/outlook.cfm*

strategic over- and under-blending by obligated parties with heterogeneous cost structures.

In fact, the RFS places the percentage blend obligations on refiners and importers of fossil fuels rather than blenders. This distinction is important since, as highlighted in [58], around 40% of US refiners in 2015 were neither integrated with blenders nor retailers[3]. The original justification for the choice of obligated party was to minimize regulated parties since the number of importers and refiners at the time was lower than the number of blenders using ethanol. However, due to the increased mandate levels under the RFS2, virtually all blenders now purchase some amount of ethanol and have therefore become regulated, though for the most part not obligated, parties.

The U.S. biofuels environment has been analyzed from a multitude of different angles. Introductions to the regulatory framework as well as the nature of RINs can be found in [52], [53] and [80][4].

A number of papers have challenged the net environmental benefits of the RFS2. [73] were the first to raise the concern that the use of corn-based ethanol might actually increase global greenhouse gas levels due to indirect land use change. A general equilibrium study by [7] cautions that the RFS in its current form increases emissions due to significant leakages in both land and fuel markets. This result is also underlined by [31] who explore unintended consequences of the RFS and find that land-use costs from erosion and habitat loss

[3]Percentage based on the number of active firms in the market rather than throughput

[4]The Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign (*http://farmdocdaily.illinois.edu*) and the Center for Agricultural and Rural Development at Iowa State University (CARD) (*https://www.card.iastate.edu/policy_briefs*) provide regular market commentary and policy briefs to discuss changes to the regulation or market environment. Similarly, Congressional Research Service reports to congress such as [72] and [84] provide valuable insights into current discussions around the RFS2.

under the RFS2 could be as high as $693 million while a cap-and-trade system would add virtually zero costs of this type.

Another source of concern is the effect of biofuels policies on food prices. [16] explore the nature of the link between crop and fuels markets which U.S. biofuels policy has shaped. They note that agricultural crops such as corn and soybeans can effectively lock on to oil prices through the fundamental pricing relationships implied by blending. These insights help to inform the debate on biofuels' potential to exacerbate famine during food price booms such as in 2007/2008 ([26], [83]).

## A.2 Nested Model

The blender and refiner cost functions are represented by constant elasticity functions. In this case, we require $\epsilon > 1$ in order to ensure convexity of the cost function, and hence concavity of the profit function. We allow for synergy effects such as benefits from shared storage infrastructure in the production of gasoline and diesel, as well as in the blending of E10 and E85 by letting both products feed into one joint cost function. Motor gasoline blending and diesel fuel blending on the other hand are treated as separate. The proposed cost functions are therefore of the form highlighted in equation A.1.

$$\begin{aligned} \text{Refiner:} \quad & (A^R_{C_G} q_G + A^R_{C_D} q_D)^{\epsilon^R_C} \\ \text{Blender (MG):} \quad & (A^B_{C_{E10}} q_{E10} + A^B_{C_{E85}} q_{E85})^{\epsilon^B_{C_{MG}}} \\ \text{Blender (DF):} \quad & A^B_{C_{DF}} q_{DF}^{\epsilon^B_{C_{DF}}} \end{aligned} \tag{A.1}$$

In total, the model consists of 25 equations in 25 unknowns:

- Nine quantities: $q_{E10}$, $q_{E85}$, $q_G$, $q_{DF}$, $q_D$, $q^R_{D4}$, $q^B_{D4}$, $q^R_{D6}$, $q^B_{D6}$
- Nine prices: $p_{E10}$, $p_{E85}$, $p_G$, $p_{DF}$, $p_D$, $p_{BD}$, $p_{D4}$, $p_{D6}$, $p_E$
- Five Lagrange multipliers: $\gamma^R_{D4}$, $\gamma^R_{D6}$, $\gamma^B_{D4}$, $\gamma^B_{D6}$, $\gamma^B_{E10}$
- Two Blend Ratios: $\theta_{E10}$, $\theta_{DF}$

The full set of behavioral equations is shown in the set of equations A.2.

## A.3 Data and Calibration Results

Throughout this paper, the symbol $\circ$ stands for a generic subscript. The capital letters $D_\circ$, $C_\circ$ and $S_\circ$ represent demand, cost and supply functions respectively. Quantities are always denoted by $q_\circ$ while $p_\circ$ represents prices. The superscripts $B$ and $R$ designate variables pertaining to the blender or refiner respectively.

Table A.1 highlights key data sources as well as the corresponding realizations for 2015. Quantities are shown in billion gallons (bGAL) and prices in USD per gallon or USD per RIN.

Most of the data is collected from the EIA Monthly Energy Review Beta website[5], while the quantity of E85 is calculated as the sum of conventional gasoline blended with fuel ethanol for blends higher than 55%[6] and renewable fuel and oxygenate plant net production of finished motor gasoline[7]. Ethanol prices were

[5]http://www.eia.gov/beta/
[6]EIA Weekly Refiner and Blender Net Production update
[7]EIA Petroleum and Other Liquids report

**First Order Conditions**

$$\begin{aligned}
&\textit{FOC Refiner} && p_g - \frac{\partial C^R}{\partial q_G} - \gamma^R_{D4}\kappa_{BBD} - \gamma^R_{D6}\kappa_{TR} = 0\\
&\textit{FOC Refiner} && p_D - \frac{\partial C^R}{\partial q_D} - \gamma^R_{D4}\kappa_{BBD} - \gamma^R_{D6}\kappa_{TR} = 0\\
&\textit{FOC Refiner} && p_{D4} - \gamma^R_{D4} - \gamma^R_{D6} = 0\\
&\textit{FOC Refiner} && p_{D6} - \gamma^R_{D6} = 0\\
&\textit{FOC Blender} && p_{E10} - t_{MG} - \frac{\partial C^{MG}_B}{\partial q_{E10}} - (1-\theta_{E10})p_g\\
& && -\theta_{E10}(p_E - \gamma^B_{D6}) = 0\\
&\textit{FOC Blender} && p_{E85} - t_{MG} - \frac{\partial C^{MG}_B}{\partial q_{E85}} - 0.26p_g - 0.74(p_E - \gamma^B_{D6}) = 0\\
&\textit{FOC Blender} && p_{DF} - t_{DF} - \frac{\partial C^{DF}_B}{\partial q_{DF}} - (1-\theta_{DF})p_D\\
& && -\theta_{DF}(p_{BD} - 1.5\gamma^B_{D4} - tc_{BD}) = 0\\
&\textit{FOC Blender} && p_{D4} - \gamma^B_{D4} = 0\\
&\textit{FOC Blender} && p_{D6} - \gamma^B_{D6} = 0\\
&\textit{FOC Blender} && q_{E10}(p_G + \gamma^B_{D6} - p_E) - \gamma^B_{E10} = 0\\
&\textit{FOC Blender} && q_{DF}(p_D + 1.5\gamma^B_{D4} - p_{BD} + tc_{BD}) = 0
\end{aligned}$$

**Market Clearing**

$$\begin{aligned}
&\textit{MC Motor Gasoline} && q_{E10} - A_{D_C}p_{E10}^{\epsilon_{D_{MG}}} = 0\\
&\textit{MC Motor Gasoline} && q_{E85} - A_{D_{FFV}}\left(\frac{1}{\lambda}p_{E85}\right)^{\epsilon_{D_{MG}}} = 0\\
&\textit{MC Diesel Fuel} && q_{DF} - A_{D_{DF}}p_{DF}^{\epsilon_{D_{DF}}} = 0\\
&\textit{MC Gasoline} && q_G - (1-\theta_{E10})q_{E10} - 0.26q_{E85} = 0\\
&\textit{MC Ethanol} && S_E(p_E) - \theta_{E10}q_{E10} - 0.74q_{E85} = 0\\
&\textit{MC Diesel} && q_D - (1-\theta_{DF})q_{DF} = 0\\
&\textit{MC Biodiesel} && S_{BD}(p_{BD}) - \theta_{DF}q_{DF} = 0\\
&\textit{MC D4 RINs} && q^B_{D4} - q^R_{D4} = 0\\
&\textit{MC D6 RINs} && q^B_{D6} - q^R_{D6} = 0
\end{aligned}$$

**Complementary Slackness**

$$\begin{aligned}
&\textit{CS Refiner} && \gamma^R_{D4}(q^R_{D4} - \kappa_{BBD}(q_G + q_D)) = 0\\
&\textit{CS Refiner} && \gamma^R_{D6}(q^R_{D4} + q^R_{D6} - \kappa_{TR}(q_G + q_D)) = 0\\
&\textit{CS Blender} && \gamma^B_{D4}(1.5\theta_{DF}q_{DF} - q^B_{D4}) = 0\\
&\textit{CS Blender} && \gamma^B_{D6}(\theta_{E10}q_{E10} + 0.74q_{E85} - q^B_{D6}) = 0\\
&\textit{CS Blender} && \gamma^B_{E10}(0.1 - \theta_{E10}) = 0
\end{aligned} \tag{A.2}$$

obtained through Bloomberg while RIN prices were purchased from OPIS. We are using the B99/B100 biodiesel prices from the DOE Alternative Fuels Data Center as a proxy for wholesale biodiesel prices.

Relevant motor gasoline and diesel fuel tax rates were obtained from the EIA. Including federal and average state taxes, the tax rates for 2015 were at 44.89 USc/GAL and 51.64 USc/GAL respectively. Based on the outlined parametric assumptions and elasticity choices we obtain the calibration results shown in table A.2.

Using these calibrated inputs, we apply a sequential quadratic programming procedure to simulate market equilibria for different mandate levels. Additional notation is summarized in table A.3.

## A.4 Relative Importance of Compliance Channels

In order to provide evidence for the existence of the fourth compliance channel and to assess its relative significance, we need to purge the simulated changes in the total quantity of gasoline and diesel consumed of the impacts of the three other compliance channels.

The first and third compliance channel represent an increase in ethanol and biodiesel blend ratios respectively. This will lead to (i) a *substitution effect* away from fossil fuels and into biofuels and (ii) a *price effect* on retail fuels, thereby changing retail fuel consumption. We take both of these effects into account in estimating the impact of channels one and three on the compliance base.

**Table A.1:** Data Sources

| Variable | Description | 2015 | Units | Source |
|---|---|---|---|---|
| $q_G$ | Gasoline in Transport. excl. Ethanol | 124.72 | bGAL | EIA |
| $q_E$ | Ethanol in Transport | 13.38 | bGAL | EIA |
| $q_{E10}$ | E10 Consumption | 138.02 | bGAL | Calculated |
| $q_{E85}$ | E85 Consumption | 0.07 | bGAL | EIA |
| $\theta_{E10}$ | Implied E10 Ethanol Content | 9.65% | Percent | Calculated |
| $q_D$ | Diesel Fuel in Transport. excl. Biodiesel | 43.17 | bGAL | EIA |
| $q_{BD}$ | Biodiesel in Transport. | 1.48 | bGAL | EIA |
| $q_{DF}$ | Distillate Fuel Oil in Transport. | 44.65 | bGAL | EIA |
| $\theta_{DF}$ | Implied Biodiesel Content | 3.31% | Percent | Calculated |
| $p_G$ | Refiner Price of Motor Gasoline for Resale | 1.72 | USD/GAL | EIA |
| $p_E$ | Prompt Month Denatured Ethanol Futures | 1.51 | USD/GAL | Bloomberg |
| $p_{E10}$ | Regular Motor Gasoline, All Areas | 2.43 | USD/GAL | EIA |
| $p_{E85}$ | E85 Prices | 1.96 | USD/GAL | e85prices.com |
| $p_D$ | Refiner Price of No. 2 Diesel Fuel for Resale | 1.66 | USD/GAL | EIA |
| $p_{BD}$ | U.S. Retail Fuel Prices B99/B100 | 3.65 | USD/GAL | DOE AFDC |
| $p_{DF}$ | On-Highway Diesel Fuel Price | 2.71 | USD/GAL | EIA |
| $\kappa_{TR}$ | Final Percentage Standards: Renewable Fuel | 9.52% | Percent | EPA |
| $\kappa_{BBD}$ | Final Percentage Standards: BBD | 1.49% | Percent | EPA |
| $p_{D6}$ | Ethanol RINs (D6) | 0.55 | USD/RIN | OPIS |
| $p_{D4}$ | Biodiesel RINs (D4) | 0.72 | USD/RIN | OPIS |

**Table A.2:** Calibration Results

| Variable | Description | Baseline | $\epsilon_{S_E} * \frac{1}{2}$ | $\epsilon_{D_{MG}} * 2$ | $\epsilon_{S_{BD}} * \frac{1}{2}$ | $\epsilon_{D_{DF}} * 2$ |
|---|---|---|---|---|---|---|
| $A^R_{C_G}$ | Cost Refiner (Gasoline) | 0.4526 | 0.4526 | 0.4526 | 0.4526 | 0.4526 |
| $A^R_{C_D}$ | Cost Refiner (Diesel) | 0.4369 | 0.4369 | 0.4369 | 0.4369 | 0.4369 |
| $A^B_{C_{E10}}$ | Cost Blender (E10) | 0.1292 | 0.1292 | 0.1292 | 0.1292 | 0.1292 |
| $A^B_{C_{E85}}$ | Cost Blender (E85) | 0.1394 | 0.1394 | 0.1394 | 0.1394 | 0.1394 |
| $A^B_{C_{DF}}$ | Cost Blender (Diesel Fuel) | 0.1644 | 0.1644 | 0.1644 | 0.1644 | 0.1644 |
| $A_{S_E}$ | Supply Ethanol | 5.8720 | **8.8631** | 5.8720 | 5.8720 | 5.8720 |
| $A_{S_{BD}}$ | Supply Biodiesel | 0.1111 | 0.1111 | 0.1111 | **0.4050** | 0.1111 |
| $A_{D_C}$ | Demand Motor Gasoline, Conventional Cars | 171.2015 | 171.2015 | **214.0276** | 171.2015 | 171.2015 |
| $A_{D_{FFV}}$ | Demand Motor Gasoline, FFVs | 1.2000 | 1.2000 | 1.2000 | 1.2000 | 1.2000 |
| $A_{D_{DF}}$ | Demand Diesel Fuel | 47.8719 | 47.8719 | 47.8719 | 47.8719 | **51.3280** |

**Table A.3:** Additional Notation

| Variable | Description |
|---|---|
| $\gamma^R_{D4}$ | Refiner Lagrange Multiplier Refiner (D4 RINs) |
| $\gamma^R_{D6}$ | Refiner Lagrange Multiplier Refiner (D4+D6 RINs) |
| $\gamma^B_{D4}$ | Blender Lagrange Multiplier Refiner (D4 RINs) |
| $\gamma^B_{D6}$ | Blender Lagrange Multiplier Refiner (D6 RINs) |
| $\gamma^B_{E10}$ | Blender Lagrange Multiplier Refiner (E10 blend ratio) |
| $t_{MG}$ | Motor Gasoline Taxes |
| $t_{DF}$ | Diesel Fuel Taxes |
| $tc_{BD}$ | Biodiesel Tax Credit |
| $\lambda$ | Energy-equivalence Factor between E10 and E85 |
| $C^R(\cdot)$ | Refiner Cost Function |
| $C^B_{MG}(\cdot)$ | Blender Motor Gasoline Cost Function |
| $C^B_{DF}(\cdot)$ | Blender Diesel Fuel Cost Function |

$$\text{Substitution Effect} \equiv \Delta\theta * q^{initial}_{retail}$$

$$\text{Price Effect} \equiv \theta^{initial} * \Delta q_{retail}$$

The second channel, an increase in E85 sales, leads to substitution away from E10 and into E85, and hence once again a substitution from fossil fuels into biofuels. In the case of E85 however, increasing mandate levels lead to lower prices and increased consumption, meaning that the price effect here partially offsets the substitution effect.

We estimate the decrease in the overall compliance base attributable to channel four as the total decrease in gasoline and diesel sales minus the estimated substitution and price effects of the first three channels. Figures 1.7 and 1.8 highlight the order in which the four compliance channels are activated as well as their relative significance under the baseline calibration. Table A.4 summarizes the sensitivity of these results to the underlying elasticity assumptions (the corresponding calibration results are shown in table A.2).

We tested the effect of doubling the motor gasoline or diesel fuel demand elasticity, or reducing the ethanol or biodiesel supply elasticity by half. All four scenarios intuitively lead to a more stringent blend wall effect as they either make consumers more price sensitive or make it more costly to secure additional biofuel supply. In line with this intuition, we find that most of these changes lead to (i) the activation of channel four at lower mandate levels, and (ii) a higher reliance on channel four at high mandate levels.

The only exception is a reduction in the motor gasoline demand elasticity. This change effects the fuel market equilibrium in multiple ways: FFV drivers become more price-responsive, increasing the effectiveness of channel 2; and conventional vehicle consumers become more price-responsive, leading to a faster reduction in E10 demand and hence in the overall compliance base since E10 has a high gasoline content. We therefore see a delayed start-up of channel four at a mandate level of 11.5%, but the reliance on channel four at the 14% mandate level is higher than for any other scenario.

## A.5 Reduced Model: Motor Gasoline Only

Refiner's Problem (linear cost of blending, $c^R$):

**Table A.4:** Sensitivity Analysis: Compliance Channel Four

| | Start-Up Point of Channel Four ($\kappa_{TR}$) | Estimated Reduction in Overall Compliance Base at Different Mandate Levels (bn GAL) | | |
|---|---|---|---|---|
| | | $\kappa_{TR} = 12\%$ | $\kappa_{TR} = 13\%$ | $\kappa_{TR} = 14\%$ |
| Baseline | 11.1% | 0.22 | 0.55 | 0.95 |
| $\epsilon_{S_E} * \frac{1}{2}$ | 10.9% | 0.27 | 0.60 | 1.00 |
| $\epsilon_{D_{MG}} * 2$ | 11.5% | 0.25 | 0.83 | 1.50 |
| $\epsilon_{S_{BD}} * \frac{1}{2}$ | 10.9% | 0.53 | 1.00 | 1.01 |
| $\epsilon_{D_{DF}} * 2$ | 10.5% | 0.42 | 0.80 | 1.25 |

$$\kappa_{TR} \leq \frac{\epsilon_{S_E}(0.1 * 0.9\lambda B_{D_C} + 0.26 * 0.74 B_{D_{FFV}}) + \frac{p_E}{q_E}\frac{q_G}{p_{D6}}(0.1^2\lambda B_{D_C} + 0.74^2 B_{D_{FFV}}) + \lambda\epsilon_{S_E}\frac{q_G}{p_{D6}}}{\epsilon_{S_E}(0.9^2\lambda B_{D_C} + 0.26^2 B_{D_{FFV}}) + \frac{p_E}{q_E}B_{D_C}B_{D_{FFV}}(0.1 * 0.26 - 0.74 * 0.9)^2} \tag{A.3}$$

$$\begin{aligned} &\max_{\{q_G,\, q_{D6}^R\}} \Pi^R = \\ &p_G q_G - c^R q_G - p_{D6} q_{D6}^R \\ s.t. \quad &q_{D6}^R \geq \kappa_{TR} q_G \end{aligned} \tag{A.4}$$

Blender's Problem (linear markups for blending, $(m_{E10}, m_{E85})$):

$$
\begin{aligned}
\max_{\{q_{E10},\ q_{E85},\ q_{D6}^B\}} \Pi^B = \quad & \\
& q_{E10}(p_{E10} - m_{E10} - t_{MG}) \\
& + q_{E85}(p_{E85} - m_{E85} - t_{MG}) \\
& + p_{D6}q_{D6}^B \\
& - (0.9q_{E10} + 0.26q_{E85})p_G \\
& - (0.1q_{E10} + 0.74q_{E85})p_E \\
s.t. \qquad & q_{D6}^B \leq \theta_{E10}q_{E10} + 0.74q_{E85}
\end{aligned}
\tag{A.5}
$$

Linear Demand Functions:

- Case 1: $p_{E85} > \lambda p_{E10}$

$$
\begin{aligned}
D_{E85} &= 0 \\
D_{E10} &= A_{D_{FFV}} + A_{D_C} \\
&\quad - (B_{D_{FFV}} + B_{D_C})p_{E10}
\end{aligned}
$$

- Case 2: $p_{E85} = \lambda p_{E10}$

$$
\begin{aligned}
D_{E85} &\in [0, A_{D_{FFV}} - B_{D_{FFV}}p_{E10}] \\
D_{E10} &= A_{D_{FFV}} + A_{D_C} \\
&\quad - (B_{D_{FFV}} + B_{D_C})p_{E10} - q_{E85}
\end{aligned}
$$

- Case 3: $p_{E85} < \lambda p_{E10}$

$$D_{E85} = A_{D_{FFV}} - \frac{B_{D_{FFV}}}{\lambda} p_{E85}$$

$$D_{E10} = A_{D_C} - B_{D_C} p_{E10}$$

Upper bound on $\kappa_{TR}$ to ensure that the price of D6 RINs rises with increasing mandate levels provided in equation A.3.

APPENDIX B

# APPENDIX FOR CHAPTER 2

## B.1 Data

### B.1.1 Emission Factors

We use lifecycle analysis emission factors provided by the EPA shown in table B.1.[1] These emission factors account for all 'well-to-wheel' and 'field-to-wheel' emissions including feedstock production and the resulting land use change, fuel production and distribution as well as emissions generated during the combustion of the finished fuel. The emission factor for corn ethanol assumes a dry mill production process using natural gas. The biodiesel emission factor represents results for a transesterification process of soybean oil.

**Table B.1:** Lifecycle Emission Factors by Fuel Type

| Fuel Type | Net Emissions (kgCO2e / mmBtu) | Net Emissions (kgCO2e / gal.) | % Saving |
|---|---|---|---|
| Gasoline | 98.2 | 11.83 | - |
| Ethanol | 77.2 | 6.54 | 21% |
| Diesel | 97.0 | 13.33 | - |
| Biodiesel | 42.2 | 5.38 | 57% |

Note: Percentage savings reflect the percent reduction in CO2 emissions compared to the petroleum-based fuel being replaced.

[1] https://www.epa.gov/fuels-registration-reporting-and-compliance-help/lifecycle-greenhouse-gas-results

### B.1.2 Calibration Results

Table A2 highlights the calibration results.

**Table A2:** Calibration Results

| Variable | Description | Value |
|---|---|---|
| $A^R_{C_G}$ | Cost Factor Refiner (Gasoline) | 2.22 |
| $A^R_{C_D}$ | Cost Factor Refiner (Diesel) | 2.85 |
| $A^B_{C_{E10}}$ | Cost Factor Blender (E10) | 0.32 |
| $A^B_{C_{E85}}$ | Cost Factor Blender (E85) | 0.28 |
| $A^B_{C_{DF}}$ | Cost Factor Blender (Diesel Fuel) | 0.53 |
| $A_{S_E}$ | Supply Factor Ethanol | 5.38 |
| $A_{S_{BD}}$ | Supply Factor Biodiesel | 0.13 |
| $A_{D_{DF}}$ | Demand Factor Diesel Fuel | 59.67 |

## B.2 Demand Function Choice

In this section, we compare the demand specifications employed in (i) [64], (ii) [43], and (iii) [54]. Our simulation results indicate that the most important factor determining the shift of the welfare losses from motor gasoline to diesel fuel consumers is the constraint on E85 demand, and hence the degree to which the blend wall is binding. We show that the demand specifications in [64] and [43] lead to very similar simulation results, while the functional form used in

[54] slightly underestimates the welfare impact on diesel fuel consumers at high mandate levels. Note that the simulation results in this section are extended to total renewable blend mandate levels of up to 16%. The three demands lead to similar simulation results for mandate levels of up to 12%.

### B.2.1 Pouliot and Babcock (2016)

The demand functions in [64] are locally weighted quadratic regression estimates of the detailed demands derived in [62]. [62] employ a choice model to estimate E85 demand by FFV drivers taking into account (i) heterogeneous preferences for the two fuels; (ii) the effort cost of finding E85 given the distribution of gas stations offering E85 relative to the location of FFVs; and (iii) constrained E85 distribution infrastructure.

### B.2.2 Korting and Just (2017)

[43] use a simplified demand specification in which FFV drivers value fuel purely based on the miles per gallon they provide, and do not have heterogeneous tastes for E10 and E85. As a result, drivers switch between E10 and E85 based purely on the relative price of the two fuels in energy equivalent terms. In line with the EPA RFS2 rule documents, we assume that a price discount of at least 22% relative to E10 is necessary for E85 to be attractive on an energy-equivalent basis. We therefore apply an energy-equivalence factor of $\lambda = 1 - 0.22$ to E10 prices to make the fuel prices comparable. FFV consumers will consume only E10 when $p_{E85} > \lambda p_{E10}$, purchase only E85 when $p_{E85} < \lambda p_{E10}$ and are indif-

ferent between the two types of fuel otherwise. The corresponding demand for E85 is shown in equation B.1.

$$D_{E85}(p_{E10}, p_{E85}) \begin{cases} = 0 & if \; p_{E85} > \lambda p_{E10} \\ \in [0, A_{D_{FFV}} p_{E10}^{\epsilon_{D_{MG}}}] & if \; p_{E85} = \lambda p_{E10} \\ = A_{D_{FFV}} \left(\frac{1}{\lambda} p_{E85}\right)^{\epsilon_{D_{MG}}} & otherwise \end{cases} \tag{B.1}$$

When prices are exactly equal in energy equivalent terms, consumers are completely indifferent between the two types of fuel. As in [43], we choose $\epsilon_{D_{MG}} = -0.25$ and $A_{D_{FFV}} = 1.2$ for our simulation results.

Similarly, equation B.2 shows the demand for E10 which is either made up of all consumers, just conventional car owners or conventional car owners and some share of FFV owners in case prices are equal in energy equivalent terms.

$$D_{E10}(p_{E10}, p_{E85}) \begin{cases} = A_{D_{FFV}} p_{E10}^{\epsilon_{D_{MG}}} + A_{D_C} p_{E10}^{\epsilon_{D_{MG}}} & if\ p_{E85} > \lambda p_{E10} \\ = A_{D_{FFV}} p_{E10}^{\epsilon_{DMG}} + A_{DC} p_{E10}^{\epsilon_{D_{MG}}} - D_{E85} & if\ p_{E85} = \lambda p_{E10} \\ = A_{DC} p_{E10}^{\epsilon_{DMG}} & otherwise \end{cases} \tag{B.2}$$

Based on calibration to 2015 market data, we use $A_{D_{FFV}} = 174.5$.

### B.2.3 Meiselman (2016)

[54] provides inverse demand functions for E10 and E85 of the form

$$p_{E10} = \alpha \phi_{MG} q_{E10}^{\frac{1-\sigma}{\sigma}} \left[ \alpha q_{E10}^{\frac{1}{\sigma}} + (1-\alpha)\left(\gamma q_{E85}\right)^{\frac{1}{\sigma}} \right]^{\left(\sigma - \frac{\sigma}{\epsilon_{D_{MG}}} - 1\right)} \tag{B.3}$$

$$p_{E85} = \gamma(1-\alpha)\phi_{MG}\left(\gamma q_{E85}\right)^{\frac{1-\sigma}{\sigma}} \left[ \alpha q_{E10}^{\frac{1}{\sigma}} + (1-\alpha)\left(\gamma q_{E85}\right)^{\frac{1}{\sigma}} \right]^{\left(\sigma - \frac{\sigma}{\epsilon_{D_{MG}}} - 1\right)} \tag{B.4}$$

These demand functions are derived from the maximization of a quasi-linear utility function subject to a budget constraint. The functional forms allow for heterogeneous fuel preferences. We recalibrate these demand functions to our 2015 data for comparison, choosing an energy equivalence factor of $\gamma = \lambda = 0.78$

and a demand elasticity of $\epsilon_{D_{MG}} = -0.25$ as before. $\sigma = 1.15$ is given in the paper. This yields $\alpha = 0.73$ and $log(\phi_{MG}) = 19.59$ which we use for our numerical simulations.

### B.2.4 Simulation Results under Different Demand Specifications

Figure B.1 plots the different E85 demand specifications.

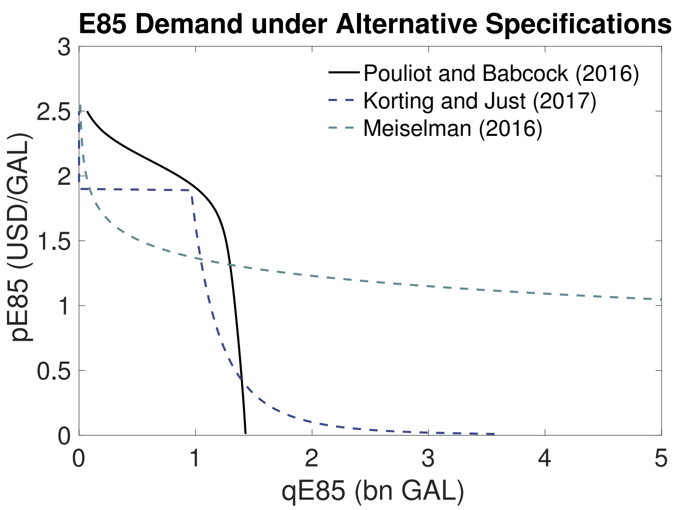


**Figure B.1:** E85 demand under alternative demand specifications

Figure B.2 provides simulation results for the quantity of E85 consumed and the diesel fuel consumer surplus under the three alternative demand specifications. As can be seen in the second panel of figure B.2, [64] and [43] generate very similar welfare results. The heterogeneous demand functions in [64] effectively act like a smoothed version of the discrete switching assumed in [43].

[54] also generates comparable welfare results up to mandate levels of 12%. For higher mandate levels, numerical simulations based on the demand func-

tion specification in [54] predict lower diesel fuel consumer surplus losses since E85 demand is far less constrained under this setup.

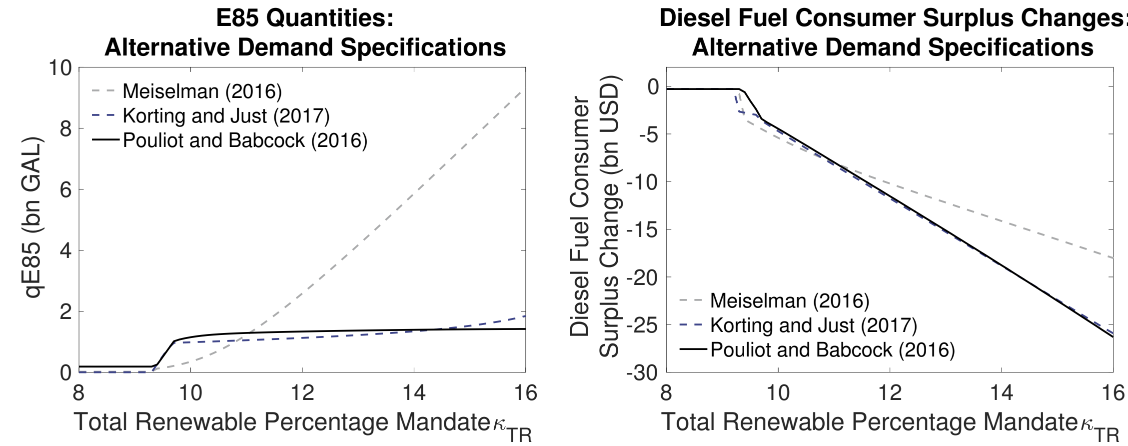


**Figure B.2:** Simulation results under alternative motor gasoline demand specifications
Notes: Numerical simulations hold the BBD blend mandate constant at its 2015 level of $\kappa_{BBD} = 1.5\%$ while varying the total renewable blend mandate level $\kappa_{TR}$. Welfare results are calculated as welfare changes compared to the no-mandate case ($\kappa_{TR} = \kappa_{BBD} = 0$).

## B.3 Behavioral Equations under the Reference Model

In total, the reference model consists of 25 equations in 25 unknowns:

- Nine quantities: $q_{E10}$, $q_{E85}$, $q_G$, $q_{DF}$, $q_D$, $q^R_{D4}$, $q^B_{D4}$, $q^R_{D6}$, $q^B_{D6}$
- Nine prices: $p_{E10}$, $p_{E85}$, $p_G$, $p_{DF}$, $p_D$, $p_{BD}$, $p_{D4}$, $p_{D6}$, $p_E$
- Five Lagrange multipliers: $\gamma^R_{D4}$, $\gamma^R_{D6}$, $\gamma^B_{D4}$, $\gamma^B_{D6}$, $\gamma^B_{E10}$
- Two Blend Ratios: $\theta_{E10}$, $\theta_{DF}$

$\gamma^B_{E10}$ represents the Lagrange multiplier on the blend wall constraint for the blender, i.e. on the constraint $\theta_{E10} \leq 10\%$. The other Lagrange multipliers apply to the compliance and RIN generation constraints for the blender and refiner respectively.

The full set of behavioral equations is shown in the system of equations below.

**First Order Conditions (FOC)**

$$\begin{aligned}
&FOC\ Refiner && p_G - \frac{\partial C^R}{\partial q_G} - \gamma^R_{D4}\kappa_{BBD} - \gamma^R_{D6}\kappa_{TR} = 0 \\
&FOC\ Refiner && p_D - \frac{\partial C^R}{\partial q_D} - \gamma^R_{D4}\kappa_{BBD} - \gamma^R_{D6}\kappa_{TR} = 0 \\
&FOC\ Refiner && p_{D4} - \gamma^R_{D4} - \gamma^R_{D6} = 0 \\
&FOC\ Refiner && p_{D6} - \gamma^R_{D6} = 0 \\
&FOC\ Blender && p_{E10} - t_G - \frac{\partial C^{MG}_B}{\partial q_{E10}} - (1-\theta_{E10})p_G - \theta_{E10}(p_E - \gamma^B_{D6}) = 0 \\
&FOC\ Blender && p_{E85} - t_G - \frac{\partial C^{MG}_B}{\partial q_{E85}} - 0.26p_G - 0.74(p_E - \gamma^B_{D6}) = 0 \\
&FOC\ Blender && p_{DF} - t_D - \frac{\partial C^{DF}_B}{\partial q_{DF}} - (1-\theta_{DF})p_D \\
& && - \theta_{DF}(p_{BD} - 1.5\gamma^B_{D4} - tc_{BD}) = 0 \\
&FOC\ Blender && p_{D4} - \gamma^B_{D4} = 0 \\
&FOC\ Blender && p_{D6} - \gamma^B_{D6} = 0 \\
&FOC\ Blender && q_{E10}(p_G + \gamma^B_{D6} - p_E) - \gamma^B_{E10} = 0 \\
&FOC\ Blender && q_{DF}(p_D + 1.5\gamma^B_{D4} - p_{BD} + tc_{BD}) = 0
\end{aligned}$$

**Market Clearing (MC)**

$$\begin{aligned}
&MC\ Motor\ Gasoline && q_{E10} - D_{E10}(p_{E10}, p_{E85}) = 0 \\
&MC\ Motor\ Gasoline && q_{E85} - D_{E85}(p_{E10}, p_{E85}) = 0 \\
&MC\ Diesel\ Fuel && q_{DF} - A_{D_{DF}} p_{DF}^{\epsilon_{D_{DF}}} = 0 \\
&MC\ Gasoline && q_G - (1-\theta_{E10})q_{E10} - 0.26q_{E85} = 0 \\
&MC\ Ethanol && S_E(p_E) - \theta_{E10}q_{E10} - 0.74q_{E85} = 0 \\
&MC\ Diesel && q_D - (1-\theta_{DF})q_{DF} = 0 \\
&MC\ Biodiesel && S_{BD}(p_{BD}) - \theta_{DF}q_{DF} = 0 \\
&MC\ D4\ RINs && q^B_{D4} - q^R_{D4} = 0 \\
&MC\ D6\ RINs && q^B_{D6} - q^R_{D6} = 0
\end{aligned}$$

**Complementary Slackness (CS)**

$$\begin{aligned}
&CS\ Refiner && \gamma^R_{D4}(q^R_{D4} - \kappa_{BBD}(q_G + q_D)) = 0 \\
&CS\ Refiner && \gamma^R_{D6}(q^R_{D4} + q^R_{D6} - \kappa_{TR}(q_G + q_D)) = 0 \\
&CS\ Blender && \gamma^B_{D4}(1.5\theta_{DF}q_{DF} - q^B_{D4}) = 0 \\
&CS\ Blender && \gamma^B_{D6}(\theta_{E10}q_{E10} + 0.74q_{E85} - q^B_{D6}) = 0 \\
&CS\ Blender && \gamma^B_{E10}(0.1 - \theta_{E10}) = 0
\end{aligned}$$

APPENDIX C

**APPENDIX FOR CHAPTER 3**

## C.1 Exclusion Criteria

**Table C.1:** 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 $\chi^2$-test ($p$-value = 0.048).

Figure C.1 shows the raw lottery valuations for the participants included in the sample, split out by the token vs. lottery frame.

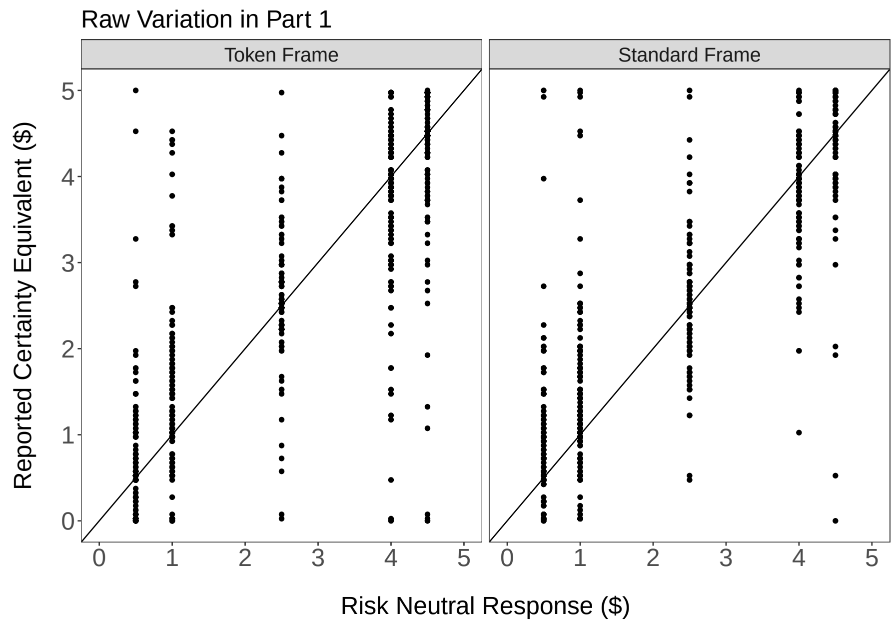


**Figure C.1:** Raw variation in lottery valuations by lottery type.

## C.2 Derivation of Theoretical Certainty Equivalents

### C.2.1 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

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

Rearranging yields

$$
\begin{aligned}
&(1-p)\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k-c_L)\right] \\
&+p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u(k-EFF_i)-\sum_{k=0}^{10}\mathbb{P}_i(k|30s)u_i(k)\right]=0
\end{aligned}
\tag{C.2}
$$

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

$$
\begin{aligned}
&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u_i(k-EFF_i)-\sum_{k=0}^{10}\mathbb{P}_i(k|s=30)u_i(k)\right] \\
=&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k)\right]\right. \\
&\left.+\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)u_i(k)-\sum_{k=0}^{10}\mathbb{P}_i(k|s=30)u_i(k)\right] \\
=&p\left[\sum_{k=0}^{10}\mathbb{P}_i(k|s=60)\left[u_i(k-EFF_i)-u_i(k)\right]+u_i(WTP_i)\right]
\end{aligned}
\tag{C.3}
$$

Equation C.2 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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s = 60)\left[(1-p)(c_L - EFF_i) + p(WTP_i - EFF_i)\right] = 0 \tag{C.4}$$

$$\Rightarrow (1-p)c_L + pWTP_i - EFF_i = 0 \tag{C.5}$$

which yields

$$EFF_i = (1-p)c_L + pWTP_i \tag{C.6}$$

*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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s = 60)u(k - EFF_i) = p\sum_{k=0}^{10} \mathbb{P}_i(k|s = 30)u_i(k) + (1-p)\sum_{k=0}^{10} \mathbb{P}_i(k|s = 60)u_i(k - c_L) \tag{C.7}$$

which again yields

$$EFF_i = (1-p)c_L + pWTP_i \tag{C.8}$$

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

$$\sum_{k=0}^{10} \mathbb{P}_i(k|s=60)u_i(k-EFF_i) = \sum_{k=0}^{10} \mathbb{P}_i(k|s=60)\left[(1-p)u_i(k-c_L)+pu_i(k-c_{i,H})\right]. \tag{C.9}$$

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

$$EFF_i = (1-p)c_L + pc_{i,H} \tag{C.10}$$

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. $EFF_i = WTP_i$.

### C.2.2 Cumulative Prospect Theory

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

$$\begin{aligned}\sum_{k=0}^{10} V_i(60)u_i(k - EFF_i) &= \pi_i(1)\sum_{k=0}^{10} V_i(30)u_i(k) \\ &\quad + \pi_i(1-p)\left[\sum_{k=0}^{10} V_i(60)u_i(k-c_L) - \sum_{k=0}^{10} V_i(30)u_i(k)\right] \\ &= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)u_i(k-c_L) \\ &\quad + (1-\pi_i(1-p))\sum_{k=0}^{10} V_i(30)u_i(k).\end{aligned} \tag{C.11}$$

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

$$
\begin{aligned}
0 &= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)\left[u_i(k-EFF_i) - u_i(k-c_L)\right]\\
&+ (1-\pi_i(1-p))\left[\sum_{k=0}^{10} V_i(60)(u_i(k-EFF_i) - u_i(k)) + \sum_{k=0}^{10} V_i(60)u_i(k) - \sum_{k=0}^{10} V_i(30)u_i(k)\right]\\
&= \pi_i(1-p)\sum_{k=0}^{10} V_i(60)\left[u_i(k-EFF_i) - u_i(k-c_L)\right]\\
&+ (1-\pi_i(1-p))\sum_{k=0}^{10} V_i(60)\left[(u_i(k-EFF_i) - u_i(k)) + u_i(WTP_i)\right].
\end{aligned}
\tag{C.12}
$$

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

$$
0 = \sum_{k=0}^{10} V_i(60)\left[\pi_i(1-p)(c_L - EFF_i) + (1-\pi_i(1-p))(WTP_i - EFF_i)\right] \tag{C.13}
$$

which yields

$$
EFF_i = \pi_i(1-p)c_L + (1-\pi_i(1-p))c_{i,H} \tag{C.14}
$$

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

The other three cases follow as shown in appendix C.2.1

## C.3 Supplementary Results

### C.3.1 Task Performance Measures for Part 2

Figure C.2 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.

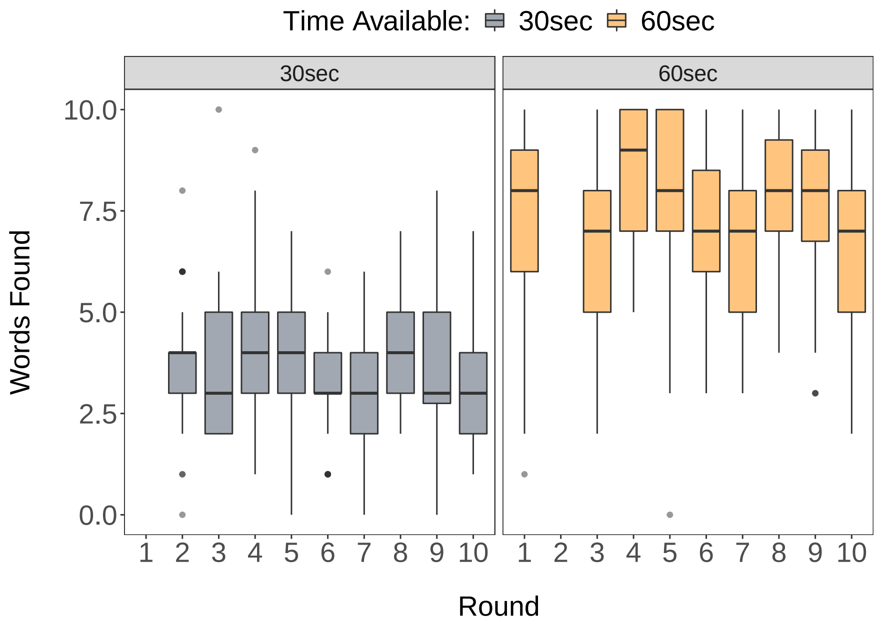


**Figure C.2:** Box plot of the number of words found per task for tasks of 30 versus 60 seconds. The box reflects the first, second and third quartiles (i.e. the 25th percentile, median, and 75th percentile). The whiskers extend to the largest value no further than 1.5 times the inter-quartile range. Data beyond the end of the whiskers are plotted individually.

Figure C.3 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. 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.

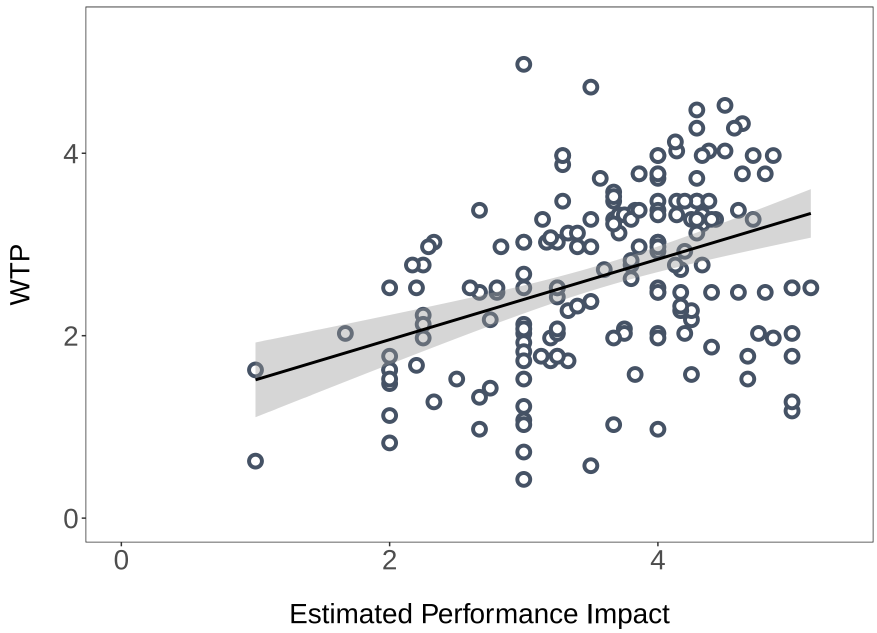


**Figure C.3:** Stated WTP to upgrade versus simple performance measure (average number of words found in 60 versus 30 seconds in Part 2). The black linear represents the linear regression of WTP on the estimated performance impact which has a positieve and statistically significant slope coefficient ($p$-value $<$ 0.001).

Figure C.4 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.

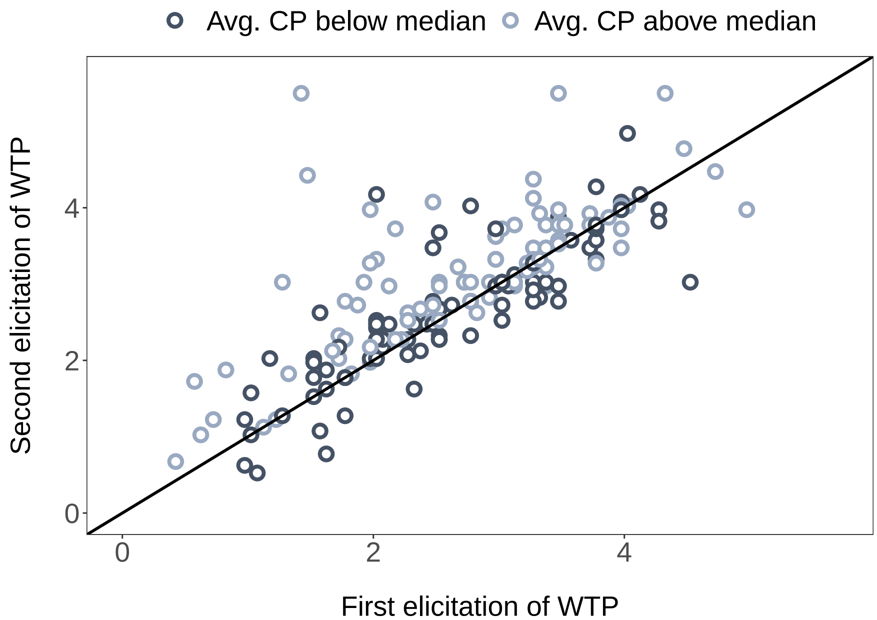


**Figure C.4:** Difference in first and second reported WTP to upgrade to 60 seconds. The color of each circle indicates whether the average control premium of the individual across the three main control lottery decisions in Part 3 was above or belove the population median.

### C.3.2 Stated EFFs

Figure C.5 visually splits out Figure 3.11 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 highlights that most of the participants whose valuations fall far below the 45 degree line frequently reported such valuations.

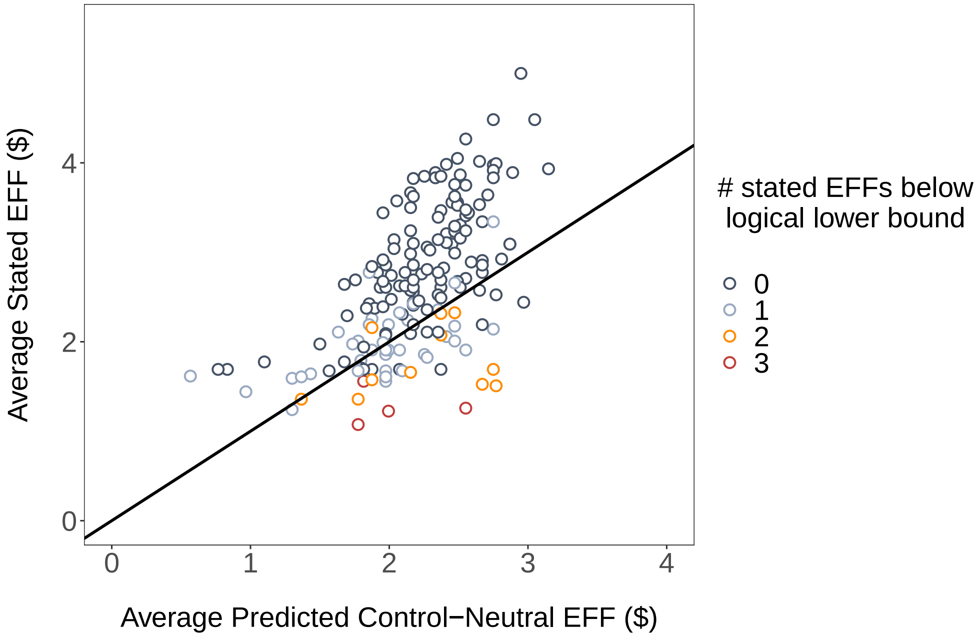


**Figure C.5:** Individual participants' average stated equivalent flat fees (EFFs) across the three main control decisions versus the implied control-neutral EFFs (accounting for population-average probability weighting). Each dot represents one participant. Points above the 45 degree line imply a positive average control premium across decisions. Colors indicate how many times a participant reported an EFF below its logical lower bound.

Figure C.6 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.

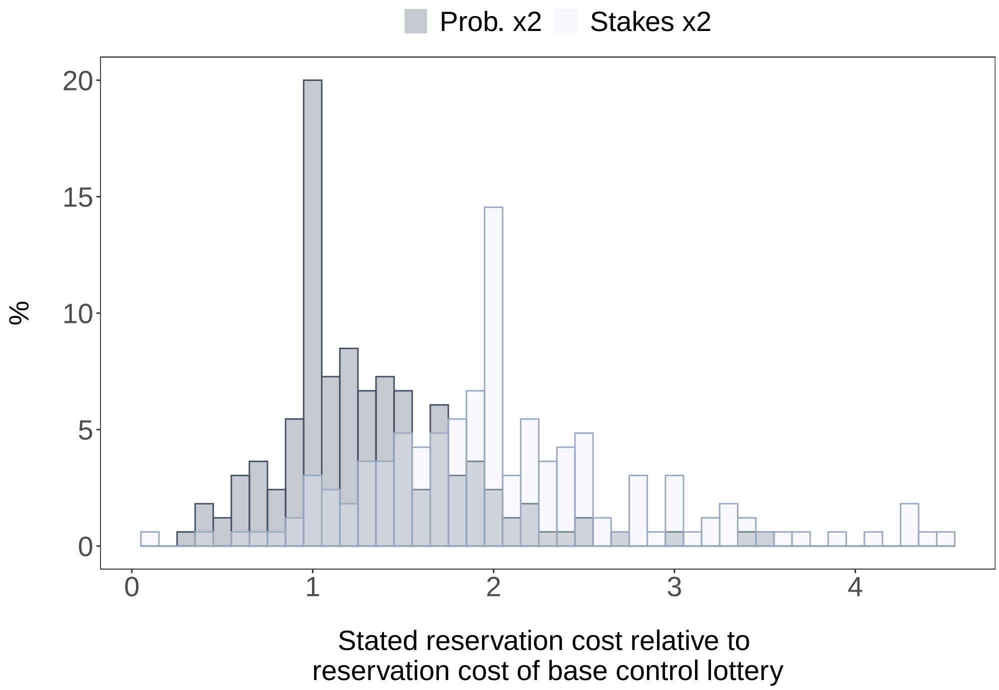


**Figure C.6:** Histogram of the ratio of reported equivalent flat fees (EFFs) for control-lotteries $\mathcal{P}$ (doubled probability of control events) and $\mathcal{S}$ (doubled stakes) relative to the base control lottery $\mathcal{B}$.

### C.3.3 Endowment Effect Questions

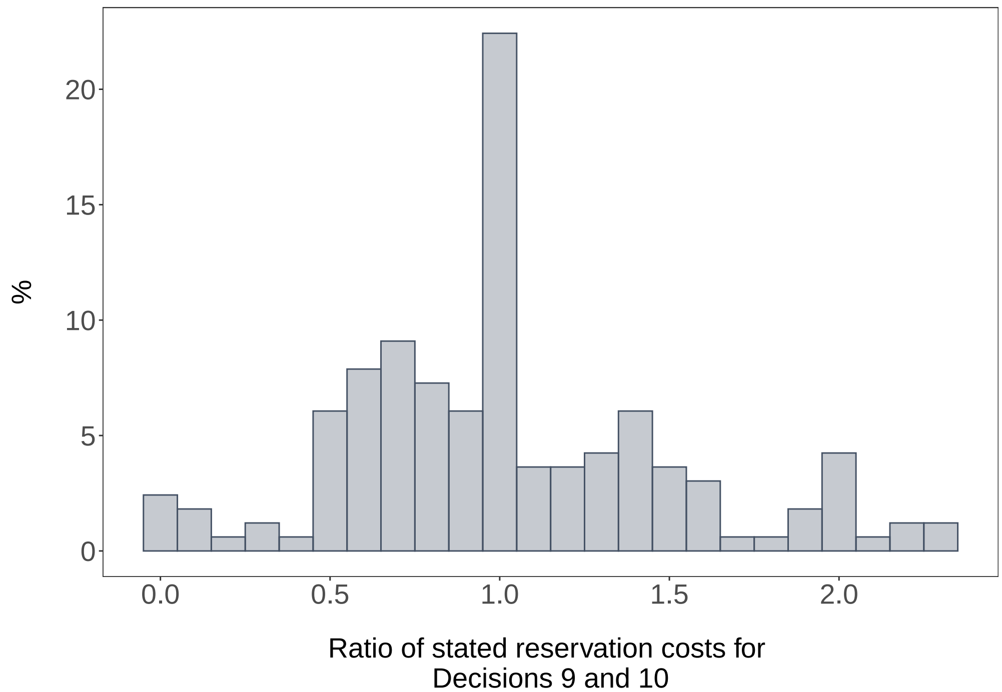


**Figure C.7:** Histogram of the ratio of reported equivalent flat fees (EFFs) for the control lotteries in decisions 9 and 10 of Part 3.

## C.4 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 [60] 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.

## C.5 Survey Screens

### C.5.1 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<br>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 |

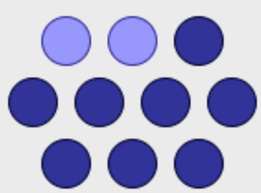

Submit

**Figure C.8:** Part 1: Lottery Questions - 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<br>Lottery: | | or | Option B<br>For sure: |
|---|---|---|---|
| 2 in 10 chance: pay $0<br>8 in 10 chance: pay $5.00<br>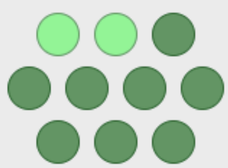 | Play Lottery ☐ | or | ☐ pay $5.00 |
| | Play Lottery ☐ | or | ☐ pay $4.50 |
| | Play Lottery ☐ | or | ☐ pay $4.00 |
| | Play Lottery ☐ | or | ☐ pay $3.50 |
| | Play Lottery ☐ | or | ☐ pay $3.00 |
| | Play Lottery ☐ | or | ☐ pay $2.50 |
| | Play Lottery ☐ | or | ☐ pay $2.00 |
| | Play Lottery ☐ | or | ☐ pay $1.50 |
| | Play Lottery ☐ | or | ☐ pay $1.00 |
| | Play Lottery ☐ | or | ☐ pay $0.50 |
| | Play Lottery ☐ | or | ☐ pay $0 |

Submit

**Figure C.9:** Part 1: Lottery Questions - Token frame

### C.5.2 Part 2: Word-Search Task

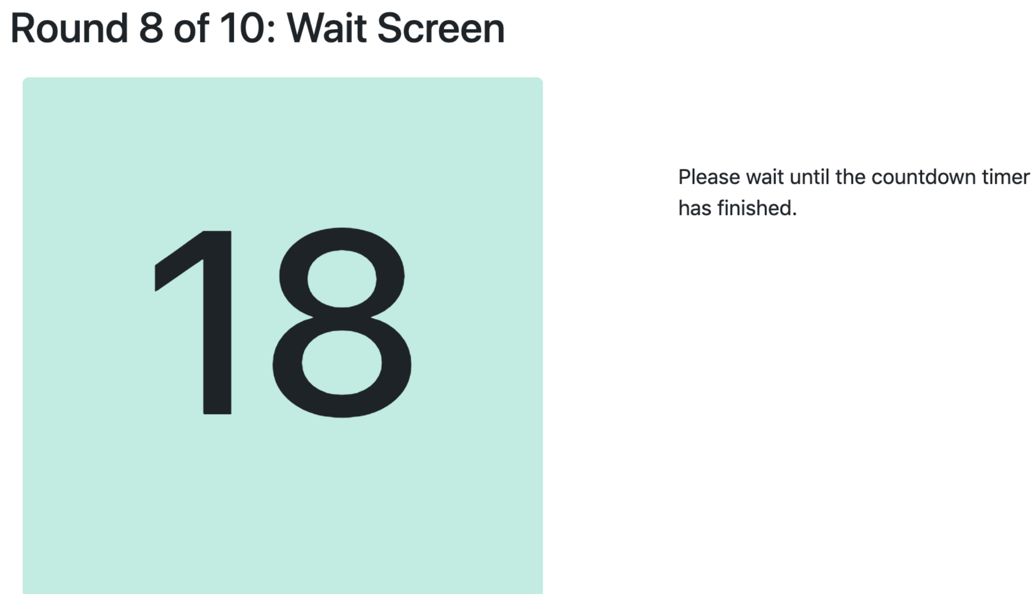


**Figure C.10:** Screenshot of wait screen during 30 second countdown

## C.5.3 Part 3: Control Lottery and Peak-Price Lottery Decisions

| Option A<br>Lottery: | | | or | | Option B<br>For sure: |
|---|---|---|---|---|---|
| 2 in 10 chance: CANNOT upgrade<br>8 in 10 chance: option to upgrade for $1.25<br>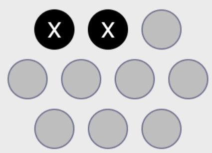 | Play Lottery | ▢ | or | ▢ | Option to upgrade for $5.50 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $5.00 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $4.50 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $4.00 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $3.50 |
| | 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 C.11:** Screenshot of base control lottery ($\mathcal{B}$)

| Option A<br>Lottery: | | | or | | Option B<br>For sure: |
|---|---|---|---|---|---|
| 4 in 10 chance: CANNOT upgrade<br>6 in 10 chance: option to upgrade for $1.25<br>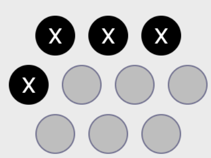 | Play Lottery | ▢ | or | ▢ | Option to upgrade for $5.50 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $5.00 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $4.50 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $4.00 |
| | Play Lottery | ▢ | or | ▢ | Option to upgrade for $3.50 |
| | 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 C.12:** Screenshot of control lottery with doubled probability ($\mathcal{P}$)

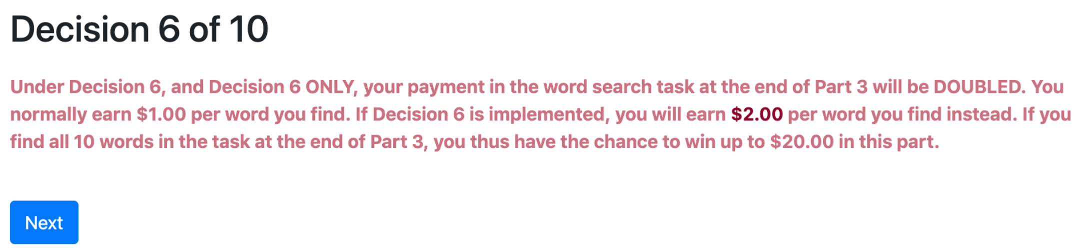

## Decision 6 of 10

**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

**Figure C.13:** Screenshot of control lottery with doubled stakes ($\mathcal{S}$) - Announcement that stakes will be doubled

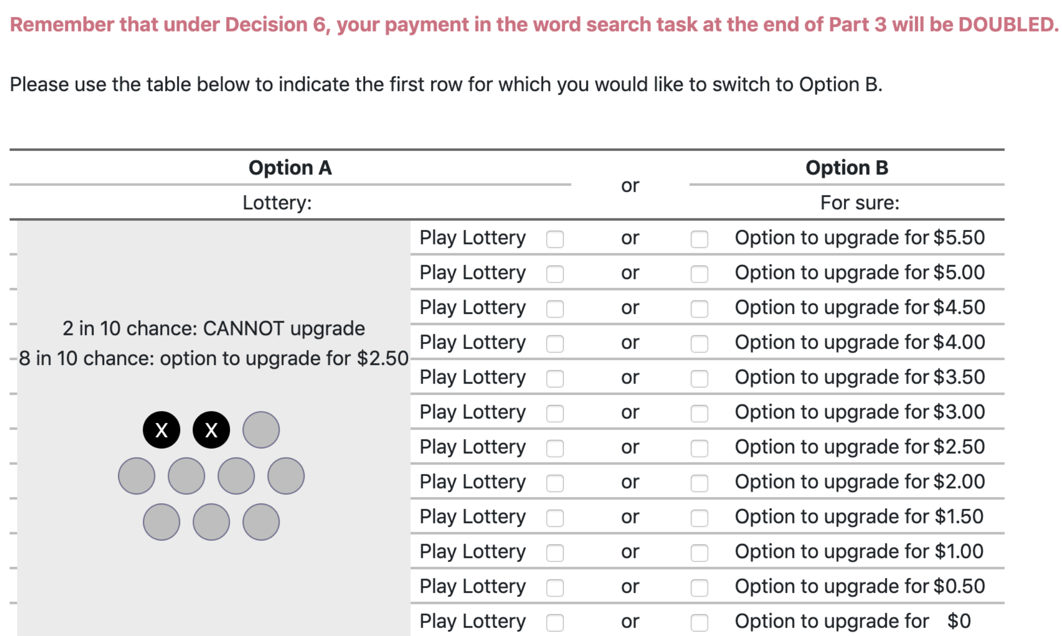

**Remember that under Decision 6, your payment in the word search task at the end of Part 3 will be DOUBLED.**

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

| **Option A** Lottery: | | or | **Option B** For sure: |
|---|---|---|---|
| 2 in 10 chance: CANNOT upgrade<br>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 |

**Figure C.14:** Screenshot of control lottery with doubled stakes ($\mathcal{S}$) - Control lottery screen

| Option A | | or | Option B |
|---|---|---|---|
| Lottery: | | | For sure: |
| 2 in 10 chance: Option to upgrade for $2.81<br>8 in 10 chance: Option to upgrade for $1.25<br>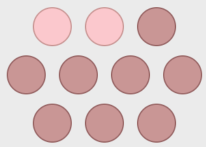 | Play Lottery | or | Option to upgrade for $5.50 |
| | Play Lottery | or | Option to upgrade for $5.00 |
| | Play Lottery | or | Option to upgrade for $4.50 |
| | Play Lottery | or | Option to upgrade for $4.00 |
| | Play Lottery | or | Option to upgrade for $3.50 |
| | Play Lottery | or | Option to 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 C.15:** Screenshot of peak-price lottery $\mathcal{CPP}$

### C.5.4 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

**Figure C.16:** Screenshot of attention check - Attention check announcement at the start of Part 3

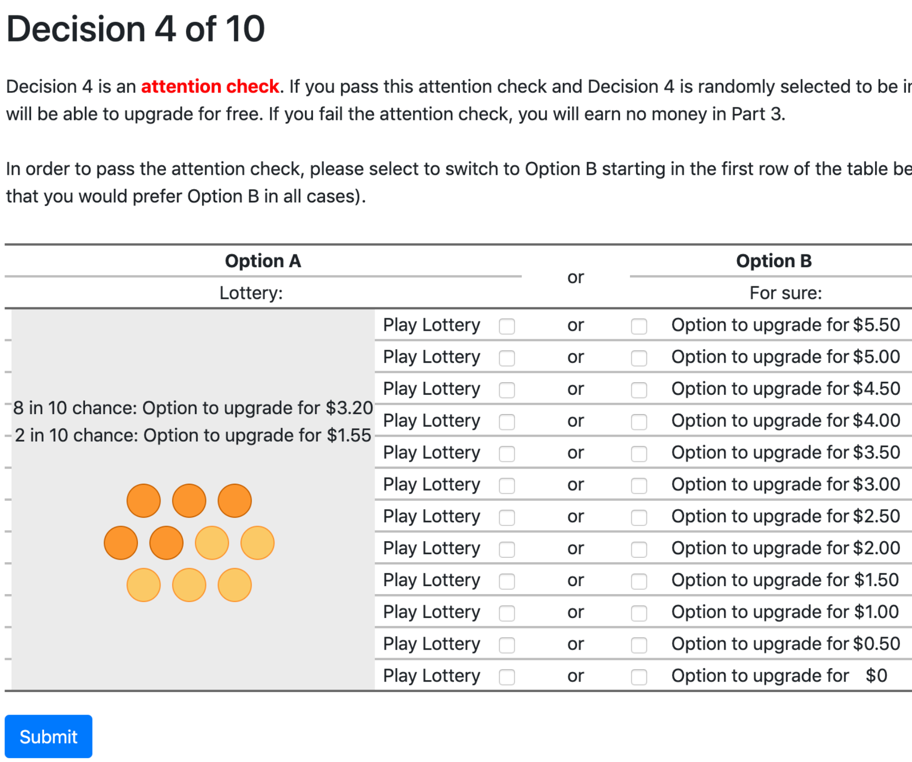

## Decision 4 of 10

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 | | or | Option B |
|---|---|---|---|
| Lottery: | | | For sure: |
| 8 in 10 chance: Option to upgrade for $3.20<br>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

**Figure C.17:** Screenshot of attention check - Attention check screen

# BIBLIOGRAPHY

[1] Steffen Andersen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström. Elicitation using multiple price list formats. *Experimental Economics*, 9(4):383–405, 2006.

[2] Soren T Anderson. The demand for ethanol as a gasoline substitute. *Journal of Environmental Economics and Management*, 63(2):151–168, 2012.

[3] API and AFPM. AFPM/API Comments on 2014-2016 RFS Proposal (Docket ID No. EPA-HQ-OAR-2015-0111), July 2015.

[4] Bruce A Babcock, Marcelo Moreira, Yixing Peng, and others. Biofuel taxes, subsidies, and mandates: Impacts on US and Brazilian markets. *CARD Staff Report*, 13-SR 108, 2013.

[5] Björn Bartling, Ernst Fehr, and Holger Herz. The intrinsic value of decision rights. *Econometrica*, 82(6):2005–2039, 2014.

[6] Gordon M Becker, Morris H DeGroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232, 1964.

[7] Antonio M. Bento, Richard Klotz, and Joel Landry. Are There Carbon Savings from US Biofuel Policies? The Critical Importance of Accounting for Leakage in Land and Fuel Markets. *Energy Journal*, 36(3):75–109, July 2015.

[8] B Douglas Bernheim and Charles Sprenger. Direct tests of cumulative prospect theory. *Available at SSRN 3350196*, 2019.

[9] Sebastian Bobadilla-Suarez, Cass R Sunstein, and Tali Sharot. 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)*, 2016.

[10] Jerry M Burger and Harris M Cooper. The desirability of control. *Motivation and emotion*, 3(4):381–393, 1979.

[11] Daniel L Chen, Martin Schonger, and Chris Wickens. otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.

[12] Adam Christensen and Sauleh Siddiqui. A mixed complementarity model for the us biofuel market with federal policy interventions. *Biofuels, Bio-products and Biorefining*, 9(4):397–411, 2015.

[13] Carol A Dahl. Measuring global gasoline and diesel price and income elasticities. *Energy Policy*, 41:2–13, 2012.

[14] Harry de Gorter and David R Just. The economics of a blend mandate for biofuels. *American Journal of Agricultural Economics*, 91(3):738–750, 2009.

[15] Harry de Gorter and David R. Just. The Social Costs and Benefits of Biofuels: The Intersection of Environmental, Energy and Agricultural Policy. *Applied Economic Perspectives and Policy*, 32(1):4–32, March 2010.

[16] Harry de Gorter, David R. Just, and Dusan Drabik. *The Economics of Biofuel Policies - Impacts on Price Volatility in Grain and Oilseed Markets*. Palgrave Macmillan, April 2015.

[17] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3):1155–1170, 2016.

[18] EIA. Annual Energy Outlook 2017, January 2017.

[19] Electric Power Research Institute. The demand-side management information directory: Interim report. 3426, 1985.

[20] EPA. 2010 Final Rule, March 2010.

[21] EPA. Renewable Fuel Standard Program: Standards for 2014, 2015, and 2016 and Biomass-Based Diesel Volume for 2017; Final Rule, December 2015.

[22] Armin Falk and Michael Kosfeld. The hidden costs of control. *American Economic Review*, 96(5):1611–1630, 2006.

[23] Michael J Fell, David Shipworth, Gesche M Huebner, and Clifford A Elwell. Exploring perceived control in domestic electricity demand-side response. *Technology Analysis & Strategic Management*, 26(10):1118–1130, 2014.

[24] Michael J Fell, David Shipworth, Gesche M Huebner, and Clifford A Elwell. Knowing me, knowing you: the role of trust, locus of control and privacy concern in acceptance of domestic electricity demand-side response. ECEEE, 2015.

[25] Stephen M Fleming and Dan Bang. Shouldering responsibility. *Science*, 361(6401):449–450, 2018.

[26] Carlisle Ford and Benjamin Senauer. How biofuels could starve the poor. *Foreign Affairs, www. foreignaffairs. org/20070501faessay86305/c-ford-runge-benjamin-senauer/how-biofuels-could-starve-the-poor. html*, 2007.

[27] Judd Hammack and Gardner Malard Brown. *Waterfowl and wetlands: toward bio-economic analysis*. Resources for the Future; distributed by

the Johns Hopkins University Press, Baltimore, [Washington], 1974. http://newcatalog.library.cornell.edu/catalog/461924.

[28] Tomas Havranek and Ondrej Kokes. Income elasticity of gasoline demand: A meta-analysis. *Energy Economics*, 47:77–86, 2015.

[29] Tristen Hayfield, Jeffrey S Racine, et al. Nonparametric econometrics: The np package. *Journal of statistical software*, 27(5):1–32, 2008.

[30] Thomas A Heberlein and Richard C Bishop. Assessing the validity of contingent valuation: Three field experiments. *Science of the Total Environment*, 56:99–107, 1986.

[31] Stephen P. Holland, Jonathan E. Hughes, Christopher R. Knittel, and Nathan C. Parker. Unintended Consequences of Carbon Policies: Transportation Fuels, Land-Use, Emissions, and Innovation. *Energy Journal*, 36(3):35–74, 2015.

[32] John K Horowitz and Kenneth E McConnell. A review of wta/wtp studies. *Journal of environmental economics and Management*, 44(3):426–447, 2002.

[33] Scott Irwin. Estimating the Biodiesel Supply Curve. *farmdoc daily*, October 2013.

[34] Scott Irwin. What's up with rins prices? *farmdoc daily*, (6):188((6):188), 2016.

[35] Fleeming Jenkin. On the principles which regulate the incidence of taxes. *Proceedings of the Royal Society of Edinburgh*, 7:618–631, 1872.

[36] Paul L. Joskow and Catherine D. Wolfram. Dynamic Pricing of Electricity. *The American Economic Review*, 102(3):381–385, 2012.

[37] Richard Just, Darell L Hueth, and Andrew Schmitz. *Applied welfare economics*. Edward Elgar Publishing, 2008.

[38] Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6):1325–1348, December 1990.

[39] Jack L. Knetsch. The Endowment Effect and Evidence of Nonreversible Indifference Curves. *The American Economic Review*, 79(5):1277–1284, 1989.

[40] Christopher R. Knittel, Ben S. Meiselman, and James H. Stock. The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard. Working Paper 21343, National Bureau of Economic Research, 2015.

[41] Christopher R Knittel, Ben S Meiselman, and James H Stock. The pass-through of rin prices to wholesale and retail fuels under the renewable fuel standard. *Journal of the Association of Environmental and Resource Economists*, 4(4):1081–1119, 2017.

[42] Christina Korting. Establishing Properties of the Control Premium. *AEA RCT Registry*, June 2019.

[43] Christina Korting and David R Just. Demystifying rins: A partial equilibrium model of us biofuel markets. *Energy Economics*, 64:353–362, 2017.

[44] David M Kreps. A representation theorem for" preference for flexibility". *Econometrica: Journal of the Econometric Society*, pages 565–577, 1979.

[45] Gabriel Lade and James Bushnell. Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard. Technical report, CARD Working Papers, 2016.

[46] Gabriel E Lade, C-Y Cynthia Lin Lawell, and Aaron Smith. Policy Shocks and Market-Based Regulations: Evidence from the Renewable Fuel Standard. Technical report, CARD Working Paper [16-WP 565], 2015.

[47] Gabriel E Lade, C-Y Cynthia Lin Lawell, and Aaron Smith. Policy shocks and market-based regulations: Evidence from the renewable fuel standard. *American Journal of Agricultural Economics*, 100(3):707–731, 2018.

[48] Ellen J Langer. The illusion of control. *Journal of personality and social psychology*, 32(2):311, 1975.

[49] Harvey Lapan and GianCarlo Moschini. Second-best biofuel policies and the welfare effects of quantity mandates and subsidies. *Journal of Environmental Economics and Management*, 63(2):224–241, 2012.

[50] Hyunok Lee and Daniel A Sumner. International trade patterns and policy for ethanol in the united states. In *Handbook of bioenergy economics and policy*, pages 327–345. Springer, 2010.

[51] Evan Markel, Burton C English, and Dayton Lambert. Thresholds and regime change in the market for renewable identification numbers. Technical report, Agricultural and Applied Economics Association, 2016.

[52] Bruce A. McCarl and Fred O. Boadu. Bioenergy and U.S. Renewable Fuels Standards: Law, Economic, Policy/Climate Change and Implementation Concerns. *Drake Journal of Agricultural Law*, 14:74, 2009.

[53] Lihong McPhail, Paul Westcott, and Heather Lutman. The renewable identification number system and US biofuel mandates. *ERS Report, USDA, November*, BIO-03, 2011.

[54] Ben S Meiselman. Breaching the Blendwall: RINs and the Market for Renewable Fuel. Technical report, University of Michigan Working Paper, 2016.

[55] Wilma Mert, Jürgen Suschek-Berger, and Wibke Tritthart. Consumer acceptance of smart appliances (report SMART-A D5.5). *Smart domestic appliances in sustainable energy systems (Smart-A)*, 2008.

[56] GianCarlo Moschini, Harvey Lapan, and Hyunseok Kim. The renewable fuel standard in competitive equilibrium: Market and welfare effects. *American Journal of Agricultural Economics*, 99(5):1117–1142, 2017.

[57] National Research Council. *Renewable fuel standard: potential economic and environmental effects of US biofuel policy*. National Academies Press, 2012.

[58] NERA. Effects of Moving the Compliance Obligation under RFS2 to Suppliers of Finished Products. Final Report, July 2015.

[59] Claudia Neri and Hendrik Rommeswinkel. Freedom, power and interference. *Available at SSRN 2485107*, 2015.

[60] David Owens, Zachary Grossman, and Ryan Fackler. The Control Premium: A Preference for Payoff Autonomy. *American Economic Journal. Microeconomics; Pittsburgh*, 6(4):138–161, November 2014.

[61] Nick Paulson. 2015 Year End RIN Update. *farmdoc daily*, (6):42((6):42), March 2016.

[62] Sebastien Pouliot and Bruce A Babcock. The demand for e85: Geographical location and retail capacity constraints. *Energy Economics*, 45:134–143, 2014.

[63] Sebastien Pouliot and Bruce A. Babcock. Compliance Path and Impact of Ethanol Mandates on Retail Fuel Market in the Short Run. *American Journal of Agricultural Economics*, 98(3):744–764, December 2015.

[64] Sébastien Pouliot and Bruce A Babcock. Compliance path and impact of ethanol mandates on retail fuel market in the short run. *American Journal of Agricultural Economics*, 98(3):744–764, 2016.

[65] Sebastien Pouliot, Aaron Smith, and James H. Stock. RIN Pass-Through at Gasoline Terminals. Working Paper, 2017.

[66] Drazen Prelec et al. The probability weighting function. *ECONOMETRICA*, 66:497–528, 1998.

[67] Brian L. Quick and Jennifer R. Considine. 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(5):483–491, 2008. PMID: 18850395.

[68] Matthew Rabin. Risk-aversion for small stakes: A calibration theorem. *Econometrica*, 68:1281–1292, 2000.

[69] Deepak Rajagopal and David Zilberman. *Review of environmental, economic and policy aspects of biofuels*, volume 4341. World Bank Publications, 2007.

[70] Katie Ryder, Rachel Reiss Buckley, and Charles Framularo. DLC Program Benchmarking: Results from the 2016 E Source Direct Load Control Program Study. Technical report, E Source, October 2017.

[71] Alberto Salvo and Cristian Huse. Build it, but will they come? Evidence from consumer choice between gasoline and sugarcane ethanol. *Journal of Environmental Economics and Management*, 66(2):251–279, 2013.

[72] Randy Schnepf and Brent D Yacobucci. Renewable Fuel Standard (RFS): overview and issues. In *CRS Report for Congress*, October 2010.

[73] Timothy Searchinger, Ralph Heimlich, Richard A Houghton, Fengxia Dong, Amani Elobeid, Jacinto Fabiosa, Simla Tokgoz, Dermot Hayes, and Tun-Hsiang Yu. Use of US croplands for biofuels increases greenhouse gases through emissions from land-use change. *Science*, 319(5867):1238–1240, 2008.

[74] O Seppanen, William J Fisk, and QH Lei. Room temperature and productivity in office work. 2006.

[75] Ellen A Skinner. A guide to constructs of control. *Journal of personality and social psychology*, 71(3):549, 1996.

[76] Chauncey Starr. Social benefit versus technological risk. *Science*, pages 1232–1238, 1969.

[77] James H. Stock. The Renewable Fuel Standard: A Path Forward. *Columbia SIPA Center on Global Energy Policy*, 2015, 2015.

[78] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, October 1992.

[79] U.S. Energy Information Administration. Annual Energy Outlook 2016, 2016.

[80] P Verleger. Renewable Identification Numbers. *Presentation to Agricultural Advisory Committee: Commodity Futures Trading Commission*, 2013, 2013.

[81] E. Glen Weyl and Michal Fabinger. Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy*, 121(3):528–583, 2013.

[82] Jarrett Whistance and Wyatt Thompson. A critical assessment of rin price behavior and the implications for corn, ethanol, and gasoline price relationships. *Applied Economic Perspectives and Policy*, page ppu012, 2014.

[83] Brian Wright. Global Biofuels: Key to the Puzzle of Grain Market Behavior. *The Journal of Economic Perspectives*, 28(1):73–97, 2014.

[84] Brent D Yacobucci. Intermediate-Level Blends of Ethanol in Gasoline, and the Ethanol "Blend Wall". In *CRS Report for Congress*, October 2010.