

Partial Identification of Welfare Effects in the Presence of Demand Frictions

Tommaso Coen *

11 · 28 · 2025

Please read the most recent version available [here](#)

Abstract

The validity of traditional welfare analysis in economics, based on the revealed preference paradigm, can be undermined by demand frictions such as default options, cognitive limitations, and limited or distorted information. This paper develops a general framework for discrete choice welfare analysis given quasi-experimental interventions that remove such frictions, which relies on relatively minimal assumptions on individual heterogeneity and overcomes key limitations of existing methodologies. It does so by mapping this problem to the literature on identification of functionals of the joint distribution of two potential outcomes, and takes a partial identification approach. I illustrate the approach in the context of product demand with non-salient taxes.

*Brown University, Department of Economics, (tommaso_coen@brown.edu); I am greatly indebted to Peter Hull, Toru Kitagawa, Jon Roth, and Neil Thakral for numerous discussions and their support throughout this project. I am also grateful to Jesse Bruhn, Pedro Dal Bó, John Friedman, Soonwoo Kwon and participants in all seminars at Brown University for their encouragement and helpful comments. All errors are mine.

1. Introduction

Welfare analysis in economics was traditionally rooted in the paradigm of revealed preference, which instructs us to infer objectives and welfare from choices [Beshears et al. \(2008\)](#), [Bernheim and Taubinsky \(2018\)](#). But choices can sometimes be influenced by factors that are not considered relevant for welfare, such as default options, cognitive limitations, or informational frictions [DellaVigna \(2009\)](#), [Bernheim and Taubinsky \(2018\)](#), [Train \(2015\)](#), [Handel et al. \(2019\)](#). This idea, which can be traced back to the distinction between a decision utility and an experienced utility [Kahneman et al. \(1997\)](#), [Mullainathan et al. \(2012\)](#), undermines the validity of the revealed-preferences approach. For this reason over the last decade standard techniques have been complemented with methods that try to provide valid results when decision utility and experienced utility diverge [Mullainathan et al. \(2012\)](#), [Bernheim and Taubinsky \(2018\)](#).

[Mullainathan et al. \(2012\)](#) show that to study the welfare consequences of taxation on agents that don't perfectly optimize with respect to taxes, it's crucial to discern two distinct demand curves. The first represents the scenario where agents optimize correctly, and the second reflects their flawed optimization. If a behavioral model suggests that agents make optimal decisions under specific conditions (for instance, in the absence of taxes or when taxes are transparently displayed in prices), then (quasi-)experimental exposure to these conditions can be used to identify the relevant demand curves. For example, [Chetty et al. \(2009\)](#) identify these alternative demands in a field experiment, randomly posting tax-inclusive prices in a grocery store chain. In a separate estimation strategy, they conduct a similar analysis for beer leveraging differences across states and over time in excise taxes (visible on listed prices) and sales taxes (applied during checkout).

While [Mullainathan et al. \(2012\)](#) underline how their framework can be used to study several problems beyond commodity taxation, including the provision of social insurance and the correction of externalities, they restrict attention to the case of homogeneous and non-stochastic behavioral biases. Relaxing this assumption, the demand curves for optimizing

and behavioral agents are no longer sufficient to identify welfare effects. More recent work by [Allcott and Taubinsky \(2015\)](#), [Spinnewijn \(2017\)](#), [Taubinsky and Rees-Jones \(2018\)](#), has shown that working under the assumption of homogeneous frictions can severely distort the results of the analysis. Some of these papers suggest the use of within-subject designs, that try to elicit willingness to pay under multiple conditions from the same subjects, to recover identification of welfare effects. One drawback of these methods is that they can be applied to a more limited range of data. Furthermore, within-subjects designs suffer from common threats that can undermine their internal validity, such as order effects ([Clark and Friesen \(2008\)](#)) and identification comes at the cost of additional assumptions.

This paper addresses these challenges in a binary choice setting and proposes a framework to conduct behavioral welfare analysis with arbitrarily heterogeneous behavioral biases, while relying exclusively on between-subject (quasi-)experimental variation. The key assumption in my framework is that I can observe the “demand curve” under both “clear” (without frictions) and “noisy” (with frictions) states. These are the same demand functions required by the original work of [Mullainathan et al. \(2012\)](#), which I use to provide bounds for welfare effects, taking a partial identification approach. A difference between this paper and the existing literature in behavioral public finance is that my bounds do not impose any parametric assumption on the demand curves, while many of the cited papers adopt a sufficient statistics approach that relies on demand linearization. I adopt instead the approach used in [Bhattacharya \(2015\)](#) to conduct nonparametric welfare analysis in discrete choice settings, extending it to the context of behavioral welfare analysis. My partial identification approach is also closely connected to the literature on distributional treatment effects [Heckman et al. \(1997\)](#), [Frandsen and Lefgren \(2021\)](#), a parallel that allows me to use ideas that originated in that literature to narrow my identified intervals. More generally, my exercise connects to a literature that uses optimal transport methods to study the partial identification of functionals of the joint distribution of potential outcomes [Fan et al. \(2017\)](#).

The main application studied in this paper concerns agents whose choices are impacted by

the salience of taxes. This is one of the leading examples studied in the growing literature on behavioral public finance Chetty et al. (2009), Taubinsky and Rees-Jones (2018), Bernheim and Taubinsky (2018). Following some of the most relevant papers in this literature, and in particular Taubinsky and Rees-Jones (2018), I focus on a simple model of binary choice where agents’ decision utility is quasilinear and agents can misperceive taxes when these are not salient. In this setting, I study the identification and estimation of the excess burden of taxation and show that this is not identified by the observable demand curves unless very stringent conditions are imposed. I re-analyze the experimental data from Taubinsky and Rees-Jones (2018) in which participants in an online platform are asked to indicate at what price they would purchase a series of items. Choice takes place in a condition in which state taxes are charged, or in a condition in which no taxes are charged. Importantly the experiment nests a between-subject design that I use for my bounds, and a within-subject design that they need in their approach to obtain point identification. I can therefore use their data to compare the results obtained with both methodologies. I find that the two approaches provide overall similar results. This suggests that my bounds can be useful to extend the analysis to those settings where we lack reliable within-subject data. At the same time, given that the bounds are based on weaker assumptions compared to point estimates, any point estimates falling outside the bounds hint at a possible violation of the additional assumptions needed for the latter.

The paper proceeds as follows: Section 2 presents my motivating example. Section 3 presents the general theoretical framework and discusses the main results. Section 4 discusses an empirical application to salience and taxation. Section 5 concludes.

2. Motivating example

The main application studied in this paper concerns agents whose choices are impacted by the salience of taxes. This is one of the leading examples studied in the growing literature on behavioral public finance Chetty et al. (2009), Taubinsky and Rees-Jones (2018), Bernheim

and Taubinsky (2018). Following some of the most relevant papers in this literature, and in particular Taubinsky and Rees-Jones (2018), I focus on a simple model of binary choice where agents' decision utility is quasilinear and agents can misperceive taxes when they are not salient, because not shown on a product's price-tag. In this setting, I will discuss the introduction of a value tax τ on the previously untaxed binary good. The welfare measures of interest will be the efficiency cost of the tax, or excess burden of taxation (EB), and the consumer loss from its introduction, as measured by equivalent variation (EV).

Agent i 's valuation for a binary product is v_i^* and, when taxes are not salient, i behaves as if the tax rate τ was actually $\theta_i\tau$ for some non-negative and bounded scalar θ_i . This setup is compatible with several frictions: inattention, incorrect beliefs, and rounding heuristics. The benchmark of correctly anticipated taxes is given by $\theta_i = 1$, which is the case for all agents if taxes are salient (e.g. posted on a product's price tag). We have therefore that with non salient taxes, i buys if his valuation is greater than the total price he anticipates to pay: $x_i(p, \tau) = \mathbb{1}\{v_i^* > p(1 + \theta_i\tau)\}$. If taxes were correctly anticipated, i would take the optimal choice $x_i^*(p, \tau) = \mathbb{1}\{v_i^* > p(1 + \tau)\}$, and maximize his utility function

$$U_i(x) = x[v_i^* - p(1 + \tau)]$$

Assuming supply is perfectly elastic, the excess burden of taxation is simply the lost surplus from discouraged transactions. The transaction of agent i is discouraged by the introduction of the tax if i would have bought the product without the tax, but does not buy it with the tax, so that $x_i(p, \tau) < x_i(p, 0)$. In this case i loses the surplus that he would have benefited from the transaction, given by $v_i^* - p$: the difference between his valuation for the product and the pre-tax price of the good. The excess burden of taxation EB is obtained aggregating

these losses over the population. ¹²

$$\begin{aligned} EB &= \mathbb{E} \left[(v_i^* - p) \times \mathbb{1}\{x_i(p, \tau) < x_i(p, 0)\} \right] \\ &= \mathbb{E} \left[(v_i^* - p) \times \mathbb{1}\{p < v_i^* \leq p(1 + \theta_i \tau)\} \right] \end{aligned} \quad (1)$$

Mullainathan et al. (2012) show that, when θ_i is homogeneous and non-stochastic in the population, observing two alternative demand curves is sufficient to study the welfare consequences of taxation. The first of these demands, that I denote $D^*(p)$, should represent a scenario in which agents optimize correctly. The second demand instead should be representative of agents' actual choices $x_i(p, \tau)$ in the setting in which taxes are present but not shown on posted prices, I denote this demand curve $D(p)$. My approach requires knowledge of these two functions, which can be obtained with any de-biasing intervention that as-good-as-randomly makes taxes salient, or removes them. For example, Chetty et al. (2009) ran a quasi-experiment with a grocery store chain, posting tax-inclusive prices on some products of some of the stores, over a limited time period, to estimate the elasticity of demand to taxes. They exploit independent variation in prices to separately estimate the price elasticity of demand, and identify the two demands assuming that these are approximately linear. Alternatively, some identify $D^*(p)$ is through variations in the tax rate. Indeed, in our model $x(p, 0) = x^*(p, 0)$, so that choices made when $\tau = 0$ are optimal. I focus on this case for expositional simplicity, since in the empirical section of the paper I will re-analyze data from Taubinsky and Rees-Jones (2018), who identify $D^*(p)$ eliciting willingness to pay with and without taxes in a randomized experiment. So I define

$$D^*(p) = Pr[x_i^*(p, 0) = 1] \quad D(p) = Pr[x_i(p, \tau) = 1] \quad (2)$$

²Without taxes, i purchases a product provided his evaluation is above the pre-tax price: $v_i^* > p$. When the tax rate rises to τ , i switches to not purchasing the same product if his evaluation is below his perceived tax inclusive price: $v_i^* \leq p(1 + \theta_i \tau)$. Therefore i 's choice is affected by the introduction of the tax if and only if his evaluation is between the pre-tax price and his perceived tax-inclusive price. Since $\theta > 0$ we can rewrite this as $p < v_i^* \leq p(1 + \theta_i \tau)$ or, equivalently, $v_i^*/(1 + \theta_i \tau) < p \leq v_i^*$.

The next session describes how we can use these demand curves to identify the excess burden of taxation when θ_i is heterogeneous in the population.

2.1. Identification

Start with the following observation: knowing the two demand curves is equivalent to knowing the marginal distributions of two potential reservation prices: one is the reservation price that individuals reveal when taxes are zero, and which coincides with their valuation v_i^* for the product. The other is the reservation price they reveal when taxes are equal to τ , this is $v_i = v_i^*/(1 + \theta_i\tau)$.

$$\begin{aligned} D^*(p) &= Pr[x_i(p, 0) = 1] & D(p) &= Pr[x_i(p, \tau) = 1] \\ &= Pr[v_i^* > p] & &= Pr[v_i^*/(1 + \theta_i\tau) > p] \\ &= 1 - F_{v^*}(p) & &= 1 - F_v(p) \end{aligned} \tag{3}$$

We can similarly rewrite EB in (1) in terms of v_i^* and v_i :

$$EB = \mathbb{E} \left[(v_i^* - p) \times \mathbb{1}\{v_i \leq p < v_i^*\} \right]. \tag{4}$$

Knowledge of the marginals F_{v^*} and F_v does not pin down EB: the expectation in (4) is taken with respect to the unobserved joint distribution of v_i and v_i^* and it cannot be separated in two expectations over the marginals. The impossibility of identifying EB arises because we don't have a way to choose the true joint distribution of v_i and v_i^* among all those that share the observed marginals. Consider a simple example where our population consists of three individuals. Then F_{v^*} is the distribution of their three reservation prices without taxes $v_1^*, v_2^*,$ and v_3^* , and F_v is the distribution of their three reservation prices

²Remember that without taxes, i purchases a product provided his evaluation is above the pre-tax price: $v_i^* > p$. When the tax rate raises to τ , i switches to not purchasing the same product if his evaluation is below his perceived tax inclusive price: $v_i^* \leq p(1 + \theta_i\tau)$. Therefore i 's choice is affected by the introduction of the tax if and only if his evaluation is between the pre-tax price and his perceived tax-inclusive price. Since $\theta > 0$ we can rewrite this as $p < v_i^* \leq p(1 + \theta_i\tau)$ or, equivalently, $v_i^*/(1 + \theta_i\tau) < p \leq v_i^*$.

with taxes v_1, v_2 , and v_3 . Figure 1 shows that these distributions (and hence these demand curves) do not identify EB because they can originate from joint distributions associated with different values of EB . In Figure 1a and 1b, the (fixed) marginal distributions of v_i^* and v_i are matched differently, determining different joint distributions. In the two resulting joint distributions, the dashed line represents the individual whose transaction is lost due to the introduction of the tax. This is the individual who has a reservation price $v_i^* > p$ before the tax is introduced, and a reservation price $v_i < p$ after the tax is introduced. This individual is different in the two distributions, and so is the surplus lost. In Figure 1a the affected individual loses a surplus of \$1, while in 1b the affected individual loses a surplus of \$3.

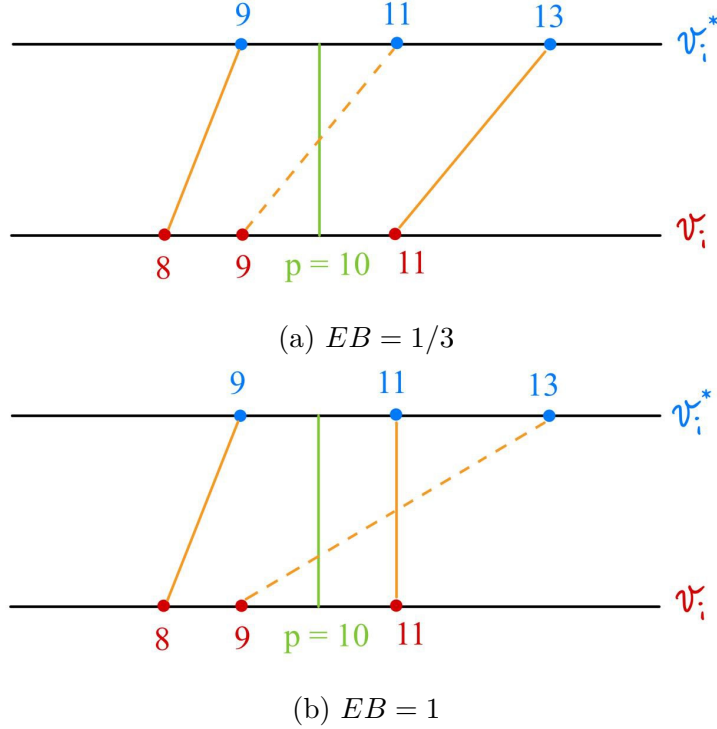


Figure 1: Different joint distributions are compatible with the given marginals of v_i and v_i^* . The value of EB would be different if the true joint distribution was the one represented in (a) or the one represented in (b).

Even if we cannot determine EB from F_{v^*} and F_v , we can use our knowledge of the marginals to bound it. In our simple example with three individuals, we know that EB

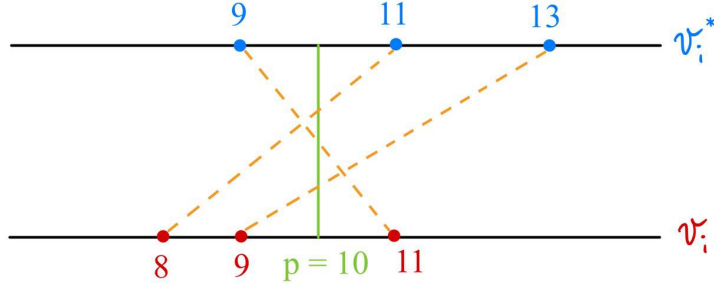


Figure 2: This matching of v_i^* and v_i gives $EB = 5/3$ but violates our assumption that $\theta_i \geq 0$, which implies that $v_i < v_i^*$ for all i . $5/3$ remains an upper bound for EB , but not a tight one.

results from some matching of their reservation prices without taxes v_1^*, v_2^* , and v_3^* , to their reservation prices with taxes v_1, v_2 , and v_3 . Table 1 reports all such matchings, together with the excess burden of the tax that they would imply. For example, the two matchings represented in figures 1a and 1b correspond to the first and the third lines of the table. All the matchings in the table 1 generate values of EB that fall between $1/3$ and $5/3$, which must therefore bound the true value of EB .

These bounds are obtained using only the constraints imposed by our knowledge of the marginal distributions F_{v^*} and F_v , and we may hope to narrow them imposing any additional assumption that we consider reasonable. For example, in our model we assumed $\theta_i > 0$, and we defined $v_i = v_i^*/(1 + \theta_i\tau)$. This implies that for every i we must have $v_i \leq v_i^*$. Some of the matchings in table 1 include pairs (v_i^*, v_i) that violate this condition, and are therefore ruled out by our assumption. The highest and lowest values of EB that come from matchings compatible with $\theta_i > 0$ are therefore $1/3$ and 1 , which are narrower than the bounds obtained without assuming $\theta_i > 0$. Figure 2 shows an example of a matching that is incompatible with $\theta_i \geq 0$.

Matching			EB	any $v_i^* < v_i$?
(v_1^*, v_1)	(v_2^*, v_2)	(v_3^*, v_3)		
(9, 8)	(11,9)	(13, 11)	1/3	no
(9, 9)	(11,8)	(13, 11)	1/3	no
(9, 8)	(11, 11)	(13,9)	1	no
(9, 9)	(11, 11)	(13,8)	1	no
(9,11)	(11,8)	(13,9)	5/3	yes
(9,11)	(11,9)	(13,8)	5/3	yes

Table 1: enumeration of all possible matchings in the simple example displayed in figures 1-2. Each line shows a possible matching $(v_1^*, v_1), (v_2^*, v_2), (v_3^*, v_3)$, the values of the excess burden of taxation associated with that matching (assuming a pre-tax price $p = 10$ as in the figures) and whether $v_i^* < v_i$ for any individual in that matching. Pairs in bold characters indicate individuals who change their choice when the tax is introduced.

3. General Framework

This section starts defining choice settings, and describing the process of making choices through individual choice functions and revealed valuations. Choice settings are categorized into two groups: *clear* and *noisy*, based on whether choices within these settings maximize a specified utility function among the available alternatives. The concept of policies is introduced as changes in choice settings. The subsequent discussion focuses on defining the welfare implications of a policy, with a specific focus on metrics such as equivalent variation and related measures. The main results of the paper are then presented. These express the equivalent variation of a policy in terms of revealed values across different scenarios, and establish a natural bridge between the distribution of welfare effects and demand functions. The identification problem is then discussed and bounds are obtained as solutions to two optimal transport problems. The section proceeds discussing assumptions under which narrower bounds can be obtained.

3.1. Choice, utility, and revealed valuations

Agent $i \in I = [0, 1]$ has a total amount of money m_i and chooses $x \in \{0, 1\}$. Choosing $x = 1$ costs p . Any money left goes into consumption of a unit cost numeraire good

y , therefore any feasible consumption bundle (x, y) satisfies with equality the budget constraint $xp + y = m_i$. To focus only on nontrivial choice problems, I assume that $m_i > p$ for all i .

Agent i 's utility from consuming (x, y) is

$$U_i(x, y) = \phi_i x + \varphi_i(y) \tag{5}$$

which satisfies the following

Assumption 1 $\varphi_i(y)$ is a strictly increasing continuous function and $\varphi_i^{-1}(\cdot)$ denotes the corresponding unique inverse.

This is the same assumption imposed on utility by [Bhattacharya \(2015\)](#). It basically requires that more of the numeraire is better, the continuity assumption is technical and guarantees that inverses are defined everywhere.

In this paper, I will think of choice as taking place in different environments. I will refer to these interchangeably using the terms environments, settings, situations, or states. I follow [Bernheim and Rangel \(2009\)](#) and define the choice environment faced by individual i as a set of feasible alternatives (the budget line, characterized by p and m_i) and ancillary conditions ζ . Ancillary conditions are features of the environment that may affect behavior, but which are not taken as relevant to a social planner's evaluation. Typical examples of ancillary conditions include the information available to i at the moment of choice, how information or alternatives are presented, the salience of a default option or of different components of price, the point in time at which a choice is made. A generic choice environment is therefore represented by the triplet (m, p, ζ) .

Optimal choice is a function of price p and monetary endowment m

$$x_i^*(m, p) = \arg \max_{x \in \{0,1\}} U_i(x, y) \quad \text{s.t. } m = xp + y$$

while i 's choice $x_i(m, p, \zeta)$ may be affected by ancillary conditions as well, and be suboptimal. I will say that the choice setting is **clear** if i chooses optimally, so that $x_i(m, p, \zeta) = x_i^*(m, p)$. When this is not the case I'll say that the setting is **noisy**.

As a regularity condition on choices, I require that individual demand for the binary good be weakly decreasing in its price.

Assumption 2 $x_i(m, p, \zeta)$ is weakly decreasing in p , for any pair (m, ζ) .

In light of assumption 2, individual choice in a given setting can be conveniently described defining i 's reservation price in that setting

$$v_i(m, \zeta) = \inf\{p : x_i(m, p, \zeta) = 0\}$$

In a clear setting, assumption 1 implies that $v_i(m, \zeta) = v_i^*(m)$, where

$$v_i^*(m) = m - \varphi_i^{-1}(\varphi_i(m) - \phi_i)$$

Therefore, assumption 2 is enough to rewrite demand functions in terms of the marginal distributions F_v and F_{v^*} as in equation (3).

3.2. policies, indirect utility, and equivalent variation

A *policy* π affects individuals' choice settings, transforming the pre-policy settings (p_i, m_i, ζ_i) into (p'_i, m'_i, ζ'_i) for all $i \in I$. To study the average welfare effects of π in the population. I focus on one of the most widespread money metrics for welfare effects: *equivalent variation*. When the policy of interest consists of the introduction of a new tax on the binary good, another measure of interest will be the efficiency cost, or *excess burden* of this tax, introduced in the motivating example. I will discuss this measure as well, and its connection to equivalent variation.

These welfare measures can be conveniently described through the indirect utility function.

$$V_i(m, p, \zeta) = U_i(x_i(m, p, \zeta), m - p \times x_i(m, p, \zeta))$$

The *equivalent variation* of policy π for individual i is defined as the change in wealth, at pre-policy conditions, that would have the same effect on i 's utility as would the policy. This can be expressed in terms of indirect utility as the smallest solution S to

$$V_i(m - S, p, \zeta) = V_i(m', p', \zeta') \quad (6)$$

Finally, following [Auerbach \(1985\)](#), I measure the *excess burden* (deadweight cost) of a tax as the amount of additional tax revenue that could be collected from the consumer i while keeping his utility constant, if the distortionary tax were replaced with a lump-sum tax. The excess burden of the tax on i is zero if i 's choice is not affected by the introduction of the tax and coincides with equivalent variation otherwise. So that results for the excess burden of taxation will easily follow from results on equivalent variation.

3.3. Results

3.3.a. Equivalent variation from a change in price.— Consider a generalized version of the problem discussed in the motivating example, where we replace the quasilinear utility specification with the expression in [5](#). In this case the ancillary condition ζ will represent the part of the total price which is not salient to individuals. For example, the total price of the product would be $q = p + \zeta$ where only p is salient (the pretax price in our previous example) while ζ is not (the total tax, in our example) and is taken into imperfect consideration. Just like before we represented imperfect consideration of the nonsalient component of the price by a parameter $\theta_i \geq 0$ so that i 's choice is

$$x_i(m, p + \zeta, \zeta) = \mathbb{1}\{\phi_i + \varphi_i(m - p - \theta_i \zeta) > \varphi_i(m)\} \quad (7)$$

while i 's optimal choice would be

$$x_i^*(m, p + \zeta) = \mathbb{1}\{\phi_i + \varphi_i(m - p - \zeta) > \varphi_i(m)\} \quad (8)$$

Choice coincides with optimal choice when the total price is salient, which in our example can happen when taxes are zero, or when tax inclusive prices are posted. We would represent this as $x_i(m, p, 0) = x_i^*(m, p)$.

The following theorem, proved in appendix A, studies the equivalent variation associated with an increase in the price of the binary good, when part of the price increase is non-salient.

Theorem 1 *Suppose that $x(m, p, \zeta)$ and $x^*(m, p)$ are as in (7) and (8), respectively. Consider a policy π that introduces a price increase, only part of which is salient. In other words, π changes the choice environment for i from $(m_i, p, 0)$ to $(m_i, p + \Delta p, \zeta)$, where $\zeta \in [0, \Delta p]$. Let EV_i denote the equivalent variation to π for i . Then*

$$\begin{array}{ll} EV_i = 0 & \text{if } v_i^*(m_i) \leq p \\ EV_i = v_i^*(m_i) - p & \text{if } v_i^*(m_i) > p \text{ and } v_i(m_i) \leq p + \Delta p \\ EV_i > [v_i^*(m_i) - p, p + \Delta p] & \text{if } p < v_i^*(m_i) < p + \Delta p \text{ and } v_i(m_i) > p + \Delta p \\ EV_i = \Delta p & \text{if } v_i^*(m_i) \geq p + \Delta p \text{ and } v_i(m_i) > p + \Delta p \end{array}$$

Some interesting results follow from this as special cases. For example [Bhattacharya \(2015\)](#) studies identification of the welfare effects of price changes in clear settings. If we set $\zeta = 0$ in theorem 1 above, we obtain one of the main results in that paper.

Corollary 2 [price change in clear setting] *If the price change is entirely salient ($\zeta = 0$), we recover an equivalent version³ of lemma 1 of [Bhattacharya \(2015\)](#). In this case, both pre-policy*

³The original one is expressed in terms of utility functions instead of reservation prices.

and post-policy settings are clear, so that $v_i(m) = v_i^*(m)$. Therefore the lemma above becomes

$$EV_i = \begin{cases} 0 & \text{if } v_i(m_i) \leq p \\ v_i(m_i) - p & \text{if } p < v_i(m_i) \leq p + \Delta p \\ \Delta p & \text{if } v_i(m_i) > p + \Delta p \end{cases}$$

A version of the first half of the main theorem in [Bhattacharya \(2015\)](#) also follows very naturally from this expression. Here EV_i is just a function of $v_i(m)$, whose distribution is identified from aggregate demand. It follows that the entire distribution of EV_i is identified by aggregate demand curve $D(\cdot)$

$$Pr\{EV_i \leq a\} = \begin{cases} 0 & \text{if } a < 0 \\ F_v(a + p) & \text{if } 0 \leq a \leq \Delta p \\ 1 & \text{if } a > \Delta p \end{cases}$$

A second corollary of the lemma is that the formula for EB_i presented in the motivating example remains valid without assuming quasilinear utility, and without assuming that the entire tax is not salient ⁴.

Corollary 3 [Excess Burden of Tax] *Suppose that $x(m, p, \zeta)$ and $x^*(m, p)$ are as in (7) and (8), respectively. Consider a policy π that introduces a tax, only part of which is salient, in a context where all previous taxes were salient. In other words, π changes the choice environment for i from $(m_i, p, 0)$ to $(m_i, p + tp, \zeta)$, where $\zeta \in [0, t]$. Let EB_i denote the excess*

⁴In our introduction we could have applied the exact same formulas to study the contextual introduction of two taxes, such as a salient excise tax and a non-salient sales tax, in a population with arbitrary utility functions.

burden of the tax for i . Then EB_i is identified by $v_i^*(m_i)$ and $v_i(m_i)$ as follows:

$$EB_i = \begin{cases} 0 & \text{if } v_i^*(m_i) \leq p \\ v_i^*(m_i) - p & \text{if } v_i^*(m_i) > p \text{ and } v_i(m_i) \leq p + t \\ 0 & \text{if } v_i(m_i) > p + t \end{cases}$$

In this case, EB_i is a function of two distinct reservation prices: $EB_i = EB(v_i^*, v_i)$ where

$$EB(v^*, v) = (v^* - p) \mathbb{1}\{p < v^*\} \mathbb{1}\{v \leq p + t\} \quad (9)$$

In this case, we would expect to be able to identify the entire distribution of EB_i from the joint distribution of v_i and v_i^* . Indeed, we can write

$$Pr\{EB_i \leq a\} = \begin{cases} 0 & \text{if } a < 0 \\ F_{v^*}(p) + 1 - F_v(p + t) & \text{if } a = 0 \\ Pr\{EB_i \leq 0\} + F_{v^*|M}(a + p) \times P_M & \text{if } a > 0 \\ 1 & \text{if } a > t + \zeta(\bar{\theta} - 1) \end{cases}$$

where $F_{v^*|M}$ is the cdf of v^* for individuals that are “marginal” in the sense that their transactions are discouraged by the tax: $F_{v^*|M}(\cdot) = Pr\{v_i^* \leq \cdot \mid v_i^*(m_i) > p, v_i(m_i) \leq p + t\}$. Similarly P_M is the fraction of marginal individuals in the population. Finally, $\bar{\theta}$ is any upper bound to the support of θ_i , possibly ∞ .

3.3.b. Equivalent variation from a generic policy, under quasilinear utility.– If we restrict to quasilinear utility, we can provide a similar result to theorem 1 for a broader class of policies:

Lemma 4 Suppose that utility is quasilinear and that under policy π , the choice environment for individual i changes from the clean setting (m, p, ζ) to (m', p', ζ') . Let EV_i denote the

equivalent variation to π for i . Then

$$EV_i = \begin{cases} -\Delta m & \text{if } v_i^*(m) \leq p \text{ and } v_i(m', \zeta') \leq p' \\ p' - \Delta m - v_i^*(m) & \text{if } v_i^*(m) \leq p \text{ and } v_i(m', \zeta') > p' \\ v_i^*(m) - p - \Delta m & \text{if } v_i^*(m) > p \text{ and } v_i(m', \zeta') \leq p' \\ \Delta p - \Delta m & \text{if } v_i^*(m) > p \text{ and } v_i(m', \zeta') > p' \end{cases}$$

where $\Delta m = (m' - m)$ and $\Delta p = (p' - p)$.

Proof. in appendix A. ■

3.4. Welfare Analysis

In the simple example considered in table 1, I obtained upper and lower bounds for EB going through all the possible joint distributions compatible with the observed marginal F_v and F_v^* . This approach can be easily generalized: let $\Pi(F_v, F_v^*)$ denote the set of all joint distributions with marginals F_v and F_v^* . We can bound EB from below and from above finding in this set the joint distributions that minimize, and maximize, the expectation of $EB(v^*, v)$. The best bounds we can obtain without making use of any information besides the one contained in the marginals of v^* and v are therefore the values of two optimal transport problems

$$EB_L = \inf_{\pi \in \Pi(F_v^*, F_v)} E_\pi[EB(v^*, v)], \quad (10)$$

and

$$EB_U = \sup_{\pi \in \Pi(F_v^*, F_v)} E_\pi[EB(v^*, v)]. \quad (11)$$

An *optimal transport* (OT) problem is any problem of the form

$$\inf_{\pi \in \Pi(F_X, F_Y)} E_\pi[f(X, Y)] \quad (12)$$

where $\Pi(F_X, F_Y)$ denotes the set of all joint distributions of X and Y with marginals F_X

and F_Y , and $E_\pi[\cdot]$ is the expectation taken with respect to the joint distribution π . If f is interpreted as a cost, problem (12) is the problem of transforming distribution F_X in F_Y as efficiently as possible, and when X and Y are continuous random variables it is the infinite-dimensional extension of the discrete problem of matching. This basic problem has several applications within mathematics and many other fields and is currently an extremely active research area both in theory and applications. Our problems in (10) and (11) are special cases of (12) where $F_X = F_{v^*}$, $F_Y = F_v$, and $f(\cdot, \cdot)$ is equal to $EB(\cdot, \cdot)$ in (9) or $-EB(\cdot, \cdot)$, respectively.

One of the few cases in which problem (12) admits a closed-form solution is when $f(\cdot, \cdot)$ is *submodular*, or *supermodular*. A function $f(\cdot, \cdot)$ is submodular if, for any $x < x'$ and $y < y'$ we have that $f(x, y) + f(x', y') \leq f(x, y') + f(x', y)$ and is *supermodular* if $-f(\cdot, \cdot)$ is submodular. The following lemma shows that $EB(\cdot, \cdot)$ is a submodular function and then applies standard results in optimal transport theory to derive a closed-form expression for the bounds, in terms of the marginal distributions F_{v^*} and F_v .

Lemma 5 *The function*

$$EB(x, y) = (x - p)\mathbb{1}\{p \leq x\}\mathbb{1}\{y \leq p + t\}$$

is submodular. Classical results in optimal transport theory then imply that EB_L and EB_U , as defined in (10) and (11), are

$$EB_L = E\left[EB\left(F_{v^*}^{-1}(U), F_v^{-1}(U)\right)\right]$$

and

$$EB_U = E\left[EB\left(F_{v^*}^{-1}(U), F_v^{-1}(1 - U)\right)\right]$$

where $U \sim \mathcal{U}(0, 1)$ and $F_{v^*}^{-1}$ and F_v^{-1} denote the quantile functions of v_i^* and v_i , defined as

$$F_{v^*}^{-1}(u) = \inf \left\{ x : F_{v^*}(x) > u \right\}$$

$$F_v^{-1}(u) = \inf \left\{ x : F_v(x) > u \right\}$$

Proof.

Let $x < x'$ and $y < y'$. We need to show that

$$EB(x, y) + EB(x', y') \leq EB(x, y') + EB(x', y).$$

Note that $EB(\cdot, y)$ is increasing for any y , therefore $EB(x, y) \leq EB(x', y)$. Similarly $EB(x, \cdot)$ is decreasing for any x , therefore $EB(x', y') \leq EB(x', y)$. This proves that $EB(\cdot, \cdot)$ is submodular. The expressions for the bounds follow then by well-known results in optimal transport theory (see e.g. Galichon (2017) Thm. 4.3), which generalize the Fréchet-Hoeffding inequalities to submodular functions. ■

The result states that the lower and upper bounds are obtained when the marginal distributions are matched via the Fréchet-Hoeffding lower and upper copulas, respectively. An intuitive interpretation of this fact in the non-salient taxes example, is that the tax deadweight loss is minimized if the tax introduction does not impact agents' ranking in terms of their reservation prices. In other words, while the bias may affect how many transactions are discouraged by the tax, it does not affect the order in which agents leave the market as the tax increases. If this happens, then the first transactions to be discouraged when a tax is introduced are those of the buyers with the lowest valuations for the product, whose surplus from the deal is minimal. In our framework, this happens for example if θ_i is homogeneous in the population. If $\theta_i = \theta$ for all i , then the revealed valuation v_i is just a rescaled version of the true valuation as $v_i = v_i^*/(1 + \theta\tau)$. Agents rankings in terms of v_i and v_i^* must therefore coincide as $v_i \geq v_j$ if and only if $v_i^* \geq v_j^*$.

The upper bound obtained via the Fréchet-Hoeffding upper copula, on the other side,

requires that individual rankings are reversed when the tax is introduced. The individual holding the highest valuation in the clear setting becomes the one with the lowest revealed valuation in the noisy setting, and so on. While this upper bound is typically much larger than the lower one, it depicts a situation that may be incompatible with our theoretical model. For example, requiring that the individual with the lowest v_i^* be also the one with the highest v_i typically implies that for this individual $\theta_i < 0$, which is ruled out by our model. We may therefore hope to obtain more informative upper bounds by restricting the set of joint distributions over which the optimization in (11) takes place to those which are also compatible with our modelling assumptions.

3.5. Narrowing bounds via additional assumptions

Identified sets based on Fréchet-Hoeffding bounds are common in the literature on distributional treatment effects, that studies the identification and estimation of functionals of the treatment effects distribution beyond the ATE. Under standard experiments and in many quasi-experimental setups, researchers can identify the marginal cdfs F_0 and F_1 of potential outcomes $Y(0)$ and $Y(1)$. For reasons analogous to the ones discussed above for welfare effects, researchers cannot identify parameters that depend on the joint distribution of potential outcomes such as the fraction of individuals harmed by a treatment, the median treatment effect, or the expected treatment effect conditional on the potential outcome in the control state. Most treatment effects parameters of interest can be expressed as $\theta_0 = E[\nu(\Delta)]$, where Δ denotes the difference between two potential outcomes $\Delta = Y(1) - Y(0)$. The problem of finding an upper bound for these parameters based uniquely on the observed marginals is, therefore, a version of problem (12) where $F_X = F_0$, $F_Y = F_1$, and the loss function $f(X, Y)$ is given by $-\nu(Y - X)$. For example, the fraction of individuals harmed by treatment can be expressed as $\theta_0 = E[\mathbb{1}(\Delta \leq 0)]$, and the α -quantile treatment effect $\theta_0 = F_\Delta^{-1}(\alpha)$ can be bounded inverting the bounds for $F_\Delta(\delta) = E[\mathbb{1}(\Delta \leq \delta)]$. In this literature, bounds based on Fréchet-Hoeffding copulas tend to be wide and preclude meaningful inference, so research has developed to come up with additional restrictions to tighten them.

Some of these assumptions impose constraints on the support of the joint distribution of the two potential outcomes, requiring that $Pr((Y(0), Y(1)) \in C) = 1$ for some closed set $C \in \mathbb{R}^2$. The general case is discussed in [Kim \(2014\)](#), but special cases such as assuming monotone treatment effects, or a Roy selection model, were among the first restrictions to be studied in the literature ([Manski \(1997\)](#), [Heckman et al. \(1997\)](#)). Other restrictions grant point identification but may be too strong to be palatable in many empirical settings. One example is the assumption of constant treatment effects or the slightly more general assumption of rank invariance ([Heckman et al. \(1997\)](#), [Chernozhukov and Hansen \(2005\)](#)) which would be equivalent to assuming a lower Fréchet bound. Similarly, assuming that gains Δ are independent of the base state $Y(0)$ implies that the distribution of $Y(1)$ is a convolution of $Y(0)$ and Δ , and deconvolution methods can be applied to point identify the distribution of Δ from F_1 and F_0 ([Heckman et al. \(1997\)](#)).

An assumption that seems to have wider applicability is studied in [Frandsen and Lefgren \(2021\)](#) and requires that the two potential outcomes be stochastically increasing in each other. Contrarily to rank invariance, this allows individuals with the same $Y(0)$ to have different values of $Y(1)$. If $Y(1)$ is stochastically increasing in $Y(0)$, though, individuals with high values of $Y(0)$ must draw their values of $Y(1)$ from more favorable (first order stochastically dominant) distributions. Other similar but weaker requirements of positive dependence between potential outcomes are also used to narrow the Fréchet bounds. Some examples are positive quadrant dependence ([Bhattacharya et al. \(2012\)](#)), stochastic dominance ([Blundell et al. \(2007\)](#)), and imposing bounds on measures of association such as Kendall's τ or Spearman's ρ ([Heckman et al. \(1997\)](#), [Fan and Soo Park \(2009\)](#)).

The difference between the problem of finding an upper bound to the excess burden of a tax and bounding distributional treatment effects is therefore the different loss function $f(X, Y)$ that these problems imply in their OT formulation [\(12\)](#). For this reason, results obtained using these restrictions in the treatment effects literature cannot be applied directly to our case, and minor modifications are required.

The assumptions that I will be using to narrow the upper bound for EB are two: the first is a version of the support assumption, while the second is the assumption that v_i is stochastically increasing in v_i^* .

Assumption 3 $Pr\left(v_i^* + \zeta(1 - \bar{\theta}) \leq v_i \leq v_i^* + \zeta\right) = 1$ for known $\bar{\theta} < \infty$.

This follows directly from the fact that in our model the parameter θ_i is non-negative and bounded.

Assumption 4 $Pr\left(v_i \leq p \mid v_i^* = x\right)$ is (weakly) decreasing in x .

This assumption requires that individuals with higher valuations v_i^* for the binary good are more likely to buy it also in the presence of taxes, and rules out an exceedingly strong positive correlation between v_i^* and θ_i . This is for example weaker than assuming that biases are independent of product valuations. In the upper bound derivation, this property will only be required to hold almost surely with respect to F_{v^*} .

Proposition 6 [upper bound] *Under Assumptions 3 and 4, $EB \leq EB_U$ where*

$$EB_U = \mathbb{E}\left[(v_i^* - p) \times \mathbb{1}\{p \leq v_i^* \leq q\}\right] \times \left(\frac{F_v(p) - F_{v^*}(p)}{F_{v^*}(q) - F_{v^*}(p)}\right)$$

where $q = p + \bar{\theta}t$.

Proof. The function $\kappa_0(x) = Pr(v_i \leq p \mid v_i^* = x)$ is sufficient to identify EB . Indeed

$$\begin{aligned} EB &= E\left[(v_i^* - p) \times \mathbb{1}(p < v_i^*) \times \mathbb{1}(v_i \leq p + t)\right] \\ &= E\left[(v_i^* - p) \times \mathbb{1}(p < v_i^*) \times E\left[\mathbb{1}(v_i \leq p + t) \mid v_i^*\right]\right] \\ &= E\left[(v_i^* - p) \times \mathbb{1}(p < v_i^*) \times \kappa_0(v_i^*)\right] \end{aligned} \tag{13}$$

Both assumptions impose constraints on this function. Under assumption 4

$$k_0(x) \leq k_0(x') \quad \text{for all } x' > x \quad (14)$$

while under assumption 3

$$\kappa_0(x) = 1 \quad \text{if } x \leq p \quad (15)$$

$$\kappa_0(x) = 0 \quad \text{if } x > q \quad (16)$$

where $q = p + \bar{\theta}t$. Always under assumption 3, I prove (claim 8 in appendix) that

$$F_v(p) - F_{v^*}(p) = \int_p^q \kappa_0(z) dF_{v^*}(z) \quad (17)$$

Therefore under assumptions 3 and 4 we know that $\kappa_0 \in \mathcal{K}$, where

$$\mathcal{K} = \left\{ \kappa : \mathbb{R} \rightarrow [0, 1] \text{ such that (14) (15) (16) and (17) hold replacing } \kappa \text{ to } \kappa_0 \right\}$$

Expression (13) implies that $EB \leq EB_U$, where

$$EB_U = \sup_{\kappa \in \mathcal{K}} E \left[(v_i^* - p) \times \mathbb{1}(p < v_i^*) \times \kappa(v_i^*) \right].$$

Since k can be seen as a decreasing weighting function (whose integral is fixed), the sup is attained as a max by the only $\kappa^* \in \mathcal{K}$ which is constant over the interval $[p, q]$. The value of the constant is pinned down by (17) and is

$$K = \frac{F_v(p) - F_{v^*}(p)}{F_{v^*}(p(1 + \bar{\theta}\tau)) - F_{v^*}(p)}$$

So we have that

$$\kappa^*(x) = \mathbb{1}(x \leq p) + \mathbb{1}(x \leq q) \times K$$

EB_U is the optimand evaluated at $\kappa^*(x)$. ■

The idea behind this result is that EB is a weighted average of $(v_i^* - p)$, where the weights are given by $\kappa(x) = Pr(v_i \leq p \mid v_i^* = x)$ so that knowledge of κ would allow us to identify EB . Assumptions 3 and 4 impose restrictions on κ , so we can construct an upper bound maximizing the weighted average expression for EB over all weighting functions that satisfy these restrictions. In particular, the assumption that v_i is stochastically increasing in v_i^* imposes non-increasing weights, while the support assumption implies zero weight on values of $v_i^* > q$. The weighting function that maximizes the expression for EB is then constant between p and q , and zero after. Finally, the value of the constant is pinned down as the only one compatible with the total amount of discouraged transactions $F_v(p) - F_{v^*}(p)$.

Proposition 7 [lower bound] *If $\theta_i > 0$ almost surely, then $EB \geq EB_L$ where*

$$EB_L = \mathbb{E} \left[(v_i^* - p) \times \mathbb{1} \left\{ p \leq v_i^* \leq F_{v^*}^{-1}(F_v(p)) \right\} \right]$$

This bound is sharp and remains such under Assumptions 3 and 4.

This is a restatement of the result for the lower bound in section 3.4. The assumption $\theta_i > 0$ is not used to derive the bound and is only necessary for the characterization of the excess burden of taxation given in (1) to hold. The bound is sharp because there is a joint distribution of v_i^* and v_i that attains it. In particular, $EB = EB_L$ when the ranking of individuals is the same regardless of whether they are sorted by v_i^* or by v_i . It remains sharp under both assumptions because this particular joint distribution is not ruled out by either.

3.6. point identification in special cases

A simple but important fact proved as claim 8 in appendix A is that, under assumption 3 we can identify how many transactions are affected by the tax, and

$$Pr \{ x_i(m, p, 0) \neq x_i(m, p, t) \} = F_v(p) - F_{v^*}(p) \quad (18)$$

In some special cases this gives us point identification of EB . For example, an immediate consequence of equation (18) is that EB is point identified when it is equal to zero ⁵. Indeed this happens only when no transactions are discouraged by the tax, which is point identified under Assumption 3. Similarly, we can point identify EB when the share of discouraged transactions is maximal. Under Assumption 3 we know that the only transactions that can be discouraged are those of individuals with $p < v_i^* \leq q$, where $q = p + \bar{\theta}t$ ⁶. The amount of individuals with v_i^* in this range is $F_{v^*}(q) - F_{v^*}(p)$ and because of (18) it coincides with the amount of discouraged transactions when $F_v(p) = F_{v^*}(p + \bar{\theta}t)$.

In this case

$$EB = \left[v_i^* - p \mid p < v_i^* \leq q \right] \times Pr\left\{ p < v_i^* \leq q \right\}$$

Two important cases that lead to the situations above are when θ is degenerate on one corner of the admissible region $[0, \bar{\theta}]$. Indeed if $\theta = 0$ with probability one, the two marginals are pointwise equal, not just at p . Similarly if $\theta = \bar{\theta}$ with probability one $F_v(x) = F_{v^*}(x(1 + \bar{\theta}\tau))$ for all values of x , not just when $x = p$.

4. Application: tax salience

The goal of the empirical section of this paper is to show how inference conducted using aggregate demand curves alone, giving up point identification of the excess burden of taxation, can be as informative as inference obtained using methods that impose higher data requirements, and stronger assumptions, with the objective of obtaining point identification.

⁵Or, equivalently, when $F_{v^*}(\cdot)$ and $F_v(\cdot)$ coincide at the product price p .

⁶To see this, note that under Assumption 3 we know that individuals with $v_i^* < p$ will not buy the product regardless of the tax, since for them $v_i \leq v_i^* + t < p + t$. Similarly, individuals with $v_i^* \geq p + \bar{\theta}t$ will buy the product in any case, since for them $v_i \geq v_i^* + t(1 - \bar{\theta}) \geq p + t$.

4.1. Data and Results

We estimate the bounds described in the previous section using data from an experiment run by [Taubinsky and Rees-Jones \(2018\)](#). For the sake of brevity, we will only describe the experiment in its fundamental aspects. We refer the reader to the original paper by [Taubinsky and Rees-Jones \(2018\)](#) for a comprehensive and detailed description.

The experiment consisted of a series of shopping decisions involving twenty common household products. For each product, consumers saw a picture and a product description drawn from Amazon.com. Consumers then used a slider to select the highest tag price at which they would be willing to purchase the product. These decisions were taken by subjects in two consecutive modules of the experiment: in the first module, consumers made shopping decisions with either a zero tax rate (no-tax arm), or a standard tax rate corresponding to their city of residence (standard-tax arm). Each consumer was randomly assigned to one of the two treatment arms. The second module of the experiment was a repetition of the first one, except that all consumers made shopping decisions with a zero tax rate.

We will not use data from the second module of the experiment to estimate the bounds for EB and construct the corresponding confidence intervals. We will only use these data to compare our results to those obtained with the methods used in [Taubinsky and Rees-Jones \(2018\)](#), where within-subject measurements and additional assumptions are required to establish point identification. We see in the weaker data requirements of our approach one of its main advantages, as data sources that enable identification of the marginal distributions of two potential outcomes (the aggregate demands) are much more common than those that allow the identification of their joint distribution.

This experiment was designed specifically to elicit from consumers the reservation prices we introduced in section [2.1](#). Indeed, letting p_1^{ik} denote the reservation price revealed in

module 1 of the experiment by agent i for product k , we have

$$p_1^{ik} = \begin{cases} v_{ik}^* & \text{if } i \text{ assigned to no-tax arm} \\ w_{ik} & \text{if } i \text{ assigned to standard-tax arm} \end{cases}$$

where we are allowing consumers' valuations and biases to vary with the products. This data is sufficient to estimate the bounds defined in the previous section and conduct inference on EB for each product k . We estimate the bounds using the following sample averages taken over the no-tax arm of the experiment (of size N_0)

$$\widehat{EB}_L^k = \frac{1}{N_0} \sum_i (p_1^{ik} - p_k) \mathbb{1}\{p_k \leq p_1^{ik} \leq \widehat{F}_{vk}^{-1}(\widehat{F}_{wk}(p))\} \quad (19)$$

$$\widehat{EB}_U^k = \frac{1}{N_0} \sum_i (p_1^{ik} - p_k) \times \mathbb{1}\{p_k \leq p_1^{ik} \leq p_k(1 + \bar{\theta}\tau)\} \times \frac{\widehat{F}_w^k(p) - \widehat{F}_v^k(p)}{\widehat{F}_v^k(p(1 + \bar{\theta}\tau)) - \widehat{F}_v^k(p)} \quad (20)$$

Where \widehat{F}_v^k and \widehat{F}_w^k are the empirical distribution functions of p_1^{ik} in the no-tax and standard-tax arms, respectively. We obtain 95% confidence intervals via bootstrapping. Our results are shown in Figure 3.

Let now p_2^{ik} be the price selected in module 2 by agent i for product k . According to our setup, we should have that for all participants to the experiment $p_2^{ik} = v_{ik}^*$ ⁷. In practice, there might be systematic differences in the values elicited from the two modules, which we refer to as *order effects*. This is a phenomenon commonly found in pricing experiments, and the control arm of the experiment was designed by Taubinsky and Rees-Jones to accommodate these effects. For simplicity suppose that these effects take the form of an additive error ϵ_{ik} so that

$$p_2^{ik} = v_{ik}^* + \epsilon_{ik}$$

⁷we would therefore be able to observe $\theta_{ik} \approx \log(p_2^{ik}) - \log(p_1^{ik})$ for individuals in the standard-tax arm.

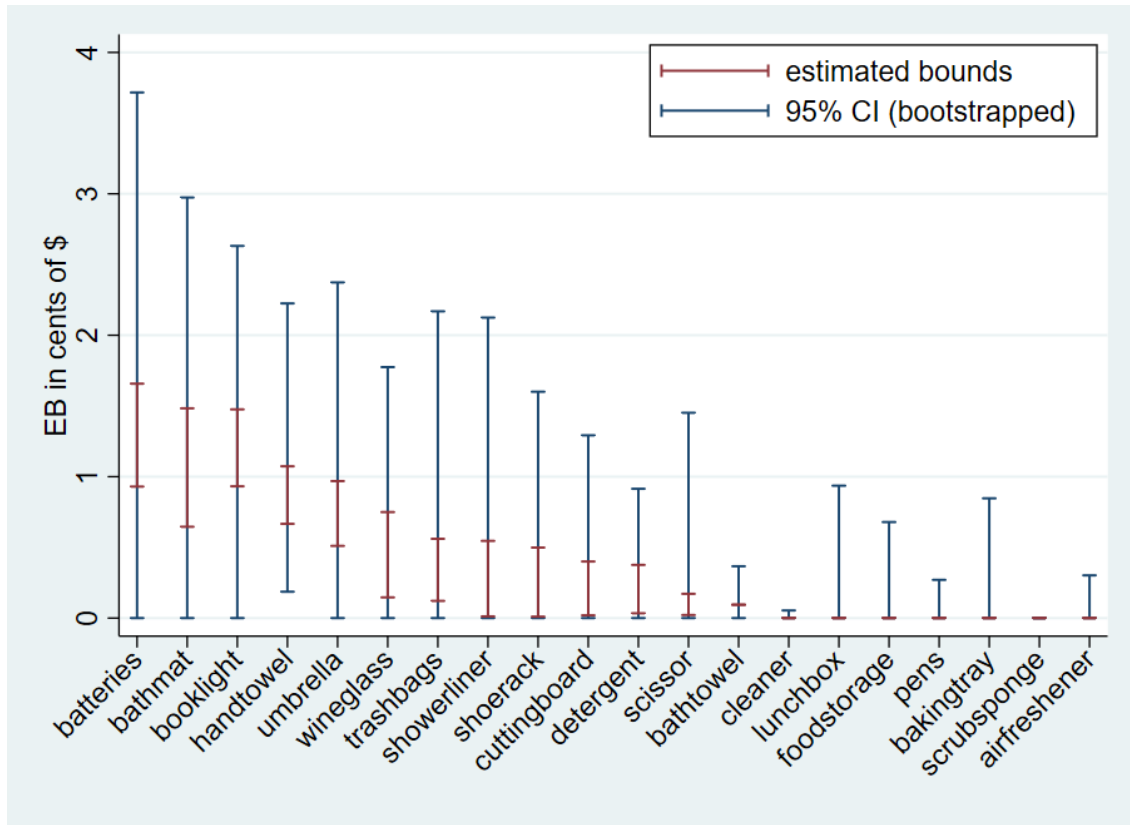


Figure 3: Identified regions for the excess burden of taxation for each of the products. 95% confidence intervals obtained by bootstrap.

4.2. Compare to Taubinsky and Rees-Jones (2018)

Taubinsky and Rees-Jones (2018) suggest an approximation to the excess burden of taxation

$$EB \approx -\frac{1}{2}t^2 \left(E[\theta_i|p, t]D_t(p, t) + Var(\theta_i|p, t)D_p(p, t) \right) \quad (21)$$

where D_t and D_p denote partial derivatives of demand with respect to p and t , and $E[\theta_i|p, t]$ and $Var[\theta_i|p, t]$ are the mean and variance of θ_i for consumers who are indifferent about purchasing the product when its price is p and taxes are t .

This is a good approximation if demand $D(p, t)$ as a function of price p and total taxes $t = p(1 + \tau)$ is approximately linear, and if θ_i is approximately independent of v_i^* when $v_i^* \in [p, p(1 + \bar{\theta}\tau)]$. In this case (21) is equivalent to

$$EB \approx EB_L + Var(\theta_i|p, t) \times EB^* \quad (22)$$

where EB_L is the lower bound for EB described above, and EB^* is the excess burden of taxation under salient taxes:

$$EB^* = \mathbb{E} \left[(v_i^* - p) \times \mathbb{1} \{ p \leq v_i^* \leq p(1 + \tau) \} \right]$$

which is point identified by F_{v^*} and can be estimated by the sample average in the no-tax arm

$$\widehat{EB^*} = \frac{1}{N} \sum_i (v_i^* - p) \times \mathbb{1} \{ p \leq v_i^* \leq p(1 + \tau) \}$$

Since $Var[\theta_i|p, t]$ is not identified by $D(p)$ and $D^*(p)$, the expression in (22) is also not identified. Anyways, it will simplify the comparison of our results to those that Taubinsky and Rees-Jones obtain under stricter assumptions and using more data.

Taubinsky and Rees-Jones provide estimators for $E[\theta_i|p, t]$ and $Var(\theta_i|p, t)$ that use data from both modules of their experiment, and exploit within-subject elicitation of both potential reservation prices. These estimators are valid under the following assumption, which fundamentally requires that order effects do not vary by experimental arm.

Assumption 5 For any vector of covariates X_{ik} , $E[\log(p_2^{ik}) - \log(p_1^{ik}) - \log(1 + \theta_{ik}\tau)|X_{ik}]$ does not depend on τ .

We want to check whether these additional data and assumptions are useful for obtaining narrower confidence intervals for EB . To do so, we use their estimator for $Var(\theta_i|p, t)$, together with the approximation in (22) and our estimators for EB_L and EB^* , to obtain a point estimate for EB . We bootstrap this estimate and compare the resulting 95% confidence intervals with the ones we obtained using only the first module of the experiment, without the need to impose linearity of the demand function or any particular property of the order effects. The results are shown in Figure 4.

4.3. Discussion

From a comparison of the confidence intervals obtained using the two alternative approaches, the first thing we notice is that for most products they seem to substantially coincide, and that only for a few products (thrashbags, shoerack, detergent) leveraging additional data and assumptions seems to bring substantially tighter intervals. The second is that for some products (in particular for booklight, but also for others) the intervals have comparable measures, but they do not overlap for a good part. We suggest that this is due to the fact that $D(p, t)$ is not well approximated by a linear function, and that this undermines the validity of the approximation that underpins the approach of Taubinsky and Rees-Jones.

The approximation in (22) is good if $F_{v*}(\cdot)$ is approximately linear over the interval $[p, p(1 + \bar{\theta}\tau)]$. For most of the products in our sample, F_{v*} displays strong nonlinearities around integer values, as shown in Figure 5, while it is approximately linear between these

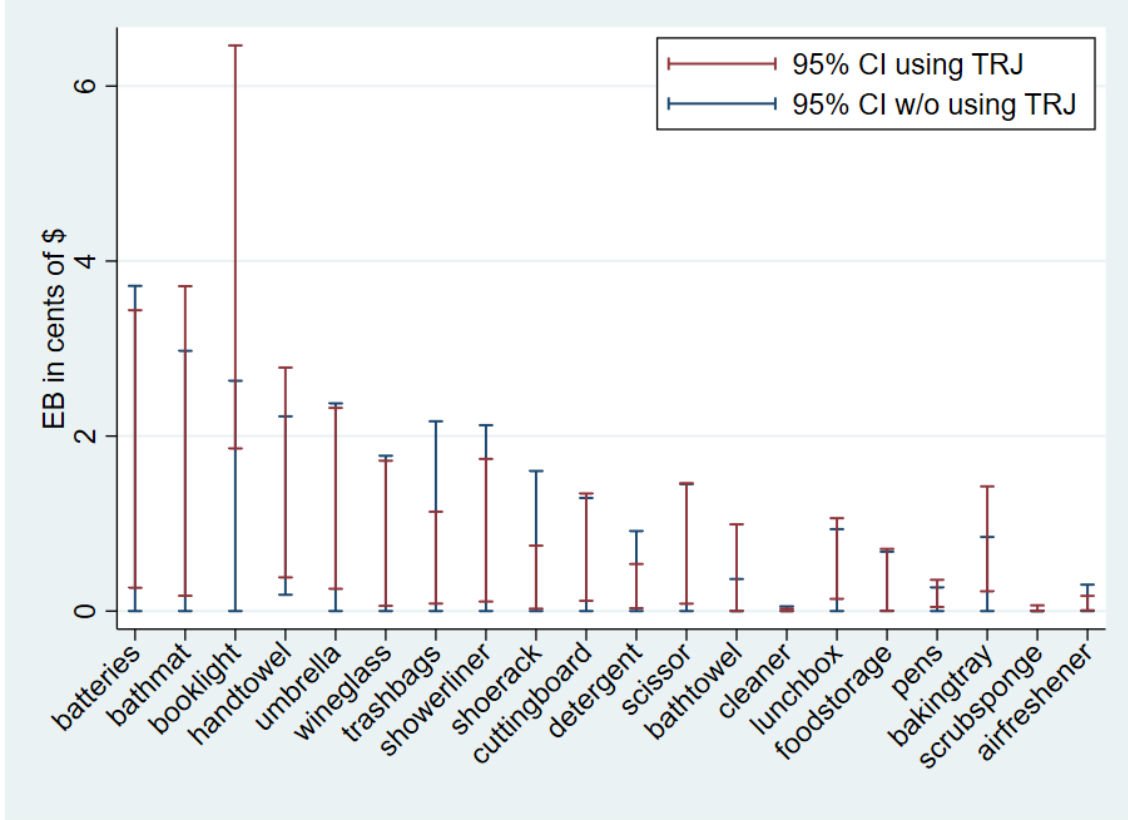


Figure 4: In blue: 95% confidence intervals for EB obtained using only data from module 1 of the experiment as in Figure 3. In red: 95% confidence intervals for EB obtained using also within-subject data (from module 2 of the experiment) and additional assumptions for point-identification of EB , as detailed in Taubinsky and Rees-Jones (2018).

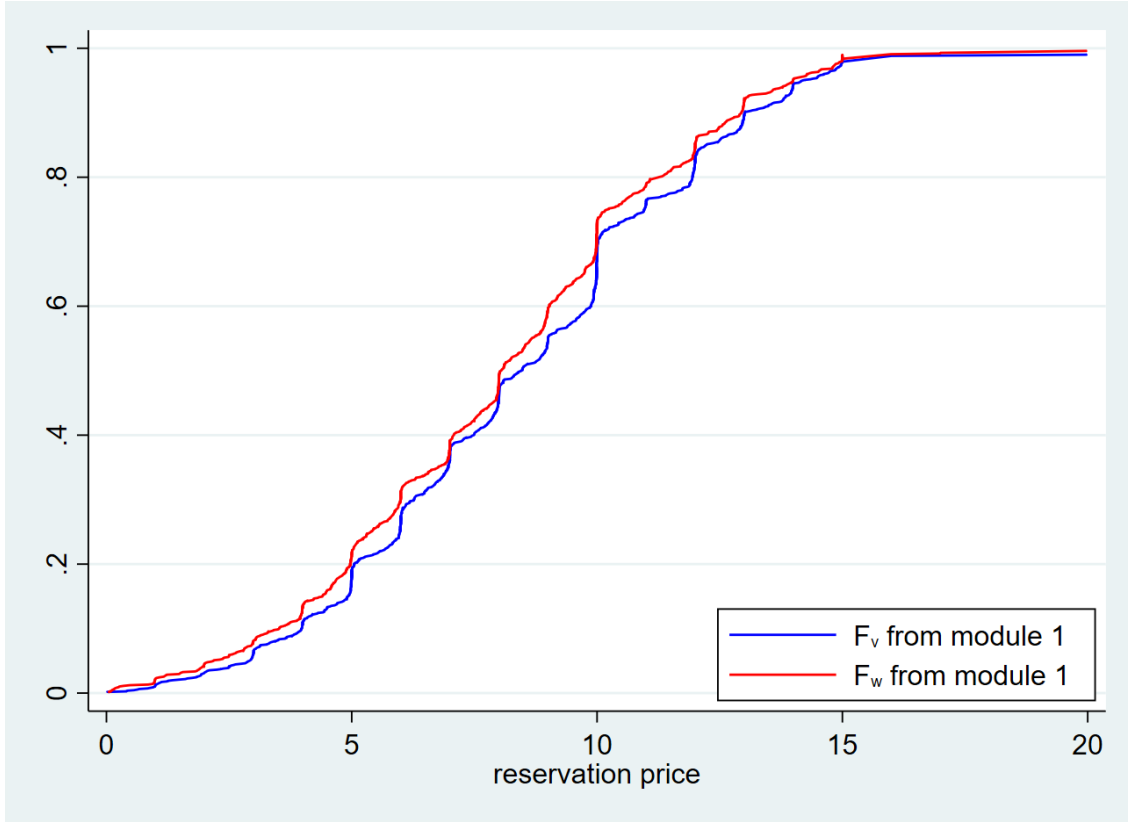


Figure 5: Empirical cdfs for v_i^* and v_i for the product "bathmat".

values. We would therefore expect the assumption to be particularly weak when a major nonlinearity falls in the interval $[p, p(1 + \bar{\theta}\tau)]$. We test this computing the confidence intervals under different product prices, chosen so that one of the main nonlinearities of F_{v^*} is included in the interval. As shown in Figure 6, this drives the confidence intervals computed following Taubinsky and Rees-Jones (2018) away from those that do not rely on the linearity assumption.

5. Summary and Conclusion

This paper underscores a fundamental insight: methods initially developed for nonparametric classical welfare analysis in discrete choice scenarios naturally extend to the realm of behavioral welfare analysis. These methods also establish a direct connection with the litera-

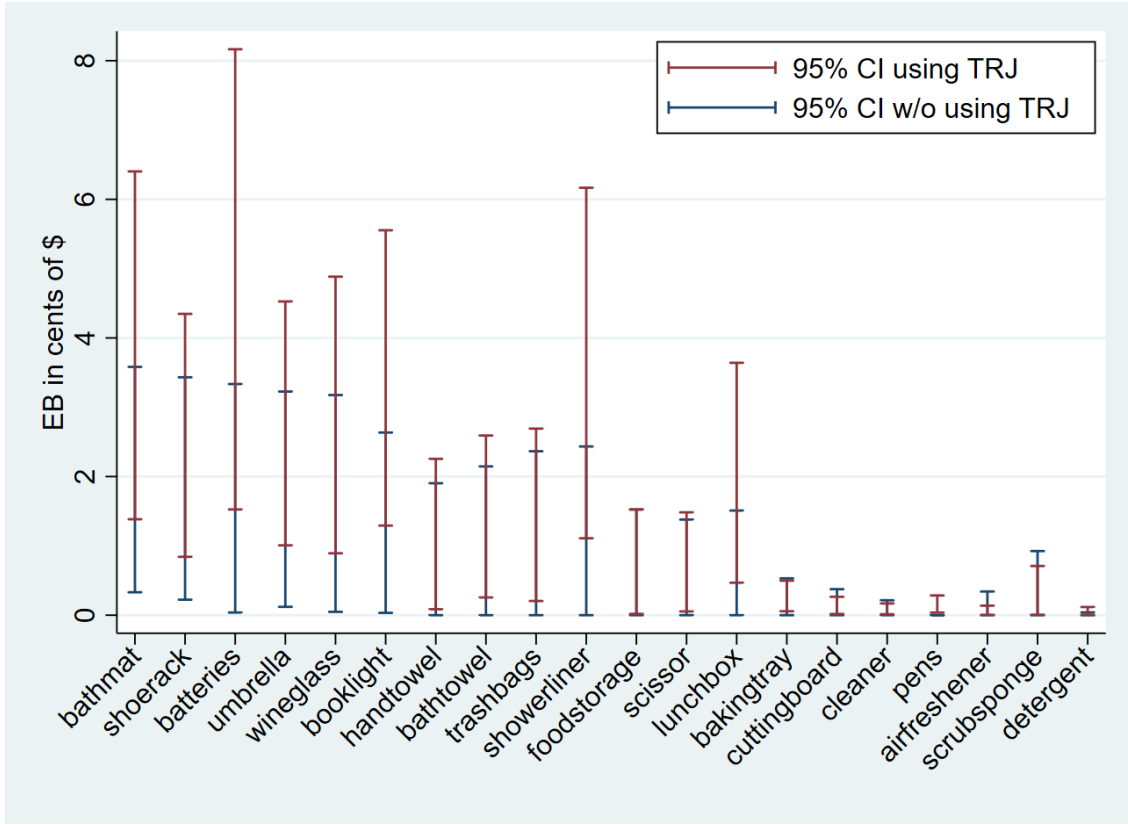


Figure 6: Replication of the results from Figure 4, but replacing actual product prices with either \$9.5 or \$4.5, depending on the shape of F_{v^*} .

ture on distributional treatment effects and, more broadly, with research employing optimal transport methods to explore partial identification. Importantly, this approach accommodates arbitrarily heterogeneous behavioral biases while exclusively relying on between-subject (quasi-)experimental variation. This differentiation sets it apart from earlier works in behavioral public finance, which assumed homogeneous behavioral biases in similar data settings. It also distinguishes it from recent research, which accommodates heterogeneous biases but necessitates within-subject variation and additional assumptions. Furthermore, the absence of parametric assumptions on demand curves in our bounds distinguishes this paper from others in behavioral public finance adopting a sufficient statistics approach reliant on demand linearization.

In conclusion, our approach offers distinctive characteristics that augment existing methods in the field of behavioral welfare analysis. Foremost, it emphasizes directness, enabling a clear and straightforward interpretation of underlying assumptions. This clarity empowers users to better understand foundations and the implications of the assumptions, and promotes greater trust and reliability in the results. Our methodology demonstrates the potential to provide informative bounds, with evidence from a validation exercise suggesting that it may yield confidence intervals comparable to those derived from existing approaches. Moreover, our approach exhibits notable flexibility, making it applicable to standard experimental or quasi-experimental data. This versatility reduces the necessity for within-subjects elicitation of willingness to pay in controlled lab settings, expanding the range of real-world applications where it can be effectively employed. Lastly, our approach aligns closely with the prevailing practices in the related literature on treatment effects distribution. By incorporating methodologies and principles that resonate with the current state of the field, our approach maintains relevance and consistency with contemporary research practices. In summary, our approach stands as a valuable contribution to the field, offering simplicity, transparency, precision, flexibility, and alignment with current research trends.

Appendix

A. Proofs

Proof of lemma 1.

First note that in this setting $v_i(m)$ is characterized by

$$\phi_i + \varphi_i(m - v_i(m) - \zeta(\theta_i - 1)) = \varphi(m)$$

therefore

$$\begin{aligned} v_i(m) &= m - \varphi_i^{-1}(\varphi_i(m) - \phi_i) + \zeta(1 - \theta_i) \\ &= v_i^*(m) + \zeta(1 - \theta_i). \end{aligned}$$

If $\theta_i \in [0, \bar{\theta}]$, this implies

$$\begin{aligned} v_i^*(m) &\leq v_i(m) + \zeta(\bar{\theta} - 1) \\ v_i(m) &\leq v_i^*(m) + \zeta \end{aligned}$$

Secondly, note that

$$\begin{aligned} V_i(m - S, p, 0) &= \max \left\{ \phi_i + \varphi_i(m - S - p), \varphi_i(m - S) \right\} \\ V_i(m, p + \Delta p, \zeta) &\in \left\{ \phi_i + \varphi_i(m - \Delta p - p), \varphi_i(m) \right\} \end{aligned}$$

So that for $S < 0$ we must have $V_i(m - S, p, 0) > V_i(m, p + \Delta p, \zeta)$.

Now, since EV_i is the smallest solution S to

$$\underbrace{V_i(m - S, p, 0)}_{LHS} = \underbrace{V_i(m, p + \Delta p, \zeta)}_{RHS} \quad (23)$$

we must have

$$EV_i \geq 0. \quad (24)$$

Finally, remember that by definition of $v_i^*(m)$, we have $\phi_i + \varphi_i(m - v_i^*(m)) = \varphi_i(m)$, therefore for any $S \in [0, v_i^*(m) - p]$

$$\begin{aligned} \phi_i + \varphi_i(m - S - p) &\geq \phi_i + \varphi_i(m - v_i^*(m)) \\ &= \varphi_i(m) \\ &\geq \varphi_i(m - S) \end{aligned}$$

where the inequalities follow from the fact that $\varphi_i(\cdot)$ is strictly increasing. This implies that

$$x_i^*(m - S, p) = 1 \quad \forall S \in [0, v_i^*(m) - p] \quad (25)$$

Now we can start analyzing the different cases enunciated in the lemma.

Case $v_i^*(m_i) \leq p$:

Since $v_i(m) \leq v_i^*(m) + \zeta$ we know that $v(m) \leq p + \Delta p$ and therefore $x_i(m_i, p + \Delta p, \zeta) = 0$. This means that RHS in (23) is equal to $\varphi_i(m)$. Also, since $v_i^*(m) < p$, we know that $x_i^*(m_i - S, p, 0) \Big|_{S=0} = 0$ and the LHS is also equal to $\varphi_i(m)$. Since $EV_i \geq 0$ by (24), we must have $EV_i = 0$.

Case $p < v_i^*(m_i)$ and $v_i(m_i) \leq p + \Delta p$:

We know from (25) that LHS of (23) is $\phi_i + \varphi_i(m_i - S - p)$ for all $S \in [0, v_i^*(m_i) - p]$. This is strictly decreasing in S and, for $S = v_i^*(m_i) - p$, it becomes $\phi_i + \varphi_i(m_i - v_i^*(m_i)) = \varphi_i(m_i)$. The RHS of (23) is also equal to $\varphi_i(m_i)$, given that $v_i(m_i) \leq p + \Delta p$. This implies that $EV_i = v_i^*(m_i) - p$.

Case $p < v_i^*(m_i) < p + \Delta p$ and $v_i(m_i) > p + \Delta p$:

We know from (25) that LHS of (23) is $\phi_i + \varphi_i(m_i - S - p)$ for all $S \in [0, v_i^*(m_i) - p]$. This is

strictly decreasing in S and, for $S = v_i^*(m_i) - p$, it becomes $\phi_i + \varphi_i(m_i - v_i^*(m_i)) = \varphi_i(m_i) > \phi_i + \varphi_i(m_i - p - \Delta p)$, which is the *RHS* of (23) given that $v_i(m_i) > p + \Delta p$. This implies that $EV_i \geq v_i^*(m_i) - p$, which proves the lower bound. Now consider $S = p + \Delta p$. We have that $V_i(m - S, p, 0) = \max\{\varphi_i(m - p - \Delta p), \phi_i + \varphi_i(m - 2p - \Delta p)\}$, where either of the terms is smaller than the *RHS*. Therefore must be $EV_i < p + \Delta p$.

Case $v_i^*(m_i) \geq p + \Delta p$ and $v_i(m_i) > p + \Delta p$:

We know from (25) that *LHS* of (23) is $\phi_i + \varphi_i(m_i - S - p)$ for all $S \in [0, \Delta p]$. This is strictly decreasing in S and, for $S = \Delta p$, it becomes $\phi_i + \varphi_i(m_i - p - \Delta p)$. The *RHS* of (23) is also equal to $\phi_i + \varphi_i(m_i - p - \Delta p)$, given that $v_i(m_i) > p + \Delta p$. This implies that $EV_i = \Delta p$. ■

Proof of Lemma 4.

We know from (6) that EV_i is the smallest solution S to

$$V_i(m - S, p, \zeta_i) = V_i(m', p', \zeta')$$

also, since utility is quasilinear, we can rewrite without loss of generality

$$U(x, y) = \phi_i + y$$

and optimal reservation prices do not depend on income: $v_i^*(m) = \phi_i$ for all $m \in \mathbb{R}_+$.

Case $v_i^*(m) \leq p$ and $v_i(m', \zeta') \leq p'$:

Since utility is quasilinear, $v_i^*(m - S) = v_i^*(m) \leq p$ and $x_i^*(m - S, p) = 0$. Since $v_i'(m', \zeta') \leq p'$, we also have $x_i(m', p', \zeta') = 0$ and can therefore rewrite (6) as

$$m - S = m'$$

which has unique solution $EV_i = -\Delta m$

Case $v_i^*(m) \leq p$ and $v_i(m', \zeta') > p'$:

Since utility is quasilinear, $v_i^*(m - S) = v_i^*(m) \leq p$ and $x_i^*(m - S, p) = 0$. Since $v_i'(m', \zeta') > p'$, we also have $x_i(m', p', \zeta') = 1$ and can therefore rewrite (6) as

$$m - S = \phi_i + m' - p'$$

which has unique solution $EV_i = p' - \Delta m - \phi_i = p' - \Delta m - v_i^*(m)$

Case $v_i^*(m) > p$ and $v_i(m', \zeta') \leq p'$:

Since utility is quasilinear, $v_i^*(m - S) = v_i^*(m) > p$ and $x_i^*(m - S, p) = 1$. Since $v_i'(m', \zeta') \leq p'$, we also have $x_i(m', p', \zeta') = 0$ and can therefore rewrite (6) as

$$\phi_i + m - S - p = m'$$

which has unique solution $EV_i = \phi_i - p - \Delta m = v_i^*(m) - p - \Delta m$

Case $v_i^*(m) > p$ and $v_i(m', \zeta') > p'$:

Since utility is quasilinear, $v_i^*(m - S) = v_i^*(m) > p$ and $x_i^*(m - S, p) = 1$. Since $v_i'(m', \zeta') > p'$, we also have $x_i(m', p', \zeta') = 1$ and can therefore rewrite (6) as

$$\phi_i + m - S - p = \phi_i + m' - p'$$

which has unique solution $EV_i = \Delta p - \Delta m$ ■

Claim 8 Under assumption 3

$$F_v(p) - F_{v^*}(p) = \int_p^q \kappa_0(z) dF_{v^*}(z)$$

Proof. The idea is to prove that under assumption 3, both sides of the expression represent the total amount of discouraged transactions. Indeed

$$\begin{aligned} \Pr\{x_i(m, p, 0) \neq x_i(m, p, t)\} &= \Pr\{v_i - t \leq p < v_i^*\} \\ &= \Pr\{v_i^* > p\} - \Pr\{v_i > p\} \\ &= F_v(p) - F_{v^*}(p) \end{aligned} \tag{26}$$

where the assumption is used in the first equality, that requires $v_i - t \leq v_i^*$ to hold. Similarly

$$\begin{aligned} \int_p^q \kappa_0(z) dF_{v^*}(z) &= E\left[\mathbb{1}(p \leq v_i^* \leq q) \times E\left[\mathbb{1}(v_i \leq p + t) \middle| v^*\right]\right] \\ &= E\left[\mathbb{1}(p \leq v_i^* \leq q) \times \mathbb{1}(v_i \leq p + t)\right] \\ &= \Pr\{x_i(p, 0) \neq x_i(p, \tau)\} \end{aligned}$$

■

References

- Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–38.
- Auerbach, A. J. (1985). Chapter 2 the theory of excess burden and optimal taxation. volume 1 of *Handbook of Public Economics*, pages 61–127. Elsevier.
- Bernheim, B. D. and Rangel, A. (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics*. *The Quarterly Journal of Economics*, 124(1):51–104.
- Bernheim, B. D. and Taubinsky, D. (2018). Behavioral public economics. In Bernheim, B. D., DellaVigna, S., and Laibson, D., editors, *Handbook of Behavioral Economics - Foundations and*

- Applications 1*, volume 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, pages 381–516. North-Holland.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics*, 92(8):1787–1794. Special Issue: Happiness and Public Economics.
- Bhattacharya, D. (2015). Nonparametric welfare analysis for discrete choice. *Econometrica*, 83(2):617–649.
- Bhattacharya, J., Shaikh, A., and Vytlacil, E. (2012). Treatment effect bounds: An application to swan–ganz catheterization. *Journal of Econometrics*, 168(2):223–243.
- Blundell, R., Gosling, A., Ichimura, H., and Meghir, C. (2007). Changes in the distribution of male and female wages accounting for employment composition using bounds. *Econometrica*, 75(2):323–363.
- Chernozhukov, V. and Hansen, C. (2005). An iv model of quantile treatment effects. *Econometrica*, 73(1):245–261.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77.
- Clark, J. and Friesen, L. (2008). The causes of order effects in contingent valuation surveys: An experimental investigation. *Journal of Environmental Economics and Management*, 56(2):195–206.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–72.
- Fan, Y., Guerre, E., and Zhu, D. (2017). Partial identification of functionals of the joint distribution of “potential outcomes”. *Journal of Econometrics*, 197(1):42–59.
- Fan, Y. and Soo Park, S. (2009). Partial identification of the distribution of treatment effects and its confidence sets. In *Nonparametric Econometric Methods*, pages 3–70. Emerald Group Publishing Limited.
- Frandsen, B. R. and Lefgren, L. J. (2021). Partial identification of the distribution of treatment effects with an application to the knowledge is power program (kipp). *Quantitative Economics*, 12(1):143–171.
- Handel, B. R., Kolstad, J. T., and Spinnewijn, J. (2019). Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets. *The Review of Economics and Statistics*, 101(2):326–340.
- Heckman, J. J., Smith, J., and Clements, N. (1997). Making the most out of programme evaluations

- and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies*, 64(4):487–535.
- Kahneman, D., Wakker, P. P., and Sarin, R. (1997). Back to Bentham? Explorations of Experienced Utility. *The Quarterly Journal of Economics*, 112(2):375–406.
- Kim, J. H. (2014). Identifying the distribution of treatment effects under support restrictions.
- Manski, C. F. (1997). Monotone treatment response. *Econometrica*, 65(6):1311–1334.
- Mullainathan, S., Schwartzstein, J., and Congdon, W. J. (2012). A reduced-form approach to behavioral public finance. *Annual Review of Economics*, 4(1):511–540.
- Spinnewijn, J. (2017). Heterogeneity, demand for insurance, and adverse selection. *American Economic Journal: Economic Policy*, 9(1):308–43.
- Taubinsky, D. and Rees-Jones, A. (2018). Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment. *The Review of Economic Studies*, 85(4):2462–2496.
- Train, K. (2015). Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples. *Journal of Choice Modelling*, 16:15–22.