Author Manuscript

Brusco Giacomo ORCID iD: 0000-0003-4398-7266

Attending to Inattention: Identification of Deadweight Loss under Non-Salient Taxes

Giacomo Brusco* and Benjamin Glass† ‡

August 28, 2019

Abstract

Recent developments in behavioral public economics have shown that heterogeneous biases prevent point identification of deadweight loss. We replicate this result for an arbitrary (closed) consumption set, whereas previous results on heterogeneous attention focused on binary choice. We find that one can bound the efficiency costs of taxation based on aggregate features of demand. When individuals have linear demand functions, the bounds for deadweight loss are easy to calculate from linear regressions.

*Economics, University of Michigan at Ann Arbor

†Economics, Pennsylvania State University at University Park

‡The work in the following pages would have been impossible without the thought-provoking conversations we had with James Hines, Joel Slemrod, Dmitry Taubinsky, Alex Rees-Jones, Jeremy Fox, Tilman Börgers, Ying Fan, Ugo Troiano, William Boning, Luis Alejos, Tejaswi Velayudhan, Aristos Hudson, Benjamin Lockwood, Johannes Spinnewijn, Sarah Moshary, Alessandro Lizzeri, Amine Ouazad, the participants of the 2017 NTA conference, the participants of the 2018 Quebec Political Economy Conference, and the participants of the Public Finance, Theory, and IO seminars at the economics department of the University of Michigan. We warmly thank all of them for their invaluable advice. Last but not least, a special thanks goes to anonymous reviewers at the Journal of Public Economic Theory for their insightful comments that reshaped and refocused the scope of this paper. Any mistakes are, of course, our own.

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/jpet.12401

This article is protected by copyright. All rights reserved.

Author Manuscript

1 Introduction

Taxing a good results in a loss of economic efficiency whenever it distorts equilibrium behavior away from the Pareto optimum. To the extent that agents do not notice a tax, the burden of the tax is exacerbated by the fact that agents cannot adjust the behavior to protect themselves from the tax. However, the burden of the tax in excess of government revenue, or deadweight loss, is mitigated when agents do not pay attention to the tax: if consumers pay the tax without noticing it, they are effectively transferring some of their income to the government in a lump sum. Chetty, Looney, and Kroft (2009, henceforth CLK) were the first to make these points. In addition to their theoretical contributions, they also showed that consumers in the U.S., where sales tax is applied at the register rather than included on the prices displayed on shelves (or, "sticker prices"), tend to under-react to sales taxes.

While CLK (2009) focuses on the case of homogeneous attention, recent work by Taubinsky and Rees-Jones (2018, henceforth TRJ) has noted that introducing the possibility of heterogeneous attention may prevent the computation of deadweight loss from aggregate data. If each person faced a different tax rate when buying a certain good, understanding welfare effects would require us to study not only aggregate demand, but the demand of every individual. Imposing a high tax on low elasticity individuals and a low tax on high elasticity individuals would have a very different effect on welfare than doing the opposite. A similar reasoning applies when all agents face the same tax rate, but perceive different tax rates. TRJ (2018) find that allowing for heterogeneous attention introduces an issue of allocative inefficiency that is normally absent from the study of the welfare effects of taxation. In a world of heterogeneous attention, there is no guarantee that the individuals who end up consuming the good are the ones who value it the most.

TRJ (2018) make these points in a binary choice model. This is well-suited to their experiment, in which people are choosing whether or not to buy a certain object, but their proof does not generalize trivially. Given the predominance of continuous choice settings in much of the literature on tax salience, including CLK's seminal paper, this motivates us to study the issue further.

We begin by developing a model of choice under misperceived prices with an arbitrary closed consumption set, and develop our welfare measure: compensating variation due to the tax, net of tax revenue. Bernheim and Rangel (2009) laid the foundations of welfare analysis with behavioral agents. Our model is similar to the models of CLK (2007, 2009) and TRJ (2018), but we slightly modify the treatment of income effects, along the lines of Gabaix's (2014) model of rational inattention. In the absence of income effects, our model of choice for an individual agent is essentially equivalent to the model in CLK (2009), except that we allow for arbitrary consumption sets. Our model is also similar to the model of Chetty (2009), but for the fact that we specify a particular way in which behavioral agents maximize utility. While this does not impose severe restrictions on behavior, it offers a useful framework when we move on to identification. We confirm that some of the major results in CLK (2009) and TRJ (2018) hold quite broadly: inattention to taxes increases the size of the loss in consumer surplus, but decreases the size of deadweight loss; attention heterogeneity amplifies deadweight loss, and invalidates CLK's sufficient statistic approach.

The main contribution of this paper is to generalize TRJ's non-identification result to an arbitrary closed choice set. We show that an econometrician who only observes aggregate consumption data can only determine the true value of aggregate deadweight loss to lie on an interval. These bounds were

This article is protected by copyright. All rights reserved.

first noted by TRJ (2018) in their proposition A.2. We find these bounds hold generally and propose to use them as a novel empirical tool.

The lower bound for deadweight loss is the calculation one would perform in the case of a representative consumer. Since the loss in efficiency is a convex function of the perceived tax rate, the calculation of deadweight loss from *one* perceived tax-inclusive price consistent with aggregate demand will generically underestimate deadweight loss. Heterogeneity in tax salience creates heterogeneity in perceived net-of-tax prices, which creates an allocative inefficiency across consumers. As the calculation with a representative consumer only accounts for inefficiency from *aggregate* foregone consumption due to the tax, it will under-estimate excess burden. However, in the case in which all agents pay the same amount of attention to the tax, there is no allocative inefficiency between consumers, and so performing the calculation as with a representative consumer yields the correct value for deadweight loss. The formula for this lower bound to deadweight loss is an extension of formulas provided by CLK (2009) and TRJ (2018).

Following TRJ (2018), we obtain an upper bound for deadweight loss by letting the econometrician assume that tax salience has support on a known bounded non-negative interval. The upper bound comes from maximizing perceived price heterogeneity, again exploiting the convexity of deadweight loss with respect to the perceived tax. This is achieved by positing that agents have either zero or maximal salience. Generalizing introduces two additional considerations in calculating the upper bound for deadweight loss. One, a distribution yielding the upper bound for deadweight loss assigns high tax salience precisely where it will "hurt" most: to those agents whose particular preferences yield maximal deadweight loss from that agent relative to the change in consumption of that agent. This distribution allocates high tax salience to those agents who have more convex demand curves, keeping the aggregate change in quantity demanded constant. Two, deadweight loss is maximized for a given aggregate demand if any agent with multiple optimal decisions at the perceived price consumes the highest amount consistent with their preference when they perceive low prices, whereas they consume the lowest amount consistent with their preferences when they perceive high prices.[1] This is because heterogeneity in perceived prices permits different equilibria with the same sticker price, tax rate, and aggregate consumption, yet yielding different values of deadweight loss due to different distributions of consumption among individuals.

Our approach to compute the upper bound forces the econometrician to impose limits on the possible values of attention – something that the empirical literature has not settled yet and might be highly context-dependent. While it might seem natural to assume that attention varies between zero and one, many papers have found evidence of salience above one – see for instance Allcott and Taubinsky (2015). In theory, one might allow for unlimited (positive) tax salience, under regularity conditions that avoid the possibility of unlimited distortion to consumer behavior.[2]

Our general results follow the work of TRJ (2018), and rely on the fact that the distribution of preferences is independent of taxes and prices, but that the distribution of salience might change depending on the value of the tax. Indeed, we show that, even assuming that the distribution of preferences is entirely known to the econometrician and invariant to observables, one cannot identify deadweight loss.

[1]TRJ (2018) do not deal with cases like this because they restrict attention to non-atomic distributions of willingness to pay, so almost every agent has a unique choice that is perceived to maximize utility.

[2]One might assume that consumer surplus is uniformly bounded to ensure that the upper bound for deadweight loss is finite even if the upper bound for tax salience is infinite.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

We then turn to the special case in which demand is linear in a relevant range where consumption is positive. This case is of particular interest for three reasons: first, its ease of applicability; second, its special relationship to the second order approximation of deadweight loss; and third, its illustrative value in the kind of problems we can face in identifying deadweight loss. In the case of binary choice, non-identification ultimately comes from the possibility that taxes and attention to taxes are not independent of each other. Although the experimental evidence in TRJ (2018) suggests that attention varies with how large the tax is, if one assumed away this possibility we could identify a full distribution of responsiveness to both taxes and sticker prices using existing models of discrete choice with random coefficients, as in Masten (2017) and Fox (2017). This estimated distribution could then be used to yield a point-estimate of deadweight loss. In the case of linear demand, instead, even assuming that attention and taxes are independent would not help identify deadweight loss with aggregate data beyond the bounds described above.

Our results complement a growing literature on tax salience. Rosen (1976) does not find evidence of limited tax salience, but besides CLK (2009) and TRJ (2018), Finkelstein (2009), Gallagher and Muehlegger (2011), and Goldin and Homonoff (2013) all find strong evidence of dramatically limited tax salience. Most of this literature looks at sales taxes in the U.S., as they lend themselves very credibly to a story about lack of salience. However, work on salience has also looked at other settings: Finkelstein (2009) studies car tolls; Weber and Schram (2016) study whether income taxes being remitted by the employer or the employee affects differently people's attitude towards public spending and the burden of the tax;[3] Morone, Nemore and Nuzzo (2018) explore a similar question in the context of a double-auction market. Blake, Moshary and Tadelis (2017) study how people react differently to back-end and upfront fees in online purchases; and Bradley and Feldman (2018) study how changes in the disclosure of ticket taxes affect the demand for airlines. As we mentioned above, most of these empirical papers, as well as other theoretically focused papers like Goldin (2015), use models where the choice set is continuous or mixed discrete-continuous.

The paper proceeds as follows. In section 2, we develop a general model of choice under misperceived prices. Once we have replicated some of the major theoretical results in CLK (2009) and TRJ (2018), we shift focus to identification of deadweight loss. Section 3 lays out the main results of our paper, establishing the non-identification result and the bounds. Section 4 focuses on the special case in which demand is linear, and provides a straightforward way to compute our bounds in the context of linear models. Section 5 concludes. Proofs and other minor results are relegated to the online appendix.

2 Choice and Deadweight Loss under Non-Salient Taxes

This section describes the theoretical model and results that underlie the rest of this paper. Many of our results here simply mirror previous literature, but we make slightly different modeling choices. The main modeling challenge in dealing with misperceived prices is to allow for the misperception of prices while keeping agents financially solvent. CLK (2007, 2009) assume that one good "absorbs" all optimization mistakes. In contrast our approach, inspired by parts of the model in Gabaix (2014), has agents conjecture a certain income such that they end up consuming on their true budget constraint

[3] Interestingly, assuming some agents face credit constraints as in Boadway, Garon, and Perrault (2018), also breaks traditional optimal tax theory.

This article is protected by copyright. All rights reserved.

Author Manuscript

when presented with the relative prices they perceive. While this framework preserves all results of interest from CLK (2009) and TRJ (2018), we find that our approach eases exposition while freeing the researcher from having to make ad-hoc assumption about which good (or goods) absorb optimization mistakes. It should be noted that while we do generalize the model to include multiple taxed and non-taxed goods in the online appendix, in the body of the paper agents will face a choice over two goods, only one of which is subject to tax.

The agent has a closed consumption set $X = X^T \times \mathbb{R}_+ \subseteq \mathbb{R}^2_+$. She also has a choice function for the taxed good, $q(\bar{p}, p^{NT}, \tau, W)$, with $(\bar{p}, p^{NT}) \in \mathbb{R}^2_{++}$, where $\bar{p}$ and p^{NT} are respectively the sticker price of the taxed and non-taxed good, $\tau \in \mathbb{R}$ is the sales tax on the taxed good, and W is the income of the agent.[4] We express taxes as if they were specific, so that $\bar{p} + \tau$ is the tax-inclusive price of the taxed good.

The agent has a continuous and strictly monotonic utility function $u(q, q^{NT})$, where q denotes generic consumption of the taxed good. The choice vector function $\boldsymbol{q}(\bar{p}, p^{NT}, \tau, W) = (q(\bar{p}, p^{NT}, \tau, W), q^{NT}(\bar{p}, p^{NT}, \tau, W)) \in X$ meets two requirements. One, the agent spends all available income:[5]

$$(\bar{p} + \tau) q(\bar{p}, p^{NT}, \tau, W) + p^{NT} q^{NT}(\bar{p}, p^{NT}, \tau, W) = W. \tag{1}$$

Two, the agent correctly optimizes in the choice of all consumption bundles when there is no tax:

$$\boldsymbol{q}(\bar{p}, p^{NT}, 0, W) \in \underset{\{\bar{p}*q + p^{NT}*q^{NT} \leq W\}}{\arg\max} u(q, q^{NT}). \tag{2}$$

It turns out that this model is quite general, in the sense that it rules out very few possible behaviors. Indeed, this model is equivalent to one in which agents pick rationally given a perceived price p^s, and conjecture themselves an income W^s so that they satisfy their true budget constraint at their perceived price. Proposition 2 and subsequent work in the online appendix shows that, under weak convexity assumptions on preferences, one can find a pair (p^s, W^s) to satisfy equations 1 and 2 for any choice function $\boldsymbol{q}(\cdot)$.

We now want some measure of the incidence of the tax on the consumer. For concreteness, we consider the compensating variation due to the tax with complete pass-through, defined as:

$$\Delta CS \equiv \inf\{\Delta W |\ u(\boldsymbol{q}(\bar{p}, p^{NT}, \tau, W + \Delta W)) \geq u(\boldsymbol{q}(\bar{p}, p^{NT}, 0, W))\}.$$

[4]We implicitly restrict consideration to sticker prices, taxes, and income such that $\mathbf{q}(\bar{\boldsymbol{p}}, \tau, W)$ is well-defined at those values.

[5]When we consider multiple non-taxed goods in the online appendix, we also have agents optimally choose q^{NT} given their choice of q.

This article is protected by copyright. All rights reserved.

Author Manuscript

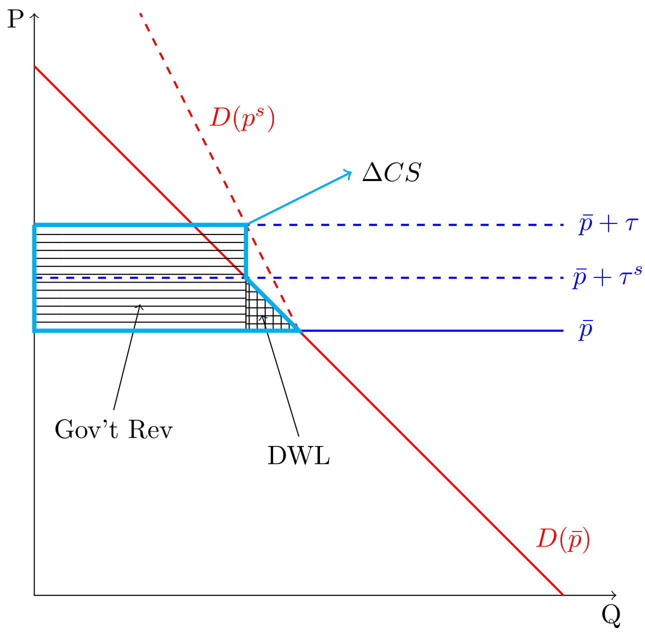

Figure 1: Welfare effects from the imposition of a non-salient tax.

In words, the change in consumer surplus is the greatest lower bound of the amount of money the agent requires to achieve the utility reached before the imposition of the tax. Online appendix proposition 3 shows that the change in consumer surplus can be written as the sum of what consumer would have to be compensated if the tax-inclusive price were *actually* p^s and the income they "lost" due to their inattention:

$$\Delta CS = \underbrace{(\bar{p}+\tau-p^s)h(p^s)}_{\text{Income lost}} + \underbrace{e(p^s)-e(\bar{p})}_{\Delta CS \text{ under } p^s}. \tag{3}$$

This representation, which is graphically illustrated in figure 1, readily gives us two interesting results. First, as noted by CLK (2007, 2009), if $p^s \in [\bar{p}, \bar{p}+\tau]$, then the consumer will be worse off than if she paid attention to the tax:

$$\begin{aligned}\Delta CS &= (\bar{p}+\tau-p^s)h(p^s) + \int_{\bar{p}}^{p^s} h(p)dp \\ &\geq \int_{p^s}^{\bar{p}+\tau} h(p)dp + \int_{\bar{p}}^{p^s} h(p)dp \\ &= \int_{\bar{p}}^{\bar{p}+\tau} h(p)dp.\end{aligned}$$

Second, a consumer can be made better off by an increase in the tax, which we do not believe previous literature to have noted. This is because a tax increase can induce inattentive agents to reduce their consumption of the taxed good to zero, improving welfare for consumers who would have already avoided consuming the taxed good had they been attentive to the tax.[6] We provide a graphical example

[6]More formally, consumption does not necessarily have to be reduced to zero for the consumer to be made better off. The loss of consumer surplus decreases in τ whenever the tax is sufficiently high such that consumption is sufficiently

This article is protected by copyright. All rights reserved.

Author Manuscript

in figure 2.

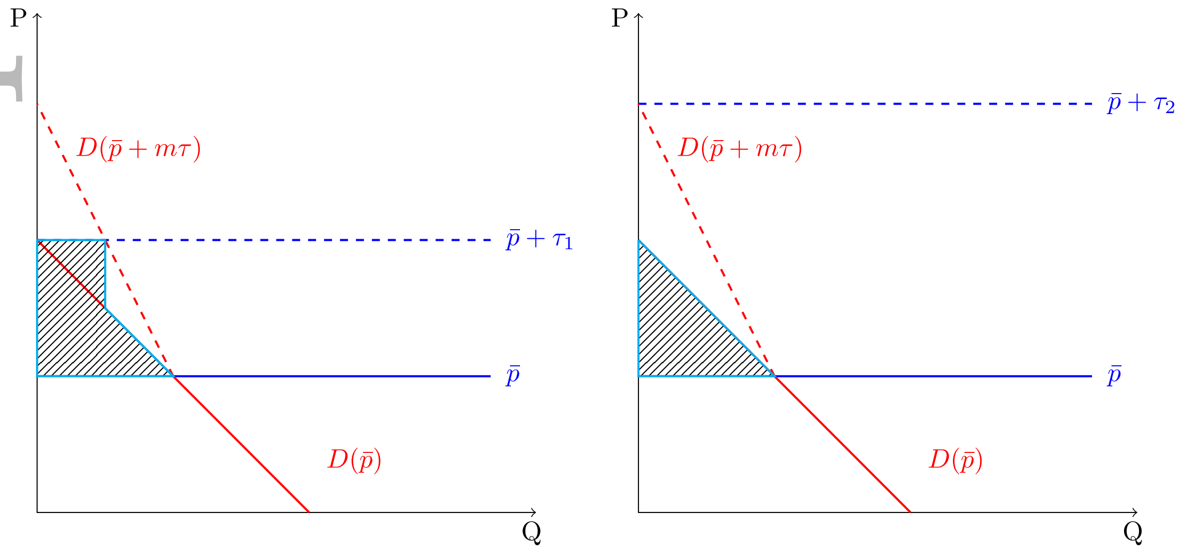

(a) Under τ_1 the consumer loses more than total CS

(b) Under τ_2 the consumer loses exactly total CS

Figure 2: The consumer in (a) is worse off than in (b), although she is subject to a lower tax

To obtain deadweight loss, we need to adjust ΔCS for the change in tax revenue:[7]

$$dwl \equiv \Delta CS - \tau q(\bar{p}, p^{NT}, \tau, W + \Delta CS) = e(p^s) - e(\bar{p}) - (p^s - \bar{p})q(\bar{p}, p^{NT}, W + \Delta CS). \qquad (4)$$

Note that deadweight loss is exactly as if the agent was correct that the tax-inclusive price were p^s.

To introduce heterogeneity, let $i \in \mathcal{I}$ index consumers. Each consumer is characterized by her perception of the price of the taxed good, p_i^s, type, θ_i, standing in for her preferences $\succeq_{\theta_i}$ and income W_{θ_i}, and tie-breaking parameter ζ_i, which we need for technical reasons.[8] These consumer-specific parameters are distributed according to $F^*_{p^s,\theta,\zeta}$. Each agent has a choice function for the taxed good, satisfying:

$$q(\bar{p}, p^{NT}, \tau, W_{\theta_i}; \theta_i, \zeta_i) = q(p_i^s, W_i^s; \theta_i, \zeta_i) \in \{q | \exists q^{NT} : (q, q^{NT}) \succeq_{\theta_i} \boldsymbol{q}' \ \forall \boldsymbol{q}' \in X : (p_i^s, p^{NT}) * \boldsymbol{q}' \leq W_i^s\},$$

where p_i^s and W_i^s are determined as in the model of section 2, with corresponding expenditure function $e(p_i^s; \theta_i)$. If demand is single-valued, letting us ignore the tie-breaking parameter ζ, total deadweight loss is:

$$DWL = \int_{p_i^s, \theta_i} [e(p_i^s; \theta_i) - e(\bar{p}; \theta_i)] - (p_i^s - \bar{p})q(p_i^s; \theta_i) dF^*_{p^s,\theta}(p_i^s, \theta_i).$$

small (but positive).

[7]We maintain the convention that deadweight loss is a generically positive value.

[8]ζ_i acts as a tie-breaker among bundles that could all have been chosen: choices do not necessarily reflect true preferences when agents misperceive prices, and agents might appear indifferent between choices that do not actually yield the same ex-post utility. This is in sharp contrast with the neo-classical model, where the actual choice that one selects among indifferent bundles has no impact on consumer surplus. But in this model of choice, different values of conjectured income might yield different choices, even given the same preferences, perceived prices, and true income.

This article is protected by copyright. All rights reserved.

Author Manuscript

This paper emphasizes continuous choice, as the binary choice case is worked out in TRJ (2018). We momentarily assume that h is continuously differentiable with respect to its own price, and that p^s is continuously differentiable with respect to τ. Then tax salience at a tax rate of zero is $m = \frac{\partial p^s}{\partial \tau}|_{(\bar{p},0)}$.[9] In line with existing deadweight loss analyses, we can consider a second-order approximation, letting us characterize our object of interest in terms of first derivatives:

$$DWL \approx -\frac{\tau^2}{2} \int_{p_i^s,\theta_i} m_i^2 \frac{\partial h(\bar{p};\theta_i)}{\partial p} dF^*_{p^s,\theta}(p_i^s,\theta_i). \tag{5}$$

This process of aggregation makes apparent two important points that confirm the analysis of TRJ (2018) extends naturally to continuous choice. First, allowing for attention heterogeneity introduces an issue of allocative inefficiency, as it is no longer guaranteed that consumers who value some units of the taxed good the most will be the ones who end up purchasing those units. It should be noted that we are assuming throughout this paper that supply is perfectly elastic, and so tax increases will be reflected one-for-one in the after-tax price. This is not really an issue of concern: as TRJ (2018) note, all that is needed to generalize this to an arbitrary supply function is to account for the change in sticker price and the change in profits to suppliers. Nonetheless, as in TRJ's work, it is interesting to deviate for a moment from this assumption, to consider what happens to aggregate deadweight loss when supply is perfectly inelastic. In that case, we can use appendix proposition 4 to show that

$$DWL \approx -\frac{\tau^2}{2} \int_{m_i,\theta_i} \left[m_i^2 \frac{\partial q_i}{\partial p} dF_{m,\theta}(m_i,\theta_i) - \frac{(\int_{m_i,\theta_i} m_i \frac{\partial q_i}{\partial p})^2}{\int_{m_i,\theta_i} \frac{\partial q_i}{\partial p} dF_{m,\theta}(m_i,\theta_i)} \right] \geq 0,$$

and $DWL = 0$ when attention is homogeneous, i.e. $m_i = m \ \forall i$. Thus, a non-salient tax may yield excess burden even without changing the equilibrium quantity, due to its effects on allocative efficiency. Second, allowing for heterogeneous attention introduces a serious problem of identification, as neither aggregate price responsiveness $\int_{p_i^s,\theta_i} \frac{\partial h(\bar{p};\theta_i)}{\partial p} dF^*_{p^s,\theta}(p_i^s,\theta_i)$ nor aggregate tax responsiveness $\int_{p_i^s,\theta_i} m_i \frac{\partial h(\bar{p};\theta_i)}{\partial p} dF^*_{p^s,\theta}(p_i^s,\theta_i)$ are sufficient statistics for the deadweight loss in equation 5. These points effectively extend some of TRJ's (2018) major results to continuous choice. In the next section, we formalize this non-identification result with an arbitrary choice function, and show how one might bound deadweight loss with mere information on aggregate parameters.

3 Non-Identification with Aggregate Data

This section discusses to what degree one can infer deadweight loss from aggregate choice data. Our results generalize TRJ's work on binary choice to an arbitrary choice set. We find this to be of interest for two reasons. First, many papers dealing with tax salience, including the seminal paper by CLK (2009), operate in a continuous choice setting. Second, while point identification is impossible using aggregate

[9] Formally, the claim is that there exist (p^s, W^s) such that $\frac{\partial p^s}{\partial \tau}$, where the derivative is taken while the consumer is being compensated, is well defined. If $\frac{\partial h}{\partial p}(\bar{p}) \neq 0$, then the Inverse Function Theorem implies that $\frac{\partial p^s}{\partial \tau} = \frac{\frac{\partial q}{\partial \tau} + \frac{\partial q}{\partial W}\frac{\partial \Delta CS}{\partial \tau}}{\frac{\partial h}{\partial p}}$. If $\frac{\partial h}{\partial p} = 0$ in a neighborhood around $\bar{p}$, then set $\frac{\partial p^s}{\partial \tau}|_{\tau=0} = 0$ and $\frac{\partial \Delta CS}{\partial \tau} = -\frac{\frac{\partial q}{\partial \tau}}{\frac{\partial q}{\partial W}}$.

This article is protected by copyright. All rights reserved.

parameters, we can still provide tight bounds based solely on aggregate (or average) quantities.

For simplicity, we assume that the econometrician already knows the distribution of preference types, but this should not be considered a limiting assumption. Consumer preferences can be identified with sticker price variation when there are no (non-salient) taxes. Regardless, even when the econometrician can fully observe the true distribution of preferences, and how much aggregate consumption there is at every tax level, she still cannot infer the exact value of deadweight loss. However, we provide a lower bound and an upper bound for deadweight loss. The lower bound is achieved by assuming that all agents perceive the same tax-inclusive price, i.e. assume there is no attention heterogeneity. The upper bound for deadweight loss is achieved by imposing maximal attention heterogeneity. Since the data do not reveal the individual variation in tax salience, one cannot point identify deadweight loss from aggregate data. Deadweight loss can take on any value between the upper and lower bounds.[10]

The results in this section are described as if all agents face the same sales tax. Also, since we are considering the problem of identification with aggregate demand, we assume there are no income effects. This is because even the standard model with fully salient taxes requires strong restrictions on income effects in order to achieve identification with aggregate data, as the same income can yield different consumption bundles whenever there are several conjectured incomes that solve the agent's problem. Suppressing income and price for the non-taxed good, we denote the consumption function for agent i with type θ_i and perceived tax-inclusive price p_i^s for the taxed good by $q(p_i^s;\theta_i,\zeta_i)$. However, all of these results follow if one reinterprets $q(p_i^s;\theta_i,\zeta_i)$ as the *compensated* choice of agent i.

To ensure integrability, we assume the econometrician knows that $F^*_{p^s}$ has support with lower bound greater than zero, and so only considers marginal distributions of subjective prices bounded above zero. The econometrician observes aggregate demand:

$$\int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF^*_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).$$

Deadweight loss for an individual i is a function of their expenditure function $e(p)$ and prices via:

$$dwl(p_i^s;\theta_i,\zeta_i) = e(p_i^s;\theta_i) - e(\bar{p};\theta_i) - (p_i^s - \bar{p})q(p_i^s;\theta_i,\zeta_i).$$

We are interested in aggregate deadweight loss:

$$DWL = \int_{p_i^s,\theta_i,\zeta_i} dwl(p_i^s;\theta_i,\zeta_i)dF^*_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).$$

The problem of identification is to find conditions for which any joint distribution $F_{p^s,\theta,\zeta}$ of (p^s,θ,ζ), as a function of observable variables $\bar{p}$ & τ, satisfying these conditions and such that aggregate demand is rationalized; that is, such that for any observed values of observable variables, any F satisfying

$$\int_{p_i^s,\theta_i\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta_i}(p_i^s,\theta_i,\zeta_i) = \int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF^*_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i), \tag{6}$$

[10] This claim follows by taking any weighted average of the distributions of parameters yielding upper and lower bounds of deadweight loss.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

also yields the same value for deadweight loss:

$$\int_{p_i^s,\theta_i,\zeta_i} dwl(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i) = DWL.$$

The main message of this section will be the failure of such a result obtaining. We show that there are at least two distributions satisfying 6 that yield different values of DWL. These two distributions also turn out to yield tight bounds to the possible values of DWL that are consistent with aggregate demand.

Finally, we impose regularity conditions throughout this section to rule out ill-defined integrals. Formally, we insist that the econometrician only consider distributions that satisfy the *integrability conditions*, described below.

Definition 1. *A distribution $F_{p^s,\theta,\zeta}$ satisfies the integrability conditions if:*

1. *q and dwl are integrable on any measurable set.*
2. *$q(p;\theta_i,z)$ is integrable on any subset of the support of θ for any $p>0$ and any z in the range of ζ.*

For instance, all distributions with a finite support of (p^s,θ,ζ) satisfy the above conditions.

3.1 Lower Bound on Deadweight Loss

Consider arbitrary $\bar{p}$, p^{NT}, and τ. For arbitrary $F_{p^s,\theta,\zeta}$ consistent with the data, we can choose a price $\hat{p}^s$ that could also rationalize the data if perceived homogeneously.

Proposition 1. *For any $F_{p^s,\theta,\zeta}$ that yields integrable aggregate demand, there exists $\hat{p}^s$ such that for some distribution $F'_{\theta,\zeta}$ such that $F'_\theta = F_\theta$:*

$$\int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i) = \int_{\theta_i,\zeta_i} q(\hat{p}^s;\theta_i,\zeta_i)dF'_{\theta,\zeta}(\theta_i,\zeta_i).$$

We can always rationalize the data with a joint distribution of (p^s,θ) in which θ has marginal distribution F^*_θ, whereas $p^s=\hat{p}^s$ with probability one. We now show that such a joint distribution provides a generic underestimate to the possible values of deadweight loss.

Theorem 1. *Consider any joint distributions $F_{p^s,\theta,\zeta}$ and $F_{\theta,\zeta}$ with corresponding value $\hat{p}^s$ such that:*

$$\int_{\theta,\zeta} q(\hat{p}^s;\theta_i,\zeta_i)dF_{\theta,\zeta}(\theta_i,\zeta_i) = \int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).$$

Then the following inequality obtains:

$$\int_{\theta_i,\zeta_i} dwl(\hat{p}^s;\theta_i,\zeta_i)dF_{\theta,\zeta}(\theta_i,\zeta_i) \le \int_{p_i^s,\theta_i,\zeta_i} dwl(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).$$

Intuitively, introducing heterogeneity in perceived prices can mute gains from trade. If someone with a higher marginal valuation for the good has a higher perceived price than someone with a lower marginal

This article is protected by copyright. All rights reserved.

Author Manuscript

valuation, they could both gain by trading with each other after making their consumption decisions. If they could exchange with each other, the one who perceived the higher price could purchase some of the good from the other agent, making both agents better off. Thus, ruling out perceived price heterogeneity by assuming a homogeneous perceived price $\hat{p}^s$ eliminates the possibility of an allocative inefficiency. Figure 3 offers graphical intuition for the lower bound.

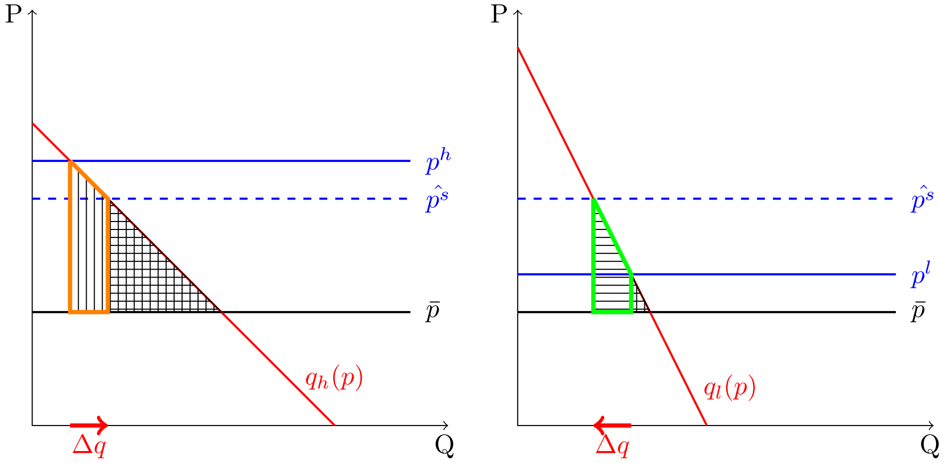

(a) High perceived price

(b) Low perceived price

Figure 3: A graphical illustration of Theorem 1. When one picks $\hat{p}^s$ as to make the change in demand equal for the consumer in (a) and in (b), the decrease in *dwl* for the consumer in (a) (orange) must be at least as large as the increase in *dwl* for the consumer in (b) (green).

Theorem 1 points out that for any distribution that rationalizes the data, i.e. that explains the observed aggregate demand, one can alternatively rationalize the data with a homogeneous perceived price that yields (weakly) less deadweight loss. From this, we can reach two conclusions. One, we generally cannot identify deadweight loss because we could always alternatively rationalize the data with a homogeneous perceived price.[11] This holds even if we already knew the distribution of preference types F^*_θ. Two, if there is a minimum value of deadweight loss that is consistent with the data, that value of deadweight loss comes from a distribution with no heterogeneity in tax salience.

3.2 Upper bound on deadweight loss

The upper bound comes from an assumption on the limits to tax salience:

Assumption 1. *There is some value $\bar{m} \geq 0$ such that p^s has support known to be contained in $\mathcal{P} \equiv [\bar{p}, \bar{p} + \bar{m}\tau]$.*

This assumption says that agents must perceive a non-negative tax τ^s no greater than fraction $\bar{m}$ of the true tax.[12] The econometrician is allowed to assume an arbitrarily large $\bar{m}$, but the gain in robustness will likely come at the expense of precision. For instance, setting $\bar{m} = 1$ would be to assume that

[11] This claim holds generically, but would not hold, for instance, if there was no heterogeneity in tax salience.
[12] This description implicitly assumes that $\tau > 0$.

This article is protected by copyright. All rights reserved.

agents never over-react to a tax rate. Imposing that $\tau^s \geq 0$ with probability one already ensures that deadweight loss is no greater than the original consumer surplus.[13] But the interval restriction implies any distribution yields no more deadweight loss than a distribution with "binary" perceived prices, i.e. where p^s can only take on values in $\{\bar{p}, \bar{p}+\bar{m}\tau\} \equiv \partial\mathcal{P}$.

The gist of the upper bound of deadweight loss is that, for any data-generating process that rationalizes observed aggregate demand, there is another data-generating process that also rationalizes observed demand, but which would yield at least as much deadweight loss. This alternative explanation of the observed demand insists that all agents pay either zero or maximal attention.

Before formally stating our main result, we demonstrate how one can always pick such a distribution of attention to rationalize observed demand, for any underlying (known) distribution of preferences. Then, we state the main result, theorem 2, providing some intuition for why such a distribution would yield a weakly higher deadweight loss. Because our model of choice may result in several choices given the same sticker prices and taxes, sometimes there can be multiple equilibria with the same level of aggregate demand, but different consequences for welfare. We briefly discuss appendix theorem 3, which deals deals with such cases.

Consider any $F_{p^s,\theta,\zeta}$ that rationalizes the data, and such that:

$$\lim_{m\to\bar{m}^-} F_{p^s}(\bar{p}+m\tau) - F_{p^s}(\bar{p}) > 0.$$

In words, the distribution assumes some positive mass of agents pay neither zero nor maximal attention, i.e. $m \in (\bar{p}, \bar{p}+\bar{m}\tau) \equiv int(\mathcal{P})$. Pick $\tilde{p}^s \in int(\mathcal{P})$ and a corresponding $p^b(p_i^s) \equiv \bar{p} + \mathbb{I}(p_i^s > \tilde{p}^s)\bar{m}\tau$ such that:

$$\begin{aligned}\int_{p_i^s\in int(\mathcal{P}),\theta_i} q(p^b(p_i^s);\theta_i,l)dF_{p^s,\theta}(p_i^s,\theta_i) &\leq \int_{p_i^s\in int(\mathcal{P}),\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i)\\ &\leq \int_{p_i^s\in int(\mathcal{P}),\theta_i} q(p^b(p_i^s);\theta_i,h)dF_{p^s,\theta}(p_i^s,\theta_i).\end{aligned}$$

In words, for any distribution that puts mass on $int(\mathcal{P})$, we pick a value $\tilde{p}^s$ that acts as a divide: those below it get assigned to a group that does not perceive the tax at all, while those above it get assigned to a group that perceives it "maximally". Since demand is monotonic in p, and given our definitions of l and h, one can always pick $\tilde{p}^s$ such that the above inequalities hold weakly. Thus, one can always find $\lambda \in [0,1]$ such that:

$$\begin{aligned}\lambda\int_{p_i^s\in int\mathcal{P},\theta_i} q(p^b(p_i^s);\theta_i,h)dF_{p^s,\theta}(p_i^s,\theta_i) + (1-\lambda)\int_{p_i^s\in int\mathcal{P},\theta_i} q(p^b(p_i^s);\theta_i,l)dF_{p^s,\theta}(p_i^s,\theta_i)\\ = \int_{p_i^s\in int\mathcal{P},\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).\end{aligned} \tag{7}$$

Equation 7 implies that we can always pick a threshold $\tilde{p}^s$ that rationalizes demand, so long as we randomly assign a fraction λ of consumers with $m \in int(\mathcal{P})$ to the tie-breaking parameter $\zeta = h$, and the remaining $1-\lambda$ to $\zeta = l$. But in turn, this implies that we have found binary distribution of p^s and

[13] Recall that deadweight loss equals its calculation as if taxes actually satisfied $\tau_i = p_i^s - \bar{p}$, so excess burden cannot exceed the original consumer surplus for any agent. One can show that if τ^s has support on negative values, then it's possible to have total deadweight loss substantially greater than the original total consumer surplus.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

ζ that rationalizes aggregate demand.

Let us call such distribution $F''_{p^s,\theta,\zeta}$. This new distribution has the same marginal distribution of preferences, $F_\theta = F''_\theta$. If a consumer perceived a price in the boundary region in the original distribution $F_{p^s,\theta,\zeta}$, then so will she in the new distribution: $F''_{p^s,\theta,\zeta|p^s\in\partial\mathcal{P}} = F_{p^s,\theta,\zeta|p^s\in\partial\mathcal{P}}$. As for those consumers who perceived a price in the interior, we propose to split them up in a manner akin to what we just did in equation 7: $F''_{p^b(p^s),\theta} = F_{p^s,\theta}$, and ζ is assigned randomly as we described above. One can quickly confirm that this distribution rationalizes the same demand as the original distribution $F_{p^s,\theta,\zeta}$:

$$\begin{aligned}\int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i) &= \int_{p_i^s\in int(\mathcal{P}),\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i)\\ &+ \int_{p_i^s\in\partial\mathcal{P},\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i)\\ &= \int_{p_i^s\in int(\mathcal{P}),\theta_i} [\lambda q(p^b(p_i^s);\theta_i,h) + (1-\lambda)q(p^b(p_i^s),\theta_i,l)]dF_{p^s,\theta}(p_i^s,\theta_i)\\ &+ \int_{p_i^s\in\partial\mathcal{P},\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF''_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i)\\ &= \int_{p_i^s,\theta_i,\zeta_i} q(p_i^s;\theta_i,\zeta_i)dF''_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).\end{aligned}$$

Furthermore, such a distribution provides a generically larger value of deadweight loss than does $F_{p^s,\theta,\zeta}$.

Theorem 2. *Under assumption 1, for any $F_{p^s,\theta,\zeta}$ and any corresponding $F''_{p^s,\theta,\zeta}$ as described above:*

$$\int_{p_i^s,\theta_i,\zeta_i} dwl(p_i^s;\theta_i,\zeta_i)dF''_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i) \geq \int_{p_i^s,\theta_i,\zeta_i} dwl(p_i^s;\theta_i,\zeta_i)dF_{p^s,\theta,\zeta}(p_i^s,\theta_i,\zeta_i).$$

We can obtain intuition in two ways. One is to note that the method of forcing binary perceived prices increases heterogeneity of perceived prices compared to $F_{p^s,\theta,\zeta}$. Another is by considering the case where $\bar{m} = 1$ and F^*_θ is known to be degenerate, so that all agents have the same preferences. For a given aggregate demand, deadweight loss is maximized under these preferences when some perceive price $p_i^s = \bar{p}$, while others correctly perceived the true tax rate $p_i^s = \bar{p} + \tau$. This is because for each individual agent, deadweight loss is convex in the perceived price. Hence, for a given aggregate demand, *aggregate* deadweight loss will be highest when it is as high as possible for some – namely, those who fully perceive the tax – while it is null for everybody else – as those who don't perceive the tax at all are effectively subject to a lump-sum tax. We provide a graphical illustration of this argument in figure 4.

This article is protected by copyright. All rights reserved.

Author Manuscript

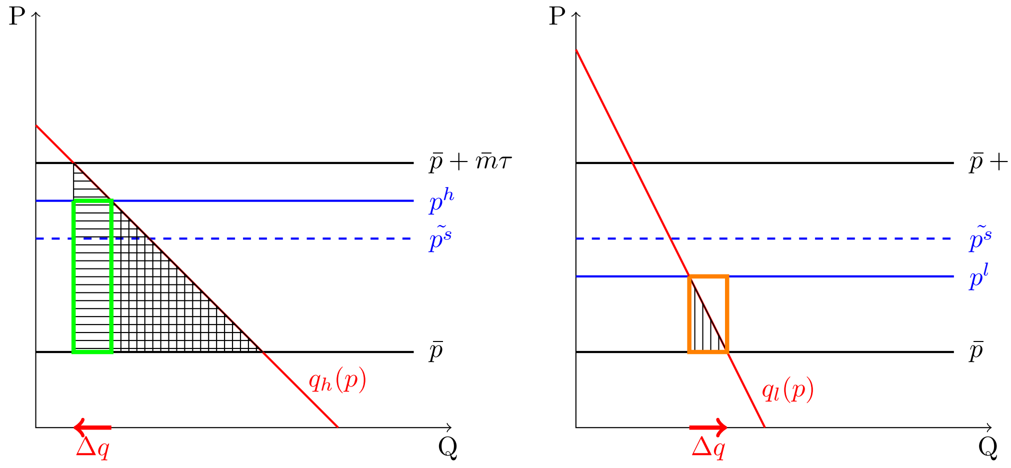

(a) High perceived price

(b) Low perceived price

Figure 4: A graphical illustration of Theorem 2. The watershed price $\tilde{p}^s$ is chosen to make the change in demand equal for the consumer in (a) and in (b). As long as we are dealing with weakly decreasing demand functions, the increase in deadweight loss for (a) is at least as big as the green box, while the decrease in deadweight loss for (b) is at most as big as the orange box. By assigning a perceived price of $\bar{p} + \bar{m}\tau$ to the consumer in (a) and $\bar{p}$ to the consumer in (b), we have increased aggregate deadweight loss holding aggregate demand constant.

Theorem 2 illustrates that for any distribution of (p^s, θ) that rationalizes the data, we can alternatively rationalize the data with a distribution with support for p_i^s on $\{\bar{p}, \bar{p} + \bar{m}\tau\}$ that yields (weakly) greater deadweight loss. Again, we see that identification of deadweight loss is not generally possible even if we knew the distribution of F_θ^*, as different marginal distributions of p^s and ζ could have different implications for deadweight loss. Also, any upper bound to the possible values of deadweight loss must be generated from a distribution with support of perceived prices on $\{\bar{p}, \bar{p} + \bar{m}\tau\}$.

However, not all distributions that have $p^s \in \partial\mathcal{P}$ with probability one yield the same value of deadweight loss, even when rationalizing the same data with the same distribution of preference types. The allocation of the good that yields the highest possible deadweight loss will also assign more consumption to agents with more convex demand curves. Theorem 3 in the online appendix spells out how to assign consumption of the good in the way it will "do the least good", and deals with cases where the tie-breaking parameter ζ is relevant, to get a general expression for the upper bound for deadweight loss consistent with aggregate demand and the distribution of preferences.

4 Linear Special Case

In this section, we discuss the special case in which q is known to be linear in $\bar{p}$ and τ (for fixed p^{NT}). We focus on this example both because of how frequently economists estimate linear models and because of its relationship to the second order approximation of deadweight loss.

We can also use the linear special case to better illustrate the general identification problem. In this subsection, we will no longer assume that the distribution of preference parameters is known; the econometrician will, as is usually the case, be able to identify preferences with exogenous price

This article is protected by copyright. All rights reserved.

variation. As in TRJ (2018), we assume the distribution of preferences does not depend on taxes (or prices). Further, we permit the econometrician to assume that the distribution of tax salience does not change as sticker prices and taxes vary. This entirely rules out any sort of endogeneity between attention and taxes – which was driving non-identification in the binary case – and yet we will get the same non-identification result.

Naturally, no demand curve can be entirely linear, for the simple reason that agents cannot consume negative quantities. But in practice, one rarely gets the privilege of such rich variation in prices when doing empirical work. What we are implicitly assuming in this section is that the values of $(\bar{p}, \tau)$ that are considered are all such that $\bar{p} > 0$, and $q_i > 0$ for all consumers in the market, so that aggregate demand is also linear at those values. In other words, one can think of our work here as modeling linear demand conditional on buying at all prices under consideration.[14]

We might recall from section 2 that one can express a second order approximation to deadweight loss as a function of derivatives. If the choice function is linear in regressors $\bar{p}$ and τ, the second order approximation is an exact calculation of deadweight loss, and our results from the previous sections apply.

Formally, each preference type θ_i takes the form $\theta_i = (\beta_i, \epsilon_i) \in \mathbb{R}^2$.[15] To maintain linearity in regressors, we also assume that tax salience m is constant with respect to τ. The choice function q then takes the form:

$$q_i = \alpha + \beta_i p_i^s + \epsilon_i = \alpha + \beta_i[\bar{p} + m_i\tau] + \epsilon_i$$

We are suppressing the tie-breaking parameter ζ because in this linear example $\theta_i, m_i, \bar{p}$, and τ always uniquely determine consumption. We have the parameter α so that we can assume without loss of generality that $\mathbb{E}[\epsilon] = 0$.

Defining $\tilde{\beta}_i \equiv m_i\beta_i$ yields:

$$q_i = \alpha + \beta_i\bar{p} + \tilde{\beta}_i\tau + \epsilon_i, \tag{8}$$

with corresponding deadweight loss per agent, from equation 4:

$$\begin{aligned} dwl_i &= \int_{\bar{p}}^{p^s} [\alpha + \beta_i p + \epsilon_i]dp - (p^s - \bar{p})[\alpha + \beta_i p^s + \epsilon_i] = \int_{\bar{p}}^{p^s} (p - p^s)\beta_i dp \\ &= \frac{1}{2}\Big[\frac{p^{s2} - \bar{p}^2}{2} - (p^s - \bar{p})p^s\Big]\beta_i = \frac{1}{2}(p^s - \bar{p})[(p^s + \bar{p}) - 2p^s]\beta_i = -\frac{1}{2}\tau^{s2}\beta_i \\ &= -\frac{1}{2}m_i^2\beta_i\tau^2. \end{aligned}$$

We assume that not just the distribution of preference parameters, but also the joint distribution of preference and salience parameters remains unaffected by the specific values of $\bar{p}$ and τ. The econometrician

[14]We thank an anonymous referee at the Journal of Public Economic Theory for helping us clarify our thinking on this matter.

[15]Agents have quasi-linear utility $u_i = \frac{q_i^2/2 - (\alpha+\epsilon_i)q_i}{\beth_i} + q_i^{NT}$. For a given p^{NT}, we define $\beta_i \equiv \frac{\beth_i}{p^{NT}}$, yielding utility representation $U_i = \frac{q_i^2/2 - (\alpha+\epsilon_i)q_i}{\beta_i} + p^{NT}q_i^{NT}$.

This article is protected by copyright. All rights reserved.

Author Manuscript

Author Manuscript

observes for various values of regressors:

$$\mathbb{E}[q|\bar{p},\tau] \equiv \int_{\beta_i,\tilde{\beta}_i,\epsilon_i} [\alpha + \beta_i \bar{p} + \tilde{\beta}_i \tau + \epsilon_i] dF^*_{\beta,\tilde{\beta},\epsilon}(\beta_i,\tilde{\beta}_i,\epsilon_i) = \alpha + \mathbb{E}[\beta]\bar{p} + \mathbb{E}[\tilde{\beta}]\tau. \tag{9}$$

where $F^*_{\beta,\tilde{\beta},\epsilon}$ is the true distribution of $(\beta, m\beta, \epsilon)$. The challenge is to use the observed values of triplets $(\bar{p},\tau,\ \mathbb{E}[q|\bar{p},\tau])$ to infer aggregate deadweight loss, which in this case is equivalent to its second order approximation around $\tau = 0$:

$$DWL = -\frac{1}{2}\int_{\beta_i,m_i} m_i^2 \beta_i dF_{\beta,m}(\beta_i,m_i)\tau^2 = -\frac{1}{2}\mathbb{E}[m^2\beta]\tau^2 = -\frac{1}{2}\mathbb{E}[m\tilde{\beta}]\tau^2.$$

The only restriction that the econometrician imposes on the distribution of tax salience m is that the support of tax salience is contained within the interval $[0,\bar{m}]$. The econometrician can also use the Compensated Law of Demand as defined in appendix lemma 2, which shows that compensated demand is always weakly decreasing, so that $\mathbb{P}[\beta \le 0] = 1$. In fact, we can permit the econometrician to know the entire distribution of $\theta = (\beta,\epsilon)$. It will not affect our results.} First, we can find a homogeneous perceived price that rationalizes the data for any τ. In particular, a linear regression of aggregate demand on sticker prices and taxes may permit identification of $\hat{\beta} \equiv \mathbb{E}[\beta]$ and $\hat{\tilde{\beta}} \equiv \mathbb{E}[\tilde{\beta}]$, respectively.[16] We define a measure of central tendency of tax salience:[17]

$$\hat{m} \equiv \frac{\hat{\tilde{\beta}}}{\hat{\beta}}.$$

Then the homogeneous perceived price that rationalizes the data is $\hat{p}^s = \bar{p} + \hat{m}\tau$. To see this, note that assuming all agents have tax salience $m_i = \hat{m}$ yields aggregate demand as in equation 9:

$$\begin{aligned}\int_{\beta_i,\epsilon_i} [\alpha + \beta_i \hat{p}^s + \epsilon_i] dF^*_{\beta,\epsilon}(\beta_i,\epsilon_i) &= \alpha + \bar{p}\int_{\beta_i,\epsilon_i} \beta_i dF^*_{\beta}(\beta_i) + \hat{m}\tau \int_{\beta_i} \beta_i dF^*_{\beta}(\beta_i) \\ &= \alpha + \hat{\beta}\bar{p} + \hat{m}\hat{\tilde{\beta}}\tau \\ &= \alpha + \hat{\beta}\ \bar{p} + \hat{\tilde{\beta}}\ \tau.\end{aligned}$$

Thus, the econometrician cannot rule out all agents perceiving the same price $\hat{p}^s$, and so cannot rule out $m_i = \hat{m}\ \forall i$. For tax τ, this would yield deadweight loss:

$$DWL_{low} = -\frac{1}{2}\hat{m}\hat{\tilde{\beta}}\tau^2.$$

By theorem 1, this is a lower bound for deadweight loss.

Alternatively, the econometrician cannot rule out the perceived tax τ^s having support in $\{0, \bar{m}\tau\}$. To see this, consider $\mathbb{P}(p^s = \bar{p} + \bar{m}\tau) = \frac{\hat{m}}{\bar{m}}$ and $\mathbb{P}(p^s = \bar{p}) = 1 - \frac{\hat{m}}{\bar{m}}$ independently of other parameters and

[16] Such identification requires exogenous and non-collinear variation in sticker prices and taxes. If the econometrician cannot identify these terms, so much the worse for identifying aggregate deadweight loss.

[17] If $\hat{\beta} = 0$, then let $\hat{m} = 0$.

This article is protected by copyright. All rights reserved.

Author Manuscript

regressors.[18] This will rationalize aggregate demand:

$$\int_{\beta_i,\epsilon_i}\Big[\alpha+\beta_i\bar{p}_i+\frac{\hat{m}}{\bar{m}}\beta_i\bar{m}\tau_i+\epsilon_i\Big]dF_{\beta,\epsilon}(\beta_i,\tilde{\beta}_i,\epsilon_i)=\alpha+\mathbb{E}[\beta]\bar{p}+\hat{m}\mathbb{E}[\beta]\tau=\alpha+\mathbb{E}[\beta]\bar{p}+\mathbb{E}[\tilde{\beta}]\tau.$$

This yields deadweight loss for tax τ:

$$DWL_{high}=-\frac{1}{2}\frac{\hat{m}}{\bar{m}}\mathbb{E}[\beta]\bar{m}^2\tau^2=-\frac{1}{2}\hat{m}\hat{\beta}\bar{m}\tau^2=-\frac{1}{2}\bar{m}\hat{\tilde{\beta}}\tau^2.$$

For instance, if $\bar{m}=1$, then the value of deadweight loss under a homogeneous perceived price is fraction $\hat{m}$ of the above calculation of deadweight loss.

Proceeding from theorem 2, we noted that there is a specific distribution of perceived prices on $\{\bar{p},\bar{p}+\bar{m}\tau\}$ that maximizes deadweight loss. We describe that distribution in theorem 3, noting that it involves assigning high or low perceived prices based on the ratio of per-person deadweight loss to the change in consumption for that individual. But in this context:

$$\frac{dwl_i}{q_i(\bar{p})-q_i(p^s)}=\frac{\tau^s}{2}.$$

Thus, the distribution of tax salience independent of all other parameters and regressors in which $\mathbb{P}(m=\bar{m})=\frac{\hat{m}}{\bar{m}}$ and $\mathbb{P}(m=0)=1-\frac{\hat{m}}{\bar{m}}$ maximizes deadweight loss. More generally, the econometrician cannot rule out this maximal value of deadweight loss so long as they cannot rule out the possibility of some distribution F with $F_\beta=F^*_\beta$ such that $supp(m)\in\{0,\bar{m}\}$ with:

$$\mathbb{P}_F(m=\bar{m})\mathbb{E}_F[\tilde{\beta}|m=\bar{m}]=\hat{m}\hat{\beta}=\hat{\tilde{\beta}}.$$

One can check that this distribution rationalizes the data,

$$\mathbb{E}[q|\bar{p},\tau]=\alpha+\mathbb{E}[\beta]\bar{p}+\mathbb{E}_F[\tilde{\beta}]\tau=\alpha+\hat{\beta}\bar{p}+\mathbb{P}_F[m=\bar{m}]\mathbb{E}_F[\tilde{\beta}|m=\bar{m}]\tau=\alpha+\hat{\beta}\bar{p}+\hat{\tilde{\beta}}\tau,$$

and yields the maximal value of deadweight loss,

$$-\frac{1}{2}\mathbb{E}_F[m^2\beta]\tau^2=-\frac{1}{2}\mathbb{P}_F[m=\bar{m}]\bar{m}\mathbb{E}_F[\tilde{\beta}|m=\bar{m}]\tau^2=-\frac{1}{2}\bar{m}\hat{\tilde{\beta}}\tau^2=DWL_{high}.$$

More intuitively, once one knows $\hat{\beta}$ and $\hat{\tilde{\beta}}$, one can rationalize the aggregate data. Since the ratio of deadweight loss to the change in quantity is constant, the relationship between tax salience and preferences doesn't matter upon attaining the observed aggregate demand.

Finally, consider a distribution with $m\perp(\beta,\epsilon)$ with $supp(m)\subseteq\{0,\hat{m},\bar{m}\}$, $\mathbb{P}(m=\hat{m})=\lambda$ and $\mathbb{P}(m=\bar{m}|m\neq\hat{m})=\frac{\hat{m}}{\bar{m}}$. Varying λ from zero to one yields:

$$DWL\in\Big[-\frac{1}{2}\hat{m}\hat{\tilde{\beta}}\tau^2,-\frac{1}{2}\bar{m}\hat{\tilde{\beta}}\tau^2\Big].$$

[18]In the true distribution, it must be that $\hat{m}\in[0,\bar{m}]$. Alternatively, one could consider checking whether $\hat{m}\in[0,\bar{m}]$ as a weak test of the null hypothesis that tax salience is bounded within that interval.

This article is protected by copyright. All rights reserved.

We can conclude from this result that one cannot even identify a second order approximation of deadweight loss with aggregate data alone.[19] Imposing structure on preferences to facilitate identification of F_θ^* still only permits interval identification. Nonetheless, we can use aggregate data to obtain bounds, or at least $\hat{m}$, which gives us a sense of the uncertainty over the possible values of deadweight loss.

Besides its illustrative value, this linear framework also gives researchers a quick and easy way to compute bounds for the deadweight loss of non-salient taxes or fees in a variety of empirical contexts. In the online appendix, we apply these findings to the framework of CLK's study of aggregate beer consumption and Goldin and Homonoff's (2013) study of cigarette consumption. Details on the estimation procedures are provided in the online appendix. In the baseline specification of the CLK (2009) data, we estimate that $\hat{m} \approx 0.31$. This estimate suggests that even assuming that salience cannot exceed one, $\bar{m} = 1$, the upper bound of deadweight loss is about three times the lower bound. These estimates, however, all seem fairly imprecise. Across the two data-sets, there is no specification in which we can reject the null hypothesis that $\hat{m} = 0$, permitting the upper bound to be arbitrarily large in proportion to the lower bound.[20] Similarly, in most specifications we cannot reject that $\hat{m} = 1$, which would imply that upper bound and lower bound are identical. This underlying uncertainty is mirrored in previous work on tax salience – e.g., TRJ (2018) find that individual differences increase excess burden by at least 200% relative to the case of homogeneous attention – but our wide confidence intervals may be the result of the specific data sets we are using here – for example, Goldin and Homonoff (2013) often cannot reject that consumers do not react at all to sales taxes. Our procedure, however, seems so straightforward to carry out that it might turn out useful in future research on tax salience.

Finally, in both the setting of CLK (2009) and Goldin and Homonoff (2013), functional form assumptions seem to matter. The limitations of the linear setting would prompt us to undergo more sophisticated and less parametric exercises, but we are dissuaded by the fact that our statistical power is already very low.

5 Conclusion

In this paper, we studied deadweight loss in a model where agents misperceive prices. We started by generalizing the theoretical results of CLK (2007, 2009) with an arbitrary closed choice set. Inattentiveness to taxes makes agents worse off while reducing deadweight loss by preventing agents from avoiding the tax.

As in the binary choice model of TRJ (2018), heterogeneous attention adds another layer of complexity to deadweight loss. In our general setting, we show that aggregate consumption can be consistent with a wide variety of co-distributions of attention and preferences, each with a different implication for deadweight loss. This is because it matters who gets the good and who doesn't: when prices are misperceived, there is no guarantee that agents who end up with some units of the good are the ones who value those units most. By minimizing and maximizing this allocative inefficiency, we show that deadweight loss can only vary between two extremes for any given aggregate demand. The lower bound holds generally, while the upper bound relies on the assumption that tax salience has support contained

[19] One can identify a first order approximation trivially: it is zero.

[20] Note that $\frac{DWL_{high}}{DWL_{low}} = \frac{\bar{m}}{\hat{m}}$, so that as $\hat{m}$ approaches zero from above, this ratio of upper to lower bounds blows up to infinity.

Author Manuscript

This article is protected by copyright. All rights reserved.

in a known non-negative interval.

Finally, we explore the special case in which demand is linear, which is of special interest due to its relationship to both the empirical literature and the second order approximation of deadweight loss. Our analysis shows that, while identification of deadweight loss under binary choice may be restated as an endogeneity problem, the same cannot be said regardless of the choice set. Indeed, when individual demand is linear, assuming independence of tax salience from taxes and prices does not change the interval of possible values of deadweight loss.

The linear model yields bounds for deadweight loss that one can easily compute from linear regression estimates. While this doesn't necessarily doom any future application in empirical work, our own applications of this method on the existing work of CLK (2009) and Goldin and Homonoff (2013) leave us without many answers about how tight these bounds might be. While some point estimates seem reasonable, they also can be imprecise and dependent on functional form specification.

Author Manuscript

This article is protected by copyright. All rights reserved.

Author Manuscript

References

Hunt Allcott and Dmitry Taubinsky. Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–38, 2015.

B Douglas Bernheim and Antonio Rangel. Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics*, 124(1):51–104, 2009.

Tom Blake, Sarah Moshary, and Steve Tadelis. Price salience and product choice. *NBER Working Paper N. 25186*, 2017.

Robin Boadway, Jean-Denis Garon, and Louis Perrault. Optimal mixed taxation, credit constraints, and the timing of income tax reporting. *Journal of Public Economic Theory*, 21:708–737, 2018.

Steven Bradley and Naomi Feldman. Hidden baggage: behavioral responses to changes in airline ticket tax disclosure. *Finance and Economics Discussion Series 2018-057. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2018.057*, 2018.

Raj Chetty. The simple economics of salience and taxation. *NBER Working Paper N. 15246*, 2009.

Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation. *NBER Working Paper N. 13330*, 2007.

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

Amy Finkelstein. Ez-tax: Tax salience and tax rates. *Quarterly Journal of Economics*, 124(3):969–1010, 2009.

Jeremy T Fox. A note on nonparametric identification of distributions of random coefficients in multinomial choice models. *NBER Working Paper No. 23621*, 2017.

Xavier Gabaix. A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*, 129 (4):1661–1710, 2014.

Kelly Sims Gallagher and Erich Muehlegger. Giving green to get green? incentives and consumer adoption of hybrid vehicle technology. *Journal of Environmental Economics and Management*, 61(1): 1–15, 2011.

Jacob Goldin. Optimal tax salience. *Journal of Public Economics*, 131:115–123, 2015.

Jacob Goldin and Tatiana Homonoff. Smoke gets in your eyes: cigarette tax salience and regressivity. *American Economic Journal: Economic Policy*, 5(1):302–336, 2013.

Matthew A Masten. Random coefficients on endogenous variables in simultaneous equations models. *The Review of Economic Studies*, 85(2):1193–1250, 2017.

Andrea Morone, Francesco Nemore, and Simone Nuzzo. Experimental evidence on tax salience and tax incidence. *Journal of Public Economic Theory*, 20(4):582–612, 2018.

This article is protected by copyright. All rights reserved.

Harvey S Rosen. Tax illusion and the labor supply of married women. *The Review of Economics and Statistics*, 58(2):167–172, 1976.

Dmitry Taubinsky and Alex Rees-Jones. Attention variation and welfare: theory and evidence from a tax salience experiment. *Review of Economic Studies*, 85(4):2461–2496, 2018.

Matthias Weber and Arthur Schram. The non-equivalence of labour market taxes: A real-effort experiment. *The Economic Journal*, 127(604), 2016.

Author Manuscript

This article is protected by copyright. All rights reserved.