Robust Bounds for Welfare Analysis[*]

Zi Yang Kang[†]   Shoshana Vasserman[‡]

May 2022

Abstract

Economists routinely make functional form assumptions about consumer demand to obtain welfare estimates. How sensitive are welfare estimates to these assumptions? We answer this question by providing bounds on welfare that hold for families of demand curves commonly considered in different literatures. We show that commonly chosen functional forms, such as linear, exponential, and CES demand, are extremal in different families: they yield either the highest or lowest welfare estimate among all demand curves in those families. To illustrate our approach, we apply our results to the welfare analysis of energy subsidies, trade tariffs, pensions, and income taxation.

JEL classification: C14, C51, D04, D61, F14, H21, Q48

Keywords: welfare analysis, welfare bounds, robustness, functional form assumptions

---

[†] Graduate School of Business, Stanford University; zykang@stanford.edu.
[‡] Graduate School of Business, Stanford University and NBER; svass@stanford.edu.

1 Introduction

Welfare analysis is central to economic policy evaluation. Raising the tax on a good, for example, leads to a gain in tax revenue but a loss in consumer surplus. But while revenue changes are often easy to measure, consumer surplus cannot be directly observed. Economic theory dating back to Marshall (1890) has established methods for extrapolating consumer surplus from observations of consumer demand: the loss in consumer surplus due to a price increase is equal to the area below the demand curve between the two prices.

However, this approach requires the entire demand curve to be observed, or at least inferred from granular data—conditions that are rarely satisfied in practice. This is especially salient in empirical studies that evaluate the welfare impact of policies on the basis of observed responses to (quasi-)experimental price shocks. The number of such studies has increased exponentially in recent years, following advances in econometric techniques (the so-called "credibility revolution") that facilitate credible and precise estimates of responses to discrete policy changes (see, for example, the surveys of Finkelstein and Hendren, 2020 and Kleven, 2021). However, practical limitations (*e.g.*, political barriers and administrative costs) often restrict the number of points along the demand curve that can be examined.

Many empirical studies therefore interpolate between points along the demand curve using standard functional forms. For instance, linear interpolations (Harberger, 1964; Feldstein, 1999; Einav, Finkelstein, and Cullen, 2010; Hackmann, Kolstad, and Kowalski, 2015; Cohen, Hahn, Hall, Levitt, and Metcalfe, 2016; Amiti, Redding, and Weinstein, 2019; Hahn and Metcalfe, 2021) and constant elasticity of substitution (CES) interpolations (Hausman, 1981; Hausman, Pakes, and Rosston, 1997; Brynjolfsson, Hu, and Smith, 2003; Fajgelbaum, Goldberg, Kennedy, and Khandelwal, 2020) are widely used across different fields of economics, including applied microeconomics, international trade, and public finance. However, as Finkelstein and Hendren acknowledge, these interpolations are often used for convenience and "ease of implementation," rather than realism or economic reasoning.

Other studies take a more conservative approach: by assuming that marginal individuals have either zero or full willingness to pay, they estimate lower and upper bounds, respectively, for the welfare impact of policies. This approach is equivalent to interpolating with either the pointwise highest or lowest demand curve that passes through the observed points, and is common in applied microeconomics and public finance (Varian, 1985; Hendren and Sprung-Keyser, 2020; Jácome, 2020; Giesecke and Jäger, 2021; Gray, Leive, Prager, Pukelis, and Zaki, 2021; Deshpande and Mueller-Smith, 2022). Yet these assumptions are often unrealistic: they imply that demand

is either perfectly elastic or perfectly inelastic over a potentially wide range of prices. These assumptions are also routinely rejected when tested: extensive empirical literatures have estimated price elasticities in various settings, such as the short-run demand for gasoline (ranging between -0.2 and -0.4; see Morris, 2014 and Kilian and Zhou, 2020) and demand for medical care (widely cited as -0.2 due to the RAND Health Insurance Experiment; see Manning, Newhouse, Duan, Keeler, and Leibowitz, 1987 and Cutler and Zeckhauser, 2000).

In this paper, we propose a complementary approach for evaluating welfare changes in settings with limited data. Rather than estimate an entire demand curve to infer the change in welfare, our approach bounds the change in welfare within different families of demand curves. Our bounds are *simple*: they can be computed in closed form using only data from before and after a policy change. Our bounds are also *robust*: they apply to any demand curve in each family that we consider. Thus, rather than estimate a demand curve that "reasonably" captures preferences exhibited in data in order to evaluate the change in welfare, we compute the smallest and largest changes in welfare that are consistent with any "reasonable" demand curve. Finally, our bounds are *sharp*: any value between the upper and lower bounds can be attained as the change in welfare for some demand curve in that family.

To fix ideas, consider the canonical example of a tax levied on a good (Harberger, 1964). There are two periods: $t = 0$ before the tax is levied, and $t = 1$ after. The market is perfectly competitive and the demand curve does not shift between the two periods. At $t = 0$, q_0 units of the good are sold at a unit price of p_0. At $t = 1$, the posted price remains unchanged, but an *ad valorem* tax τ is introduced, yielding an effective price of $p_1 = (1 + \tau) p_0$ and a new equilibrium quantity q_1. To evaluate the net impact of the tax, the researcher faces the problem depicted in Figure 1(a).
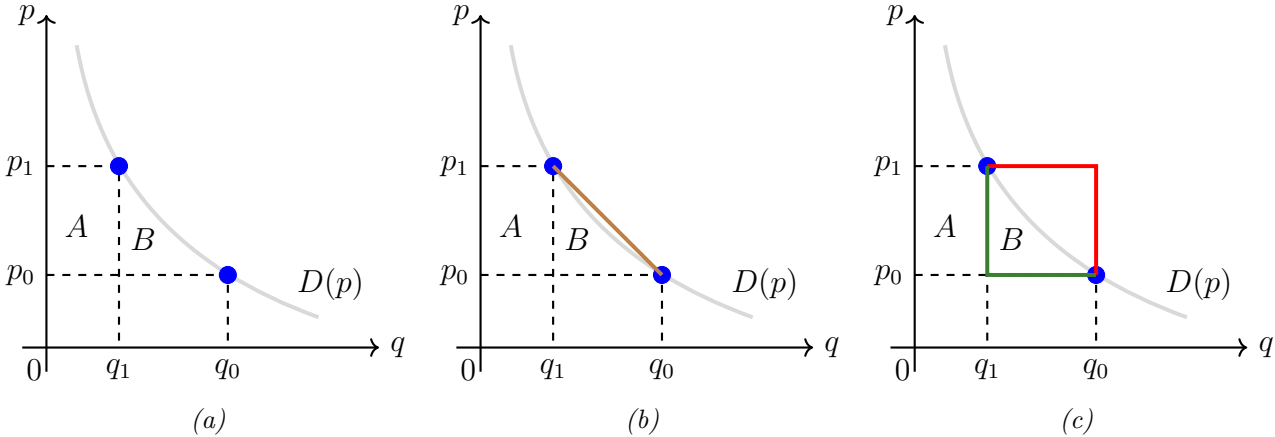

Figure 1: *Illustration of how the change in consumer surplus from a price increase can be estimated.*

She observes the points (p_0, q_0) and (p_1, q_1), but not the demand curve $D(p)$ that connects them. While computing the revenue gain from the tax (area A) is easy, computing the loss in consumer surplus (the sum of areas A and B) requires integrating over $D(p)$, which is unknown.

To resolve this, Harberger interpolated between (p_0, q_0) and (p_1, q_1) with a line, as depicted in Figure 1(b)—an approach that remains popular to date, as we argued above. When the tax τ is small, such an interpolation is justified via Taylor's theorem; but no such guarantee holds when τ is large—which is often the case in practice.

By contrast, the approach that we take in this paper bounds the change in welfare under two types of assumptions, which we formally discuss in Section 2. A special case of our approach is the conservative one that many empirical studies already take, depicted in Figure 1(c). These conservative bounds are attained by the extremal demand curves that decrease from (p_1, q_1) to (p_0, q_0): between p_0 and p_1, the lower extremal demand curve (in green) is constant at q_1, while the upper extremal demand curve (in red) is constant at q_0. These bounds are intuitive and robust: they hold for any downward-sloping demand curve and do not depend on supply-side assumptions. We generalize the conservative approach with additional restrictions on the demand curve that can capture more realistic distributions of willingness to pay while preserving the intuitiveness and robustness of the conservative bounds.

The first type of assumption is a restriction on the price elasticity of demand. For instance, the researcher might reason that elasticities are unlikely to vary unboundedly as the conservative extremal demand curves shown in Figure 1(c) do. In this case, it might make sense to restrict attention to demand curves for which elasticities are contained in a finite interval $[\underline{\varepsilon}, \overline{\varepsilon}]$ between p_0 and p_1. The researcher might have measured the elasticities at p_0 and p_1, and reason that the elasticity at any intermediate price lies in between; or the researcher might have obtained a range of elasticities measured in comparable settings, and posit that elasticities cannot deviate too far from that range. Our first main result shows that under this restriction, bounds on the change in consumer surplus are attained by demand curves with a surprisingly simple structure: each extremal demand curve consists of two pieces, such that the elasticity throughout each piece is equal to either $\underline{\varepsilon}$ or $\overline{\varepsilon}$. In the limiting case when $\underline{\varepsilon} = -\infty$ and $\overline{\varepsilon} = 0$, our bounds are identical to the conservative bounds.

The second type of assumption consists of restrictions on the curvature of demand. Such restrictions are commonly imposed to connect empirical observations to theoretical models across different economic fields. An example is Marshall's second law—demand is more elastic at higher prices—which is often assumed in models of international trade in order to guarantee that tougher

3

competition between firms leads to higher markups.[1] Other examples include the assumptions of decreasing marginal revenue and log-concavity, which guarantee the existence (and sometimes uniqueness) of equilibrium in models of industrial organization (Caplin and Nalebuff, 1991a), as well as concavity, ρ-concavity, convexity, log-convexity, and ρ-convexity. Our second main result derives the largest and smallest possible changes in consumer surplus under each of these assumptions, as well as the extremal demand curves that attain them.

Our results uncover a key insight: commonly used functional forms are often extremal. For example, our analysis implies that the *smallest* possible loss in consumer surplus is attained by a CES demand curve under Marshall's second law. Consequently, a CES functional form assumption never produces a "representative" welfare estimate in this family of demand curves—rather, it produces the most conservative welfare estimate. Similarly, a linear functional form assumption produces the largest possible welfare estimate under the assumption that demand is convex.

As we show in Section 3, our framework can be extended in various dimensions motivated by empirical applications. For example, we derive bounds for the case that the researcher observes fewer points (*e.g.*, when evaluating the welfare impact of a counterfactual policy, the researcher does not observe the quantity of the good sold at the counterfactual price) or more points (*e.g.*, when there are multiple policy changes) along the demand curve. Motivated by the fact that points on the demand curve are almost never observed perfectly in empirical applications, we also demonstrate how sampling error can be incorporated into our framework. Finally, we show how our results extend to other measures of welfare, including deadweight loss, compensating variation, equivalent variation, supply-side welfare measures, and welfare measures that incorporates equity through social welfare weights (Harberger, 1978; Saez and Stantcheva, 2016).

A natural concern is that the bounds that we derive may be too wide to be informative in practice. To assuage this concern, we apply our approach to four empirical settings in Section 4. In each setting, we demonstrate how our bounds compare to and complement existing approaches to welfare evaluation:

(i) **Energy subsidies.** We apply our results to Hahn and Metcalfe (2021), who employ a large field experiment to evaluate the welfare impact of energy subsidies under the California Alternative Rates for Energy (CARE) program. While they show that the net welfare impact of CARE is negative under the functional form assumption that demand is linear,

---

[1] A notable example is Krugman (1979), who invoked Marshall's second law "without apology" given that the assumption "seems to be necessary if [his] model is to yield reasonable results." This assumption seemed so intuitive that Marshall (1890) required it as part of his definition of a demand curve. Melitz (2018) discusses the role of Marshall's second law in the international trade literature as well as supporting empirical evidence.

our results imply that this result is robust: it holds not only for linear demand curves, but also for a large set of demand curves that are consistent with their elasticity estimates.

(ii) **Trade tariffs.** We apply our results to Amiti et al.'s (2019) analysis of the deadweight loss due to trade tariffs imposed by the United States on foreign imports between 2018 and 2019. Our results quantify the sensitivity of Amiti et al.'s estimates under a linear demand model to alternative specifications from the same literature, such as a CES demand model, and demonstrate different conditions under which their estimate might be too high or too low.

(iii) **Pensions.** We show how our results apply to the marginal value of public funds (MVPF) literature by building on Giesecke and Jäger's (2021) analysis of old-age pensions in the UK. While Giesecke and Jäger estimate the gain in worker surplus using a conservative bound, our results show how heterogeneity in income and disutility from working leads to higher estimates of surplus gain, which in turn imply higher MVPF estimates.

(iv) **Income taxes.** We apply our results to Feldstein's (1995) analysis of the excess burden of income taxation. While Feldstein employs a linear interpolation to estimate the excess burden, our results show how narrow bounds can be obtained using estimated elasticities of taxable income, even when no functional form is assumed.

Finally, Section 5 concludes. We provide closed-form expressions for our bounds in Appendix A, and demonstrate how our main results can be viewed through the lens of information design in Appendix B. In Appendix C, we discuss the relationship between different assumptions on the curvature of demand, while Appendix D collects omitted proofs and some additional discussion.

1.1   Related literature

Conceptually, our paper is closely related to the literature on "sufficient statistics" for welfare analysis—a phrase coined by Chetty (2009). Since Chetty's article, a growing body of work has embraced this approach for policy evaluation, as recently surveyed by Kleven (2021). The sufficient statistics approach stipulates that the welfare impacts of small policy changes can be well approximated without specifying a comprehensive model for the determinants of market equilibria. Instead, carefully deployed envelope conditions facilitate simple formulas that can be computed with local measurements and reduced form elasticity estimates. As we discuss in Section 4.4, this approach has many conceptual features in common with ours. Like ours, papers adopting the sufficient statistics approach shy away from specifying a particular demand curve, instead focusing

on inferences around an exogenous policy change. In this sense, sufficient statistic estimates are also robust to parametric assumptions. However whereas sufficient statistic estimators require the changes being analyzed to be sufficiently small that a local approximation suffices, our approach does not. Instead, one can think of our approach as an intermediate between the sufficient statistics approach and a fully structural one: by allowing researchers to specify conditions on the shape of demand, our bounds expand the set of policies that can be studied without requiring a commitment to a particular parametric model of the demand curve itself.

There is an extensive literature on welfare analysis in economics,[2] and the insight that a linear demand curve provides the upper bound for the change in consumer surplus when demand is convex is considered lore by many economists (Hausman, 2003). However, beyond this particular one-sided bound, linear interpolations are often described as "*ad hoc* at best" and other assumptions are generally not considered (Hausman, 1996).

Complementary to our approach, there is a large literature on demand estimation (see, for example, Berry and Haile, 2021 and Gandhi and Nevo, 2021). Demand estimation allows welfare to be computed in a straightforward manner, but requires granular data for credible identification. Our approach is motivated by empirical applications (such as those in Section 4) where the data are not granular enough for structural modeling. In these cases, we show that meaningful bounds on welfare can nonetheless often be obtained.

Our paper also relates to the literature on set identification: like other papers in this literature, our exercise of deriving bounds under relatively weak restrictions aims to provide transparency for the mapping between modeling assumptions and subsequent conclusions about welfare (Tamer, 2010). Closest to our paper, Manski (1997) derives sharp bounds on the distribution of demand curves in a cross-section of markets under a monotonicity assumption similar to Varian (1985) as well as a concavity assumption similar to our Assumption (**CA4**). However, while Manski set-identifies demand curves themselves, our approach set-identifies the welfare measures that are consistent with feasible demand curves. Closer to this idea, a burgeoning literature studies identification for counterfactual equilibria that may arise from perturbations of model primitives estimated in data.[3] We view these papers, which study different types of games and often require more data and structure in order to consider counterfactual outcomes, as complementary to ours.

Theoretically, our work relates to Bulow and Pfleiderer's (1983) critique of an empirical study by Sumner (1981) that used observed changes in marginal costs to estimate demand elasticities. Bulow and Pfleiderer point out that even though two demand curves might be close to each other,

---

[2] See, for instance, Willig (1976), Varian (1982, 1985), and Hausman (1981).
[3] See, for instance, Aguirregabiria (2010), Reguant (2016), and Kalouptsidi, Scott, and Souza-Rodrigues (2021).

their elasticities might be very different—just as the derivatives of two "similar-looking" functions can be far apart. The main result of our paper leverages the converse of this observation: even though two functions might be different, their *integrals* might be close. Our bounds on welfare—which are obtained by integrating different demand curves—may thus be quite narrow even if the family of demand curves that they account for is very broad.

Our work is also connected with Anderson and Renault (2003), Tsitsiklis and Xu (2014), and Condorelli and Szentes (2020), who bound welfare measures in Cournot competition; while our paper shares a similar theoretical objective, our bounds do not rely on supply-side assumptions, but rather elasticity and curvature restrictions arising from different empirical applications. Our work is also related in spirit to the theoretical literature on optimal pricing with limited knowledge of demand, such as Bergemann and Schlag (2011), Cohen, Perakis, and Pindyck (2021), and Bergemann, Castro, and Weintraub (2022); rather than optimal pricing, however, we focus on deriving sharp bounds on welfare measures.

From a methodological point of view, our analysis exploits a novel connection between welfare analysis and problems in mechanism and information design. This connection allows us to leverage tools from recent work on mechanism and information design: the proof of our main results builds on an insight of Kang and Vondrák (2019), who solve a constrained optimization problem in mechanism design by showing that the underlying change in welfare is monotone with respect to the convex partial order. This approach has also been used in information design by Gentzkow and Kamenica (2016), and was recently generalized by Kleiner, Moldovanu, and Strack (2021), who present a general framework for solving infinite-dimensional optimization problems with a majorization constraint.

2 Robust bounds for welfare analysis

In this section, we derive our main theoretical results: tight welfare bounds under distinct sets of assumptions that are widely maintained in empirical applications. We consider two types of assumptions: restrictions on price elasticities of demand (*e.g.*, using treatment effect estimates from field experiments) and restrictions on the curvature of demand (*e.g.*, implied by observed comparative statics).

We present our results in the context of a basic model, so as to emphasize the key ideas and intuition behind our approach. Limitations of this basic model are discussed at the end of this section, and Section 3 shows how our results extend to more general environments.

2.1 Basic model

We study a market with a consumption good and a numéraire ("money"). The market consists of a continuum of consumers. We make the following standard assumption:

(**A1**) Consumers have quasilinear utility in money.

This assumption can be viewed as an approximation that the market in question is a sufficiently small part of the economy, so that income effects can be ignored (Vives, 1987). Moreover, it is widely maintained in many empirical studies, and implies that the market has a downward-sloping Marshallian demand curve, which we denote by $D(\cdot)$. In Section 3, we show how this assumption can be relaxed to accommodate income effects.

The good is exposed to an exogenous price shock (*e.g.*, a policy change, such as a new *ad valorem* tax). For simplicity, we assume that there are two time periods: before the policy change ($t = 0$) and after ($t = 1$). The price increases from p_0 to p_1; correspondingly, the quantities of the good sold in each period are denoted by $q_0 = D(p_0)$ and $q_1 = D(p_1)$.

We evaluate the change in consumer surplus arising from the price increase. As we show later, a similar analysis applies to other welfare measures. By (**A1**), the change in consumer surplus is equal to the area below the demand curve between p_0 and p_1:

$$\Delta \text{CS} = \int_{p_0}^{p_1} D(p) \; \mathrm{d}p. \tag{1}$$

Typically, in empirical applications, not all of $D(\cdot)$ is observed: data about the quantities sold are available at the realized prices p_0 and p_1, but not at any other price. In this basic model, we impose a standard regularity condition on $D(\cdot)$ and assume that $D(p_0)$ and $D(p_1)$ can be perfectly observed by the researcher:

(**A2**) $D(\cdot)$ is absolutely continuous and passes through the points (p_0, q_0) and (p_1, q_1).

The first part of (**A2**)—that $D(\cdot)$ is absolutely continuous—is relatively innocuous: absolutely continuous functions are dense (in the L^1 norm) in the space of integrable functions defined on the interval $[p_0, p_1]$; thus ΔCS under *any* demand curve is well-approximated by ΔCS under an absolutely continuous demand curve. Furthermore, absolute continuity guarantees that $D(\cdot)$ is differentiable almost everywhere, so that its price elasticity is well-defined almost everywhere.

The second part of (**A2**)—that $D(\cdot)$ passes through both (p_0, q_0) and (p_1, q_1)—is, in our view, the simplest assumption that captures the data limitations faced by empirical researchers and

policymakers. Its sole purpose is to serve as a benchmark for our analysis. Later, we relax this assumption in various directions by studying settings with fewer or more observed points, as well as settings with error so that $D(\cdot)$ does not pass through the observed points exactly.

2.2 Welfare bounds with elasticity restrictions

As we argued in the introduction, imposing (**A1**) and (**A2**) as the only assumptions on $D(\cdot)$ allows the price elasticity of demand to be $-\infty$ or 0 almost everywhere, which seems unrealistic and overly conservative. In many cases, such extreme values may even be ruled out *a priori* via elasticity estimates from field experiments or quasi-experimental designs.

In particular, elasticity estimates from field experiments or quasi-experimental designs allow the researcher to restrict the price elasticity of demand to a range of values of interest. We consider the simplest such setting by imposing the following elasticity assumption (abbreviated by "EA"):

(**EA**) The price elasticity of demand $\varepsilon(\cdot)$ between p_0 and p_1 lies between $\underline{\varepsilon}$ and $\overline{\varepsilon}$, so that

$$\underline{\varepsilon} \leq \varepsilon(p) := \frac{pD'(p)}{D(p)} \leq \overline{\varepsilon} \quad \text{for any } p \in [p_0, p_1].$$

To avoid trivialities, we further assume that the average elasticity between p_0 and p_1 is bounded by $\underline{\varepsilon}$ and $\overline{\varepsilon}$, which is a necessary and sufficient condition for the existence of a demand curve that satisfies (**A1**), (**A2**), and (**EA**):

$$\underline{\varepsilon} \leq \frac{\log(q_1/q_0)}{\log(p_1/p_0)} \leq \overline{\varepsilon}.$$

The extreme elasticities $\underline{\varepsilon}$ and $\overline{\varepsilon}$ in (**EA**) flexibly parameterize how much *a priori* knowledge about $D(\cdot)$ is assumed. As a special case, the researcher might not have any *a priori* knowledge: $\underline{\varepsilon} = -\infty$ and $\overline{\varepsilon} = 0$. As we show in Section 4, comparative statics with respect to $\underline{\varepsilon}$ and $\overline{\varepsilon}$ allow for a precise quantification of how robust welfare estimates are to different elasticity assumptions.

To determine appropriate values of $\underline{\varepsilon}$ and $\overline{\varepsilon}$, the researcher might appeal to elasticity estimates from field experiments and quasi-experimental designs. In these environments, the price elasticity of demand can typically be estimated at either p_0 or p_1—or both. The variability of elasticity between p_0 and p_1 can then be parameterized via $\underline{\varepsilon}$ and $\overline{\varepsilon}$: for example, the researcher might require that $\underline{\varepsilon}$ and $\overline{\varepsilon}$ are within 10% of the elasticities $\varepsilon(p_0)$ and $\varepsilon(p_1)$. We illustrate this approach in Section 4 with a variety of empirical applications.

Alternatively, the researcher might also obtain priors on $\underline{\varepsilon}$ and $\overline{\varepsilon}$ from institutional knowledge and surveys of related studies. For example, Andreyeva, Long, and Brownell (2010) summarize

elasticities for food and beverage in the U.S. from 160 empirical studies and determine that all lie between -3.18 and -0.01; they also provide narrower ranges for each distinct food category. A researcher studying the welfare impact of a sugar tax might thus compute conservative bounds by considering all demand functions that satisfy $-3.18 \leq \varepsilon(p) \leq -0.01$.

Unlike functional form assumptions, (**A1**), (**A2**), and (**EA**) do not jointly determine a unique demand curve in general. By contrast, (**A1**), (**A2**), and (**EA**) are typically satisfied by a *family* of demand curves, which we denote by $\mathcal{D}$:

$$\mathcal{D} := \{D : [p_0, p_1] \to \mathbb{R} \text{ is decreasing and satisfies } (\mathbf{A2}) \text{ and } (\mathbf{EA})\}.$$

For every demand curve in $\mathcal{D}$, the formula (1) can be applied to give the corresponding change in consumer surplus *for that demand curve*. Our goal is to find the range of all possible changes in consumer surplus. We denote the largest and smallest possible changes in consumer surplus as follows:

$$\begin{cases} \overline{\Delta \mathrm{CS}} := \max_{D \in \mathcal{D}} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p, \\ \underline{\Delta \mathrm{CS}} := \min_{D \in \mathcal{D}} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p. \end{cases} \tag{2}$$

Proposition 1. Under (**A1**), (**A2**), and (**EA**), the identified set $[\underline{\Delta \mathrm{CS}}, \overline{\Delta \mathrm{CS}}]$ is sharp: for any possible change in consumer surplus that lies within this set, there exists a demand curve $D(\cdot) \in \mathcal{D}$ that generates it.

Proof. Because $\mathcal{D}$ is convex and $\Delta \mathrm{CS}$ is a linear map of $D(\cdot)$ according to the formula (1), the set of possible changes in consumer surplus is a convex subset of $\mathbb{R}$. Therefore, it is an interval and can be equivalently characterized by its endpoints. $\qquad \square$

In general, the family of demand curves $\mathcal{D}$ is very large. To be concrete, we begin by considering a demand curve that we know with certainty to be in $\mathcal{D}$: the constant elasticity of substitution (CES) demand curve that connects the points (p_0, q_0) and (p_1, q_1),

$$D_{\mathrm{CES}}(p) := q_0 \cdot \left(\frac{p}{p_0} \right)^{\frac{\log(q_1/q_0)}{\log(p_1/p_0)}}.$$

The elasticity of $D_{\mathrm{CES}}(\cdot)$ at any price is equal to the average elasticity between p_0 and p_1: $\log(q_1/q_0)/\log(p_1/p_0)$. We call $D_{\mathrm{CES}}(\cdot)$ a *1-piece CES interpolation* between (p_0, q_0) and (p_1, q_1). Under a 1-piece CES interpolation, the change in consumer surplus can be computed exactly by

using the formula (1):

$$\Delta\mathrm{CS}_{\mathrm{CES}} = \frac{p_0 q_0}{1 + \frac{\log(q_1/q_0)}{\log(p_1/p_0)}} \cdot \left[\left(\frac{p_1}{p_0}\right)^{1 + \frac{\log(q_1/q_0)}{\log(p_1/p_0)}} - 1 \right].$$

Is there a demand curve in $\mathcal{D}$ that yields a higher change in consumer surplus than a 1-piece CES interpolation? The answer is yes: rather than interpolate between (p_0, q_0) and (p_1, q_1) with a CES demand curve, one can construct an auxiliary point, say (p^*, q^*), that lies slightly above $D_{\mathrm{CES}}(\cdot)$, as shown in Figure 2. One can then interpolate between the three points with CES demand curves (each of which has a different elasticity). We call any such demand curve a *2-piece CES interpolation* between (p_0, q_0) and (p_1, q_1).
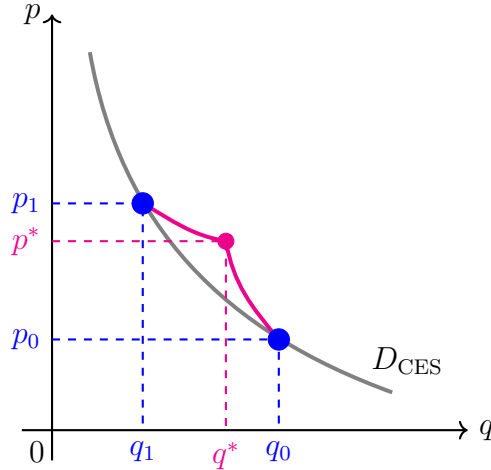
Figure 2: *Example of a demand curve (in red) that yields a larger ΔCS than $D_{\mathrm{CES}}(\cdot)$.*

Unlike a 1-piece CES interpolation, 2-piece CES interpolations between (p_0, q_0) and (p_1, q_1) are not unique. Thus a natural question is which 2-piece CES interpolation, call it $D^*(\cdot)$, maximizes ΔCS. A relatively simple solution is to interpolate between (p_0, q_0) and (p^*, q^*) with the most inelastic demand curve possible (*i.e.*, with elasticity $\bar{\varepsilon}$), and then interpolate between (p^*, q^*) and (p_1, q_1) with the most elastic demand curve possible (*i.e.*, with elasticity $\underline{\varepsilon}$). These conditions on elasticity uniquely determine the auxiliary point (p^*, q^*) for $D^*(\cdot)$:

$$\bar{\varepsilon} = \frac{\log(q^*/q_0)}{\log(p^*/p_0)} \quad \text{and} \quad \underline{\varepsilon} = \frac{\log(q_1/q^*)}{\log(p_1/p^*)}.$$

11

A symmetric argument implies that a 2-piece CES interpolation can also yield a smaller ΔCS than $D_{\mathrm{CES}}(\cdot)$. The 2-piece CES interpolation that minimizes ΔCS, call it $D_*(\cdot)$, interpolates between (p_0, q_0) and an auxiliary point (p_*, q_*) with the most elastic demand curve possible, and then interpolates between (p_*, q_*) and (p_1, q_1) with the most inelastic demand curve possible:

$$\underline{\varepsilon} = \frac{\log(q_*/q_0)}{\log(p_*/p_0)} \quad \text{and} \quad \overline{\varepsilon} = \frac{\log(q_1/q_*)}{\log(p_1/p_*)}.$$

It might be tempting to extend this argument in a variety of ways. Perhaps 3-piece CES interpolations yield an even wider range of possible changes in consumer surplus? How about linear interpolations rather than CES interpolations? Our first main result indicates that these attempts will ultimately fail.

Theorem 1. Under (**A1**), (**A2**), and (**EA**), the largest and smallest possible changes in consumer surplus between p_0 and p_1, $\overline{\Delta\mathrm{CS}}$ and $\underline{\Delta\mathrm{CS}}$, are attained by 2-piece CES interpolations:

$$\int_{p_0}^{p_1} D_*(p)\ \mathrm{d}p = \underline{\Delta\mathrm{CS}} \leq \int_{p_0}^{p_1} D(p)\ \mathrm{d}p \leq \overline{\Delta\mathrm{CS}} = \int_{p_0}^{p_1} D^*(p)\ \mathrm{d}p \quad \text{for any } D \in \mathcal{D}.$$

Here, $D_*(\cdot)$ and $D^*(\cdot)$ are defined by

$$D_*(p) := \begin{cases} q_1 \left(\dfrac{p_1}{p}\right)^{\overline{\varepsilon}} & \text{if } p > \left(\dfrac{q_0 p_1^{\overline{\varepsilon}}}{q_1 p_0^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon}-\underline{\varepsilon}}}, \\[12pt] q_0 \left(\dfrac{p}{p_0}\right)^{\underline{\varepsilon}} & \text{if } p \leq \left(\dfrac{q_0 p_1^{\overline{\varepsilon}}}{q_1 p_0^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon}-\underline{\varepsilon}}}, \end{cases} \quad \text{and} \quad D^*(p) := \begin{cases} q_1 \left(\dfrac{p_1}{p}\right)^{\underline{\varepsilon}} & \text{if } p > \left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon}-\underline{\varepsilon}}}, \\[12pt] q_0 \left(\dfrac{p}{p_0}\right)^{\overline{\varepsilon}} & \text{if } p \leq \left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon}-\underline{\varepsilon}}}. \end{cases}$$

Closed-form expressions for $\underline{\Delta\mathrm{CS}}$ and $\overline{\Delta\mathrm{CS}}$ are provided in Appendix A.

We offer two proofs of Theorem 1. The first, which we present below, is simple and intuitive due to its geometric nature. Its generalizability, however, is rather limited: in more complicated environments that we consider later, this simple geometric approach no longer applies. The second, which we present in Appendix B, is more technical and relies on a fortuitous connection between our problem (2) and Bayesian persuasion problems that have been considered by the theoretical literature stemming from Kamenica and Gentzkow (2011). While our second approach is less straightforward, it generalizes more easily to later environments.

We now explain the geometric proof of Theorem 1 with the help of Figure 3. We begin with a change of variables: rather than plot prices against quantities, we depict a demand curve by
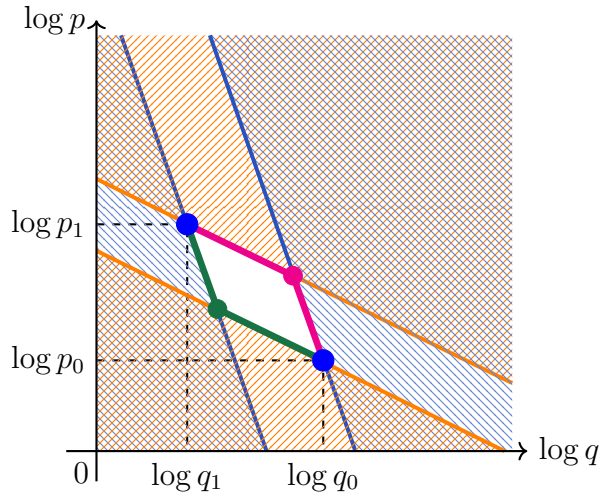
Figure 3: Sketch of the proof of Theorem 1.

plotting *log-prices* against *log-quantities*. The key insight is that, because the logarithm is a monotone transformation, this change of variables does not qualitatively alter our problem: as before, our goal is to find decreasing curves that pass through two points, $(\log p_0, \log q_0)$ and $(\log p_1, \log q_1)$, that respectively maximize and minimize the areas under the curves that are bounded between $\log p_0$ and $\log p_1$.

Although it does not qualitatively alter the problem, this change of variables enables a more natural interpretation of **(EA)**. Notice that demand curves with constant elasticity correspond to *linear* curves when we plot log-prices against log-quantities. To rule out demand curves with elasticities higher than $\bar{\varepsilon}$, we draw two (blue) straight lines such that one passes through $(\log p_0, \log q_0)$ and the other, $(\log p_1, \log q_1)$; each of these lines has a gradient of $\bar{\varepsilon}$. Any demand curve in $\mathcal{D}$ must therefore correspond to a curve that lies between these two lines—and not in the (blue) shaded regions. Similarly, to rule out demand curves with elasticities lower than $\underline{\varepsilon}$, we draw two (orange) straight lines with gradients equal to $\underline{\varepsilon}$. Any demand curve in $\mathcal{D}$ must also correspond to a curve that lies between these two lines—and not in the (orange) shaded regions.

The next step is to find the curve within the unshaded parallelogram-shaped region that maximizes the area under it between $\log p_0$ and $\log p_1$ (as ΔCS must vary monotonically with this area). This curve can be read directly off the diagram: it must be the top boundary of the parallelogram (depicted by the red curve). Likewise, the bottom boundary of the parallelogram (depicted by the green curve) minimizes the area under it between $\log p_0$ and $\log p_1$.

It follows that the demand curves corresponding to each of these curves must respectively maximize and minimize the change in consumer surplus. Reversing the change of variables, the

top (red) curve corresponds to a 2-piece CES interpolation that has elasticity $\bar{\bar{\varepsilon}}$ between (p_0, q_0) and an auxiliary point, and elasticity $\underline{\varepsilon}$ between the auxiliary point and (p_1, q_1), which is the demand curve $D^*(\cdot)$. Likewise, the bottom (green) curve corresponds to the demand curve $D_*(\cdot)$.

One minor technicality is that, while any demand curve in $\mathcal{D}$ must correspond to a curve within the unshaded parallelogram-shaped region, the converse is not true: not any curve within this region can be mapped back into a demand curve in $\mathcal{D}$. Nevertheless, it is readily verified that the resulting demand curves, $D^*(\cdot)$ and $D_*(\cdot)$, are both in $\mathcal{D}$.

Although minor, this technicality highlights the somewhat "coincidental" simplicity of this proof. The heart of the geometric argument relies on finding the constraints implied by (**EA**) that are not only necessary, but also sufficient *a posteriori*. In more complicated environments, the binding constraints are not as easily determined, nor do they necessary take on such a simple form. In turn, this explains why the geometric argument fails to generalize to later environments, necessitating the information design approach that we undertake in Appendix B.

2.3   Welfare bounds with curvature restrictions

We now turn to a different type of assumption than (**EA**), namely restrictions on the curvature of the demand curve, and show that analogs of Proposition 1 and Theorem 1 hold in this setting.

Restrictions on the curvature of demand have a long history of precedence in economics. For instance, Marshall (1890) went so far as to *define* a demand curve as a decreasing function whose elasticity also decreases with price, while Robinson (1933) suggested that demand curves, ought to be convex lest the monopoly output rises when price discrimination causes prices to rise. These intuitions underlie the standard textbook depiction of a demand curve as a convex function.

Such restrictions are imposed for a variety of reasons, arguably the most important of which is to ensure that comparative statics predicted by models are consistent with observed data. For example, Marshall's assumption (now more commonly known as Marshall second law) was maintained "without apology" by Krugman (1979) so that his model would produce "reasonable results." Melitz (2018) argues that Marshall second law—which he also imposes in his model—is "equivalent to the property that more productive firms (or alternatively lower cost) set higher markups," and that violations would "directly contradict the [empirical] evidence on markups and pass-through."

Different literatures in economics employ a variety of assumptions on the curvature of demand that capture other intuitions pertaining to their fields of interest. To be comprehensive, we consider a range of assumptions that are considered standard in different fields. Each curvature assumption
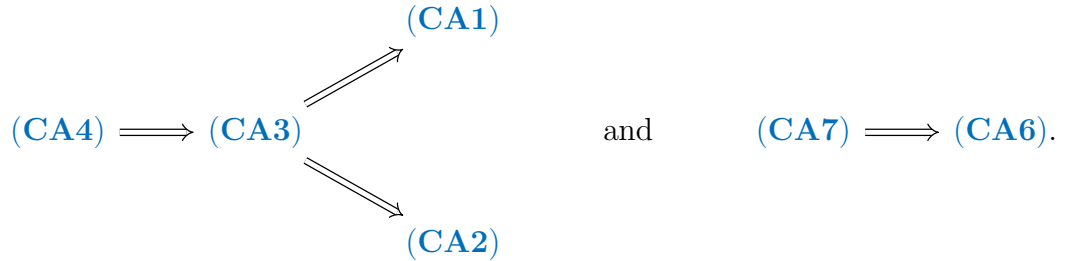
(abbreviated by "CA") restricts ΔCS in a different way. We detail these assumptions below and provide some examples of how they are invoked in different fields.

(**CA1**) **Marshall's second law.** Demand is said to satisfy Marshall's second law if its price elasticity $\varepsilon(p) = pD'(p)/D(p)$ is decreasing in p. This was introduced by Marshall (1890) and is widely used in international trade, macroeconomics, and microeconomics, including by Krugman (1979), Bishop (1968), Johnson (2017), and Melitz (2018), who also provides some empirical justification for this assumption in the context of trade models.

(**CA2**) **Decreasing marginal revenue.** Let $P(q) := D^{-1}(q)$ denote the inverse demand curve. Demand exhibits decreasing marginal revenue if marginal revenue $\mathrm{MR}(q) := P(q) + qP'(q)$ is decreasing in q. This assumption is standard in microeconomics (see Robinson, 1933, for example) and ensures that a profit-maximizing price exists for a monopolist who faces a convex cost function.

(**CA3**) **Log-concave demand.** Demand is log-concave if $D'(p)/D(p)$ is decreasing in p. The comprehensive surveys of Bagnoli and Bergstrom (2005) and An (1998) demonstrate that many common demand curves are log-concave. Log-concave demand also has a simple economic interpretation, as Amir, Maret, and Troege (2004) show: the pass-through rate of a change in a monopolist's marginal cost is less than one if and only if demand is log-concave (see also Weyl and Fabinger, 2013). It is also well-known that log-concavity is a sufficient condition for a unique equilibrium to exist in common models of Cournot competition (Dixit, 1986) and differentiated products Bertrand competition (Caplin and Nalebuff, 1991a).

(**CA4**) **Concave demand.** Demand is concave if $D'(p)$ is decreasing in p. Robinson (1933) shows that concave demand has a simple economic interpretation: total output increases when monopolistic price discrimination causes prices to rise in markets with concave demands (see also Malueg, 1994 and Aguirre, Cowan, and Vickers, 2010).

(**CA5**) ρ**-concave demand.** For a given real number ρ, demand is ρ-concave if $D'(p)\left[D(p)\right]^{\rho-1}$ is decreasing in p. Based on the work of Prékopa (1973), this assumption was introduced to the economics literature by Caplin and Nalebuff (1991a,b) as a generalization of log-concavity ($\rho = 0$) and concavity ($\rho = 1$). Different values of ρ parametrize the restrictiveness of this assumption: a ρ'-concave demand curve is ρ''-concave for any $\rho'' < \rho'$.

(**CA6**) **Convex demand.** Demand is convex if $D'(p)$ is increasing in p. Similar to concave demand (**CA4**), Robinson (1933) shows that total output increases when monopolistic price discrimination causes prices to fall in markets with convex demands.

(**CA7**) **Log-convex demand.** Demand is log-convex if $D'(p)/D(p)$ is increasing in p. Similar to log-concave demand (**CA3**), Amir et al. (2004) show that the pass-through rate of a change in a monopolist's marginal cost is more than one if and only if demand is log-convex.

(**CA8**) **ρ-convex demand.** For a given real number ρ, demand is ρ-convex if $D'(p)\left[D(p)\right]^{\rho-1}$ is increasing in p. Similar to ρ-concave demand (**CA5**), ρ-convexity generalizes convexity ($\rho = 1$) and log-convexity ($\rho = 0$); a ρ'-convex demand curve is ρ''-convex for any $\rho'' > \rho'$.

These assumptions can be divided into two categories: *concave-like* assumptions (**CA1**)–(**CA5**) and *convex-like* assumptions (**CA6**)–(**CA8**). Concave-like and convex-like assumptions bound the curvature of the demand curve from above and from below, respectively.

These assumptions are not mutually disjoint. For example, it is well-known that concave demand curves are log-concave, and that log-convex demand curves are convex. In fact:

$$(\mathbf{CA4}) \implies (\mathbf{CA3}) \begin{array}{c} \nearrow (\mathbf{CA1}) \\ \\ \searrow (\mathbf{CA2}) \end{array} \qquad \text{and} \qquad (\mathbf{CA7}) \implies (\mathbf{CA6}).$$

For reference, these relationships are proven in Appendix C, where we also provide examples of common demand curves that satisfy each assumption.

Analogous to the demand family $\mathcal{D}$, we define the following families of demand curves $\mathcal{D}_i$ that correspond to the different curvature assumptions for each $i = 1, \ldots, 8$:

$$\mathcal{D}_i := \{D : [p_0, p_1] \to \mathbb{R} \text{ is decreasing and satisfies } (\mathbf{A2}) \text{ and } (\mathbf{CA}i)\}.$$

As before, we find the largest and smallest possible changes in consumer surplus:

$$\begin{cases} \overline{\Delta\mathrm{CS}}_i := \max_{D \in \mathcal{D}_i} \int_{p_0}^{p_1} D(p)\,\mathrm{d}p, \\ \underline{\Delta\mathrm{CS}}_i := \min_{D \in \mathcal{D}_i} \int_{p_0}^{p_1} D(p)\,\mathrm{d}p. \end{cases}$$

16

We can now state our second main result:

Theorem 2. Under (**A1**) and (**A2**), the following hold:

(a) For concave-like assumptions (**CA1**)–(**CA5**), $\overline{\Delta \mathrm{CS}}_i$ is attained by a 2-piece interpolation and $\underline{\Delta \mathrm{CS}}_i$ is attained by a 1-piece interpolation.

(b) For convex-like assumptions (**CA6**)–(**CA8**), $\overline{\Delta \mathrm{CS}}_i$ is attained by a 1-piece interpolation and $\underline{\Delta \mathrm{CS}}_i$ is attained by a 2-piece interpolation.

Specifically, the following bounds on changes in consumer surplus hold:

(i) Under (**A1**), (**A2**), and (**CA1**), $\overline{\Delta \mathrm{CS}}_1$ is attained by a 2-piece CES interpolation and $\underline{\Delta \mathrm{CS}}_1$ is attained by a 1-piece CES interpolation.

(ii) Under (**A1**), (**A2**), and (**CA2**), $\overline{\Delta \mathrm{CS}}_2$ is attained by a 2-piece constant marginal revenue interpolation and $\underline{\Delta \mathrm{CS}}_2$ is attained by a 1-piece constant marginal revenue interpolation.

(iii) Under (**A1**), (**A2**), and (**CA3**), $\overline{\Delta \mathrm{CS}}_3$ is attained by a 2-piece exponential interpolation and $\underline{\Delta \mathrm{CS}}_3$ is attained by a 1-piece exponential interpolation.

(iv) Under (**A1**), (**A2**), and (**CA4**), $\overline{\Delta \mathrm{CS}}_4$ is attained by a 2-piece linear interpolation and $\underline{\Delta \mathrm{CS}}_4$ is attained by a 1-piece linear interpolation.

(v) Under (**A1**), (**A2**), and (**CA5**), $\overline{\Delta \mathrm{CS}}_5$ is attained by a 2-piece ρ-linear interpolation and $\underline{\Delta \mathrm{CS}}_5$ is attained by a 1-piece ρ-linear interpolation.

(vi) Under (**A1**), (**A2**), and (**CA6**), $\overline{\Delta \mathrm{CS}}_6$ is attained by a 1-piece linear interpolation and $\underline{\Delta \mathrm{CS}}_6$ is attained by a 2-piece linear interpolation.

(vii) Under (**A1**), (**A2**), and (**CA7**), $\overline{\Delta \mathrm{CS}}_7$ is attained by a 1-piece exponential interpolation and $\underline{\Delta \mathrm{CS}}_7$ is attained by a 2-piece exponential interpolation.

(viii) Under (**A1**), (**A2**), and (**CA8**), $\overline{\Delta \mathrm{CS}}_8$ is attained by a 1-piece ρ-linear interpolation and $\underline{\Delta \mathrm{CS}}_8$ is attained by a 2-piece ρ-linear interpolation.

Closed-form expressions for $\underline{\Delta \mathrm{CS}}$ and $\overline{\Delta \mathrm{CS}}$ are provided in Appendix A.

As the proof is an extension of the geometric proof of Theorem 1, we defer it to Appendix D. Instead, we emphasize a key implication of Theorem 2 for empirical applications here: common demand curves often attain extremal values of ΔCS. For instance, Theorem 2 implies that a CES demand curve achieves the smallest possible value of ΔCS among all demand curves in $\mathcal{D}_1$ that satisfy Marshall's second law. Similarly, a linear demand curve achieves the largest possible value of ΔCS among all convex demand curves in $\mathcal{D}_6$. Thus, Theorem 2 provides a formal sense in which interpolations commonly used by practitioners in different fields—such as the CES interpolation in international trade and the linear interpolation in applied microeconomics—are extremal.

Moreover, although Theorem 2 states bounds separately for each $\mathcal{D}_i$, curvature assumptions can also be combined: for instance, a convex demand curve may also satisfy Marshall's second law. Such a demand curve belongs to both $\mathcal{D}_1$ and $\mathcal{D}_6$. In that case, $\overline{\Delta\text{CS}}$ will be attained by a 1-piece linear demand curve, whereas $\underline{\Delta\text{CS}}$ will be attained by a 1-piece CES demand curve.

A natural question is whether an analogue of Proposition 1 holds: is it true that the identified sets remain sharp under (**CA1**)–(**CA8**) rather than (**EA**)? We show that the answer is yes:

Proposition 2. Under (**A1**), (**A2**), and (**CA**i), the identified set $[\underline{\Delta\text{CS}}_i, \overline{\Delta\text{CS}}_i]$ is sharp for each $i = 1, \ldots, 8$: for any possible change in consumer surplus that lies within this set, there must be a demand curve $D(\cdot) \in \mathcal{D}_i$ that generates it.

The proof of Proposition 2 is more complicated than the proof of Proposition 1. The key challenge is that, unlike $\mathcal{D}$, many of the families $\mathcal{D}_i$ are no longer convex. In Appendix D, we show how our information design approach yields a simple proof of Proposition 2. The key step in our proof expresses the set of all possible changes in consumer surplus as the image of a convex set under a continuous real-valued function. We then establish a one-to-one map between this convex set and the demand family $\mathcal{D}_i$. This guarantees that the set of all possible changes in consumer surplus is convex, and hence exactly equal to the interval $[\underline{\Delta\text{CS}}_i, \overline{\Delta\text{CS}}_i]$.

2.4   Discussion

Before moving on to extensions, we comment on the modeling choices and assumptions imposed in the basic version of the model, and how they affect the interpretation of our results.

First, as Theorem 1 demonstrates, the extremality of 2-piece CES interpolations arises from our restriction (**EA**) on the range of admissible elasticities. As we argued earlier, this restriction seems natural because empirical researchers and policymakers already use ranges of elasticities and treatment effect estimates to reason about policy and welfare evaluation. However, in some

applications, it might be more suitable to instead assume that the range of revenues is known. For example, a monopolist who wishes to evaluate the effect of a price increase on consumers might have more experience with likely revenues, rather than elasticities, over the interval of prices in question. The analog of Theorem 1 (which can be shown by appropriately modifying the geometric proof) then states that the largest and smallest possible changes in consumer surplus are attained by 2-piece constant marginal revenue interpolations.

Second, instead of replacing (**EA**) with a curvature assumption (**CA**$_i$) as in Theorem 2, we can combine (**CA**$_i$) with information about elasticities at p_0 and p_1 that the researcher might have obtained from estimated treatment effects. In Appendix A, we show how an extension of Theorem 2 yields simple bounds for the case when the researcher observes $\varepsilon(p_0)$ and $\varepsilon(p_1)$, and when the demand curve satisfies (**A1**), (**A2**), and (**CA**$_i$).

Third, in many empirical applications, the researcher is interested in counterfactual policies, rather than retrospective ones—in which case only one price-quantity pair is observed, leaving the quantity at the counterfactual price unknown. Extending Theorem 1 to this case requires "extrapolating" from fewer observations to the case of two observations. Before applying Theorem 1, we first establish bounds on what the unobserved point on the demand curve could possibly be. We then apply Theorem 1 to every possible pair of points and find the extreme points of the resulting set of possible ranges for ΔCS. We show how this can be done in Section 3.1 using the same analytical tools introduced above, and apply this to an empirical setting based on the work of Hahn and Metcalfe (2021) in Section 4.1.

Fourth, in other empirical applications, more than two points on the same demand curve are observed. For instance, a tax might impact different markets or a series of price shocks may be introduced sequentially over time. While our basic approach can be applied directly to each market or price shock separately, a researcher may wish to refine the set of feasible demand curves by imposing consistency with all of the observed data points. With infinite variation along the demand curve, the true demand curve $D(\cdot)$ is non-parametrically identified and may be directly recovered. In that limit, our bounds converge to the actual change in consumer surplus. Section 3.2 analyzes the intermediate case of arbitrarily (but finitely) many observations.

Fifth, an important assumption that we have made throughout this section is that the points on the demand curve are perfectly observed without sampling error or structural noise. While this simplifies the exposition, we show in Section 3.3 that our analysis extends to these cases by deriving confidence intervals for each of our bounds.

Sixth, while our basic model focuses on bounding the change in consumer surplus, our approach can be extended to other welfare measures. In Section 3.4, we discuss how deadweight loss can be

bounded. By relaxing **(A1)**, we allow income effects and show how our approach extends to the compensating and equivalent variations. We further incorporate equity into our welfare measures through the use of social welfare weights (Harberger, 1978; Saez and Stantcheva, 2016), and discuss how our results extend symmetrically to supply-side measures of welfare such as producer surplus.

Seventh, and finally, we focus on the case where there is a single good in the market. Our results extend immediately to markets where the good in consideration is independent of all other goods. By continuity, this argument also applies when the elasticity of substitution between this good and other goods in the market is small.[4]

3   Extensions

Motivated by the discussion at the end of the last section, we now consider different variations of the basic model. As we show below, the analytical tools introduced in the basic model extend naturally when more complexity is allowed for. We illustrate this by imposing the elasticity assumption **(EA)** throughout this section; similar results can be obtained under the curvature assumptions **(CA1)**–**(CA8)**.

3.1   Counterfactual extrapolation from fewer points

The basic model assumes that two points on the demand curve are observed: (p_0, q_0) and (p_1, q_1). However, in counterfactual exercises such as our application in Section 4.1, the quantity that would be demanded at p_1 is not known. Instead of **(A2)**, these applications call for a weaker assumption:

(A3) $D(\cdot)$ is absolutely continuous and passes through the point (p_0, q_0).

Let $\mathcal{D}'$ denote the family of demand curves that satisfy **(A1)**, **(A3)**, and **(EA)**:

$$\mathcal{D}' := \{D : [p_0, p_1] \to \mathbb{R} \text{ is decreasing and satisfies } \textbf{(A3)} \text{ and } \textbf{(EA)}\}.$$

As before, we find the largest and smallest possible changes in consumer surplus within $\mathcal{D}'$:

$$\begin{cases} \overline{\Delta \mathrm{CS}}' := \max_{D \in \mathcal{D}'} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p, \\ \underline{\Delta \mathrm{CS}}' := \min_{D \in \mathcal{D}'} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p. \end{cases}$$

---

[4] In ongoing work, we study how our results extend to more general markets with multiple goods.
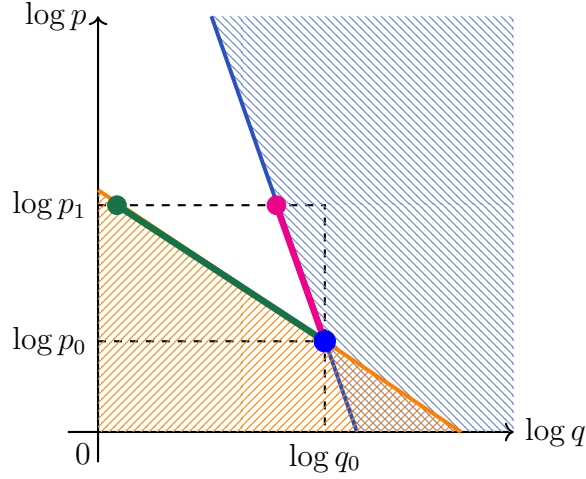
Figure 4: Illustration of bounds when only (p_0, q_0) and p_1 are observed.

A general procedure for finding $\overline{\Delta\mathrm{CS}}'$ and $\underline{\Delta\mathrm{CS}}'$ is to decompose the problem into three steps. First, we characterize the set of possible values of q_1 that are consistent with (**A3**) and (**EA**). For each possible value of q_1, we then apply Theorem 1 to compute $\overline{\Delta\mathrm{CS}}$ and $\underline{\Delta\mathrm{CS}}$ for that q_1. Finally, $\overline{\Delta\mathrm{CS}}'$ is equal to the maximal $\overline{\Delta\mathrm{CS}}$ over all possible q_1, whereas $\underline{\Delta\mathrm{CS}}'$ is equal to the minimal $\underline{\Delta\mathrm{CS}}$ over the same set. Because $\overline{\Delta\mathrm{CS}}'$ and $\underline{\Delta\mathrm{CS}}'$ obtain at some (generally different) values of q_1, Theorem 1 implies that these bounds are attained by 2-piece CES interpolations.

Actually, more can be said about $\overline{\Delta\mathrm{CS}}'$ and $\underline{\Delta\mathrm{CS}}'$ than this procedure might suggest by using our earlier geometric argument, illustrated in Figure 4. The largest possible value of $\log q_1$ that is consistent with (**A3**) and (**EA**) can be found by drawing the (blue) straight line with gradient $\overline{\varepsilon}$ that passes through the point $(\log p_0, \log q_0)$, and then finding the (red) point on the line at $\log p_1$. It is clear that this value of q_1 must also yield the maximal $\overline{\Delta\mathrm{CS}}$; hence $\overline{\Delta\mathrm{CS}}'$ must be attained by a *1-piece* CES interpolation (corresponding to the red curve). A symmetric argument shows that $\underline{\Delta\mathrm{CS}}'$ must also be attained by a 1-piece CES interpolation (corresponding to the green curve).

Theorem 3. Under (**A1**), (**A3**), and (**EA**), the largest and smallest changes in consumer surplus, $\overline{\Delta\mathrm{CS}}'$ and $\underline{\Delta\mathrm{CS}}'$, are attained by 1-piece CES interpolations (with elasticities $\overline{\varepsilon}$ and $\underline{\varepsilon}$, respectively).

From this decomposition procedure, we can deduce that Proposition 1 must hold in this setting: the identified set $[\underline{\Delta\mathrm{CS}}', \overline{\Delta\mathrm{CS}}']$ is sharp. Moreover, this decomposition procedure allows us to find bounds for more complex models to which the geometric argument does not fully extend. In particular, the geometric argument works for Theorem 3 only because our welfare measure $\Delta\mathrm{CS}$ is monotone in q_1. As we show in Section 4.1, when the researcher is interested in welfare measures

that take into account costs (that depend on q_1), the geometric argument is insufficient and this decomposition procedure is required.

3.2 Interpolating with more observations

In many empirical applications, more than two points on the same demand curve are observed. These cases require a direct generalization of the basic model to an arbitrary (finite) number of observations. To model this, we replace (**A2**) with:

(**A4**) $D(\cdot)$ is absolutely continuous and passes through the points $(p_0, q_0), \ldots, (p_{n-1}, q_{n-1})$.

Thus (**A2**) is simply a special case of (**A4**) by setting $n = 2$. As such, our first result for this subsection directly generalizes Theorem 1:

Theorem 4. Under (**A1**), (**A4**), and (**EA**), the largest and smallest possible changes in consumer surplus are attained by $(2n - 2)$-piece CES interpolations.

Theorem 4 follows by applying Theorem 1 between every two adjacent points; Proposition 1 implies that the resulting identified set is sharp. Figure 5 illustrates the argument for the case of $n = 3$ observations, where both the largest (in red) and smallest (in green) possible changes in consumer surplus are depicted.

A more challenging question is how our bounds change when we impose curvature assumptions in a model with more than two observations. For concreteness, we consider the case of Marshall's second law: what are the largest and smallest possible changes in consumer surplus under (**A1**), (**A4**), (**EA**), and (**CA1**)? Clearly, the answer must be different than Theorem 4: even in the case of $n = 3$ observations, both of the extremal demand curves (corresponding to the red and green curves) do not satisfy Marshall's second law.

An intuitive solution to this problem is to "iron" the extremal demand curves of Theorem 4 whenever Marshall's second law is violated—that is, whenever elasticity increases, rather than decreases, with price for two adjacent pieces in the CES interpolation. Ironing combines these two adjacent pieces and replaces them with a single, larger piece. A simple count of how many times Marshall's second law is violated yields the form of the extremal demand curves to this problem. By Proposition 2, the resulting identified set must also be sharp.

Theorem 5. Under (**A1**), (**A4**), (**EA**), and (**CA1**), the largest possible change in consumer surplus is attained by an n-piece CES interpolation, and the smallest possible change in consumer surplus is attained by an $(n - 1)$-piece CES interpolation.
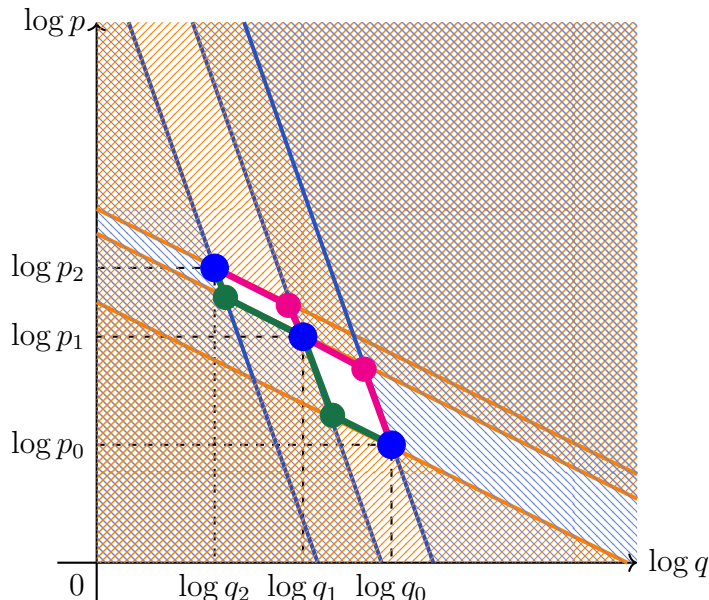
Figure 5: Illustration of bounds with $n = 3$ observations.

3.3   Sampling error and inference

In many empirical applications, observations of the quantities demanded involve sampling error: rather than precise measurements of q_0 and q_1 for the entire market, the researcher observes only a noisy sample thereof (*e.g.*, within the treatment group). When this is the case, our bounds need to account for an additional source of uncertainty: not only uncertainty about the shape of the demand curve, but also uncertainty about the points that the demand curve passes through.

In this subsection, we show that our bounds can be extended to account for uncertainty due to sampling error. We consider an extension of our basic model in which q_0 and q_1 are observed as averages of noisy observations across individuals—as would be given by the treatment effect estimate in a randomized control trial. Formally, there are N_0 individuals who are offered the good at a price of p_0 ("control") and N_1 individuals who are offered the good at a price of p_1 ("treatment"). Individuals are assumed to share the same underlying demand curve $D(\cdot)$, but the quantity demanded by each individual i is also affected by an idiosyncratic shock e_{it}:

$$q_{it} = D(p_t) + e_{it} \quad \text{for } t = 0, 1.$$

We make the standard assumption that $e_{1t}, \ldots, e_{N_t t}$ have zero mean and σ^2 variance, and are independently and identically distributed across individuals.

For ease of exposition, we assume that the control group is large but the treatment group is small, so that there is meaningful sampling error only for q_1, and not q_0. Our approach extends straightforwardly when both q_0 and q_1 have sampling error. We replace (**A2**) with the following:

(**A5**) $D(\cdot)$ is absolutely continuous, passes through the point (p_0, q_0), and satisfies

$$q_1 = D(p_1) + e,$$

where e has an asymptotic normal distribution, $\mathcal{N}(0, \sigma^2/N_1)$.

Clearly, (**A2**) is obtained as a special case of (**A5**) in the limit where the treatment group is also large so that q_1 is precisely measured.

We now show how our bounds can accommodate sampling error by determining standard confidence intervals for $\overline{\Delta\text{CS}}$ and $\underline{\Delta\text{CS}}$. With a slight abuse of notation, let $\overline{\Delta\text{CS}}(q_0, q_1)$ and $\underline{\Delta\text{CS}}(q_0, q_1)$ respectively denote the upper and lower bounds of ΔCS obtained in Theorem 1 as explicit functions of q_0 and q_1, when the demand curve passes through both (p_0, q_0) and (p_1, q_1).

Theorem 6. Let z_τ denote the τ^{th} quantile of the standard normal distribution $\mathcal{N}(0, 1)$. Under (**A1**), (**A5**), and (**EA**), the standard $100 \cdot \alpha\%$ confidence intervals for $\underline{\Delta\text{CS}}$ and $\overline{\Delta\text{CS}}$ are

$$\begin{cases} \text{CI}_\alpha \text{ for } \overline{\Delta\text{CS}} &= \left[\overline{\Delta\text{CS}}\left(q_0, q_1 + z_{(\alpha+1)/2} \cdot \frac{\sigma}{\sqrt{N_1}} \right), \overline{\Delta\text{CS}}\left(q_0, q_1 - z_{(\alpha+1)/2} \cdot \frac{\sigma}{\sqrt{N_1}} \right) \right], \\ \text{CI}_\alpha \text{ for } \underline{\Delta\text{CS}} &= \left[\underline{\Delta\text{CS}}\left(q_0, q_1 + z_{(\alpha+1)/2} \cdot \frac{\sigma}{\sqrt{N_1}} \right), \underline{\Delta\text{CS}}\left(q_0, q_1 - z_{(\alpha+1)/2} \cdot \frac{\sigma}{\sqrt{N_1}} \right) \right]. \end{cases}$$

Theorem 6 follows immediately from the observation that $\overline{\Delta\text{CS}}(q_0, \cdot)$ and $\underline{\Delta\text{CS}}(q_0, \cdot)$ are decreasing for any choice of q_0; hence we omit its proof. This approach for deriving standard confidence intervals holds generally beyond sampling error in q_1: given that our bounds admit closed-form expressions (see Appendix A), monotonicity of these bounds with respect to different variables of interest can be easily verified—which allows us to derive standard confidence intervals to account for uncertainty in these variables.

Although we focus on sampling error in this subsection, our approach can also be extended to incorporate other sources of error. For example, additive market shocks and other types of structural error can be accommodated when the researcher has a prior over their distribution. In this case, the asymptotic normal distribution in our analysis above would be replaced with the distribution in question, yielding analogous confidence intervals to those in Theorem 6.

3.4 Other welfare measures

We conclude this section by discussing how our analysis extends to welfare measures other than consumer surplus. We consider four different types of extensions: *(i)* deadweight loss; *(ii)* weighted consumer surplus using social welfare weights; *(iii)* compensating and equivalent variations when income effects are allowed; and *(iv)* supply-side welfare measures such as producer surplus.

Our results extend most straightforwardly when the welfare measure of interest is deadweight loss—as is often the case in research on tariffs (Section 4.2) and the excess burden of taxation (Section 4.4)—rather than consumer surplus. This is because the change in deadweight loss can be written as

$$\Delta\mathrm{DWL} = \int_{p_0}^{p_1} D(p)\ \mathrm{d}p - (p_1 - p_0)\, q_1 = \Delta\mathrm{CS} - (p_1 - p_0)\, q_1.$$

Since (p_0, q_0) and (p_1, q_1) are observed, maximizing or minimizing $\Delta\mathrm{DWL}$ over any demand family is equivalent to maximizing or minimizing $\Delta\mathrm{CS}$ over that demand family.[5] As such, Theorems 1 and 2—as well as Propositions 1 and 2—continue to hold.

Our results also extend naturally to the case where consumer surplus is weighted by each consumer's social welfare weight (*i.e.*, the marginal increase in social welfare from giving the consumer an additional dollar), a common objective in public finance (Saez and Stantcheva, 2016; Finkelstein and Hendren, 2020). To see how, consider a setting in which each consumer has unit demand for the good. Demand $D(p)$ at a price p can then be interpreted as the mass of consumers whose willingness to pay for the good exceeds p. Let $\lambda(p) \geq 0$ be the average social welfare weight of consumers with a willingness to pay exactly equal to p. The change in weighted consumer surplus is equal to

$$\int_{p_0}^{p_1} \lambda(p)D(p)\ \mathrm{d}p.$$

Crucially, this is a linear functional of $D(\cdot)$, which is sufficient for our information design approach (see Appendix B) to hold. Consequently, Theorems 1 and 2 and Propositions 1 and 2 also continue to hold in the case of weighted consumer surplus.

While many empirical papers focus on the case of quasilinear utility (**A1**), our analysis can also be extended to settings with income effects where the researcher is interested in changes in the compensating and equivalent variations. The key observation that enables our analysis to hold is that the compensating and equivalent variations can be expressed as the area under the Hicksian—rather than Marshallian—demand curve. Correspondingly, the elasticity assumption

---

[5] If only (p_0, q_0) is observed (as in the case of counterfactual extrapolation), then a similar approach to the one we take in Section 3.1 allows us to bound $\Delta\mathrm{DWL}$.

(**EA**) would have to be re-interpreted as a restriction on compensated price elasticities, while the curvature assumptions (**CA1**)–(**CA8**) would have to be re-interpreted as restrictions on the curvature of the Hicksian demand curve.

Finally, although we have focused on welfare analysis based on the demand curve, a similar analysis extends to welfare measures based on the supply curve instead. By combining both demand-side and supply-side welfare measures, our approach also enables us to derive bounds on the change in total surplus.

4 Empirical applications

Our approach to constructing robust welfare bounds can be applied to a number of settings. In this section, we present four applications drawn from different fields, namely applied microeconomics, international trade, and public finance.

Each application builds on an existing paper in the literature; we focus on these papers because they exploit exogenous demand shocks in their respective settings and adopt relatively simple models. This allows us to apply our framework directly and show how we can obtain meaningful bounds on welfare even in the absence of a more complicated structural model.

4.1 Energy subsidies

Empirical researchers are often interested in assessing the *counterfactual* welfare impact from a price that has not been observed. A standard approach is to impose a functional form assumption on demand that can be used to extrapolate from observed price-quantity pairs to any given price. But how robust are welfare estimates to such functional form assumptions? In this subsection, we show how our results provide an alternative approach that accounts for both the uncertainty about the true demand curve and about the counterfactual quantities that would be realized at unobserved price points.

We apply our framework to an application evaluating the energy subsidies offered through the California Alternate Rates for Energy (CARE) program, following the study by Hahn and Metcalfe (2021). The CARE program offers wholesale discounts on unit prices for gas and electricity to eligible low-income households. In Hahn and Metcalfe's sample, CARE households receive a 20% discount on marginal rates, from an average price of \$0.95 to \$0.75, per therm of gas.

Hahn and Metcalfe estimate the net welfare impact of CARE and find that the program results in a net loss of \$4.8 million. They show that, while CARE clearly benefits eligible households,

it also imposes costs through three different channels. First, discounts for eligible households are subsidized by higher-income households, who shoulder a higher cost to compensate for the difference in revenues. Second, lower gas prices encourage higher gas consumption, which harms the environment. Finally, the CARE program entails administrative costs of about \$7 million.

To reach their \$4.8 million estimate of welfare loss, Hahn and Metcalfe project the amount of gas that would be consumed both by households that are enrolled in CARE and those that are not under a linear model of demand. For CARE households, they estimate a local elasticity of consumption at the subsidized price using a LATE research design with randomized nudges for eligible households to sign up and receive the discounted rate. For non-CARE households, they use the local elasticity of consumption estimated by Auffhammer and Rubin (2018). In each case, they extrapolate from the local elasticity estimates by assuming that demand curves are linear. Under this functional form assumption, the observed price-quantity pair and local elasticity estimate pin down the entire demand curve for each type of household, allowing Hahn and Metcalfe to project a counterfactual quantity and to integrate under each demand curve to compute the changes in total surplus.

The mechanics of Hahn and Metcalfe's welfare computation are depicted in Figure 6. For CARE households, the counterfactual unit price p^* is higher than the discounted CARE price p_C,

(a) CARE households   (b) Non-CARE households
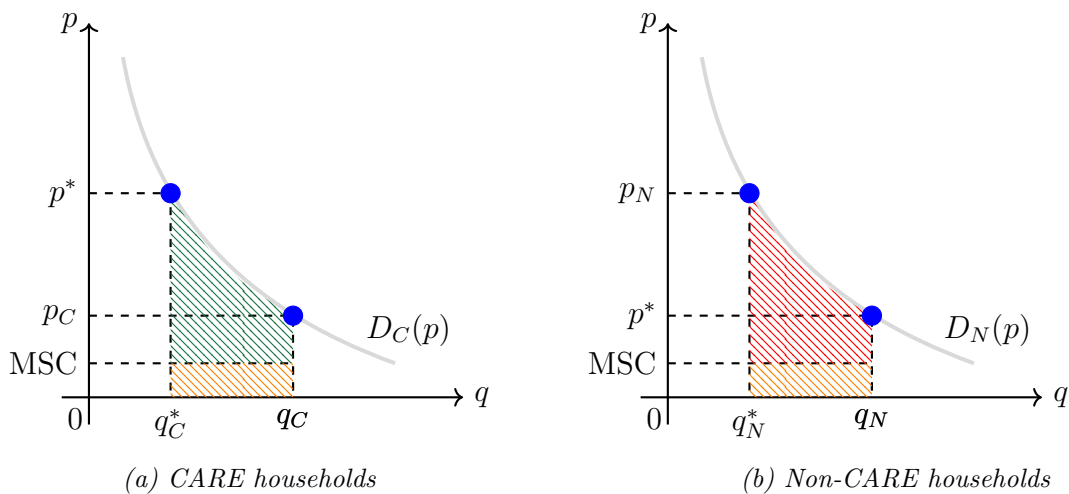
Figure 6: The change in total surplus (excluding the fixed administrative cost) from the CARE program based on Hahn and Metcalfe (2021).

Note: *Prices and quantities are not drawn to scale; the demand curves for CARE and non-CARE households are not directly related in any way. The demand curves $D_C(\cdot)$ and $D_N(\cdot)$, and counterfactual quantities q_C^* and q_N^*, are unknown to the researcher and must be inferred.*

27

and so the counterfactual quantity q_C^* is lower than the observed quantity q_C.[6] For non-CARE households, the opposite is true: $p_N > p^*$ and $q_N < q_N^*$. To compute the change in total surplus (*i.e.*, the sum of consumer and producer surplus), Hahn and Metcalfe integrate under the inverse demand curve for each group. In addition, they account for environmental costs by subtracting the change in quantities consumed multiplied by the marginal social cost (MSC), assessed at $0.68 per therm. Net of environmental costs (in orange), the gain in total surplus (in green) for a representative CARE household is shown in Figure 6(a), while the loss in total surplus (in red) for a representative non-CARE household is shown in Figure 6(b). The net change in total surplus from CARE is thus given by the difference between the green area, multiplied by the number of CARE households, and the red area, multiplied by the number of non-CARE households, minus the fixed administrative cost for the program.

Hahn and Metcalfe find that CARE households (average elasticity $= -0.35$) are substantially more elastic than non-CARE households (average elasticity $= -0.14$). This suggests that the more price-sensitive CARE households may benefit more from the subsidy than non-CARE households are harmed by it. Indeed, under their linear extrapolation, Hahn and Metcalfe estimate a total surplus gain of $5.1 million for CARE households, which outweighs a total surplus loss of $3.1 million for non-CARE households. However, the net change in total surplus for the CARE program becomes negative once the $7 million fixed administrative costs are taken into account.

Hahn and Metcalfe show that their result is robust to a number of sensitivity analyses, including different accounting formulas for the counterfactual price and varying the CARE or non-CARE elasticity over a neighborhood around their estimated values. However, these sensitivity analyses maintain the functional form assumption that the demand curve is linear.

How sensitive is Hahn and Metcalfe's result to the assumption that the demand curve is linear? By Theorem 2, the linear demand curve attains the *upper* bound on the surplus loss from non-CARE households among all convex demand curves for gas. But by the same logic, the linear demand curve also attains the *lower* bound on the surplus gain from CARE households. It is thus challenging to intuit how the linearity assumption affects their estimate barring further numerical analysis.

To overcome this challenge, we apply our results from Section 3.1 to provide an alternative approach that avoids making functional form assumptions (summarized in Figure 7). To begin, we extend Hahn and Metcalfe's sensitivity analysis to account for uncertainty—not only with respect

---

[6] Hahn and Metcalfe derive p^* using an accounting identity that equalizes status quo transfers under CARE. See their Section 4.1.2 for a detailed discussion on the derivation, its robustness to alternative specifications, and its relationship with existing policy.
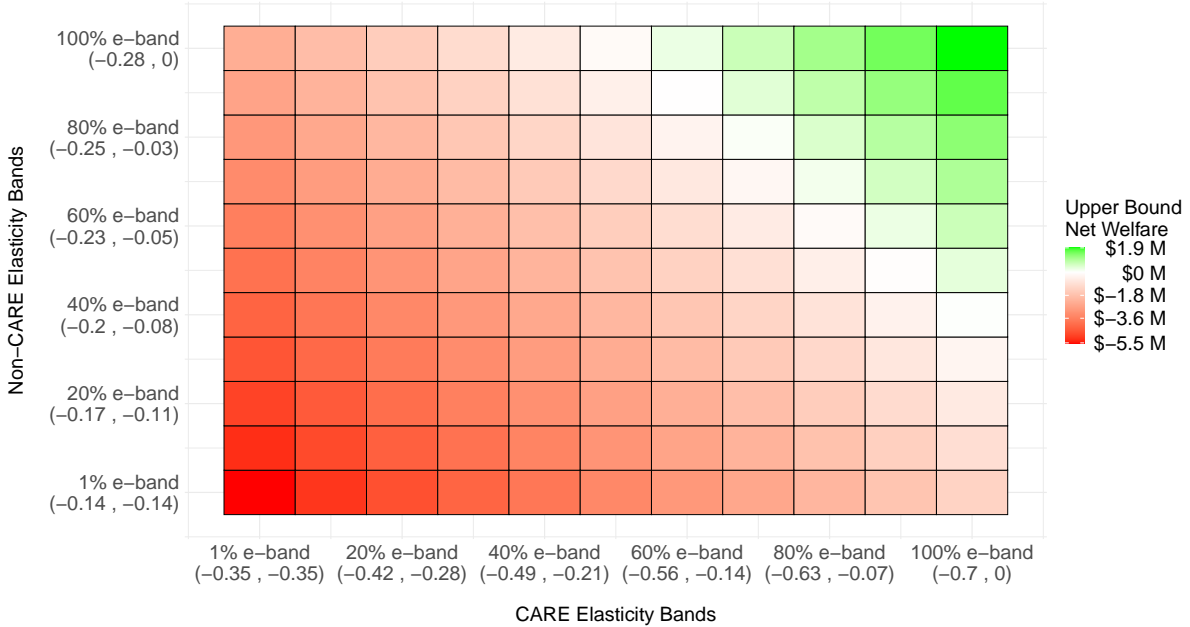
Figure 7: Upper bound of the net change in total surplus without functional form assumptions.

to the elasticity at the observed price-quantity pair for each household group, but also with respect to the counterfactual quantity at the price p^* and the elasticities along the demand curve between the observed and counterfactual price. We then plot the upper bound of the net change in total surplus under increasing elasticity bands around the Hahn and Metcalfe's estimates.[7] In the lower left corner, we assume that each group's elasticity is nearly constant (within 1% of their estimate) throughout its demand curve. In the upper right corner, we assume that each group's elasticity can vary as much as 100% of their estimate.

In each case, we apply the procedure outlined in Section 3.1 to compute our bounds. We first characterize the minimum and maximum values that the counterfactual quantity for each group could take, given the allowable range of elasticities. We then apply Theorem 1 to compute the upper and lower bounds on net welfare (including the environmental impact of counterfactual consumption) for each group and extremal quantity. Finally, we combine the group bounds to obtain bounds on the net change in total surplus.

Our analysis makes clear that Hahn and Metcalfe's result is not only robust in the various dimensions discussed in their paper, but also with respect to their functional form assumption.

---

[7] An $\alpha \cdot 100\%$ band around the elasticity estimate $\hat{\varepsilon}$ means that any elasticity between p_0 and p_1 must fall in the interval $[\underline{\varepsilon}, \overline{\varepsilon}]$ as in (EA), where $\underline{\varepsilon} = (1-\alpha) \times \hat{\varepsilon}$ and $\overline{\varepsilon} = (1+\alpha) \times \hat{\varepsilon}$. This encodes both uncertainty about the elasticity estimates at the observed price and fluctuations in the elasticities along the demand curve.

Even without the linearity assumption, fluctuations of over 50% of the estimated elasticities cannot possibly rationalize a positive net welfare impact of CARE. But Figure 7 also demonstrates how uncertainty about the shape of demand can change the interpretation of the result. Hahn and Metcalfe's estimate lies at the lower left corner of the figure. As larger elasticity fluctuations are allowed, the upper bound on net welfare increases and ultimately becomes positive. If the fixed administrative cost of CARE were lower, the range of elasticities that could rationalize a net benefit from CARE (while the linear interpolation still implies a net loss) might be narrower and more reasonable. In this sense, our approach demonstrates not only whether the particular functional form assumption that is employed is robust to alternative specifications, but also what minimal assumptions are needed to maintain the qualitative conclusions of the exercise.

4.2 Trade tariffs

Between 2018 and 2019, the United States imposed an unprecedented wave of escalating import tariffs on a large set of product sectors and major trading partners. This "return to protectionism" inspired a number of academic studies assessing the welfare impact of the new tariffs (Amiti et al., 2019; Fajgelbaum et al., 2020; Cavallo, Gopinath, Neiman, and Tang, 2021). All of these studies document the same fundamental patterns: *(i)* quantities consumed fell in sectors targeted by the tariffs; *(ii)* foreign producer prices did not change significantly in the short run; and *(iii)* the net domestic impact of the tariffs was ultimately negative.[8] But the modeling choices and empirical techniques employed in each study are slightly different—for instance, Amiti et al. assume linear demand curves, while Fajgelbaum et al. and Cavallo et al. assume CES demand curves. As result, while all of the estimates are similar, it is difficult to discern the extent to which their differences stem from substantive modeling choices (*e.g.*, accounting for substitution across product sectors) rather than different parametrizations.

In this subsection, we demonstrate how our framework can be applied to generate robust bounds for the welfare impact of import tariffs, using the 2018–2019 tariffs as a case study. For expositional simplicity, we focus our analysis on bounding the deadweight loss due to the tariffs, building on the data and approach taken by Amiti et al. For each product (defined as a 10-digit Harmonized Tariff Schedule product code) hit by a tariff, we compare two periods: a period before the trade war started (*e.g.*, March 2017), which we denote by $t = 0$, and a comparable period after tariffs were imposed (*e.g.*, March 2018), which we denote by $t = 1$. As Amiti et al. find

---

[8] By 2020, the Wall Street Journal editorial board had written about the "piling" evidence of net economic harm from tariffs in an article titled "How Many Tariff Studies Are Enough?"

near-complete tariff pass-through to consumers,[9] we assume that pre-tariff prices did not change during the trade war. Thus, a product j that was priced at $p_{j,0}$ at $t = 0$ would be priced at $p_{j,1} = (1 + \tau) p_{j,0}$ at $t = 1$, where τ is the *ad valorem* tariff imposed on good j at $t = 1$.

Furthermore, as this implies that producer prices did not change in response to the tariffs, we follow Amiti et al. in assuming that the producer supply curves are flat, and so the tariffs did not incur any losses to producer surplus.[10] In this case, computing the deadweight loss from the tariff on a given good (in a given month) is equivalent to the Harberger exercise discussed in the introduction. The deadweight loss is given by the area B in Figure 1(a): the total change in consumer surplus (given by the integral of the unobserved demand curve between the two price points), less the earned tariff revenues. Amiti et al. impute a demand curve through a linear interpolation and compute their deadweight loss estimate by calculating the area of the resulting Harberger triangle.

In order to compute bounds, we consider observations of price and quantity, $(p_{jmc,0}, q_{jmc,0})$ and $(p_{jmc,1}, q_{jmc,1})$, for each triple of product j, month m, and country c for which a tariff was introduced in 2018. To obtain these price-quantity observations, we draw from the US Customs data report following the replication code provided by Amiti et al. We close the model by assuming (as Amiti et al. do) that product sales are independent of each other, so that tariffs on one set of products do not impact sales on another set that is not yet affected. This allows us to treat each product's demand curve independently, and to aggregate the deadweight losses across all affected products for a total amount.

As the 2018 tariffs affected a variety of very different goods, it may be difficult to conjure an informative prior on the range of feasible elasticities. Instead, we begin by considering the total estimated deadweight loss that is implied by four of the interpolations described in Section 2.3: CES, constant marginal revenue (CMR), exponential, and linear. Figure 8(a) plots the monthly deadweight loss across all affected products from February through December 2018 for each demand curve. By Theorem 2, these estimates correspond to one-sided bounds for different families of demand curves. For instance, as CES demand is the lower extremum for the family of demand curves that satisfy Marshall's second law (**CA1**), the CES estimates in Figure 8(a)

---

[9]   This is consistent with the findings of Fajgelbaum et al. and Cavallo et al., who employ different estimation methodologies than Amiti et al.

[10]   As Fajgelbaum et al. point out, a more conservative conclusion based on the finding that producer prices did not change in response to the tariffs is that a flat supply curve is possible, and cannot be ruled out. If this is the case, our bounds are still valid, but they capture only the deadweight loss incurred by consumers. It would be possible to do a "doubly robust" version of our bounds to account for the deadweight loss incurred by producers, but as this does not feature in the papers we are working off of, we omit it for the sake of brevity.
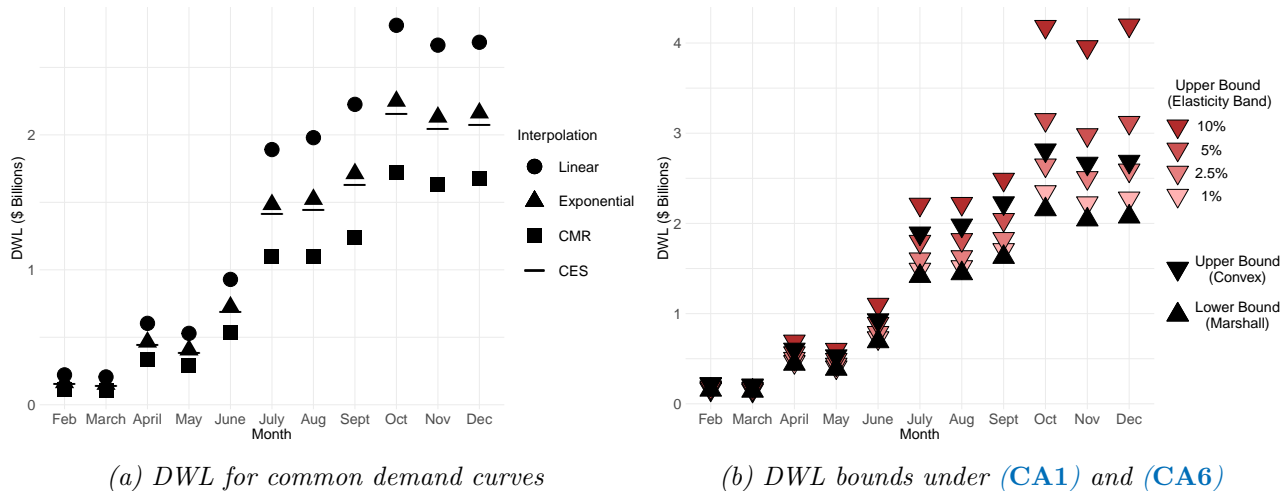
(a) DWL for common demand curves · *(b) DWL bounds under (CA1) and (CA6)*

Figure 8: Comparison of DWL bounds under different assumptions.

provide the lower bound on deadweight loss consistent with Marshall's second law. Similarly, the CMR estimates provide the lower bound consistent with decreasing marginal revenue (**CA2**), the exponential estimates provide the lower bound consistent with log-concave demand (**CA3**), and the linear estimates provide the lower bound consistent with concave demand (**CA4**).

Notably, while the estimates differ in magnitude in each month—reflecting differences in seasonal demand as well as the gradual addition of new tariffs—their relative ordering does not change. This is unsurprising: as we argue in Section 2.3 and show in Appendix C, concavity implies log-concavity, which implies decreasing marginal revenue and Marshall's second law. Thus, the lower bound in the family of concave demand curves must be at least as high as the lower bound in the family of log-concave demand curves, and so on.

To interpret the estimates in Figure 8(a) as bounds, it is necessary to take a stance on which assumptions regarding demand are most appropriate. For instance, while the linear estimate for deadweight loss provides a *lower* bound consistent with *concave* demand (**CA4**), it could instead be interpreted as the *upper* bound consistent with *convex* demand (**CA6**). Indeed, whereas Figure 8(a) demonstrates the range of extremal deadweight loss estimates that may be consistent with common families of demand, one might instead wish to know the tightest bounds that are consistent with the set of restrictions that pertain to the international trade setting.

In Figure 8(b), we demonstrate how this can be done. Given the prominence of Marshall's second law in the trade literature, we focus on (**CA1**) as a base assumption. By Theorem 2, (**CA1**) implies that the lower bound on deadweight loss is given by a 1-piece CES interpolation. The

corresponding estimate for each month drawn in a black upward-facing triangle in Figure 8(b). As we note in Section 2.3, absent any other data or restrictions, the conservative (box) upper bound is admissible and no tighter bound can be guaranteed. However, if there are other assumptions that fit the setting—whether additional curvature restrictions or elasticity restrictions—they may be combined with Marshall's second law.

To illustrate this, the black downward-facing triangles in Figure 8(b) correspond to the upper bound on deadweight loss if demand is also assumed to be convex (**CA6**)—the 1-piece linear interpolation estimates also represented by dots in Figure 8(a). Taken together, the black triangles in Figure 8(b) present a set of bounds under both (**CA1**) and (**CA6**). When the total deadweight loss is small—as in the early months of the trade war when the size and the number of tariffs was small—the bounds are very close together as the linear demand curve is a good (Taylor) approximation to the actual demand curve. However, for months with more tariffs and larger distortions, the range of admissible deadweight loss estimates is substantial. Summing across all months, we find that the total deadweight loss from tariffs was at least \$12.6 billion and at most \$16.8 billion under both (**CA1**) and (**CA6**). As Amiti et al. note, the US government internalized \$15.6 billion in revenues from import tariffs over the course of 2018, meaning that the cost to consumers and importers ranged from \$28.2 billion to \$32.4 billion.

The assumptions that are imposed on demand affect the magnitude of these estimates. For example, when the convexity assumption (**CA6**) is not appropriate, we might consider how the upper bound on deadweight loss might change under different elasticity restrictions. Since we do not have a domain-specific prior over what ranges of elasticities may be reasonable, we consider symmetric bands around the average elasticity observed in each month.[11] The colored downward-facing triangles in Figure 8(b) present the upper bound for total monthly deadweight loss under (**EA**) if elasticities can be within 1%, 2.5%, 5%, and 10% of the observed average. While the majority of the band estimates fall below the upper bound corresponding to convexity, even the 5% band bound exceeds the convexity bound starting in October. Thus, even demand curves with modestly decreasing elasticities may be consistent with total deadweight loss estimates that exceed \$16.8 billion.

---

[11] The average elasticity can be inferred from (p_0, q_0) and (p_1, q_1) alone: $\varepsilon_{\text{avg}} = \frac{\log(q_1) - \log(q_0)}{\log(p_1) - \log(p_0)}$. As in Section 4.1, an $\alpha \cdot 100\%$-band around the average elasticity ε_{avg} means that $\underline{\varepsilon} = (1 - \Delta) \times \varepsilon_{\text{avg}}$ and $\overline{\varepsilon} = (1 + \Delta) \times \varepsilon_{\text{avg}}$.

4.3 Old-age pensions

Recent work in public finance has heeded calls to "harness the gains of the 'credibility revolution' for the goal of welfare analysis" (Finkelstein and Hendren, 2020) by relating carefully estimated policy impacts to measures of willingness to pay that can be composed into a "marginal value of public funds" (MVPF) and compared across different outlets of government spending. Many of these studies evaluate programs that give a discrete subsidy to a specific group of individuals (*e.g.*, Pell Grants, supplemental security income, and pensions).[12] To evaluate welfare, the recipients of the subsidy are typically split into two groups: marginal recipients who are determined to have switched their behavior on the basis of the subsidy (*e.g.*, enrolled in college or retired), and inframarginal recipients who benefit from the subsidy with no change in behavior. Since they do not exhibit a demand response by definition, inframarginal recipients are presumed to internalize the full benefit of every dollar of subsidy. However, assessing the surplus gained by marginal recipients requires a stronger assumption—and often, studies assume that the marginal recipients internalize no benefit whatsoever from the subsidy.

In this subsection, we consider an example of the MVPF approach through a study on the impact of the 1908 Old-Age Pension Act (OPA) in the UK by Giesecke and Jäger (2021). The OPA launched the first universal pension for low-income workers in the UK. Giesecke and Jäger use individual level data from the UK census in 1891, 1901, and 1911 to assess the impact that the introduction of the OPA had on labor force participation by eligible workers. Using a regression discontinuity design around the minimum age of eligibility (70) immediately after the introduction of the OPA, Giesecke and Jäger find that labor participation for eligible workers dropped from 46% to 40%—nearly all due to workers who retired. Assuming that the labor force participation rate would have continued to evolve after the cutoff age (71+) at the same rate as for people aged 65–69 absent the OPA, they conclude that the OPA caused the total number of eligible workers who retired to grow from $557,505$ to $613,873$: a change from 52% to 58% of all eligible workers.

To estimate the welfare impact of the OPA, we consider a model of retirement on the basis of foregone wages and pension benefits. Giesecke and Jäger report that the overwhelming majority of pension recipients received the maximum pension of 260 shillings per year (22% of average earnings). However, there is little data on the distribution of incomes, other than the fact that the total annual income of most workers was no higher than 420 shillings a year (36% of average earnings) in order to qualify for the maximum pension.

---

[12] See https://policyimpacts.org/policy-impacts-library for a running list of studies that the MVPF has been computed for.
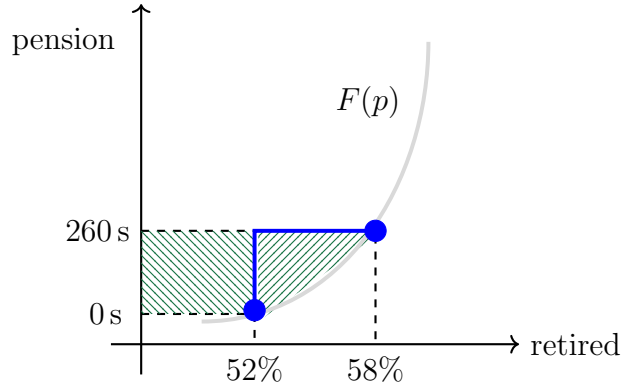
Figure 9: Illustration of how probability of retirement varies with pension offered.

As such, we apply our framework to bound the welfare impact of the OPA with respect to uncertainty about the distribution of incomes among eligible workers. We consider each worker's probability of retirement as a function of the pension that is available to him. Each eligible worker i derives a utility from working w_i that includes his wage, taste for working, and distaste for retirement. For simplicity, we assume that w_i is drawn independently from the same distribution F for each eligible worker. Given a pension amount p, the worker retires if and only if $p \geq w_i$. The probability that a given worker retires is thus $F(p)$, and by the central limit theorem (as in Section 3.3), the proportion of workers who retire is asymptotically distributed according to a Gaussian distribution with mean $F(p)$ and variance $F(p) \left[1 - F(p)\right] / N$.[13]

The change in worker surplus can now be expressed as a linear functional of the distribution F, just as how the change in consumer surplus is a linear functional of the demand curve in (1). As shown in Figure 9, the distribution F can be thought of as a supply curve: $F(p)$ is the proportion of workers who would retire at a pension of p. Thus, when the OPA increased pensions from $p_0 = 0$ shillings to $p_1 = 260$ shillings, worker surplus increased accordingly by

$$\Delta W = \int_{p_0}^{p_1} F(p) \ \mathrm{d}p.$$

As F is not observed, Giesecke and Jäger adopt a conservative approach by assuming that marginal workers who retired because of the OPA were indifferent between working and receiving a pension. By this assumption, only inframarginal retirees benefited from their pensions, and the

---

[13] Here N refers to the number of eligible workers; at the time of the OPA, there were $1,068,486$ such people. We apply the central limit theorem to the Bernoulli random variable $\mathbf{1}_{p \geq w_i}$, which indicates if worker i retires when he receives a pension amount of p.

total surplus gained by workers is proportional to the fraction of inframarginal pension recipients: 0.91 shillings of surplus for every shilling distributed. Giesecke and Jäger's assumption is equivalent to assuming that the willingness to pay for retirement is constant between $p_0 = 0$ shillings and $p_1 = 260$ shillings for eligible workers (depicted by the blue curve in Figure 9).

While this bound on ΔW is conservative, it would only be attained if all marginal workers required a pension of 260 shillings to be willing to retire—and nothing less. This is unlikely to be the case as Giesecke and Jäger describe substantial heterogeneity in hours worked prior to retirement—suggesting nontrivial curvature in the distribution of w_i between $p_0 = 0$ shillings and $p_1 = 260$ shillings. Commonly assumed functional forms for income are log-concave (McDonald and Ransom, 1979), and can be parameterized via assumptions on ρ-concavity and ρ-convexity (Caplin and Nalebuff, 1991a). Figure 10(a) illustrates bounds implied by our Theorem 2 when ρ varies. When F is ρ-concave, the value plotted in Figure 10(a) corresponds to a lower bound; and when F is ρ-convex, the value corresponds to an upper bound.[14] In the limiting case of $\rho \to -\infty$, our bounds coincide with the conservative bounds (where Giesecke and Jäger's bound is the lower bound implied by a $(-\infty)$-concave, or quasi-concave, distribution F). As in the discussion following Theorem 2, these values can be combined in some cases as well. For example, when F is log-concave and convex, then Figure 10(a) implies an upper bound of 0.954 and a lower bound of 0.953.

Beyond curvature, we can also parametrize uncertainty in the distribution F through the standard deviation of w_i. The conservative bounds (the lower of which is Giesecke and Jäger's

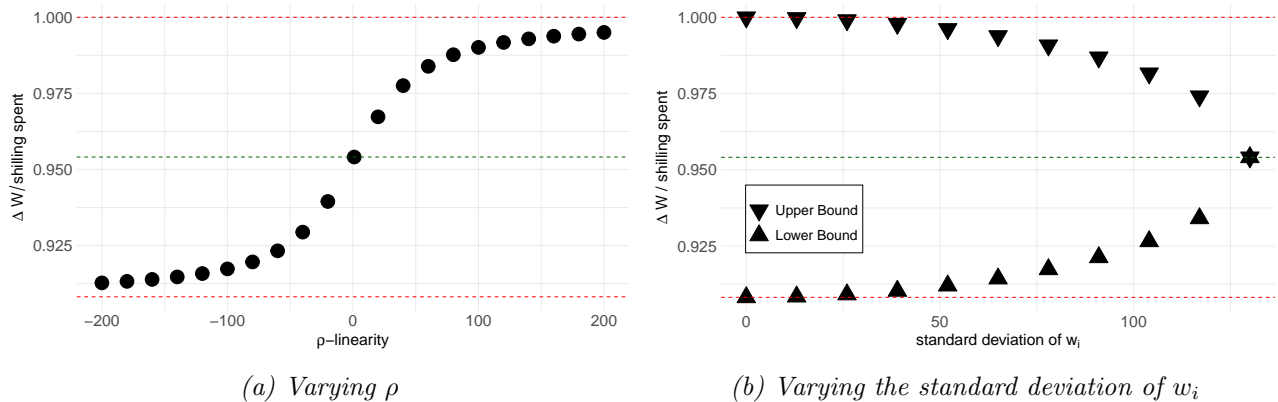(a) Varying ρ (b) Varying the standard deviation of w_i

Figure 10: Gain in worker surplus due to the OPA based on Giesecke and Jäger (2021).

---

[14] When $\rho = 1$, this value corresponds to the upper (resp. lower) bound for a convex (resp. concave) F; and when $\rho = 0$, it corresponds to the lower (resp. upper) bound for a log-concave (resp. log-convex) F.

bound) are attained only when the standard deviation of each w_i is exactly zero between $p_0 = 0$ shillings and $p_1 = 260$ shillings. Narrower bounds prevail when the standard deviation of w_i is not zero, as we illustrate in Figure 10(b). As we show in Appendix D, these bounds can be derived in a similar manner to Theorems 1 and 2 using the approach of Appendix B.

Finally, our bounds on ΔW translate into bounds on the MVPF of pensions in the context of the OPA. Giesecke and Jäger argue that each £1 of pension payments cost the government £1.13, so that bounds on the MVPF can be obtained via the formula MVPF $= \Delta W/1.13$. For example, when the standard deviation of w_i between $p_0 = 0$ shillings and $p_1 = 260$ shillings is 50 shillings, we obtain a lower bound of 0.81 and an upper bound of 0.88 on the MVPF.

4.4 Income taxation

In this subsection, we show how our framework may be used to study such fiscal externalities when the extent to which taxpayers would reduce their taxable income in response to different tax rates is uncertain. A popular approach to this question has focused on establishing "sufficient statistics" for welfare changes around marginal increases in taxes.[15] Invoking the envelope theorem to argue that the secondary effects of a small tax increase are of second-order importance when consumers are optimizing, a typical analysis decomposes the marginal welfare gain $(\mathrm{d}W/\mathrm{d}\tau)$ into a simple function of marginal consumption utilities that can be inferred by revealed preference from data (under context-appropriate assumptions). As Chetty (2009) demonstrates, this approach can be applied to generate empirical measurements of welfare effects for a broad set of policy settings.

However, an estimate of the marginal change in welfare may be less meaningful when the change in the tax rate is substantial. In this case, estimating the change in welfare requires taking a stance on how to integrate $\mathrm{d}W/\mathrm{d}\tau$ and parametric interpolations are often invoked.[16] When policy changes are small, Taylor's theorem ensures that parametric approximations work well. But there is no such guarantee when policy changes are large—which is often the case in applications. For these cases, our approach allows the researcher to assess global effects through robust bounds. The bounds are estimated using the same empirical moments that are considered under "sufficient statistics"—local elasticity measurements with respect to the policy variables. But instead of relying on local approximation assumptions, they characterize the extremal welfare changes that are consistent with prior knowledge like estimates from the contemporary literature.

---

[15] See Chetty (2009) and Kleven (2021) for thorough and accessible overviews of the sufficient statistics approach to welfare analysis and how it may be applied to different settings in public economics.

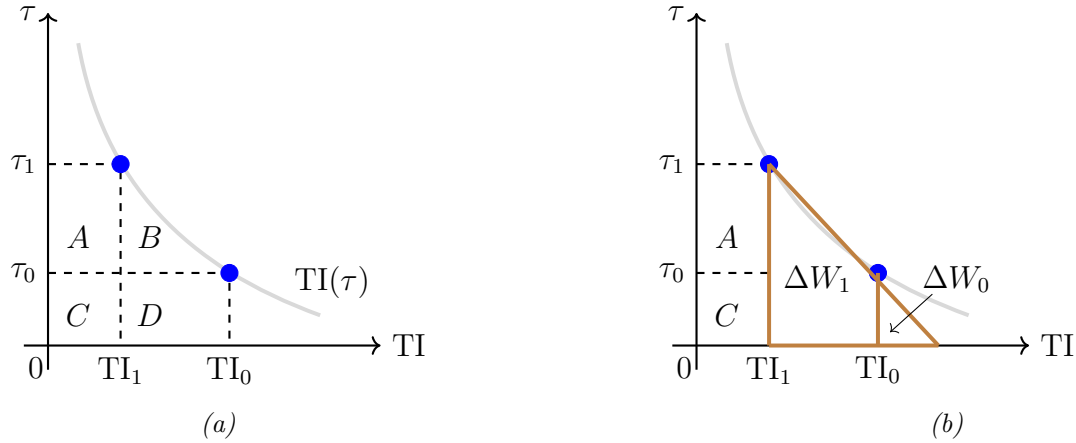[16] Kleven (2021) suggests including higher-order terms as an alternative to parametric interpolations.

Figure 11: Feldstein's sufficient statistics approach to estimating the excess burden of income taxation.

To demonstrate, we consider an exercise based on Feldstein (1999)—one of the first papers to apply sufficient statistics to study optimal income tax policy. Feldstein considers the excess burden of income taxes given equilibrium income production and (costly) income sheltering. Individuals choose how much labor to supply (across different sources with different costs and wages), how much of the earned income to shelter from tax authorities (given the cost of sheltering different amounts of money) and how to distribute the income left after taxation across consumption. Despite the many moving parts of this model, Feldstein invokes an envelope theorem argument to show that the excess burden of a marginal increase from a tax rate t is proportional to the marginal decrease in taxable income (TI):[17]

$$\frac{\mathrm{d} W(\tau)}{\mathrm{d} \tau} = \tau \cdot \frac{\mathrm{d}\, \mathrm{TI}(\tau)}{\mathrm{d} \tau}.$$

That is, although it accounts for income effects in labor decisions, the (local) first-order estimate of deadweight loss relies only on a local measurement of the elasticity of taxable income.

In Figures 11(a) and 11(b), we demonstrate Feldstein's approach in relation to our Harberger example in Section 1. We consider the change in taxable income $\mathrm{TI}(\tau)$ as the effective tax rate changes from τ_0 to τ_1. By the envelope theorem, the marginal increase in excess burden at any point τ along the curve is approximated by $\tau \cdot \mathrm{d}\, \mathrm{TI}(\tau)/\mathrm{d}\tau$. The net change in excess burden

---

[17] To make clear the connection between the sufficient statistics approach and ours, our discussion of Feldstein's analysis follows the expositional logic laid out by Chetty (2009), who gives a complete derivation of $\mathrm{d} W(\tau)/\mathrm{d}\tau$.

between τ_0 and τ_1 is therefore given by the sum of areas B and D:

$$\Delta W = W(\tau_1) - W(\tau_0) = \int_{\tau_0}^{\tau_1} \tau \cdot \text{TI}'(\tau) \, d\tau$$
$$= [\tau_1 \, \text{TI}(\tau_1) - \tau_0 \, \text{TI}(\tau_0)] - \int_{\tau_0}^{\tau_1} \text{TI}(\tau) \, d\tau$$
$$= -(\text{area } B + \text{area } D).$$

As in the Harberger example, area D is simple to evaluate without further assumptions: it is given by the product of τ_0 and $\text{TI}(\tau_0) - \text{TI}(\tau_1)$. However, as the curve $\text{TI}(\tau)$ is not observed, evaluating area B requires further assumptions.

To resolve this, Feldstein proposes a formula that approximates the change in welfare by the difference between the areas of triangle in Figure 11(b):

$$\Delta W \approx \Delta W_1 - \Delta W_0.$$

When τ_0 and τ_1 are small, each triangle can be thought of as a Taylor approximation around zero, independently of the shape of $\text{TI}(\tau)$. However, for most applications, the triangles may be better thought of as (unconditional) 1-piece linear interpolations. The interpretation of Feldstein's result therefore depends on the family of plausible $\text{TI}(\tau)$ curves from which the observed points were drawn.

To illustrate how our robust bounds approach relates to Feldstein's formula, we consider a numerical example discussed in his paper. Feldstein considers a taxpayer with an income of $\$180,000$, who would have been subject to a marginal tax rate increase from 31% to 38.9% in 1993, and been predicted to reduce their taxable income by $\$21,340$ in response. To evaluate the excess burden for this individual, Feldstein imputes a linear $\text{TI}(\tau)$ function using a taxable income elasticity of $\epsilon_\tau = -1.04$, based on a difference-in-differences exercise by Feldstein (1995). However, he notes that contemporary papers had estimated elasticities ranging from -1.33 to -0.55 using different methodologies. Under a linear interpolation, Feldstein estimates an excess burden of around $\$7,458$ for this taxpayer.

However, this estimate is not without loss of generality. Applying our robust bounds approach, conservative bounds for the taxpayer's excess burden—which make no assumptions on curvature or elasticities—suggest that it is between $\$6,615$ and $\$8,301$. Restricting elasticities to range between -1.33 and -0.55 shrinks the feasible magnitudes of the excess burden to between $\$7,400$ and $\$7,418$ by applying our Theorem 1. Feldstein's estimate of $\$7,458$ lies outside of these bounds,

however. This implies that a linear interpolation is incompatible with Feldstein's range of elasticity estimates—a fact that we can verify directly as well.

We can apply a version of Feldstein's exercise to evaluate the excess burden from taxes levied to fund the 1908 OPA in the UK, which we studied in the previous subsection. UK legislators increased marginal tax rates for the top 1% of earners by 1 to 5 percentage points (with the highest increases for the wealthiest bracket). If high earners were elastic in their earning behavior, it may be that the OPA tax increases led to distorted incomes and substantial deadweight loss.

To evaluate the excess burden from taxes levied to fund the OPA, Giesecke and Jäger propose a range of taxable income elasticities that can be used to evaluate the policy: a "mid-range estimate from the empirical literature" obtained from Diamond and Saez (2011) of -0.25, and an "extreme case" elasticity of -5. Following Giesecke and Jäger, we consider the wealthiest bracket of tax payers, who earned at least 5,000 shillings and were hit with the largest increase in marginal tax rates: from 3.8% to 8.3%. Although we do not know how the income that these taxpayers earned changed after the OPA was introduced, we can bound it using the feasible elasticity range. The resulting excess burden for an individual at the margin of the top tax bracket is then bounded between £50.50 and £227.

5 Conclusion

The rapid growth of academic articles on welfare analysis in the last decade (shown in Figure 1 of Kleven, 2021, for example) is testament to its importance and relevance to policy. While the welfare analysis of small policy changes is well understood, traditional approaches to the analysis of large policy changes often employ functional form assumptions. These approaches have played a crucial role in the understanding economists have gained about policy in different markets, but a question of robustness arises: how much do these intuitions depend on particular choices of functional forms?

In this paper, we answer this question by developing a different approach. We provide bounds on welfare that hold under various families of assumptions that are commonly made in different literatures, and illustrate our theoretical framework in a series of empirical applications. Our results demonstrate a serendipitous connection between information and mechanism design and welfare analysis in empirical work, which we view as a promising area for future research.

References

AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100, 1601–15.

AGUIRREGABIRIA, V. (2010): "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior," *Journal of Business & Economic Statistics*, 28, 201–218.

AMIR, R., I. MARET, AND M. TROEGE (2004): "On Taxation Pass-Through for a Monopoly Firm," *Annales d'Économie et de Statistique*, Jul.–Dec., 155–172.

AMITI, M., S. J. REDDING, AND D. E. WEINSTEIN (2019): "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 33, 187–210.

AN, M. Y. (1998): "Logconcavity versus Logconvexity: A Complete Characterization," *Journal of Economic Theory*, 80, 350–369.

ANDERSON, S. P. AND R. RENAULT (2003): "Efficiency and Surplus Bounds in Cournot Competition," *Journal of Economic Theory*, 113, 253–264.

ANDREYEVA, T., M. W. LONG, AND K. D. BROWNELL (2010): "The Impact of Food Prices on Consumption: A Systematic Review of Research on the Price Elasticity of Demand for Food," *American Journal of Public Health*, 100, 216–222.

AUFFHAMMER, M. AND E. RUBIN (2018): "Natural Gas Price Elasticities and Optimal Cost Recovery Under Consumer Heterogeneity: Evidence From 300 Million Natural Gas Bills," *Working paper*.

BAGNOLI, M. AND T. BERGSTROM (2005): "Log-Concave Probability and Its Applications," *Economic Theory*, 26, 445–469.

BERGEMANN, D., F. CASTRO, AND G. WEINTRAUB (2022): "Third-Degree Price Discrimination Versus Uniform Pricing," *Games and Economic Behavior*, 131, 275–291.

BERGEMANN, D. AND K. SCHLAG (2011): "Robust Monopoly Pricing," *Journal of Economic Theory*, 146, 2527–2543.

BERRY, S. T. AND P. A. HAILE (2021): "Foundations of Demand Estimation," in *Handbook of Industrial Organization* Volume 4: Elsevier, 1–62.

BISHOP, R. L. (1968): "The Effects of Specific and Ad Valorem Taxes," *Quarterly Journal of Economics*, 82, 198–218.

BRYNJOLFSSON, E., Y. HU, AND M. D. SMITH (2003): "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49, 1580–1596.

BULOW, J. I. AND P. PFLEIDERER (1983): "A Note on the Effect of Cost Changes on Prices," *Journal of Political Economy*, 91, 182–185.

CAPLIN, A. AND B. NALEBUFF (1991a): "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59, 25–59.

——— (1991b): "Aggregation and Social Choice: A Mean Voter Theorem," *Econometrica*, 59, 1–23.

CAVALLO, A., G. GOPINATH, B. NEIMAN, AND J. TANG (2021): "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy," *American Economic Review: Insights*, 3, 19–34.

CHETTY, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488.

COHEN, M. C., G. PERAKIS, AND R. S. PINDYCK (2021): "A Simple Rule for Pricing with Limited Knowledge of Demand," *Management Science*, 67, 1608–1621.

COHEN, P., R. HAHN, J. HALL, S. LEVITT, AND R. METCALFE (2016): "Using Big Data to Estimate Consumer Surplus: The Case of Uber," *Working paper*, available at https://www.nber.org/papers/w22627.

CONDORELLI, D. AND B. SZENTES (2020): "Surplus Bounds in Cournot Monopoly and Competition," *Working paper*.

CUTLER, D. M. AND R. J. ZECKHAUSER (2000): "The Anatomy of Health Insurance," in *Handbook of Health Economics* Volume 1: Elsevier, 563–643.

DESHPANDE, M. AND M. G. MUELLER-SMITH (2022): "Does Welfare Prevent Crime? The Criminal Justice Outcomes of Youth Removed from SSI," *Working paper*.

DIAMOND, P. AND E. SAEZ (2011): "The Case for a Progressive Tax: From Basic Research to Policy Recommendations," *Journal of Economic Perspectives*, 25, 165–90.

DIXIT, A. (1986): "Comparative Statics for Oligopoly," *International Economic Review*, 27, 107–122.

EINAV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 125, 877–921.

FAJGELBAUM, P. D., P. K. GOLDBERG, P. J. KENNEDY, AND A. K. KHANDELWAL (2020): "The Return to Protectionism," *Quarterly Journal of Economics*, 135, 1–55.

FELDSTEIN, M. (1995): "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act," *Journal of Political Economy*, 103, 551–572.

——— (1999): "Tax Avoidance and the Deadweight Loss of the Income Tax," *Review of Economics and Statistics*, 81, 674–680.

FINKELSTEIN, A. AND N. HENDREN (2020): "Welfare Analysis Meets Causal Inference," *Journal of Economic Perspectives*, 34, 146–67.

GANDHI, A. AND A. NEVO (2021): "Empirical Models of Demand and Supply in Differentiated Products Industries," in *Handbook of Industrial Organization* Volume 4: Elsevier, 63–139.

GENTZKOW, M. AND E. KAMENICA (2016): "A Rothschild–Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 106, 597–601.

GIESECKE, M. AND P. JÄGER (2021): "Pension Incentives and Labor Supply: Evidence from the Introduction of Universal Old-Age Assistance in the UK," *Journal of Public Economics*, 203, 104516.

GRAY, C., A. LEIVE, E. PRAGER, K. B. PUKELIS, AND M. ZAKI (2021): "Employed in a SNAP? The Impact of Work Requirements on Program Participation and Labor Supply," *Working paper*.

HACKMANN, M. B., J. T. KOLSTAD, AND A. E. KOWALSKI (2015): "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *American Economic Review*, 105, 1030–66.

HAHN, R. W. AND R. D. METCALFE (2021): "Efficiency and Equity Impacts of Energy Subsidies," *American Economic Review*, 111, 1658–88.

HARBERGER, A. C. (1954): "Monopoly and Resource Allocation," *American Economic Review: Papers & Proceedings*, 44, 77–87.

——— (1964): "The Measurement of Waste," *American Economic Review: Papers & Proceedings*, 54, 58–76.

——— (1978): "On the Use of Distributional Weights in Social Cost-Benefit Analysis," *Journal of Political Economy*, 86, S87–S120.

HAUSMAN, J. A. (1981): "Exact Consumer's Surplus and Deadweight Loss," *American Economic Review*, 71, 662–676.

——— (1996): "Valuation of New Goods under Perfect and Imperfect Competition," in *The Economics of New Goods* ed. by Bresnahan, T. F. and Gordon, R. J. Chicago, IL: University of Chicago Press, 207–248.

——— (2003): "Sources of Bias and Solutions to Bias in the Consumer Price Index," *Journal of Economic Perspectives*, 17, 23–44.

HAUSMAN, J. A., A. PAKES, AND G. L. ROSSTON (1997): "Valuing the Effect of Regulation on New Services in Telecommunications," *Brookings Papers on Economic Activity, Microeconomics*, 1997, 1–54.

HENDREN, N. AND B. SPRUNG-KEYSER (2020): "A Unified Welfare Analysis of Government Policies," *Quarterly Journal of Economics*, 135, 1209–1318.

JÁCOME, E. (2020): "Mental Health and Criminal Involvement: Evidence from Losing Medicaid Eligibility," *Working paper*.

JOHNSON, J. P. (2017): "The Agency Model and MFN Clauses," *Review of Economic Studies*, 84, 1151–1185.

KALOUPTSIDI, M., P. T. SCOTT, AND E. SOUZA-RODRIGUES (2021): "Identification of Counterfactuals in Dynamic Discrete Choice Models," *Quantitative Economics*, 12, 351–403.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KANG, Z. Y. AND J. VONDRÁK (2019): "Fixed-Price Approximations to Optimal Efficiency in Bilateral Trade," *Working paper*.

KILIAN, L. AND X. ZHOU (2020): "Gasoline Demand More Responsive to Price Changes than Economists Once Thought," https://www.dallasfed.org:443/research/economics/2020/0616.

KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89, 1557–1593.

KLEVEN, H. J. (2021): "Sufficient Statistics Revisited," *Annual Review of Economics*, 13.

KRUGMAN, P. R. (1979): "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, 9, 469–479.

MALUEG, D. A. (1994): "Monopoly Output and Welfare: The Role of Curvature of the Demand Function," *Journal of Economic Education*, 25, 235–250.

MANNING, W. G., J. P. NEWHOUSE, N. DUAN, E. B. KEELER, AND A. LEIBOWITZ (1987): "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *American Economic Review*, 251–277.

MANSKI, C. F. (1997): "Monotone Treatment Response," *Econometrica*, 65, 1311–1334.

MARSHALL, A. (1890): *Principles of Economics*, London, UK: Macmillan and Co.

MCDONALD, J. B. AND M. R. RANSOM (1979): "Functional Forms, Estimation Techniques and the Distribution of Income," *Econometrica*, 47, 1513–1525.

MELITZ, M. J. (2018): "Competitive Effects of Trade: Theory and Measurement," *Review of World Economics*, 154, 1–13.

MORRIS, M. (2014): "Gasoline Prices Tend to Have Little Effect on Demand for Car Travel," https://www.eia.gov/todayinenergy/detail.php?id=19191.

MÜLLER, A. AND M. SCARSINI (2006): "Stochastic Order Relations and Lattices of Probability Measures," *SIAM Journal on Optimization*, 16, 1024–1043.

PRÉKOPA, A. (1973): "On Logarithmic Concave Measures and Functions," *Acta Scientiarum Mathematicarum*, 34, 335–343.

REGUANT, M. (2016): "Bounding Outcomes in Counterfactual Analysis," *Working paper*.

ROBINSON, J. (1933): *The Economics of Imperfect Competition*, London, UK: Macmillan.

SAEZ, E. AND S. STANTCHEVA (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45.

SUMNER, D. A. (1981): "Measurement of Monopoly Behavior: An Application to the Cigarette Industry," *Journal of Political Economy*, 89, 1010–1019.

TAMER, E. (2010): "Partial Identification in Econometrics," *Annual Review of Economics*, 2, 167–195.

TSITSIKLIS, J. N. AND Y. XU (2014): "Efficiency Loss in a Cournot Oligopoly with Convex Market Demand," *Journal of Mathematical Economics*, 53, 46–58.

VARIAN, H. R. (1982): "The Nonparametric Approach to Demand Analysis," *Econometrica*, 945–973.

——— (1985): "Price Discrimination and Social Welfare," *American Economic Review*, 75, 870–875.

VIVES, X. (1987): "Small Income Effects: A Marshallian Theory of Consumer Surplus and Downward Sloping Demand," *Review of Economic Studies*, 54, 87–103.

WALL STREET JOURNAL (2020): "How Many Tariff Studies Are Enough?," https://www.wsj.com/articles/how-many-tariff-studies-are-enough-11579556389.

WEYL, E. G. AND M. FABINGER (2013): "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 121, 528–583.

WILLIG, R. D. (1976): "Consumer's Surplus Without Appology," *American Economic Review*, 66, 589–597.

Appendix A   Closed-form expressions for welfare bounds

In this appendix, we present closed-form expressions for the welfare bounds that we characterized in the paper. We then show how these expressions generalize to an extension where the researcher observes both $\varepsilon(p_0)$ and $\varepsilon(p_1)$.

A.1   Welfare bounds from Section 2

Our characterization of welfare bounds under the elasticity assumption (**EA**) in Theorem 1 and the curvature assumptions (**CA1**)–(**CA8**) in Theorem 2 imply closed-form expressions for $\underline{\Delta CS}$ and $\overline{\Delta CS}$. These closed-form expressions are provided in Table 1.

A.2   Welfare bounds combining elasticity and curvature restrictions

Next, we show how our closed-form expressions for welfare bounds can be extended to incorporate information about price elasticities at p_0 and p_1—in addition to the curvature assumptions (**CA1**)–(**CA8**)—which might sometimes be available from treatment effects estimates.

Formally, denote by $\mathcal{D}_i(\varepsilon_0, \varepsilon_1)$ the family of consistent demand curves that, in addition to satisfying (**CA**$_i$), have elasticities at p_0 and p_1 given by $\varepsilon(p_0) = \varepsilon_0$ and $\varepsilon(p_1) = \varepsilon_1$. We thus derive

$$
\begin{aligned}
&\text{upper bound:} \quad \max_{D(\cdot) \in \mathcal{D}_i(\varepsilon_0, \varepsilon_1)} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p, \\
&\text{lower bound:} \quad \min_{D(\cdot) \in \mathcal{D}_i(\varepsilon_0, \varepsilon_1)} \int_{p_0}^{p_1} D(p) \, \mathrm{d}p.
\end{aligned}
\tag{$\dagger$}
$$

Each family $\mathcal{D}_i(\varepsilon_0, \varepsilon_1)$ is clearly more restrictive than its corresponding $\mathcal{D}_i$. Thus, the welfare bounds with elasticity information are narrower than those without.

The problem of deriving welfare bounds with elasticity information and curvature assumptions ($\dagger$) is a generalization of the problem of deriving welfare bounds with only curvature assumptions considered in Section 2.3. However, elasticity information can be straightforwardly incorporated into the proof of Theorem 2; Theorem 2 continues to hold in this setting. In turn, Theorem 2 implies closed-form expressions for welfare bounds, which are provided in Table 2.

Demand family	Lower bound	Upper bound
$\mathcal{D}$ (elasticity restrictions)	$\dfrac{p_0 q_0 \left[\left(\frac{q_0 p_1^{\overline{\varepsilon}}}{q_1 p_0^{\overline{\varepsilon}}} \right)^{\frac{1+\underline{\varepsilon}}{\overline{\varepsilon}-\underline{\varepsilon}}} - 1 \right]}{1 + \underline{\varepsilon}} + \dfrac{p_1 q_1 \left[1 - \left(\frac{q_0 p_1^{\underline{\varepsilon}}}{q_1 p_0^{\underline{\varepsilon}}} \right)^{\frac{1-\overline{\varepsilon}}{\overline{\varepsilon}-\underline{\varepsilon}}} \right]}{1 - \overline{\varepsilon}}$	$\dfrac{p_0 q_0 \left[\left(\frac{q_1 p_0^{\underline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}} \right)^{\frac{1+\overline{\varepsilon}}{\overline{\varepsilon}-\underline{\varepsilon}}} - 1 \right]}{1 + \overline{\varepsilon}} + \dfrac{p_1 q_1 \left[1 - \left(\frac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\overline{\varepsilon}}} \right)^{\frac{1-\underline{\varepsilon}}{\overline{\varepsilon}-\underline{\varepsilon}}} \right]}{1 - \underline{\varepsilon}}$
$\mathcal{D}_1$ (Marshall's second law)	$\dfrac{(p_1 q_1 - p_0 q_0) \log(p_1/p_0)}{\log(q_1/q_0) + \log(p_1/p_0)}$	$q_0 (p_1 - p_0)$
$\mathcal{D}_2$ (decreasing MR)	$\dfrac{(p_1 - p_0) q_0 q_1 \log(q_0/q_1)}{q_0 - q_1}$	$q_0 (p_1 - p_0)$
$\mathcal{D}_3$ (log-concave demand)	$\dfrac{(p_1 - p_0)(q_0 - q_1)}{\log(q_0/q_1)}$	$q_0 (p_1 - p_0)$
$\mathcal{D}_4$ (concave demand)	$\dfrac{(p_1 - p_0)(q_1 + q_0)}{2}$	$q_0 (p_1 - p_0)$
$\mathcal{D}_5$ (ρ-concave demand)	$\dfrac{\rho (p_1 - p_0) \left(q_0^{1+\rho} - q_1^{1+\rho} \right)}{(\rho + 1)(q_0^{\rho} - q_1^{\rho})}$	$q_0 (p_1 - p_0)$
$\mathcal{D}_6$ (convex demand)	$q_1 (p_1 - p_0)$	$\dfrac{(p_1 - p_0)(q_1 + q_0)}{2}$
$\mathcal{D}_7$ (log-convex demand)	$q_1 (p_1 - p_0)$	$\dfrac{(p_1 - p_0)(q_0 - q_1)}{\log(q_0/q_1)}$
$\mathcal{D}_8$ (ρ-convex demand)	$q_1 (p_1 - p_0)$	$\dfrac{\rho (p_1 - p_0) \left(q_0^{1+\rho} - q_1^{1+\rho} \right)}{(\rho + 1)(q_0^{\rho} - q_1^{\rho})}$

Table 1: Welfare bounds for ΔCS implied by Theorems 1 and 2.

$\mathcal{D}_i(\varepsilon_0,\varepsilon_1)$	**Lower bound**	**Upper bound**	p_*	q_*
$\mathcal{D}_1(\varepsilon_0,\varepsilon_1)$	$\underbrace{\dfrac{(p_1q_1 - p_0q_0)\log(p_1/p_0)}{\log(q_1/q_0) + \log(p_1/p_0)}}_{\equiv L_1(p_0,p_1,q_0,q_1)}$	$L_1(p_0,p_*,q_0,q_*) +$ $L_1(p_*,p_1,q_*,q_1)$	$\left(\dfrac{q_0 p_0^{-\varepsilon_0}}{q_1 p_1^{-\varepsilon_1}}\right)^{\frac{1}{\varepsilon_1-\varepsilon_0}}$	$q_0 p_0^{-\varepsilon_0} p_*^{\varepsilon_0}$
$\mathcal{D}_2(\varepsilon_0,\varepsilon_1)$	$\underbrace{\dfrac{(p_1 - p_0)\,q_0 q_1 \log(q_0/q_1)}{q_0 - q_1}}_{\equiv L_2(p_0,p_1,q_0,q_1)}$	$L_2(p_0,p_*,q_0,q_*) +$ $L_2(p_*,p_1,q_*,q_1)$	$p_0 p_1 \times$ $\dfrac{q_0(1+\varepsilon_1) - q_1(1+\varepsilon_0)}{p_0 q_0 \varepsilon_1 - p_1 q_1 \varepsilon_0}$	$\dfrac{q_0}{1+\varepsilon_0 - \frac{\varepsilon_0 p_*}{p_0}}$
$\mathcal{D}_3(\varepsilon_0,\varepsilon_1)$	$\underbrace{\dfrac{(p_1 - p_0)(q_0 - q_1)}{\log(q_0/q_1)}}_{\equiv L_3(p_0,p_1,q_0,q_1)}$	$L_3(p_0,p_*,q_0,q_*) +$ $L_3(p_*,p_1,q_*,q_1)$	$\left[\varepsilon_1 - \varepsilon_0 + \log\left(\dfrac{q_0}{q_1}\right)\right] \times$ $\dfrac{p_1\varepsilon_0 - p_0\varepsilon_1}{\varepsilon_1 - \varepsilon_0}$	$q_0 e^{\varepsilon_0\left(\frac{p_*}{p_0}-1\right)}$
$\mathcal{D}_4(\varepsilon_0,\varepsilon_1)$	$\underbrace{\dfrac{(p_1 - p_0)(q_1 + q_0)}{2}}_{\equiv L_4(p_0,p_1,q_0,q_1)}$	$L_4(p_0,p_*,q_0,q_*) +$ $L_4(p_*,p_1,q_*,q_1)$	$p_0 p_1 \times$ $\dfrac{q_0(1-\varepsilon_0) - q_1(1-\varepsilon_1)}{p_0 q_1 \varepsilon_1 - p_1 q_0 \varepsilon_0}$	$q_0(1-\varepsilon_0) +$ $\dfrac{\varepsilon_0 q_0}{p_0} p_*$
$\mathcal{D}_5(\varepsilon_0,\varepsilon_1)$	$\underbrace{\dfrac{\rho(p_1 - p_0)\left(q_0^{1+\rho} - q_1^{1+\rho}\right)}{(\rho+1)(q_0^\rho - q_1^\rho)}}_{\equiv L_5(p_0,p_1,q_0,q_1)}$	$L_5(p_0,p_*,q_0,q_*) +$ $L_5(p_*,p_1,q_*,q_1)$	$p_0 p_1 \times$ $\dfrac{q_0^\rho(1-\rho\varepsilon_0) - q_1^\rho(1-\rho\varepsilon_1)}{\rho(p_0 q_1^\rho \varepsilon_1 - p_1 q_0^\rho \varepsilon_0)}$	$q_0 \times$ $\left(1 - \rho\varepsilon_0 + \dfrac{\rho\varepsilon_0 p_*}{p_0}\right)^{\frac{1}{\rho}}$
$\mathcal{D}_6(\varepsilon_0,\varepsilon_1)$	$U_6(p_0,p_*,q_0,q_*) +$ $U_6(p_*,p_1,q_*,q_1)$	$\underbrace{\dfrac{(p_1 - p_0)(q_1 + q_0)}{2}}_{\equiv U_6(p_0,p_1,q_0,q_1)}$	$p_0 p_1 \times$ $\dfrac{q_0(1-\varepsilon_0) - q_1(1-\varepsilon_1)}{p_0 q_1 \varepsilon_1 - p_1 q_0 \varepsilon_0}$	$q_0(1-\varepsilon_0) +$ $\dfrac{\varepsilon_0 q_0}{p_0} p_*$
$\mathcal{D}_7(\varepsilon_0,\varepsilon_1)$	$U_7(p_0,p_*,q_0,q_*) +$ $U_7(p_*,p_1,q_*,q_1)$	$\underbrace{\dfrac{(p_1 - p_0)(q_0 - q_1)}{\log(q_0/q_1)}}_{\equiv U_7(p_0,p_1,q_0,q_1)}$	$\left[\varepsilon_1 - \varepsilon_0 + \log\left(\dfrac{q_0}{q_1}\right)\right] \times$ $\dfrac{p_1\varepsilon_0 - p_0\varepsilon_1}{\varepsilon_1 - \varepsilon_0}$	$q_0 e^{\varepsilon_0\left(\frac{p_*}{p_0}-1\right)}$
$\mathcal{D}_8(\varepsilon_0,\varepsilon_1)$	$U_8(p_0,p_*,q_0,q_*) +$ $U_8(p_*,p_1,q_*,q_1)$	$\underbrace{\dfrac{\rho(p_1 - p_0)\left(q_0^{1+\rho} - q_1^{1+\rho}\right)}{(\rho+1)(q_0^\rho - q_1^\rho)}}_{\equiv U_8(p_0,p_1,q_0,q_1)}$	$p_0 p_1 \times$ $\dfrac{q_0^\rho(1-\rho\varepsilon_0) - q_1^\rho(1-\rho\varepsilon_1)}{\rho(p_0 q_1^\rho \varepsilon_1 - p_1 q_0^\rho \varepsilon_0)}$	$q_0 \times$ $\left(1 - \rho\varepsilon_0 + \dfrac{\rho\varepsilon_0 p_*}{p_0}\right)^{\frac{1}{\rho}}$

Table 2: *Welfare bounds for* $\Delta\mathrm{CS}$ *for each* $\mathcal{D}_i(\varepsilon_0,\varepsilon_1)$ *implied by Theorem 2.*

Appendix B   Alternative proof of Theorem 1

In this appendix, we provide an alternative proof of Theorem 1 that differs from the geometric proof presented in the paper.

B.1   Proof of Theorem 1

This alternative proof highlights a connection between our problem (2) and Bayesian persuasion problems that have been considered by the theoretical literature stemming from Kamenica and Gentzkow (2011). We divide the proof into three steps: *(i)* employing a change of variables to map the problem into an appropriate functional space; *(ii)* endowing this space with a partial order and characterizing its extremal functions; and *(iii)* mapping the solution back to the original problem.

Step #1: Changing variables

We begin by employing a change of variables. Instead of choosing a demand curve to maximize or minimize ΔCS as in problem (2), we choose the elasticity function $\eta(\cdot)$ *expressed as a function of log-price*, rather than price:

$$\eta(\pi) := \varepsilon(e^\pi).$$

Given $\eta(\cdot)$, the demand curve $D(\cdot)$ is completely determined, and vice versa:

$$\eta(\pi) = \frac{e^\pi D'(e^\pi)}{D(e^\pi)} \iff D(p) = q_0 \exp\left[\int_{\log p_0}^{\log p} \eta(\pi)\, d\pi\right] \quad \text{for any } p \in [p_0, p_1]. \tag{3}$$

Analogous to the family of demand curves $\mathcal{D}$, we define the set of elasticity functions that are consistent with (**A2**) and (**EA**):

$$\mathcal{E} := \left\{\eta : [\log p_0, \log p_1] \to [\underline{\varepsilon}, \overline{\varepsilon}] \text{ s.t. } \int_{\log p_0}^{\log p_1} \eta(\pi)\, d\pi = \log\left(\frac{q_1}{q_0}\right)\right\}.$$

Thus we arrive at the equivalent problem:

$$\begin{cases} \overline{\Delta\text{CS}} = q_0 \cdot \max_{\eta \in \mathcal{E}} \int_{p_0}^{p_1} \exp\left[\int_{\log p_0}^{\log p} \eta(\pi)\, d\pi\right]\, dp, \\[2mm] \underline{\Delta\text{CS}} = q_0 \cdot \min_{\eta \in \mathcal{E}} \int_{p_0}^{p_1} \exp\left[\int_{\log p_0}^{\log p} \eta(\pi)\, d\pi\right]\, dp. \end{cases} \tag{4}$$

Step #2: Characterizing the set $\mathcal{E}$

We now endow the set $\mathcal{E}$ with a partial order. Formally, for any two functions $\eta_1, \eta_2 \in \mathcal{E}$, we write

$$\eta_1 \succeq \eta_2 \iff \int_{\log p_0}^{\log p} \eta_1(\pi) \, \mathrm{d}\pi \geq \int_{\log p_0}^{\log p} \eta_2(\pi) \, \mathrm{d}\pi \quad \text{for every } p \in [p_0, p_1].$$

This partial order is motivated by the definition of second-order stochastic dominance, but with a few differences: η is not necessarily a monotone function, nor is $\eta(\log p_0)$ or $\eta(\log p_1)$ fixed. For these reasons, η cannot be interpreted as a cumulative distribution function (CDF), making the above definition slightly different from second-order stochastic dominance.

Nevertheless, a familiar mathematical property of second-order stochastic dominance holds in this environment. Just as the second-order stochastic dominance order defines a lattice structure on the set of all CDFs with the same mean, the partial order $\succeq$ defines a lattice structure on the set $\mathcal{E}$.

Lemma 1. Any function $\eta \in \mathcal{E}$ satisfies $\eta^* \succeq \eta \succeq \eta_*$, where

$$\eta^*(\pi) := \begin{cases} \underline{\varepsilon} & \text{if } \pi > \dfrac{1}{\overline{\varepsilon} - \underline{\varepsilon}} \cdot \log\left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right), \\[2mm] \overline{\varepsilon} & \text{if } \pi \leq \dfrac{1}{\overline{\varepsilon} - \underline{\varepsilon}} \cdot \log\left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right), \end{cases} \quad \text{and} \quad \eta_*(\pi) := \begin{cases} \overline{\varepsilon} & \text{if } \pi > \dfrac{1}{\overline{\varepsilon} - \underline{\varepsilon}} \cdot \log\left(\dfrac{q_0 p_1^{\overline{\varepsilon}}}{q_1 p_0^{\underline{\varepsilon}}}\right), \\[2mm] \underline{\varepsilon} & \text{if } \pi \leq \dfrac{1}{\overline{\varepsilon} - \underline{\varepsilon}} \cdot \log\left(\dfrac{q_0 p_1^{\overline{\varepsilon}}}{q_1 p_0^{\underline{\varepsilon}}}\right). \end{cases}$$

Proof. To see that $\eta^* \succeq \eta$ for any $\eta \in \mathcal{E}$, observe that

$$\int_{\log p_0}^{\log p} \eta^*(\pi) \, \mathrm{d}\pi = \begin{cases} \log\left(\dfrac{p}{p_0}\right) \cdot \overline{\varepsilon} & \geq \int_{\log p_0}^{\log p} \eta(\pi) \, \mathrm{d}\pi \quad \text{for any } p_0 \leq p \leq \left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon} - \underline{\varepsilon}}}, \\[3mm] \log\left(\dfrac{q_1}{q_0}\right) - \log\left(\dfrac{p_1}{p}\right) \cdot \underline{\varepsilon} & \geq \int_{\log p_0}^{\log p} \eta(\pi) \, \mathrm{d}\pi \quad \text{for any } \left(\dfrac{q_1 p_0^{\overline{\varepsilon}}}{q_0 p_1^{\underline{\varepsilon}}}\right)^{\frac{1}{\overline{\varepsilon} - \underline{\varepsilon}}} < p \leq p_1. \end{cases}$$

The inequalities follow from the fact that $\operatorname{im} \eta \subset [\underline{\eta}, \overline{\eta}]$ for any $\eta \in \mathcal{E}$. A similar argument shows that $\eta \succeq \eta_*$ for any $\eta \in \mathcal{E}$. $\qquad\square$

It is easy to check that $\eta^*, \eta_* \in \mathcal{E}$. Therefore, Lemma 1 characterizes the largest and smallest elements of the partially ordered set $(\mathcal{E}, \succeq)$. With more work, one can show that $(\mathcal{E}, \succeq)$ is a lattice (cf. Theorem 3.3 of Müller and Scarsini, 2006); however, as the lattice property is not important for our purposes, we do not pursue that here.

Step #3: Mapping back to the original problem

Having characterized the largest and smallest elements of $(\mathcal{E}, \succeq)$, it remains to map these back to the original problem. To this end, we define the functional $\Delta\mathrm{CS} : \mathcal{E} \to \mathbb{R}$ by

$$\Delta\mathrm{CS}(\eta) := q_0 \int_{p_0}^{p_1} \exp\left[\int_{\log p_0}^{\log p} \eta(\pi)\, \mathrm{d}\pi\right] \mathrm{d}p.$$

Our problem (4) is equivalent to maximizing and minimizing this functional over the family $\mathcal{E}$. The following lemma shows that this can be done with the aid of the partial order $\succeq$ defined in our previous step:

Lemma 2. The functional $\Delta\mathrm{CS}(\cdot)$ is increasing in the partial order $\succeq$:

$$q_0 \int_{p_0}^{p_1} \exp\left[\int_{\log p_0}^{\log p} \eta_1(\pi)\, \mathrm{d}\pi\right] \mathrm{d}p \geq q_0 \int_{p_0}^{p_1} \exp\left[\int_{\log p_0}^{\log p} \eta_2(\pi)\, \mathrm{d}\pi\right] \mathrm{d}p \quad \text{for any } \eta_1 \succeq \eta_2.$$

Proof. Since $\eta_1 \succeq \eta_2$, it follows from the monotonicity of the exponential function that

$$\exp\left[\int_{\log p_0}^{\log p} \eta_1(\pi)\, \mathrm{d}\pi\right] \geq \exp\left[\int_{\log p_0}^{\log p} \eta_1(\pi)\, \mathrm{d}\pi\right] \quad \text{for any } p_0 \leq p \leq p_1.$$

The result thus follows from a pointwise comparison of the two integrands. $\square$

Together, Lemmas 1 and 2 imply that:

Proposition 3. The functional $\Delta\mathrm{CS}(\cdot)$ is maximized at η^* and minimized at η_*:

$$\overline{\Delta\mathrm{CS}} = \Delta\mathrm{CS}(\eta^*) \quad \text{and} \quad \underline{\Delta\mathrm{CS}} = \Delta\mathrm{CS}(\eta_*).$$

To complete the alternative proof of Theorem 1, it remains to show that η^* and η_* correspond to the demand curves D^* and D_* respectively (defined in Section 2) via the relation (3). This is readily verified by straightforward computation and omitted for the sake of brevity.

B.2   Discussion

We conclude this appendix with a few remarks on how this alternative proof compares with the proof of Theorem 1 presented in the paper, and with some notes on its connections to similar problems in the information design literature.

As noted in Section 2, while this proof is more complex than our geometric proof, it has the advantage of being easily generalizable. Although additional constraints on the demand family—such as the ones that we consider in Section 3—might not have a simple geometric interpretation, they can be accommodated as constraints on the functional space $\mathcal{E}$. Notice also that Lemma 2 does not depend on how $\mathcal{E}$ is defined. Therefore, different constraints on $\mathcal{E}$ only require determining the analog of Lemma 1 for the constrained problem—that is, finding the largest and smallest elements of the partially ordered set $(\mathcal{E}, \succeq)$.

It is worth pointing out that the structure of $\mathcal{E}$ is reminiscent of Bayesian persuasion problems stemming from the work of Kamenica and Gentzkow (2011). If $-\eta$ could be interpreted as a posterior belief, then the mean constraint

$$\int_{\log p_0}^{\log p_1} \eta(\pi) \, d\pi = \log\left(\frac{q_1}{q_0}\right)$$

could be interpreted as a Bayes plausibility constraint, where $-\log(q_1/q_0)$ is the mean of the prior belief. This analogy breaks down for the sole reason that $-\eta$ *cannot* be interpreted as a posterior belief: $-\eta$ is not monotone and hence cannot be a cumulative distribution function.

Yet this observation also indicates that there is an *exact* equivalence between such Bayesian persuasion problems and an extension we consider in Section 2.3, rather than our basic model. Precisely, the analogy holds when we instead consider the problem of finding welfare bounds under Marshall's second law in addition to (**A2**) and (**EA**). Marshall's second law implies that η must be decreasing; hence $-\eta$ is increasing and, with appropriate rescaling, can be interpreted as a cumulative distribution function representing the posterior belief.
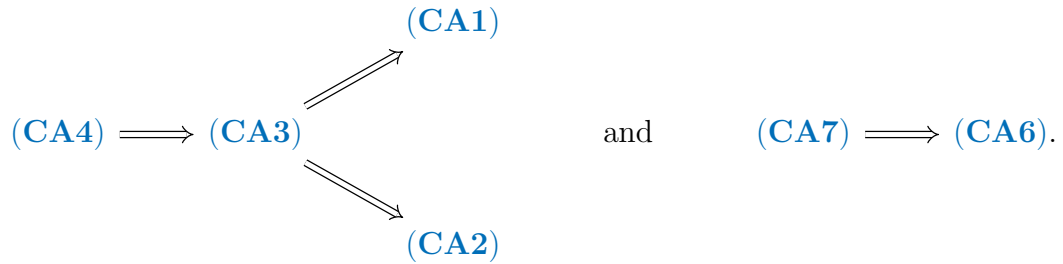
The fortuitous connection between our problem of bounding welfare in different families of demand and Bayesian persuasion problems implies that tools developed for *constrained* information design problems can potentially also be used to evaluate robust welfare bounds. From a technical point of view, our approach (in the alternative proof presented above) is based on the proof strategy of Kang and Vondrák (2019), who solve an infinite-dimensional optimization problem by showing that the objective functional is monotone with respect to the convex partial order. For the convex partial order in particular, Kleiner, Moldovanu, and Strack (2021) recently develop an approach based on a characterization of extreme points, which yields a general solution to similar problems—even when the objective function is not monotone with respect to the convex partial order. While their method can be applied to many problems in information and mechanism design, our discussion here suggests potential applications also to robust welfare bounds.

Appendix C Assumptions on the curvature of demand

In this appendix, we demonstrate the relationship between the different assumptions (**CA1**)–(**CA8**) and review some common demand curves that satisfy these assumptions.

C.1 Relationship between assumptions

We begin by showing that

$$(\mathbf{CA4}) \implies (\mathbf{CA3}) \nearrow \overset{(\mathbf{CA1})}{\underset{(\mathbf{CA2})}{}} \searrow \qquad \text{and} \qquad (\mathbf{CA7}) \implies (\mathbf{CA6}).$$

(CA4) $\implies$ (CA3)

Proof. Given a concave demand curve $D(\cdot)$, suppose on the contrary that there exist $p_H > p_L$ such that

$$\frac{D'(p_H)}{D(p_H)} > \frac{D'(p_L)}{D(p_L)} \implies D(p_L)D'(p_H) > D(p_H)D'(p_L).$$

Since $D(\cdot)$ is concave, $D'(p_H) \leq D'(p_L)$; since $D(\cdot)$ is decreasing, $D'(\cdot) \leq 0$ and $D(p_L) \geq D(p_H)$. Thus

$$D(p_L)D'(p_H) \leq D(p_H)D'(p_H) \leq D(p_H)D'(p_L).$$

This is a contradiction. Hence $D(\cdot)$ is log-concave. $\qquad\square$

(CA3) $\implies$ (CA1)

Proof. For any $p_H > p_L$, log-concavity implies that

$$\frac{D'(p_H)}{D(p_H)} \leq \frac{D'(p_L)}{D(p_L)} \implies \frac{p_H D'(p_H)}{D(p_H)} \leq \frac{p_L D'(p_H)}{D(p_H)} \leq \frac{p_L D'(p_L)}{D(p_L)}.$$

Here, we have used the fact that $D'(\cdot) \leq 0$ as $D(\cdot)$ is decreasing. Since the above inequalities hold for any $p_H > p_L$, it follows that $D(\cdot)$ satisfies Marshall's second law. $\qquad\square$

(CA3) $\implies$ (CA2)

Proof. For any $p_H > p_L$, log-concavity implies that

$$\frac{D'(p_H)}{D(p_H)} \leq \frac{D'(p_L)}{D(p_L)} \implies p_H + \frac{D(p_H)}{D'(p_H)} \geq p_L + \frac{D(p_L)}{D'(p_L)}.$$

Since this holds for any $p_H > p_L$, it follows that $D(\cdot)$ has a decreasing marginal revenue curve. $\square$

(CA7) $\implies$ (CA6)

Proof. For any $p_H > p_L$, log-convexity implies that

$$\frac{D'(p_H)}{D(p_H)} \geq \frac{D'(p_L)}{D(p_L)} \implies D(p_L)D'(p_H) \geq D(p_H)D'(p_L).$$

Since $D(\cdot)$ is decreasing, $D'(\cdot) \leq 0$ and $D(p_L) \geq D(p_H)$. Thus

$$D(p_H)D'(p_H) \geq D(p_L)D'(p_H) \geq D(p_H)D'(p_L) \implies D'(p_H) \geq D'(p_L).$$

Since this holds for any $p_H > p_L$, it follows that $D(\cdot)$ is convex. $\square$

C.2 Common demand curves

We now review some common demand curves that satisfy these assumptions. These demand curves play a crucial role in our analysis in Section 2.3 (cf. Theorem 2).

(i) *CES demand curves.* Each CES demand curve is parametrized by its elasticity $\varepsilon \leq 0$:

$$D(p) = q_0 \left(\frac{p}{p_0}\right)^\varepsilon.$$

Because elasticity is constant, it must also be trivially decreasing. Hence any CES demand curve satisfies Marshall's second law **(CA1)**.

(ii) *Constant marginal revenue demand curve.* Analogous to a CES demand curve, each constant marginal revenue demand curve is parametrized by its marginal revenue $0 \leq \mu < p_0$:

$$D(p) = \frac{q_0 (p_0 - \mu)}{p - \mu}.$$

Because marginal revenue is constant, it must also be trivially decreasing. Hence each constant marginal revenue demand curve exhibits decreasing marginal revenue (**CA2**).

(iii) Exponential demand curves. Each exponential demand curve is parametrized by $\lambda \geq 0$:

$$D(p) = q_0 \exp\left[-\lambda\left(p - p_0\right)\right].$$

Observe that the logarithm of any exponential demand curve is linear in p:

$$\log D(p) = \log q_0 - \lambda\left(p - p_0\right).$$

Hence each exponential demand curve is both log-concave (**CA3**) and log-convex (**CA7**).

(iv) Linear demand curves. Each linear demand curve is parametrized by $\lambda \geq 0$:

$$D(p) = q_0 - \lambda\left(p - p_0\right).$$

Each linear demand curve is both concave (**CA4**) and convex (**CA6**).

(v) ρ-linear demand curves. Each ρ-linear demand curve is parametrized by $\lambda \geq 0$:

$$D(p) = \left[q_0 - \lambda\left(p - p_0\right)\right]^{1/\rho}.$$

Each ρ-linear demand curve is both ρ-concave (**CA5**) and ρ-convex (**CA8**).

Appendix D   Other proofs and additional discussion

This appendix collects all other omitted proofs and includes some additional discussion.

D.1   Proof of Theorem 2

In this section, we present the geometric proof of Theorem 2 for the case of Marshall's second law (**CA1**). For brevity, we omit the proofs for the other curvature assumptions (**CA2**)–(**CA8**) as they are similar.
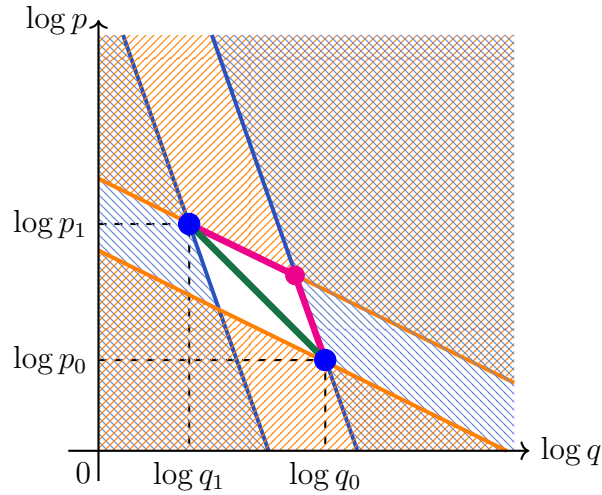

Figure D.1: Sketch of the proof of Theorem 2.

Figure D.1 summarizes our geometric proof. To begin, observe that the upper bound (in red) remains unchanged from Theorem 1, except that—in the absence of (**A2**)—we set $\underline{\varepsilon} = -\infty$ and $\bar{\varepsilon} = 0$. That is, because the upper bound from Theorem 1 satisfies Marshall's second law (**CA1**), it remains the upper bound among all demand curves that satisfy Marshall's second law (**CA1**).

It remains to show that the lower bound (in green) is attained by linearly interpolating between the two points on the log-price–log-quantity plot. To this end, notice that Marshall's second law (**CA1**) implies that log-price is *concave* in log-quantity. The pointwise smallest concave, decreasing curve that passes through the two points is precisely the straight line that connects them. This corresponds to a CES demand curve with elasticity equal to the average elasticity implied by the two points:

$$\varepsilon = \frac{\log(q_1/q_0)}{\log(p_1/p_0)}.$$

This concludes our proof of Theorem 2.

D.2 Additional discussion for Theorem 2

While Theorem 2 states our bounds by *replacing* (**EA**) with curvature assumptions, we can also derive bounds by imposing *both* (**EA**) *and* each curvature assumption. This is the case, for example, in some empirical applications that we consider in Section 4, where the researcher might observe one or both elasticities and is also willing to make additional curvature assumptions, such as Marshall's second law (**CA1**).

To give a concrete example, consider the problem of deriving bounds on ΔCS under the assumptions (**A2**), (**EA**), and (**CA1**). It can be readily verified that the geometric proof of Theorem 2 supplied in Appendix D.1 continues to hold: the upper bound is attained by a 2-piece CES interpolation, whereas the lower bound is attained by a 1-piece CES interpolation.

Things are not as straightforward, however, if we were to impose a slightly different set of assumptions, such as (**A2**), (**EA**), and (**CA2**). Crucially, the geometric picture (*e.g.*, Figure D.1) no longer captures all the relevant binding constraints except for special choices of $\underline{\varepsilon}$ and $\overline{\varepsilon}$. Instead, to derive bounds, one would have to formalize the problem using the approach of Appendix B, and then solve the resulting optimization problem.

D.3 Proof of Proposition 2

We prove Proposition 2 for the case of Marshall's second law (**CA1**); for brevity, we omit the proofs for the other curvature assumptions (**CA2**)–(**CA8**) as they are similar. To this end, we adapt the proof of Theorem 1 in Appendix B by replacing (**EA**) with (**CA1**). We express the elasticity function $\eta(\cdot)$ as a function of log-price, rather than price:

$$\eta(\pi) := \varepsilon(e^{\pi}).$$

We also define the set of elasticity functions that are consistent with (**A2**) and (**CA1**):

$$\mathcal{E} := \left\{ \eta : [\log p_0, \log p_1] \to [\underline{\varepsilon}, \overline{\varepsilon}] \text{ is decreasing and } \int_{\log p_0}^{\log p_1} \eta(\pi) \, d\pi = \log \left(\frac{q_1}{q_0} \right) \right\}.$$

Clearly, $\mathcal{E}$ is convex, and as a convex set (in the real vector space L^1), $\mathcal{E}$ is connected. Now, define the functional ΔCS $: \mathcal{E} \to \mathbb{R}$ by

$$\Delta\text{CS}(\eta) := q_0 \int_{p_0}^{p_1} \exp \left[\int_{\log p_0}^{\log p} \eta(\pi) \, d\pi \right] \, dp.$$

It is routine to check that $\Delta\mathrm{CS}$ is a continuous map (in the L^1 topology) from $\mathcal{E}$ to $\mathbb{R}$; by continuity, its image must be connected in $\mathbb{R}$ and hence an interval. Thus the identified set $[\underline{\Delta\mathrm{CS}}_1, \overline{\Delta\mathrm{CS}}_1]$ is sharp.

D.4 Proof of Theorem 5

To prove Theorem 5, we build on our proof of Theorem 2. We begin by *fixing* the elasticities at the points $(p_0, q_0), \ldots, (p_{n-1}, q_{n-1})$ at $\varepsilon_0 < \cdots < \varepsilon_{n-1}$. Our proof of Theorem 2 then implies that the upper bound between any two adjacent points is a 2-piece CES interpolation, whereas the lower bound between any two adjacent points is a 1-piece CES interpolation. This immediately implies that the smallest possible change in consumer surplus between (p_0, q_0) and (p_{n-1}, q_{n-1}) is attained by an $(n-1)$-point CES interpolation; hence it remains to prove the corresponding statement for the largest possible change in consumer surplus.

To this end, we make the additional observation that the 2-piece CES interpolation that yields the upper bound between any two adjacent points must have elasticities equal to the elasticities *at* the two points. That is, for any $i = 1, \ldots, n-2$, the 2-piece CES interpolation that attains the largest change in consumer surplus between (p_i, q_i) and (p_{i+1}, q_{i+1}) consists of a piece with elasticity ε_i and another piece with elasticity ε_{i+1}. Thus, although the statement of Theorem 2 indicates that (at most) $2n - 2$ pieces are required, this argument shows that only n pieces are required as adjacent pieces that join at the points $(p_1, q_1), \ldots, (p_{n-2}, q_{n-2})$ are actually part of the same CES demand curve. Since this argument holds for any $\varepsilon_0 < \cdots < \varepsilon_1$, the largest possible change in consumer surplus must be attained by an n-piece CES interpolation, as claimed.

D.5 Technical details for Sections 4.2 and 4.3

For each application in Section 4, we aimed to present the clearest discussion of the paper whose exercise we followed with a focus on how our bounding approach applies. To this effect, we followed the data work and framing that the authors of the original papers used as much as possible. In several cases, the original paper's exercise included a few setting-specific subtleties that we felt would distract from our main point. These subtleties introduce small discrepancies between the bounds we obtain and the welfare estimates that they report. For readers who wish to compare our numbers directly to the numbers in the original papers, we include a brief description of these discrepancies below.

D.5.1 Technical details for Section 4.2

In order to compute our bounds for the example in Section 4.2, we follow Amiti et al.'s (2019) data appendix to obtain a comprehensive dataset of products hit by new tariffs during 2018. Products are denoted by a ten-digit Harmonized Tariff Schedule (HTS10) product code and by country or origin. The dataset contains a unit quantity and total import value for each product, along with a tariff amount for each month. Unit prices are derived by dividing the total import value by the unit quantity. For each product and country, we compare the unit price in each month during 2018 against the unit price of the same product/country in the same month during 2017.

Amiti et al. take several further standardization steps to this effect. First, they impute 2018 prices using the regression of log import quantity changes against log tariff changes described in column (3) of their Table 1. Although this step is not necessary for our exercise, we follow it in order to make our results most comparable to their exercise. Second, as they explain in footnote 9 (pp. 199–200), they make use of a second Taylor approximation in computing deadweight loss:

$$-\log(m_1/m_0) \approx (m_0 - m_1)/m_1,$$

where m_t is the total import value of a product in year t. In general, it can be shown that this approximation will underestimate deadweight loss:

$$-\log z \leq \frac{1}{z} - 1 \quad \text{for any } z \in \mathbb{R}.$$

As the magnitudes of the tariffs are substantial, we find that this approximation shrinks the deadweight loss estimates substantially and makes the comparison across assumptions more difficult to interpret. As such, we skip this approximation step in our calculations and instead present the deadweight loss estimates from linear (and other) interpolations using just the quantities and prices produced in their first step.

D.5.2 Technical details for Section 4.3

Giesecke and Jäger (2021) apply two separate strategies to predict the change in retirement following the OPA. Their main exercise applies a regression discontinuity design around the cutoff age for pensions (70 years old) and finds a marginal change of 6%. Giesecke and Jäger report that the total number of eligible 70-year-old workers at the time of the OPA was 140, 288—8, 459 of whom were thus marginal. In addition, another 928, 198 workers above the

age of 70 years old were also eligible at the same time. For this group, Giesecke and Jäger assume that the linear trend of retirement would have continued from the 65–69 age group absent the OPA. Comparing this hypothetical retirement rate against the realized retirement rate for workers above the age of 70 years old after the OPA, they estimate a 5.1% increase in retirement rate. To compute the total change in welfare, we follow Giesecke and Jäger in combining these two groups. The total number of eligible workers is thus $140,288 + 928,198 = 1,068,486$; the total number of marginal workers is $8,459 + 47,909 = 56,368$; and the total number of retired workers following the OPA is $613,873$. We use the fractions $\frac{557,505}{1,068,486} \approx 0.52$ and $\frac{613,873}{1,068,486} \approx 0.58$ as the average propensities to retire before and after the OPA, respectively.

To supplement our bounds given curvature assumptions based on Theorem 2 in Figure 10(a), we presented bounds given the standard deviation of w_i in Figure 10(b). Formally, these bounds solve:

$$\begin{cases} \overline{\Delta W}(\sigma) := \max_{F \in \mathcal{F}} \left\{ \int_{p_0}^{p_1} F(p) \, \mathrm{d}p : \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) - \left[\int_{p_0}^{p_1} p \, \mathrm{d}F(p) \right]^2 = \sigma^2 \right\}, \\ \underline{\Delta W}(\sigma) := \min_{F \in \mathcal{F}} \left\{ \int_{p_0}^{p_1} F(p) \, \mathrm{d}p : \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) - \left[\int_{p_0}^{p_1} p \, \mathrm{d}F(p) \right]^2 = \sigma^2 \right\}, \end{cases}$$

where

$$\mathcal{F} := \left\{ F : [p_0, p_1] \to [0, 1] \text{ is increasing and satisfies } F(p_t) = q_t \text{ for } t = 0, 1 \right\}.$$

Integrating by parts, we get

$$\mu := \int_{p_0}^{p_1} p \, \mathrm{d}F(p) = p_1 q_1 - p_0 q_0 - \int_{p_0}^{p_1} F(p) \, \mathrm{d}p.$$

Thus our problem is equivalent to

$$\begin{cases} \overline{\Delta W}(\sigma) := \max_{F \in \mathcal{F}} \left\{ p_1 q_1 - p_0 q_0 - \mu : \int_{p_0}^{p_1} p \, \mathrm{d}F(p) = \mu, \ \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) = \mu^2 + \sigma^2 \right\}, \\ \underline{\Delta W}(\sigma) := \min_{F \in \mathcal{F}} \left\{ p_1 q_1 - p_0 q_0 - \mu : \int_{p_0}^{p_1} p \, \mathrm{d}F(p) = \mu, \ \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) = \mu^2 + \sigma^2 \right\}. \end{cases}$$

We now characterize extremal distributions $F^*(\cdot; \sigma)$ and $F_*(\cdot; \sigma)$ that respectively attain the bounds $\overline{\Delta W}(\sigma)$ and $\underline{\Delta W}(\sigma)$. Removing constants in the objective function that do not depend on F, we see that

$$F^*(\cdot; \sigma) \in \underset{F \in \mathcal{F}}{\arg\min} \left\{ \mu : \int_{p_0}^{p_1} p \, \mathrm{d}F(p) = \mu, \ \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) = \mu^2 + \sigma^2 \right\}.$$

In turn, this can be rewritten as

$$F^*(\cdot;\sigma) \in \underset{F \in \mathcal{F},\, \mu \in [0,1]}{\arg\min} \left\{ \int_{p_0}^{p_1} p^2 \, \mathrm{d}F(p) : \int_{p_0}^{p_1} p \, \mathrm{d}F(p) = \mu \right\}.$$

It can be easily verified that the approach of Appendix B applies to this transformed problem for each feasible μ. In particular, our approach implies that $F^*(\cdot;\sigma)$ consists of two atoms in the interval $[p_0, p_1]$, one at p_0 and p_1. This reduces our *a priori* infinite-dimensional to a unidimensional problem (where the only choice variable is the size of the atom at p_0, as their sum must be $q_1 - q_0$). A similar logic holds for $F_*(\cdot;\sigma)$. We solve these unidimensional problems numerically to obtain the bounds $\overline{\Delta W}(\sigma)$ and $\underline{\Delta W}(\sigma)$ presented in Figure 10(b).