

Robust Bounds for Welfare Analysis*

Zi Yang Kang[†] Shoshana Vasserman[‡]

This version: [January 2022](#)

First version: [September 2021](#)

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 trade tariffs, income taxation, and energy subsidies.

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

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

* This paper is dedicated to Eddie Lazear, whose enthusiasm and support we benefited from until his untimely passing. We are also grateful for enlightening discussions with Nageeb Ali, Susan Athey, Matt Backus, Lanier Benkard, Jeremy Bulow, Bharat Chandar, Rebecca Diamond, Liran Einav, Matt Gentzkow, Bob Hahn, Ben Handel, Ravi Jagadeesan, Chad Jones, Myrto Kalouptsidi, Robin Lee, Neale Mahoney, Rob Metcalfe, Paul Milgrom, Mike Ostrovsky, Steve Redding, Peter Reiss, Andy Skrzypacz, Chad Syverson, Chris Tonetti, Hal Varian, Heidi Williams, and Ali Yurukoglu throughout this project. Noah Jussila and Andy Myers provided excellent research assistance. Finally, we thank Daniel Feenberg for his patient help with the TAXSIM system.

[†] 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. For instance, when consumers have quasilinear utility, the loss in consumer surplus due to a price increase is equal to the area below the demand curve between the two prices.

While this approach yields an exact measurement of the loss in consumer surplus, it supposes that the entire demand curve can be observed—or at least inferred from granular data. In many empirical settings, these conditions are difficult to satisfy. When granular information about the demand curve is not available, practitioners often interpolate between the points along the curve that *are* observed under “standard” functional form assumptions ([Chetty, 2009](#)). These standards—like the logit demand model in industrial organization and the constant elasticity of substitution (CES) demand model in international trade—facilitate tractable estimation in finite samples and foster domain-specific intuitions. But they also entail a loss of generality that is often difficult to interpret: How do common functional form assumptions relate to economic fundamentals? How much do welfare estimates depend on the particular assumptions that are used in modeling demand?

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 *tight*: for each demand family, the corresponding upper and lower bounds are attained by curves that belong to that family.

To fix ideas, consider the simple example of a tax levied on a consumption good. 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




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

evaluate the net impact of the tax, a practitioner faces the problem depicted in Figure 1(a). She observes the points (p_0, q_0) and (p_1, q_1) , but does not observe the demand curve $D(p)$ that connects them. While computing the revenue gains 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.

In order to resolve this, two approaches have been taken in the literature. The first—and most common—is to interpolate between the points (p_0, q_0) and (p_1, q_1) in a parametric way. This approach dates back to Harberger (1954), who assumes a linear interpolation between the observed points, as depicted in Figure 1(b). 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.

An alternative approach, which we take in this paper, is to bound the change in welfare, rather than compute an exact point estimate for it. Closest to our exercise, Varian (1985) derives simple “box” bounds for the loss in consumer surplus associated with a price increase, without any additional assumption on the demand curve. Figure 1(c) demonstrates how these bounds apply to our setting. The 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 .

On the one hand, these box bounds are not only intuitive, but also robust: they hold for any demand curve and do not depend on supply-side assumptions. On the other hand, these extremal demand curves imply unrealistic distributions of consumer willingness to pay: both extremal demand curves are attained only when consumers have either zero or infinite elasticities between the prices p_0 and p_1 , which seems unrealistic. A natural question is whether realistic restrictions on the shape of demand might yield narrower bounds on changes in welfare, without losing the simplicity or generality of the box bounds.

In this paper, we show that the answer is yes. Our basic framework, which we present formally in Section 2, imagines a practitioner who, as before, observes only the points (p_0, q_0) and (p_1, q_1) along the demand curve. However, the practitioner is able to draw on experiments, institutional knowledge, and field expertise to restrict the domain of demand curves to consider. For instance, the practitioner might reason that even though elasticities are unlikely to vary unboundedly (as in Varian’s box bounds), it would be reasonable to restrict attention to demand curves for which elasticities are contained in a finite interval $[\underline{\varepsilon}, \bar{\varepsilon}]$ between p_0 and p_1 . The practitioner 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 practitioner 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 bounds the change in consumer surplus by maximizing and minimizing over all possible demand curves $D(p)$ that pass through the observed points (p_0, q_0) and (p_1, q_1) , such that elasticities between p_0 and p_1 lie within the interval $[\underline{\varepsilon}, \bar{\varepsilon}]$. These bounds are attained at demand curves with a surprisingly simple structure: each extremal demand curve consists of two pieces, such that elasticity throughout each piece is extremal, that is, equal to either $\underline{\varepsilon}$ or $\bar{\varepsilon}$. In the limiting case when $\underline{\varepsilon} = -\infty$ and $\bar{\varepsilon} = 0$, the bounds we obtain are identical to Varian’s box bounds. In general, bounds on possible elasticities restrict the possible changes in consumer surplus that are consistent with the data.

Beyond restrictions on elasticities, practitioners may also wish to impose more theoretically motivated assumptions on consumer demand. This idea also dates back to Marshall, who required—as part of the *definition* of a demand curve—that demand is more elastic at higher prices. This assumption, now known as Marshall’s second law, is maintained in a wide array of economic models; in its absence, unrealistic comparative statics may arise, such as tougher competition between firms leading to higher markups.¹ In many models, practitioners may even wish to impose stronger assumptions on the curvature of demand in order to guarantee theoretical properties like equilibrium existence (Caplin and Nalebuff, 1991a). How do these assumptions affect the range of welfare estimates that are consistent with the data?

To answer this question, we enrich our basic framework with a variety of assumptions that are frequently made in different literatures. Precisely, we consider eight common assumptions on the curvature of demand in Section 3: Marshall’s second law, decreasing marginal revenue, log-concavity, concavity, ρ -concavity, convexity, log-convexity, and ρ -convexity. Our second main

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.” Melitz (2018) discusses the role of Marshall’s second law in the international trade literature as well as empirical evidence.

result characterizes the extremal demand curves that attain the largest and smallest possible changes in consumer surplus under each assumption.

This result yields a key insight: commonly used functional forms are often extremal. In our motivating example of a new tax, 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.

Our results can be readily and flexibly applied to a number of different empirical settings. In [Section 4](#), we demonstrate applications to three different literatures: international trade, public finance, and applied microeconomics. Our applications vary in context, policy, and data granularity. In each case, we build on the data and framework applied by the original authors of an established paper and compare the robust bounds induced by different sets of economic assumptions. We show, for example, that estimates of the deadweight loss incurred due to the 2018 trade tariffs are sensitive to assumptions about the shape of demand: the range of estimates that is consistent with convexity and [Marshall's](#) second law—standard assumptions in the trade literature—spans a third of the estimate generated by a CES demand curve. By contrast, we show that the conclusions in a recent quasi-experimental paper on energy subsidies are largely robust to the linear parametrization of demand that the paper employs. In this way, we show not only how our approach can be applied to compute robust bounds in different settings, but also how our results can be used to interpret existing results built on standard functional form assumptions.

Our work is motivated by the critique of [Bulow and Pfleiderer \(1983\)](#) in response to an empirical study by [Sumner \(1981\)](#), who used observed changes in marginal costs to estimate local demand elasticities. [Bulow and Pfleiderer](#) point out that these elasticity estimates are sensitive to functional form assumptions. The key insight is that even though two demand curves might be close to each other, their elasticities might be very different. Mathematically, this insight stems from the observation that the derivatives of two “similar-looking” functions (that are close to each other under the supremum norm, for example) 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.

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 4.2](#), 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, there are important differences as well. Whereas sufficient statistic estimators require the changes being analyzed to be sufficiently small that a local approximation suffices, our approach stipulates no such requirement. 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.

Our paper is also closely related to the literature on price indices.² The insight that a linear curve provides the upper bound for the consumer surplus change of a population with convex demand is considered lore by many economists and is referenced in passing in surveys such as [Hausman \(2003\)](#). However, beyond this particular one-sided bound (which is typically offered as intuition), linear interpolations are cited as “*ad hoc* at best” and other assumptions are generally not considered ([Hausman, 1996](#)). Instead, most papers in this literature apply a demand curve-fitting approach ranging from hedonic regression (typically, a flexible projection of log-quantity on log-price with fixed effects) to parametrizations of CES or logit demand.³ By contrast, our paper proposes a bounding strategy as a stand-alone method for estimating welfare and demonstrates how bounds can be derived for a wide range of empirical and theoretical restrictions on the shape of demand. Notably, while flexible approaches to fitting demand curves (parametric and non-parametric) typically require granular datasets and heavy computation for precise estimation, our bounding approach can accommodate low data environments and often allows for closed-form formulas. In this sense, our approach can be seen as complementary to curve-fitting—either as a first step or a robustness check.

² For the canon of this literature, see, for instance, [Willig \(1976\)](#), [Varian \(1982, 1985\)](#), and [Hausman \(1981\)](#).

³ See, for instance, [Hausman \(1996\)](#), [Hausman, Pakes, and Rosston \(1997\)](#), [Petrin \(2002\)](#), and [Nevo \(2003\)](#).

Our paper also relates to the literatures on set identification and the identification of counterfactual outcomes. The thesis of the set (or partial) identification literature is that functional forms for structural objects and parametric distributions trade off precision against generality and credibility. Instead, the literature driven by Charles Manski and his co-authors has pushed “empirical economists to be cautious of assumption-driven conclusions” and to apply a strategy of iteratively computing worst-case bounds under strengthening assumptions (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.⁴ 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.

There has also been some related theoretical work on optimal pricing with limited knowledge of demand. For example, Bergemann and Schlag (2011) characterize the optimal pricing rules that either maximize the seller’s worst-case profit or minimize the seller’s regret; and Cohen, Perakis, and Pindyck (2021) show that a simple (linear) pricing rule achieves a good approximation of the optimal profit under various curvature assumptions. Similar to these papers, we analyze an environment with limited knowledge of demand. However, while they analyze pricing policies, for which the elasticity (and hence gradient) of the demand curve is key, we derive robust bounds on welfare estimates, which requires integrating the demand curve.

Finally, from a methodological point of view, the characterization of extremal demand curves is generally a nontrivial task because each family of admissible demand curves is infinite-dimensional. Our analysis exploits a novel connection between welfare analysis and problems in mechanism and information design. To illustrate this, consider Varian’s problem of choosing a decreasing demand curve to maximize or minimize the change in welfare. Mathematically, this problem is equivalent to that of a mechanism designer choosing a monotonic allocation function—or an information designer choosing a monotonic posterior distribution of means. The equivalence no longer holds exactly when additional assumptions are imposed, such as a restriction on demand elasticity or the curvature of demand. Nevertheless, this connection allows us to leverage tools from recent work on mechanism and information design. The proof of our main results relies on an insight of Kang

⁴ See, for instance, Aguirregabiria (2010), Reguant (2016), and Kalouptsi, Scott, and Souza-Rodrigues (2021).

and Vondrák (2019): by defining a partial order on a family of demand functions and showing that the underlying change in welfare is monotone with respect to this partial order, the solution to the constrained optimization problem is attained at the extremal elements of the demand family.⁵

2 Basic model

Consider a market with a downward-sloping demand curve $D(\cdot)$ that 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 ($t = 1$)—and that the price increases from p_0 to p_1 . Correspondingly, the quantities of the good that are sold in each period are denoted by $q_0 = D(p_0)$ and $q_1 = D(p_1)$.

In this basic version of the model, we evaluate the change in consumer surplus arising from the price increase; as we show later, a similar analysis applies to other welfare measures. Throughout our analysis, we maintain that $D(\cdot)$ is generated by consumer preferences that are quasilinear in money. Thus the change in consumer surplus is equal to the area below the demand curve between p_0 and p_1 :

$$\Delta\text{CS} = \int_{p_0}^{p_1} D(p) \, dp. \quad (1)$$

Typically in empirical applications, not all of $D(p)$ is observed. Data about the quantities sold are available at the realized prices p_0 and p_1 , but not at any other price. As such, it is not possible to compute consumer surplus exactly using the formula above without imposing strong assumptions. Nevertheless, the observations of (p_0, q_0) and (p_1, q_1) along the demand curve restrict the possible values of ΔCS . Geometrically, this means that $D(p)$ must pass through (p_0, q_0) and (p_1, q_1) . As we showed in the introductory example, even monotonicity between (p_0, q_0) and (p_1, q_1) can narrow the range of feasible ΔCS values.

2.1 Assumptions

As we argued in the introduction, imposing monotonicity as the only assumption on $D(\cdot)$ allows the elasticity of demand to be infinite or zero. In fact, the demand curves that attain the box bounds for ΔCS have elasticities of either $-\infty$ or 0 almost everywhere. In many cases, however,

⁵ Our approach is also connected to work by Kleiner, Moldovanu, and Strack (2021), who present a general framework for solving infinite-dimensional optimization problems with majorization constraints: the partial order that we define on each family of demand curves is equivalent to majorization.

such extreme values may be ruled out *a priori*. This may allow a researcher to obtain more meaningful bounds on ΔCS without without any restrictions on the functional form or curvature of $D(\cdot)$. To demonstrate this, we make the following two minimal assumptions in our basic model:

(A1) The demand curve passes through the points (p_0, q_0) and (p_1, q_1) .

(A2) The price elasticity of demand $\varepsilon(p)$ between p_0 and p_1 lies between $\underline{\varepsilon}$ and $\bar{\varepsilon}$, where:

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

To avoid trivialities, we further assume that the *average* elasticity between p_0 and p_1 is bounded by $\underline{\varepsilon}$ and $\bar{\varepsilon}$:

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

This is a necessary and sufficient condition for the existence of a demand curve that satisfies both (A1) and (A2).

Assumption (A1) 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 both directions by studying a setting where only p_0 , p_1 , and q_1 are observed, and a setting where more points along the demand curve are observed.

Assumption (A2) might seem a bit more substantive in that it requires *a priori* knowledge about the demand curve. However, we emphasize that $\underline{\varepsilon}$ and $\bar{\varepsilon}$ should rather be interpreted as a way to flexibly parameterize how much *a priori* knowledge is required. This includes as a special case the extreme where no *a priori* knowledge is assumed: $\underline{\varepsilon} = -\infty$ and $\bar{\varepsilon} = 0$.

Perhaps more realistically, empirical researchers and policymakers might derive priors on the bounds of admissible elasticities from combinations of institutional knowledge, surveys of related studies and results from (quasi-)experiments in their setting. 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 $-3.18 \leq \varepsilon(p) \leq -0.01$. Of course, while the usual concerns of external validity apply when exporting elasticity estimates from one empirical context to another, $\underline{\varepsilon}$ and $\bar{\varepsilon}$ can incorporate this uncertainty quantitatively by allowing for a wider range of elasticities than that found in the academic literature. As we show

in Section 4, comparative statics with respect to $\underline{\varepsilon}$ and $\bar{\varepsilon}$ allow for a precise quantification of how “robust” these welfare results are to different elasticity assumptions.

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

$$\mathcal{D} := \left\{ D : [p_0, p_1] \rightarrow \mathbb{R} \text{ is decreasing s.t. } D(p_0) = q_0, D(p_1) = q_1, \underline{\varepsilon} \leq \frac{pD'(p)}{D(p)} \leq \bar{\varepsilon} \quad \forall p \in [p_0, p_1] \right\}.$$

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. Equivalently, we find the largest and smallest possible changes in consumer surplus:⁶

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

2.2 Analysis

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) ,

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

The elasticity of $D_{\text{CES}}(\cdot)$ 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 $D_{\text{CES}}(\cdot)$ 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):

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

⁶ Because \mathcal{D} is convex and ΔCS is a linear map of D using 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.

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 $D_{\text{CES}}(\cdot)$, 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 $D_{\text{CES}}(\cdot)$.

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 $D^*(\cdot)$:

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

Although we have described how a 2-piece CES interpolation can yield a larger ΔCS than $D_{\text{CES}}(\cdot)$, it is easy to see that a symmetric argument implies that a 2-piece CES interpolation can also yield a smaller ΔCS than $D_{\text{CES}}(\cdot)$. 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:

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

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? However, the main result of this section indicates that these attempts will ultimately fail.

Theorem 1. The largest and smallest possible changes in consumer surplus, $\overline{\Delta CS}$ and $\underline{\Delta CS}$, are attained by 2-piece CES interpolations:

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

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 extensions that we consider later, this simple geometric approach no longer applies. The second, which we present in [Appendix A](#) as it might be of independent interest, 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 this approach is less straightforward, the proof technique also generalizes easily to later extensions.




Figure 3: Sketch of the proof of Theorem 1.

We now explain the geometric proof of Theorem 1 with the help of Figure 3. The proof begins with a change of variables: rather than plot prices against quantities, we depict a demand curve by 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 much more natural interpretation of (A2). 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 $\log p_1$. This 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 that each of these curves correspond to must respectively maximize and minimize the change in consumer surplus. Reversing the change of variables, both of these curves clearly correspond to 2-piece CES interpolations between (p_0, q_0) and (p_1, q_1) . The top (red) curve corresponds to a 2-piece CES interpolation that has elasticity $\bar{\varepsilon}$ between (p_0, q_0) and an auxiliary point, and elasticity $\underline{\varepsilon}$ between the auxiliary point and (p_1, q_1) . It is easy to see that this curve corresponds to the demand curve $D^*(\cdot)$ defined above. 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 (A2) that are not only necessary, but also sufficient *a posteriori*. In more complicated extensions, the binding constraints are not as easily determined, nor do they necessarily take on such a simple form. In turn, this explains why the geometric argument fails to generalize to later extensions, necessitating the alternative, more technical approach that we undertake in [Appendix A](#).

2.3 Discussion

Before moving on to the 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.

As our geometric proof of Theorem 1 clearly demonstrates, the extremality of 2-piece *CES* interpolations arises from our assumption (A2) that the range of *elasticities* is known. As we argued earlier, this assumption seems natural in our view because empirical researchers and policymakers already use ranges of elasticities and treatment effects 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 better *a priori* knowledge about the likely revenues, rather than elasticities, over the interval of prices in question. The analogue 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 revenue interpolations.

A related, but different, point is that the empirical researcher or policymaker might sometimes have more *a priori* knowledge than (A2) supposes. In international trade, for example, demand is often assumed to satisfy [Marshall’s](#) second law: demand becomes more elastic as price increases. While this assumption is often necessary for trade models to yield reasonable comparative statics (cf. [Krugman, 1979](#)), it is also justified by some empirical evidence, as summarized by [Melitz \(2018\)](#). However, [Marshall’s](#) second law rules out the demand curve $D_*(\cdot)$ that generates the lower bound in Theorem 1. We examine how this, as well as other assumptions on the curvature of demand, affects our results in Section 3.1.

In many 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. Of course, with infinite variation along the demand curve, the true demand curve $D(\cdot)$ is non-parametrically identified and may be directly recovered.⁷ In that limit, our bounds converge to the actual change in consumer surplus. Section 3.3 analyzes the intermediate case of arbitrarily (but finitely) many observations.

In other cases, fewer data might be available. For instance, if the welfare question at hand is *counterfactual* rather than *retrospective*, only one price and quantity pair may actually be observed, leaving the counterfactual quantity 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 3.2 using the same analytical tools introduced above.

Finally, a subtler assumption that we have made throughout this section is that the points on the demand curve are perfectly observed (*i.e.*, without noise or measurement error). While this simplifies the exposition, the assumption is typically indefensible from an econometric standpoint. However, as we show in Section 3.4, our analysis extends straightforwardly to this case.

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—whether through more assumptions on $D(\cdot)$, more data or more uncertainty—is allowed for.

3.1 Assumptions on curvature of demand

Economic models often impose restrictions on the curvature of demand for a variety of reasons. Historically, Marshall (1890) went so far as to *define* a demand curve as a decreasing function whose elasticity also decreases with price; 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 textbook depiction of a demand curve as a convex function.

Different literatures in economics have since introduced a variety of assumptions on the curvature of demand that capture other intuitions pertaining to their fields of interest. In this

⁷ See, for instance, Fox and Gandhi (2016).

subsection, we study how different assumptions on the curvature of demand affect our bounds for ΔCS . To do so, we continue maintaining (A1), but replace (A2) with an assumption on the curvature of demand.

In order to be comprehensive, we consider a range of assumptions that are considered standard in different fields. Each assumption (abbreviated by “CA” for “curvature assumption”) 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 and macroeconomics, notably “without apology” by Krugman (1979). Melitz (2018) provides some empirical justification for this assumption.⁸
- (CA2) **Decreasing marginal revenue.** Let $P(q) := D^{-1}(q)$ denote the inverse demand curve. Demand exhibits decreasing marginal revenue if marginal revenue $\text{MR}(q) := P(q) + qP'(q)$ 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).
- (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 $D'(p) [D(p)]^{\rho-1}$ 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

⁸ Marshall’s second law is also used in microeconomics, such as in the work of Bishop (1968) and Johnson (2017).

($\rho = 0$) and concavity ($\rho = 1$). Different values of ρ parametrize the restrictiveness of this assumption: a ρ' -concave demand curve is ρ'' -concave for any $\rho'' < \rho'$.

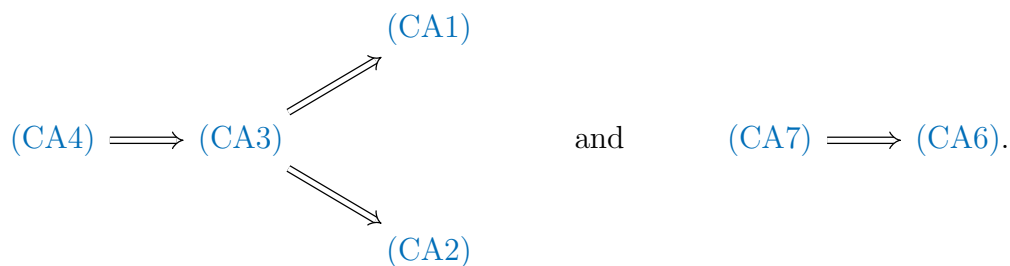
(CA6) **Convex demand.** Demand is convex if $D'(p)$ is increasing in p . Similar to concave demand (CA4), Robinson (1933) shows that total output increases when monopolistic price discrimination causes prices to fall in markets with convex demands.

(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 $D'(p)[D(p)]^{\rho-1}$ 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:



For reference, these relationships are proven in Appendix B, 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 $i = 1, \dots, 8$:

$$\mathcal{D}_i := \{D : [p_0, p_1] \rightarrow \mathbb{R} \text{ is decreasing s.t. } D(p_0) = q_0, D(p_1) = q_1, D \text{ satisfies (CA}_i\text{)}\}.$$

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

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

Theorem 2. The following bounds on changes in consumer surplus hold:

- (i) For concave-like assumptions (CA1)–(CA5), $\overline{\Delta CS}_i$ is attained by a 2-piece interpolation and $\underline{\Delta CS}_i$ is attained by a 1-piece interpolation.
- (ii) For convex-like assumptions (CA6)–(CA8), $\overline{\Delta CS}_i$ is attained by a 1-piece interpolation and $\underline{\Delta CS}_i$ is attained by a 2-piece interpolation.

Moreover, these interpolations are

- CES under Marshall’s second law (CA1);
- constant marginal revenue when marginal revenue is decreasing (CA2);
- exponential when demand is log-concave (CA3) or log-convex (CA7);
- linear when demand is concave (CA4) or convex (CA6); and
- ρ -linear when demand is ρ -concave (CA5) or ρ -convex (CA8)

As the proof is an extension of the geometric proof of Theorem 1, we defer it to [Appendix C](#). 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, $\overline{\Delta CS}$ will be attained by a 1-piece linear demand curve, whereas $\underline{\Delta CS}$ will be attained by a 1-piece CES demand curve.

Finally, instead of replacing (A2) with a curvature assumption, we can also combine each curvature assumption with (A2). In Appendix C, we show how an extension of the geometric proof yields simple bounds for the case when the demand curve satisfies (A1), (A2), and (CA1). However, when curvature assumptions other than (CA1) are combined with (A1) and (A2), the geometric proof no longer extends. In these cases, we show how our alternative proof of Theorem 1 (presented in Appendix A) extends instead.

3.2 Extrapolating from fewer observations

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 4.3, the quantity that would be demanded at p_1 is not known. Instead of (A1), these applications call for a weaker assumption:

(A1') The demand curve passes through the point (p_0, q_0) .

Let \mathcal{D}' denote the family of demand curves that satisfy (A1') and (A2):

$$\mathcal{D}' := \left\{ D : [p_0, p_1] \rightarrow \mathbb{R} \text{ is decreasing s.t. } D(p_0) = q_0, \underline{\varepsilon} \leq \frac{pD'(p)}{D(p)} \leq \bar{\varepsilon} \quad \forall p \in [p_0, p_1] \right\}.$$

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

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

A general procedure for finding $\overline{\Delta CS}'$ and $\underline{\Delta CS}'$ is to decompose the problem into three steps. First, we characterize the set of possible values of q_1 that are consistent with (A1') and (A2). For each possible value of q_1 , we then apply Theorem 1 to compute $\overline{\Delta CS}$ and $\underline{\Delta CS}$ for that q_1 . Finally, $\overline{\Delta CS}'$ is equal to the maximal $\overline{\Delta CS}$ over all possible q_1 , whereas $\underline{\Delta CS}'$ is equal to the minimal $\underline{\Delta CS}$ over the same set. Because $\overline{\Delta CS}'$ and $\underline{\Delta CS}'$ 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 $\overline{\Delta CS}'$ and $\underline{\Delta CS}'$ 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 (A1') and (A2) can be found by drawing the (blue) straight line with gradient $\bar{\varepsilon}$ that passes through the point $(\log p_0, \log q_0)$, and then finding the (red) point on the line at $\log p_1$.




Figure 4: Illustration of bounds when only (p_0, q_0) and p_1 are observed.

It is clear that this value of q_1 must also yield the maximal $\overline{\Delta CS}$; hence $\overline{\Delta CS}'$ must be attained by a 1-piece CES interpolation (corresponding to the red curve). A symmetric argument shows that $\underline{\Delta CS}'$ must also be attained by a 1-piece CES interpolation (corresponding to the green curve).

Theorem 3. The largest and smallest changes in consumer surplus, $\overline{\Delta CS}'$ and $\underline{\Delta CS}'$, are attained by 1-piece CES interpolations (with elasticities $\bar{\varepsilon}$ and $\underline{\varepsilon}$, respectively).

We conclude this subsection by noting that 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 precisely because our welfare measure ΔCS is monotone in q_1 . As we show in Section 4.3, when different welfare measures than ΔCS are used, such as welfare measures that take into account costs (that depend on q_1), the geometric argument is insufficient and this decomposition procedure is required.

3.3 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 (A1) with:

(A1'') The demand curve passes through the points $(p_0, q_0), \dots, (p_{n-1}, q_{n-1})$.

Thus (A1) is simply a special case of (A1'') by setting $n = 2$. As such, our first result for this subsection directly generalizes Theorem 1:




Figure 5: Illustration of bounds with $n = 3$ observations.

Theorem 4. Under (A1'') and (A2), the largest and smallest possible changes in consumer surplus are attained by $(2n - 2)$ -point CES interpolations.

Theorem 4 follows by applying Theorem 1 between every two adjacent points. 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''), (A2), 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.

Theorem 5. Under (A1’), (A2), and (CA1), the largest possible change in consumer surplus is attained by an n -point CES interpolation, and the smallest possible change in consumer surplus is attained by an $(n - 1)$ -point CES interpolation.

3.4 Noisy observations

The basic model assumes that prices and quantities are perfectly observed. However, most empirical observations entail measurement error of some sort.⁹ To model this, we suppose that (p_0, q_0) and (p_1, q_1) do not have to lie on the *true* demand curve $D(\cdot)$. Instead, we allow for error in the measurement of quantities, up to some known amount $\delta > 0$:

$$|D(p_0) - q_0|, |D(p_1) - q_1| \leq \delta.$$

While the inclusion of error allows the model to account for empirical noise, it does not change our analysis. To see why, consider the set \mathcal{S} of all possible quantity pairs that lie on the *true* demand curve:

$$\mathcal{S} := \{(\hat{q}_0, \hat{q}_1) : |\hat{q}_0 - q_0|, |\hat{q}_1 - q_1| \leq \delta\}.$$

For each pair $(\hat{q}_0, \hat{q}_1) \in \mathcal{S}$, we can proceed exactly as we did in the basic model. Theorem 1 implies that, when the points (p_0, \hat{q}_0) and (p_1, \hat{q}_1) lie on the true demand curve, the largest and smallest possible changes in consumer surplus are attained by 2-point CES interpolations. With slightly more work (similar to the argument in Section 3.2), we can show that $\overline{\Delta CS}$ is attained at $(\hat{q}_0, \hat{q}_1) = (q_0 + \delta, q_1 + \delta)$, whereas $\underline{\Delta CS}$ is attained at $(\hat{q}_0, \hat{q}_1) = (q_0 - \delta, q_1 - \delta)$.

Theorem 6. The largest possible change in consumer surplus is attained by a 2-piece CES interpolation between $(p_0, q_0 + \delta)$ and $(p_1, q_1 + \delta)$, and the smallest possible change in consumer surplus is attained by a 2-piece CES interpolation between $(p_0, q_0 - \delta)$ and $(p_1, q_1 - \delta)$.

Although we have focused on the case of “worst-case” error here, this analysis can be easily extended to the more intuitive form of “confidence intervals” over quantities: that is, if the error can be assumed to take the form

$$\Pr [|D(p_0) - q_0|, |D(p_1) - q_1| \leq \delta] \geq \pi \quad \text{for some } \pi \in [0, 1].$$

⁹ Our basic model assumes that observations lie on the same demand curve, which excludes the possibility of market shocks. One way to relax this assumption is to model possible market shocks as structural noise following the approach laid out in this subsection.

In that case, while a similar result to Theorem 6 holds, the interpretation is that ΔCS must fall within the resulting interval $[\underline{\Delta\text{CS}}, \overline{\Delta\text{CS}}]$ with probability at least π .

4 Applications

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

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 Trade tariffs

The domestic impact of trade wars is central to modern economic policy. Import tariffs subsidize domestic producers and exert pressure on foreign countries, but this comes at a cost for domestic consumers, who face higher prices for goods across the supply chain. While the equilibrium effects of standing trade frictions are typically difficult to disentangle, politically motivated trade wars provide a compelling exogenous price shock, around which welfare impacts can be measured. In this way, data from escalations in trade wars enable not only an evaluation of the trade war itself, but also a signal of domestic vulnerability to trade frictions more generally.

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. While each of these studies employs a slightly different modeling strategy and set of empirical techniques, they all 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.¹⁰

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,

¹⁰ 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?”](#)

building on the data and approach taken in [Amiti, Redding, and Weinstein \(2019\)](#). For each product (defined as a 10-digit Harmonized System 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 near-complete tariff pass-through to consumers,¹¹ 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.¹² 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.¹³

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 U.S. 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 3.1](#):

¹¹ This is consistent with the findings of [Fajgelbaum, Goldberg, Kennedy, and Khandelwal \(2020\)](#) and [Cavallo, Gopinath, Neiman, and Tang \(2021\)](#), who use different estimation methodologies than [Amiti et al.](#)

¹² 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 only capture the loss incurred by consumer surplus. It would be possible to do a “doubly robust” version of our bounds to account for producer surplus loss, but as this does not feature in the papers we are working off of, we omit it for the sake of brevity.

¹³ While [Fajgelbaum et al.](#) allow for a richer depiction of the economy (they account for losses from exports hit by retaliatory tariffs in addition to deadweight loss), they assume a type of CES interpolation for import demand.

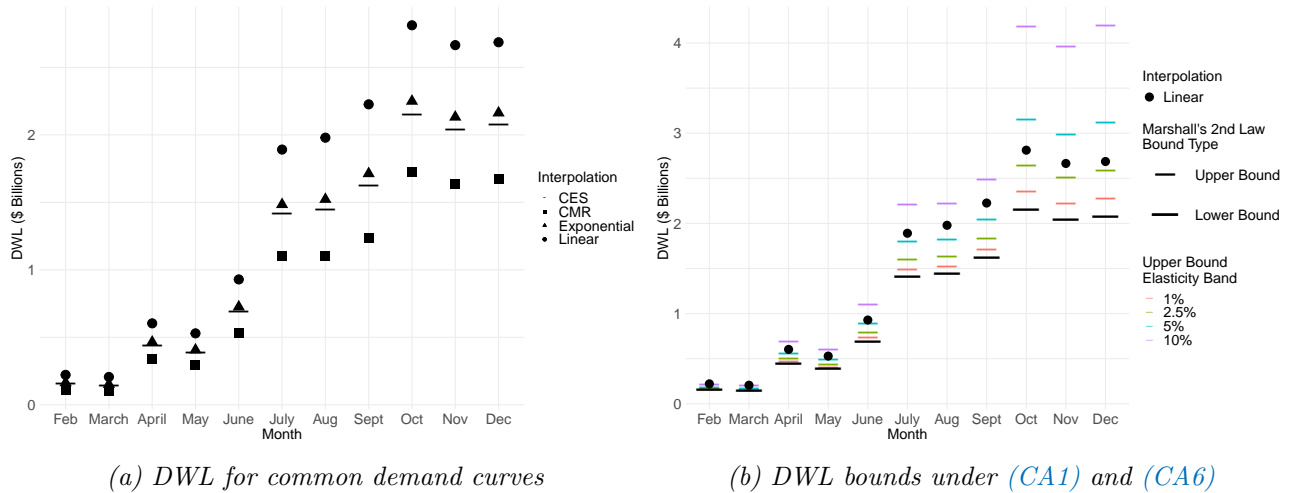


Figure 6: Comparing DWL bounds under different assumptions.

CES, constant marginal revenue (CMR), exponential, and linear. Figure 6(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 functions. 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 6(a) 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 is fixed. This is unsurprising: as we argue in Section 3.1 and show formally in Appendix B, 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 6(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 6(a) demonstrates the range of extremal deadweight loss measurements 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 6(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 long black dash in Figure 6(b). As we note in Section 3.1, absent any other data or restrictions, the Varian (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.

The black dots in Figure 6(b) correspond to the upper bound on deadweight loss if demand is also assumed to be convex (CA6)—given by the 1-piece linear estimate as in Figure 6(a). Taken together, the black dashes and dots in Figure 6(b) present a set of comprehensive 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. In these cases, the Taylor approximations of any demand function in this family would be almost the same. 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 \$12.6B and at most \$16.8B.

The assumptions that are imposed on demand may shift the interpretation of policy implications. As Amiti et al. note, the U.S. government internalized \$15.6B in revenues from import tariffs over the course of 2018. While this falls short of the upper bound on deadweight loss under (CA6), the discrepancy (\$1.2B) is lower than most estimates offered in the literature. This, in part, reflects the conservative nature of our deadweight loss analysis. However, it also reflects the additional imposition of a convexity restriction on demand—which may not be appropriate. As an alternative, we consider how the upper bound on deadweight loss might change with different priors on elasticity bounds. 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.¹⁴ The colored dashes in Figure 6(b) present the upper bound for total monthly deadweight loss under (A2) 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 exceed the convexity bound starting

¹⁴ The average elasticity can be inferred from price and quantity pairs alone: $\varepsilon_{\text{avg}} = \frac{\log(q_1) - \log(q_0)}{\log(p_1) - \log(p_0)}$. A Δ -symmetric band around the average elasticity yields $\underline{\varepsilon} = (1 - \Delta) \times \varepsilon_{\text{avg}}$ and $\bar{\varepsilon} = (1 + \Delta) \times \varepsilon_{\text{avg}}$.

in October. Thus, even demand functions with modestly decreasing elasticities may be consistent with total deadweight loss estimates that exceed \$16.8B.

4.2 Income taxation

The question of how to optimally tax productive labor in a modern economy has played a major role in economic thought since Ramsey (1927) and Mirrlees (1971). On one hand, tax revenue is necessary for funding government services: in the United States, about 50% of federal revenue came from individual tax income in 2019.¹⁵ On the other hand, the imposition of income taxes distorts individual incentives to generate income and may cause a decrease in net productivity. As with tariffs, analysts seeking to balance these two considerations must weigh the increase in revenue that would be earned from a higher tax rate against the deadweight loss from lost economic transactions that stem from the added distortion. As modeling the full scope of relevant interactions throughout the economy imposes heavy requirements on both data availability and on assumptions for tractability, a popular approach in the public finance literature has instead focused on establishing “sufficient statistics” for welfare changes around marginal increases in taxes.

The key insight of the sufficient statistics approach is two-fold.¹⁶ First, it notes—as we do—that while estimating the full curve describing the relationship between welfare (W) and tax rates (τ) is very demanding, estimating a local elasticity around a marginal change in taxes is both easier and more robust to *ad hoc* parametric assumptions. Second, it invokes the envelope theorem to argue that the secondary effects of a tax increase—for instance, changes in the distribution of an individual’s consumption of different goods that result from a change in effective income—are of second-order importance when the tax increase is small and when consumers are optimizing. Combining these two ideas, a typical analysis decomposes the marginal welfare gain ($dW/d\tau$) into a simple function of marginal consumption utilities, which 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

¹⁵ See, for example, data from the Center on Budget and Policy Priorities, available at <https://www.cbpp.org/research/federal-tax/where-do-federal-tax-revenues-come-from>.

¹⁶ 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.

taking a stance on how to integrate $dW/d\tau$ and parametric interpolations are often invoked.¹⁷ These approaches are local: when policy changes are small, Taylor’s theorem ensures that these approximations work well, but there is no such guarantee when policy changes are large—which are often precisely the cases of interest. In this subsection, we demonstrate how our framework applies to such settings using 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):¹⁸

$$\frac{dW(\tau)}{d\tau} = \tau \cdot \frac{d\text{TI}(\tau)}{d\tau}.$$

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.

As an application, [Feldstein](#) estimates the welfare impact of major changes to the tax schedule in 1986 and 1993. To do this, he relies on three empirical objects: *(i)* the baseline marginal rate of taxation, *(ii)* the level of taxable income at the baseline (and in the counterfactual), and *(iii)* an estimate of the average elasticity of taxable income relative to the net-of-tax income share. These numbers allow him to evaluate the local welfare derivative at each individual’s observed taxable income, and to extrapolate to estimate the total excess burden of taxation under each counterfactual scenario.

In [Figures 7\(a\) and 7\(b\)](#), we demonstrate [Feldstein’s](#) approach in relation to our [Harberger](#) example in [Section 1](#). We consider the change in taxable income $\text{TI}(\tau)$ 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 $\tau \cdot d\text{TI}(\tau)/d\tau$. The net change in excess burden

¹⁷ [Kleven \(2021\)](#) suggests including higher-order terms as an alternative to parametric interpolations.

¹⁸ To make clear the connection between the sufficient statistics approach and ours, our discussion of [Feldstein \(1999\)](#) follows the expositional logic laid out by [Chetty \(2009\)](#), who gives a complete derivation of $dW(\tau)/d\tau$.




Figure 7: *Feldstein's sufficient statistics approach to estimating the excess burden of income taxation.*

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

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

As in the [Harberger](#) example, area D is simple to evaluate without further assumptions: it is given by the product of τ_0 and $TI(\tau_0) - TI(\tau_1)$. However, as the curve $TI(\tau)$ 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 7(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 $TI(\tau)$. 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 $TI(\tau)$ curves from which the observed points were drawn.




Figure 8: Deadweight loss bounds based on Theorem 1 for the *Feldstein* example.

To illustrate how robust bounds interact with *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.¹⁹ To evaluate the deadweight loss for this individual, *Feldstein* relies on estimates of taxable income elasticities from previous work. His main estimate, $\epsilon_\tau = -1.04$, stems from a difference-in-differences exercise in *Feldstein* (1995). But *Feldstein* (1995) finds different estimates with different cuts of his data, and he cites a number of contemporary papers that estimate a range of elasticities from -1.33 to -0.55 .²⁰ In discussing his estimates, *Feldstein* suggests replicating his welfare results with $\epsilon_\tau = -0.5$ and $\epsilon_\tau = -1.33$ as extremal cases for robustness.

In Figure 8, we take these extremal cases as bounds on the range of feasible elasticities and compute the bounds on ΔW using Theorems 1 and 2. On the left (“No Assumption”), we compare the elasticity bounds (attained by 2-piece CES interpolations) with no curvature restrictions against box bounds (attained by 2-piece linear interpolations). Although even box bounds are fairly narrow in this example—the lower bound on ΔW is \$6.6K, and the upper bound is \$8.3K—the elasticity bounds further narrow the range of welfare estimates substantially. On the right, we zoom in on the elasticity bounds and add concave-like curvature

¹⁹ See *Feldstein* (1999) for a derivation of the counterfactual taxable income imputation.

²⁰ *Feldstein* estimates average elasticities of taxable income by comparing the taxable income amounts and marginal tax rates for a panel of taxpayers affected by the Tax Reform Act of 1986. His different estimates correspond to comparing different subsets of his panel.

restrictions following Theorem 2. Assuming Marshall’s second law yields a lower bound (attained at a 1-piece CES interpolation) of \$7.41K. Meanwhile, the linear demand assumption that Feldstein (1999) employs—which attains the lower bound of ΔW under concavity—yields an estimate of \$7.46K that is just under the elasticity-band upper bound of \$7.49K.

While this example uses projected counterfactual taxable income, it demonstrates how observed quantities (or levels of taxable income, in this case) can greatly narrow the set of possible welfare impacts. Even with no additional information, the welfare loss is bounded within a range that may be sufficiently tight for policy discussions; with elasticity restrictions spanning estimates across the literature, the welfare loss becomes pinned down to several hundreds of dollars. Furthermore, while Figure 8 makes clear that the linearity assumption is not by itself innocuous—absent elasticity restrictions, the deadweight loss estimate implied by linearity bisects the box bound—it also shows that in this case, the Taylor approximation may be appropriate.

4.3 Energy subsidies

Empirical researchers often observe only one price-quantity pair and are interested in assessing the *counterfactual* welfare impact from an alternative price. To do so, a standard approach is to impose a functional form assumption on demand, which can be used to extrapolate from the observed price-quantity pair 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 the uncertainty about counterfactual quantities that stems from uncertainty about the true demand function. To illustrate this, we consider an application to evaluating energy subsidies studied by Hahn and Metcalfe (2021).

Since 1989, the state of California has offered a rate assistance program known as for energy consumption to qualifying low-income individuals. Through the California Alternate Rates for Energy (CARE) program, these individuals receive wholesale discounts on their unit prices for gas and electricity. In the gas market that Hahn and Metcalfe study, CARE participants receive a 20% discount on marginal rates, from an average price of \$0.95 to \$0.75 per therm of gas.

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. Hahn and Metcalfe evaluate the net welfare impact of CARE and find that the program ultimately results in a net loss of \$4.8 million.

To assess CARE, [Hahn and Metcalfe](#) derive a counterfactual unit price p^* and project the amount of gas that would be consumed both by households that are enrolled in CARE and those that are not.²¹ 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.




Figure 9: 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.

The mechanics of [Hahn and Metcalfe](#)'s welfare computation are depicted in Figure 9. For CARE households, the counterfactual unit price p^* is higher than the discounted CARE price p_C , and so the counterfactual quantity q_C^* is lower than the observed quantity q_C . 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

²¹ [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.

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 9(a), while the loss in total surplus (in red) for a representative non-CARE household is shown in Figure 9(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 would their result change if the true demand curve is not linear? As we show in Section 3.1, a linear interpolation is meaningful beyond simplicity. By Theorem 2, the linear demand curve yields the *upper* bound on the surplus loss from non-CARE households among all convex demand curves for gas. By the same logic, the linear demand curve also yields the *lower* bound on the surplus gain from CARE households. In this sense, it is challenging to intuit how the linearity assumption affects their numerical estimate barring further numerical analysis.

To overcome this challenge, we apply our results from Section 3.2 to provide an alternative approach that avoids making functional form assumptions. Our approach is summarized in Figure 10. To begin, we extend Hahn and Metcalfe’s sensitivity analysis to account for uncertainty—not only with respect to the elasticity at the observed price-quantity pair for each household group, but also in 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 estimates used by Hahn and Metcalfe. 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.




Figure 10: Upper bound of net change in total surplus without functional form assumptions.

In each case, we apply the procedure outlined in Section 3.2 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 numerical 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. Even without the linearity assumption, fluctuations of over 50% of the estimated elasticities cannot possibly rationalize a positive net welfare impact of CARE.²² But Figure 10 also demonstrates how uncertainty about the shape of demand can change the interpretation of the result. Hahn and Metcalfe’s estimate of net welfare lies near 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

²² Note that the elasticity bands encode both uncertainty about the elasticity estimates at the observed price/quantity and fluctuations in the elasticities along the demand curve.

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.

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 have focused on extrapolating from small policy changes using functional form assumptions and adopting complicated structural models. These approaches have played a crucial role in the understanding economists have gained about policy in different markets. But they also beg the question of how much these intuitions depend on the functional form assumptions that these approaches rely upon.

In this paper, we present a different approach that does not depend on any one functional form. Instead, we provide bounds on welfare that hold under various families of assumptions that are commonly made in different literatures, and illustrate our framework in a series of 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.
- 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. AND K. SCHLAG (2011): “Robust Monopoly Pricing,” *Journal of Economic Theory*, 146, 2527–2543.
- BISHOP, R. L. (1968): “The Effects of Specific and Ad Valorem Taxes,” *Quarterly Journal of Economics*, 82, 198–218.
- 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.
- 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.
- FOX, J. T. AND A. GANDHI (2016): “Nonparametric Identification and Estimation of Random Coefficients in Multinomial Choice Models,” *RAND Journal of Economics*, 47, 118–139.
- 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.
- 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 T. F. Bresnahan and R. J. Gordon, 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.
- JOHNSON, J. P. (2017): “The Agency Model and MFN Clauses,” *Review of Economic Studies*, 84, 1151–1185.
- KALOUPSIDIS, 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*.
- 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.
- MANSKI, C. F. (1997): “Monotone Treatment Response,” *Econometrica*, 65, 1311–1334.
- MARSHALL, A. (1890): *Principles of Economics*, London, UK: Macmillan and Co.
- MELITZ, M. J. (2018): “Competitive Effects of Trade: Theory and Measurement,” *Review of World Economics*, 154, 1–13.
- MIRLLEES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175–208.
- MÜLLER, A. AND M. SCARSINI (2006): “Stochastic Order Relations and Lattices of Probability Measures,” *SIAM Journal on Optimization*, 16, 1024–1043.

- NEVO, A. (2003): “New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand Systems,” *Review of Economics and Statistics*, 85, 266–275.
- PETRIN, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110, 705–729.
- PRÉKOPA, A. (1973): “On Logarithmic Concave Measures and Functions,” *Acta Scientiarum Mathematicarum*, 34, 335–343.
- RAMSEY, F. P. (1927): “A Contribution to the Theory of Taxation,” *Economic Journal*, 37, 47–61.
- REGUANT, M. (2016): “Bounding Outcomes in Counterfactual Analysis,” *Working paper*.
- ROBINSON, J. (1933): *The Economics of Imperfect Competition*, London, UK: Macmillan.
- 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.
- 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.
- WALL STREET JOURNAL (2020): “How Many Tariff Studies Are Enough?” Available at <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 Apology,” *American Economic Review*, 66, 589–597.

Appendix A 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.

A.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:

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

Analogous to the family of demand curves \mathcal{D} , we define the set of elasticity functions that are consistent with (A1) and (A2):

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

Thus we arrive at the equivalent problem:

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

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

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

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

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

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

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

The inequalities follow from the fact that $\text{im } \eta \subset [\underline{\eta}, \bar{\eta}]$ for any $\eta \in \mathcal{E}$. A similar argument shows that $\eta \succeq \eta_*$ for any $\eta \in \mathcal{E}$. \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 $\Delta\text{CS} : \mathcal{E} \rightarrow \mathbb{R}$ by

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

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 $\Delta\text{CS}(\cdot)$ is increasing in the partial order \succeq :

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

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

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

The result thus follows from a pointwise comparison of the two integrands. □

Together, Lemmas 1 and 2 imply that:

Proposition 1. The functional $\Delta\text{CS}(\cdot)$ is maximized at η^* and minimized at η_* :

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

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.

A.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 3.1, 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 (A1) and (A2). [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.

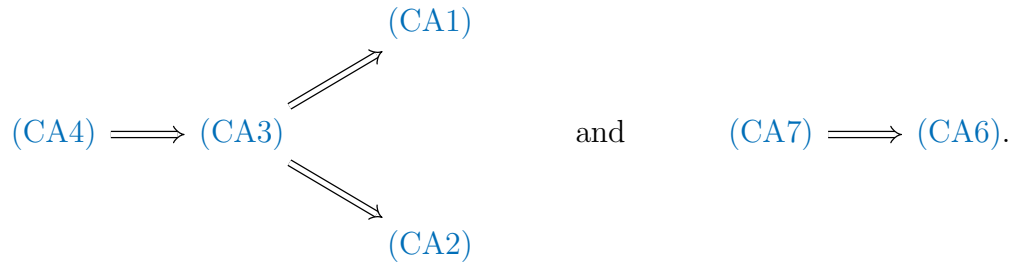
The fortuitous connection between our problem of bounding welfare in different families of demand and Bayesian persuasion problems implies that tools developed for *constrained* information design problems can potentially also be used to evaluate robust welfare bounds. From a technical point of view, 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 B 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.

B.1 Relationship between assumptions

We begin by showing that



(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. \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. \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

B.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 3.1 (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 $\lambda \geq 0$:

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

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

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

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 $\lambda \geq 0$:

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

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

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

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

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

Appendix C Other proofs and additional discussion

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

C.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 C.1: Sketch of the proof of Theorem 2.

Figure C.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 (A1)—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.

C.2 Additional discussion for Theorem 2

While Theorem 2 states our bounds by *replacing* (A2) with curvature assumptions, we can also derive bounds by imposing *both* (A2) *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 (A1), (A2), and (CA1). It can be readily verified that the geometric proof of Theorem 2 supplied in Appendix C.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 (A1), (A2), and (CA2). Crucially, the geometric picture (*e.g.*, Figure C.1) no longer captures all the relevant binding constraints except for special choices of $\underline{\varepsilon}$ and $\bar{\varepsilon}$. Instead, to derive bounds, one would have to formalize the problem using the approach of Appendix A, and then solve the resulting optimization problem.

C.3 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 $(p_0, q_0), \dots, (p_{n-1}, q_{n-1})$ at $\varepsilon_0 < \dots < \varepsilon_{n-1}$. 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 $i = 1, \dots, n - 2$, 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 $(p_1, q_1), \dots, (p_{n-2}, q_{n-2})$ are actually part of the same CES demand curve. Since this argument holds for any $\varepsilon_0 < \dots < \varepsilon_1$, the largest possible change in consumer surplus must be attained by an n -piece CES interpolation, as claimed.