

Boundedly Rational Demand*

Pavel Kocourek

University of Duisburg-Essen

Jakub Steiner

University of Zurich, CERGE-EI, and CTS

Colin Stewart

University of Toronto

May 1, 2024

Abstract

Evidence suggests that consumers do not perfectly optimize, contrary to a critical assumption of classical consumer theory. We propose a model in which consumer types can vary in both their preferences and their choice behavior. Given data on demand and the distribution of prices, we identify the set of possible values of the consumer surplus based on minimal rationality conditions: every type of consumer must be no worse off than if they either always bought the good or never did. We develop a procedure to narrow the set of surplus values using richer datasets and provide bounds on counterfactual demands.

Keywords: Behavioral welfare, bounded rationality, consumer theory, revealed preference

JEL codes: D11, D6, D9

*We circulated an earlier draft of this paper under the title “Demand in the Dark”. We have benefited from the comments of Sandro Ambuehl, Tibor Heumann, Zi Yang Kang, Anton Kolotilin, Nick Netzer, Marek Pycia, Andy Zapechelnyuk, various seminar and workshop audiences, the co-editor, Todd Sarver, and several anonymous referees. Stewart is grateful to PSE for their hospitality. This work was supported by ERC grant 770652, grant GAČR 24-10145S, DFG grant 450175676, and by the Social Sciences and Humanities Research Council of Canada.

1 Introduction

A key assumption of the standard approach to analyzing consumer demand and welfare is that consumers perfectly optimize. Yet it is clear from a number of empirical studies—if not from introspection alone—that this assumption does not generally hold in practice. For example, simply changing the way that prices are presented to consumers can have significant effects on demand (Chetty, Looney, and Kroft, 2009; Finkelstein, 2009).¹

Failures to optimize perfectly can occur for a variety of reasons. The consumer may not be fully attentive to the price of a good they buy, perhaps because it is a habitual purchase, or because prices involve complexities that require effort to understand. Similarly, consumers may not always be aware of the price of a good they do *not* buy. Alternatively, they may simply make stochastic errors, as is typically assumed in random utility models. Varying salience of certain features of the product, such as whether it is on sale, may also influence consumers’ choices.

How are we to make inferences about a consumer’s preferences and welfare in the absence of optimal choices? We propose an approach grounded in minimal assumptions about the consumer’s rationality. We focus on a setting with unit demand in which the data consist of a demand curve indicating the probability of purchase at each price together with a distribution of prices. Our rationality conditions require only that no type of consumer would be better off if they switched to either always buying or never buying the good.²

We consider an analyst who seeks to rationalize the data with a *model* describing a distribution of types of the consumer, with each type specifying the consumer’s value together with their demand curve. A model rationalizes the data if, at each price, the observed purchase probability is equal to the expected demand across the types in the model. We study two main questions. First, what can the analyst infer

¹See also Ito (2014) for empirical evidence that consumers do not correctly account for marginal electricity pricing and Feldman, Katusčák, and Kawano (2016) for related evidence regarding marginal tax rates. Dickson and Sawyer (1990) document that more than half of the supermarket shoppers they surveyed were unable to accurately report the price of an item immediately after placing it in their cart. Taubinsky and Rees-Jones (2018) and Tipoe (2021) find that there is significant heterogeneity in attention to prices.

²We evaluate expected utilities with respect to the distribution of prices observed by the analyst, making an implicit identifying assumption that this distribution agrees with the consumer’s subjective belief or experience.

from the observed data about the surplus the consumer receives from participating in the market for this good? Our approach can help to quantify the uncertainty due to bounded rationality in estimated measures of consumer surplus contributed by new products.³ Second, what can the analyst predict about the counterfactual demand if the consumer were able to fully optimize? This question could be of interest to a regulator or a monopolist considering fixing the price in the market, thereby eliminating any errors in choice arising from inattention to fluctuations in prices. For each of these questions, we first obtain bounds using simple datasets then show how to narrow these bounds with richer data.

There are two formally equivalent interpretations of our framework. Under one interpretation—which is the one we use in our description—there is a single consumer with a stochastic type. Under the other interpretation, there is a continuum of consumers, and each type describes an individual consumer whose preferences and behavior are fixed.

The consumer’s type captures both fluctuations in value (as in a random utility model) and varying attention that could be correlated with the value. For example, it could be that when the consumer’s value is high, she pays little attention to the price whereas when it is low she checks more carefully. Alternatively, her attention could vary due to external factors such as time pressure. We think of a type’s demand as resulting from a combination of her value for the good and her attention (which is unobserved and unmodelled).

One special case of our model of particular interest arises when each type of the consumer is Bayes rational and imperfections in optimization are due to incomplete information about prices. In that case, each type is described by a value for the good together with an information structure about prices, with their demand following from optimization of their expected payoff. Each such type of consumer satisfies our minimal rationality conditions since the strategies of always or never buying the good can be implemented by ignoring any information they have about prices. One might think, then, that this case would lead to narrower bounds on surplus and counterfactual demand than the ones resulting from our general rationality conditions.

³For example, Cohen et al. (2016) estimates the total consumer surplus due to Uber, Goolsbee and Petrin (2004) that due to direct broadcast satellites, and Petrin (2002) that from the introduction of the minivan.

It turns out, however, that the bounds are identical: any model satisfying our general rationality conditions can be obtained with Bayes rational consumer types for appropriately chosen information structures.

Relative to our setting, the standard assumption of optimal choice (coupled with quasilinear preferences) simplifies the analysis in two ways. First, each type has a threshold demand—buying the good if and only if the price is below that type’s value for the good—and thus the value can be directly inferred from its demand. Second, any demand curve admits a unique decomposition into threshold demands of individual consumer types; thus the distribution of values can be directly inferred from the demand observed in the data. In contrast, in our setting, types with the same value may differ in their demand, corresponding to differences in attention or sophistication. Thus types’ values cannot be directly inferred from their demands, and their demands need not take a simple threshold form.

There are generally many different models that can rationalize the data. First, the analyst must consider various decompositions of the overall demand into demands of individual types. Second, for each type, given its demand, there is a range of incentive-compatible values, i.e., values for which our rationality conditions are satisfied. In light of this flexibility, it is not possible to pin down the surplus exactly. For instance, the analyst can assume perfect optimization and rationalize the demand in the standard way to obtain the usual consumer surplus. At the opposite extreme, the analyst can attribute all stochasticity in behavior to errors by assuming a single type whose demand matches the observed demand. Or the analyst can employ a richer model with many types that may be optimizing to varying degrees.

We characterize the levels of consumer surplus (and counterfactual demands) across all rationalizations of the data. The levels of surplus consistent with the data comprise an interval ranging from 0 to an upper bound that has a simple mathematical structure akin to that of the standard consumer surplus. Just as the standard surplus is the area between the price line and the inverse demand up to the quantity demanded, the upper bound is the area between the price line and an “elevated” inverse demand up to the quantity demanded. As the name suggests, this elevated demand—which depends on both the observed demand and the price distribution—lies above the observed demand.

At first blush, it may be surprising that the consumer surplus could be higher than the standard surplus under perfect optimization. To see how this may happen, consider a consumer who believes that the price of the good is normally too high for it to be worth purchasing. She therefore ignores it and checks the price only if she notices that it is on sale. If sales are announced at prices below some threshold that is lower than the consumer’s value for the good, an analyst who treats the consumer as if she is fully optimizing will infer that the threshold price for a sale is the consumer’s value, thereby underestimating her surplus.

Each feasible level of surplus can be obtained with a simple model: as in the standard approach, the observed demand is decomposed into threshold demands of individual types. However, the value of a type that does not perfectly optimize need not be equal to the price at the corresponding threshold. Our rationality conditions imply bounds on the value that a type with a given demand threshold could have: the value can be no greater than the expected price conditional on not buying the good, and no lower than the expected price conditional on buying the good. The full range of levels of consumer surplus consistent with the data can be attained by varying the types’ values within these bounds.

It turns out to be useful to focus on the value a randomly chosen type assigns to the good, which we refer to as the *stochastic value*. (Thus a stochastic value only partially describes a model in that it does not specify the demand of each type.) Our bounds on surplus and counterfactual demand are based on bounds on the stochastic value with respect to various stochastic orders. For the upper bound on consumer surplus, we make use of the increasing convex order (ICX);⁴ for the lower bound on surplus, we use second-order stochastic dominance (SOSD); and for the bounds on counterfactual demand, we use first-order stochastic dominance (FOSD).

Bounds on the stochastic value are particularly useful with richer data. In section 6, we consider datasets comprising two or more market regimes that may differ in the distribution of prices and/or the consumer’s behavior at any given price. For example, it could be that, as in Chetty, Looney, and Kroft (2009), sales taxes are included in the posted price in one regime but not included in the other. The analyst

⁴The ICX order can be viewed as the analogue of second-order stochastic dominance for a decision-maker who is risk loving instead of risk averse. Thus whereas second-order stochastic dominance favors higher means and smaller spreads, ICX favors higher means and larger spreads.

considers all rationalizations of the datasets in which the value of each type is fixed across regimes (though its demand may vary, for example due to changes in salience or attention); in other words, the stochastic value is held constant across regimes.

We propose a simple procedure for narrowing the bounds on surplus or counterfactual demand within each regime using the data from the other regimes. This procedure involves taking the collection of bounds on the stochastic value across regimes and combining them to obtain a common tighter bound. (In particular, this procedure can give rise to a nonzero lower bound.) To compute this combined bound, we exploit a mapping between stochastic values and convex functions and identify a common bound on these convex functions.

In a similar spirit to Bernheim and Rangel (2009), we propose a revealed-preference approach to measuring the welfare of a decision-maker who may not be perfectly rational. Empirical studies of behavioral welfare typically assume, either explicitly or implicitly, that certain observed choices reveal the decision-maker’s true preferences.⁵ These preferences can then be used to assess the welfare associated with other choices that could be suboptimal. For example, Chetty, Looney, and Kroft (2009) and Taubinsky and Rees-Jones (2018) recover true preferences from consumer choices when a tax is made salient and use these to measure the welfare loss arising from mistakes that occur when the taxes are not salient; Bronnenberg, Dubé, Gentzkow, and Shapiro (2015) identify the true preferences from experts’ choices and use these to evaluate non-experts’ welfare. An alternative approach, taken by Gruber and Köszegi (2001), is to use a structural model that relates true preferences to choice behavior. Relative to these empirical studies, we make much weaker assumptions about the extent to which true preferences can be inferred from the data, requiring only that behavior satisfies certain minimal rationality conditions.

In the special case of our model in which errors are due to imperfect information, our work can be viewed as combining revealed preference with information design, where the design has the goal of maximizing or minimizing the surplus or counterfactual demand consistent with the observed data.⁶ Bergemann, Brooks, and Morris (2022) identify bounds on counterfactual behavior in abstract games. Theorem 4 in

⁵See Bernheim and Taubinsky (2018) for a survey.

⁶While information design problems typically place no restrictions on the information structure, we impose an implicit restriction to ensure that each type has monotone demand.

the present paper concerns counterfactual behavior in a more specific setting, but unlike in Bergemann, Brooks, and Morris (2022), the distribution of preferences in our model is not known to the analyst. Bergemann, Brooks, and Morris (2015, 2017) identify the range of surplus values that can be attained for given preferences as information varies in a monopolistic market or a first-price auction. Condorelli and Szentes (2020, 2022) characterize the range of surplus values consistent with partial knowledge of demand in settings with market power on the supply side. Regarding revealed preference, we are closest to the branch of the literature that uses choice data to jointly identify preferences and information, as in Masatlioglu, Nakajima, and Ozbay (2012) and Manzini and Mariotti (2014).

When considering bounds on surplus using data from multiple regimes, we represent random variables as convex functions to construct bounds with respect to the ICX or SOSD order. A similar technique has been used in Bayesian persuasion problems by Gentzkow and Kamenica (2016) and Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017). Müller and Scarsini (2006) establish lattice properties of these orders using the same transformation. This technique has a natural interpretation in our context: the convex function that represents a given stochastic value maps each price to the consumer surplus that would arise under that stochastic value under perfect optimization. Kleiner, Moldovanu, and Strack (2021) characterize the extreme points of a set of distributions bounded by a random variable with respect to the ICX or SOSD order.⁷ In contrast, we identify random variables that provide bounds with respect to these orders on a set satisfying certain constraints.

Several other papers study complementary problems of identifying bounds on consumer surplus. Varian (1985) considers issues of measurement due to gaps in observed demand. Kang and Vasserman (2022) analyze how functional form restrictions narrow these bounds. Sandomirskiy and Ushchev (2022) identify bounds on consumer welfare across possible disaggregations of an observed aggregate demand curve without quasilinear utility. Allen and Rehbeck (2021) provide bounds on surplus from finite data sets for consumers who approximately optimize; in contrast, we focus on idealized infinite data with different bounded rationality assumptions that are not nested with theirs.

⁷Relatedly, Yang and Zentefis (2023) characterize extreme points of a set of distributions bounded with respect to the FOSD order.

2 Setup

An analyst observes *data* (Q, F) describing the stochastic purchasing behavior of a consumer with unit demand together with the distribution of prices. The demand function $Q : [\underline{p}, \bar{p}] \rightarrow [0, 1]$, which we assume is non-increasing, specifies the probability $Q(p)$ of purchase at each price p ; we denote by $P(q)$ the inverse demand associated with $Q(p)$.⁸ Prices are distributed according to the continuous distribution $F(p)$ with support $[\underline{p}, \bar{p}]$, where $\underline{p} \geq 0$. As is standard when measuring consumer welfare, we assume that the analyst observes the choke price, i.e., $Q(\bar{p}) = 0$; similarly, we assume that $Q(\underline{p}) = 1$.

The demand Q is an aggregation of many choices made by the consumer across which both her value for the good and her behavior may vary. In each such choice, the consumer faces a take-it-or-leave-it offer at a random price \mathbf{p} drawn according to F . (We denote random variables in bold and their realizations with the corresponding non-bold symbol; all probabilities and expectations are evaluated with respect to these bold variables.) The consumer has a stochastic type \mathbf{i} with support $I \subset \mathbb{R}$. Each type i specifies the consumer's value v_i for the good together with a non-increasing demand $Q_i(p)$.

We interpret the data as describing the choices of a single individual whose value for the good may be changing and whose behavior also varies due to unobserved factors such as attention or salience. We allow for the possibility that these factors are related to the value since, for example, the consumer may be more attentive to the price when their value is lower. An alternative interpretation of the data is that they combine choices made by a large population of consumers, with each type corresponding to a distinct individual whose value and behavior are fixed.

We assume that the consumer's type is independent of the price. This assumption is natural for individual-level data, such as scanner data, where the individual is negligible from the perspective of the seller. If consumers' values vary systematically over time—as, for instance, would be expected for a seasonal product—then the analyst should split the data into periods within which independence of values and prices is plausible. (The tools developed in section 6 can be applied to the separated

⁸That is, $P(q) := \inf\{p : Q(p) \leq q\}$. Analogously, given any inverse demand function \tilde{P} , we define the corresponding demand function $\tilde{Q}(p) = \inf\{q : \tilde{P}(q) \leq p\}$.

data.) If, however, the analyst fails to account for temporal variation that affects the seller’s pricing strategy, then our results may not apply. Similarly, if the seller engages in third-degree price discrimination—or otherwise screens consumers by offering different distributions of prices depending on attributes correlated with the consumer’s value—the analyst should group the data according to the same attributes; our results may not apply to the original combined data.

The analyst assumes that all types of the consumer satisfy minimal rationality restrictions: we require that no type can be worse off than she would be if she never bought the good (regardless of its realized price), nor can she be worse off than if she always bought the good. Letting $s_i = E[(v_i - \mathbf{p}) Q_i(\mathbf{p})]$ denote type i ’s expected surplus, this restriction corresponds to the pair of incentive compatibility constraints

$$s_i \geq 0 \tag{1}$$

$$\text{and} \quad s_i \geq v_i - E[\mathbf{p}]. \tag{2}$$

This requirement imposes discipline on the relationship between a type’s value and its demand: substituting the definition of s_i and isolating v_i yields

$$\frac{E[\mathbf{p}(1 - Q_i(\mathbf{p}))]}{E[1 - Q_i(\mathbf{p})]} \geq v_i \geq \frac{E[\mathbf{p}Q_i(\mathbf{p})]}{E[Q_i(\mathbf{p})]}. \tag{3}$$

There are several reasons to expect that a consumer’s demand may not perfectly reflect her value. It could be that she does not always check the price of the good, or does so only if she notices that it is on sale; the posted price may not include taxes that the consumer does not accurately compute; or the consumer may make random errors in assessing the value of the good, as in a random utility model. In each of these cases, the consumer’s attentiveness and likelihood of making a mistake could depend on the current value. The analyst therefore allows for types’ demands to vary along with their values in a general way, imposing only that no type makes systematic errors such that they would be better off either always buying or never buying.

The analyst seeks to explain the observed choices with a *model* that consists of a distribution M of types $i \in I$ together with a specification $(v_i, Q_i)_{i \in I}$ of values and non-increasing demand functions for each type satisfying (1) and (2). We say that a given model *rationalizes* data (Q, F) if $Q(p) = E[Q_i(p)]$ for all p . Given a model that

rationalizes the data, the (ex ante) *consumer surplus* is $s = \mathbb{E}[s_i]$. In general, data can be rationalized by many different models which in turn yield different values of surplus. We say that surplus $s \in \mathbb{R}$ is *consistent with the data* if there exists a model that rationalizes (Q, F) and generates surplus s .

A special case of our environment that may be of particular interest arises when the consumer is Bayesian but observes only a noisy signal of the price of the good. The noise in this signal could result from inattention; for example, the consumer may assume that the price of the good exceeds her willingness to pay unless she notices that it is on sale (in which case she checks the price). In this case, a type of the consumer can be described by a value and an information structure. The type's demand is then determined by the condition that she buys precisely when her value exceeds her posterior expectation of the price. Bayesian optimality implies that (1) and (2) are satisfied for each type. Conversely, for each type (v_i, Q_i) satisfying (1) and (2), there exists an information structure for which a Bayesian type with value v_i would have demand Q_i . Namely, given a type (v_i, Q_i) , take the binary information structure that generates a “buy” signal with probability $Q_i(p)$ and an “abstain” signal otherwise. By (1) and (2), following the action recommended by the signal is incentive compatible. Consequently, all of our results apply as written to this special case; in particular, Bayesian optimality does not narrow the bounds on consumer surplus relative to those obtained with our minimal rationality assumptions.

Example 1. The analyst observes the linear demand function $Q(p) = 1 - p$ and prices uniformly distributed on $[0, 1]$. There are many possible models that rationalize this data. For instance, it could be that, as in the standard analysis, the consumer always makes the optimal decision: her value v_i is uniformly distributed on $[0, 1]$ and each type i demands the good precisely when $p < v_i$. For any realized price p , this consumer receives surplus $(1 - p)^2/2$ (corresponding to the area between the demand curve and the price). The expected consumer surplus for this model is therefore $s = \mathbb{E}[(1 - p)^2/2] = 1/6$.

Alternatively, the data can be rationalized by a model with stochastic choices. Perhaps the simplest such rationalization features a consumer with a single type. For each price realization p , the consumer purchases the good with probability $Q(p)$, which trivially generates the observed aggregate demand. The inequalities in (3)

place limits on this type's value, v : it must be at least $\underline{v} = E[\mathbf{p}Q(\mathbf{p})] / E[Q(\mathbf{p})] = 1/3$ to ensure that she does not prefer to abstain from buying, and at most $\bar{v} = E[\mathbf{p}(1 - Q(\mathbf{p}))] / E[1 - Q(\mathbf{p})] = 2/3$ to ensure that she does not prefer to always buy. Taking $v = \bar{v}$ leads to a surplus of $E[(\bar{v} - \mathbf{p})Q(\mathbf{p})] = 1/6$; taking $v = \underline{v}$ leads to a surplus of $E[(\underline{v} - \mathbf{p})Q(\mathbf{p})] = 0$. Using values of v in between these two extremes, any surplus in $[0, 1/6]$ can be obtained.⁹

More complex models can yield additional values of the surplus. Consider two equally likely types, 1 and 2, with respective demands

$$Q_1(p) = \begin{cases} 2(1-p) & \text{if } p \geq 1/2 \\ 1 & \text{otherwise} \end{cases} \quad \text{and} \quad Q_2(p) = \begin{cases} 0 & \text{if } p \geq 1/2 \\ 1-2p & \text{otherwise.} \end{cases}$$

Since $(Q_1 + Q_2)/2 = Q$, these types' demands together generate the observed total demand. Given each type's demand Q_i , (3) places restrictions on the values of the form $v_i \in [\underline{v}_i, \bar{v}_i]$. Using the maximal values gives surplus $\frac{1}{2} \sum_i \bar{v}_i E[Q_i(\mathbf{p})] - E[\mathbf{p}Q(\mathbf{p})] = 2/9$ (whereas the minimal values again give a surplus of 0).¹⁰ \triangle

We see from this example that it is possible to obtain values of the surplus consistent with the data that exceed the standard consumer surplus. This observation may seem surprising since the consumer in the standard model perfectly optimizes while in other models she does not. While it is true that introducing imperfections in decision-making can only lower the surplus *given the consumer's preferences*, allowing for these imperfections expands the range of preferences that can rationalize the data.

In the example, by splitting the aggregate demand across two types, we obtain higher levels of surplus than can be obtained with only one type. The question remains, however, as to whether further disaggregation of the two demands—or some other splitting—can expand the range of attainable surpluses. As we show in the next sections, it turns out that a “maximal” disaggregation of the demand can give rise to the highest value of the surplus (which is $1/4$ for Example 1).

⁹That the upper bound of $1/6$ is equal to the standard surplus is a coincidence that does not generally hold outside of this example. On the other hand, 0 is a tight lower bound regardless of the data as there are always models in which each type of the consumer is indifferent between choosing according to her demand and never buying the good.

¹⁰Type 1 buys with probability $3/4$ and has maximal value $\bar{v}_1 = 5/6$ and type 2 buys with probability $1/4$ and has maximal value $\bar{v}_2 = 11/18$, giving a surplus of $1/2 \cdot 5/6 \cdot 3/4 + 1/2 \cdot 11/18 \cdot 1/4 - 1/6 = 2/9$.

3 Bounds on Consumer Surplus with One Dataset

We identify tight bounds on the consumer surplus consistent with the observed data. To formulate the result, we define, for arbitrary demand \tilde{Q} and (possibly unrelated) inverse demand \hat{P} , the functional

$$\mathcal{CS}(\tilde{Q}, \hat{P}; p) := \int_0^{\tilde{Q}(p)} (\hat{P}(q) - p) dq. \quad (4)$$

When applied to the observed demand Q and its inverse demand P , $\mathcal{CS}(Q, P; p)$ returns the standard consumer surplus. In the standard case, when the consumer always chooses optimally, the inverse demand is equal to the marginal benefit of consumption at each q . If the consumer does not always choose optimally, the inverse demand is not generally equal to the marginal benefit. Nonetheless, if \tilde{Q} is the demand and \hat{P} the marginal benefit of consumption, then $\mathcal{CS}(\tilde{Q}, \hat{P}; p)$ is the consumer surplus (at price p).

For any data (Q, F) , we provide tight bounds on the consumer surplus using $\mathcal{CS}(Q, \hat{P}; p)$ for appropriate choices of \hat{P} . Accordingly, define the *elevated* and *lowered* inverse demands to be

$$\begin{aligned} \overline{P}(q) &:= \mathbb{E}[\mathbf{p} \mid \mathbf{p} \geq P(q)] \\ \text{and} \quad \underline{P}(q) &:= \mathbb{E}[\mathbf{p} \mid \mathbf{p} \leq P(q)], \end{aligned}$$

respectively. These two functions are non-increasing and satisfy $\overline{P}(q) \geq P(q) \geq \underline{P}(q)$ for all q ; see Figure 1 for an illustration.

Theorem 1. *Consumer surplus s is consistent with data (Q, F) if and only if*

$$0 = \mathbb{E}[\mathcal{CS}(Q, \underline{P}; \mathbf{p})] \leq s \leq \mathbb{E}[\mathcal{CS}(Q, \overline{P}; \mathbf{p})].$$

As explained above, $\mathcal{CS}(\tilde{Q}, \hat{P}; p)$ is the surplus obtained by a consumer who demands $\tilde{Q}(p)$ at each p and receives marginal benefit $\hat{P}(q)$ at each q . The upper bound in Theorem 1 corresponds to the “highest possible” marginal benefit function consistent with the data, in a sense that is made precise in the proof of the theorem in section 5. To obtain the marginal benefit function \overline{P} , we construct a model




Figure 1: Elevated and lowered inverse demands for a particular inverse demand $P(q)$ with uniformly distributed prices. In this case, $\bar{P}(q)$ is the midpoint between $P(q)$ and \bar{p} . Likewise, $\underline{P}(q)$ is the midpoint between $P(q)$ and \underline{p} .

rationalizing the data in which the analyst associates every marginal increase in demand with a different consumer type and takes each type's value to be the largest one that is consistent with its own demand. A similar construction with lower values can generate any surplus between 0 and the upper bound.¹¹

Example 2. To illustrate the theorem, consider the data from Example 1. The empirical inverse demand is $P(q) = 1 - q$, the elevated inverse demand is $\bar{P}(q) = (1 + P(q))/2 = 1 - q/2$, and $E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})] = 1/4$. Theorem 1 therefore indicates that the surplus levels consistent with the data are precisely those in the interval $[0, 1/4]$.

Theorem 1 also shows how the upper bound depends on both the demand and the distribution of prices in the data. By fixing the observed demand at $Q(p) = 1 - p$ and varying the distribution of prices among those with support $[0, 1]$, one can obtain an upper bound arbitrarily close to the standard consumer surplus at a price p_0 if the probability that the price is below p_0 is small and the density at prices above p_0 decreases quickly, so that for each $p \in (p_0, \bar{p})$, $E[\mathbf{p} \mid \mathbf{p} > p]$ is close to p . \triangle

4 Stochastic values

The distribution of the consumer's value for the good plays a central role in our analysis. In this section, we explain how, in viewing the value as a random variable,

¹¹In subsection 6.2, we show how positive lower bounds can be obtained using richer data.

comparisons of counterfactual demand or consumer surplus correspond to comparisons of random variables with respect to a relevant stochastic order. The proofs of our main results exploit this connection.

Given a model, let $\mathbf{v} := v_i$ be the consumer's *stochastic value* of the good. Thus \mathbf{v} is a random variable partially describing the model, disregarding types' demands. Let $Q^s(p; \mathbf{v}) := \Pr(\mathbf{v} > p)$. For any price p , $Q^s(p; \mathbf{v})$ is the probability with which the consumer would buy the good in the standard model (except possibly at atoms of \mathbf{v}).¹² Therefore, we refer to $Q^s(p; \mathbf{v})$ as the *standard demand function* for \mathbf{v} . We use the superscript s throughout to indicate elements relating to the standard model.

Note that the standard demand function is the complementary distribution function of \mathbf{v} . Likewise, the standard inverse demand function $P^s(q; \mathbf{v})$ —which is the inverse to the demand $Q^s(p; \mathbf{v})$ —is the complementary quantile function of \mathbf{v} .

In light of the connection between the distribution of the stochastic value and the demand, first-order stochastic dominance comparisons of \mathbf{v} correspond to rankings of the associated standard demands. Indeed, the following statements are equivalent: (i) \mathbf{v}' first-order stochastically dominates \mathbf{v} ; (ii) $Q^s(p; \mathbf{v}') \geq Q^s(p; \mathbf{v})$ for all p ; and (iii) $P^s(q; \mathbf{v}') \geq P^s(q; \mathbf{v})$ for all q .¹³

Our bounds on consumer surplus make use of the so-called convex order. Following Rothschild and Stiglitz (1970), mapping random variables to convex functions is useful for making comparisons with respect to this order. Define the function $CS^s(\cdot; \mathbf{v}) : \mathbb{R} \rightarrow \mathbb{R}$ by

$$CS^s(p; \mathbf{v}) := \int_p^\infty Q^s(p'; \mathbf{v}) dp'. \quad (5)$$

This mapping has a natural interpretation in our context: it is the standard consumer surplus at price p for a consumer with stochastic value \mathbf{v} .¹⁴ Observe that $CS^s(p; \mathbf{v})$ is convex in p because Q^s is downward-sloping.

In addition to its economic interpretation, the function $CS^s(\cdot; \mathbf{v})$ characterizes the convex order on \mathbf{v} . Given two real-valued random variables \mathbf{x} and \mathbf{y} , \mathbf{y} *dominates* \mathbf{x}

¹²If \mathbf{v} has an atom at p , then the demand in the standard model lies in the closed interval between the left and right limits of $Q^s(\cdot; \mathbf{v})$ at p .

¹³We use the first-order stochastic dominance order in section 7 where we study bounds on counterfactual demand that would arise in the absence of imperfections in choice.

¹⁴Gentzkow and Kamenica (2016) and Kolotilin et al. (2017) map each random variable to the integral of the *lower* tail of its distribution function. For our purposes, the relevant integral is over the *upper* tail of the *complementary* distribution function.

in the increasing convex order, denoted by $\mathbf{y} \succeq_{\text{icx}} \mathbf{x}$, if there exists a random variable \mathbf{z} such that \mathbf{z} first-order stochastically dominates \mathbf{x} and \mathbf{y} is a mean-preserving spread of \mathbf{z} .

Lemma 1. *For any \mathbf{v}' and \mathbf{v} , $\mathbf{v}' \succeq_{\text{icx}} \mathbf{v}$ if and only if $CS^s(p; \mathbf{v}') \geq CS^s(p; \mathbf{v})$ for every price p .*

Proof. The result follows from Theorem 4.A.2 of Shaked and Shanthikumar (2007) together with the fact that $Q^s(p; \mathbf{v})$ is the complementary distribution function of \mathbf{v} . \square

The increasing convex order is closely related to second-order stochastic dominance, denoted here by \succeq_{sosd} .¹⁵ (Indeed, $\mathbf{y} \succeq_{\text{icx}} \mathbf{x}$ if and only if $-\mathbf{x} \succeq_{\text{sosd}} -\mathbf{y}$.) Roughly speaking, both orders favor higher values, but the increasing convex order favors spreads while second-order stochastic dominance disfavors them. Lemma 1 is essentially the analogue for the increasing convex order of the usual characterization of SOSD in terms of integrals of the lower tails of distribution functions.

Given a stochastic value \mathbf{v} and demand $\tilde{Q}(p)$, a special role in our analysis is played by $CS\left(\tilde{Q}, P^s(\cdot; \mathbf{v}); p\right)$. This quantity is the highest possible surplus a consumer with value \mathbf{v} and demand $\tilde{Q}(p)$ can achieve, which is attained when the measure $\tilde{Q}(p)$ of the highest types are the ones that buy the good. We abuse notation and write $CS(\tilde{Q}, \mathbf{v}; p)$ for $CS\left(\tilde{Q}, P^s(\cdot; \mathbf{v}); p\right)$ throughout.

The next result suggests how the increasing convex order can be useful outside of the standard model, for a consumer who makes imperfect choices.

Lemma 2. *For every demand function \tilde{Q} and any price p , $CS(\tilde{Q}, \mathbf{v}; p)$ is nondecreasing in \mathbf{v} with respect to the increasing convex order.*

Clearly, a first-order stochastic dominance increase of values increases the consumer surplus $CS(\tilde{Q}, \mathbf{v}; p)$. A mean-preserving spread of the values also increases this surplus because it is computed under the assumption that, for each p , it is the measure $\tilde{Q}(p)$ of types with the highest values that buy. The gross surplus at a given price is therefore proportional to the mean value conditional on being among these buying

¹⁵Recall that \mathbf{y} second-order stochastically dominates \mathbf{x} if there exists \mathbf{z} such that \mathbf{x} is a mean-preserving spread of \mathbf{z} and \mathbf{y} first-order stochastically dominates \mathbf{z} .

types, which increases with a mean-preserving spread. The proof of this lemma—and those of other results not proved in the main text—may be found in the appendix.

5 Proof of Theorem 1

We begin by identifying the set of values consistent with the demand of a given type. We say that a value v_i is *consistent with demand* Q_i if v_i together with Q_i satisfy inequalities (1) and (2). Given Q_i , let

$$\begin{aligned} \underline{v}_i &:= \mathbb{E}[\mathbf{p} \mid Q_i(\mathbf{p}) \geq \mathbf{q}] = \mathbb{E}[\mathbf{p} \mid \mathbf{p} \leq P_i(\mathbf{q})] \\ \text{and} \quad \bar{v}_i &:= \mathbb{E}[\mathbf{p} \mid Q_i(\mathbf{p}) \leq \mathbf{q}] = \mathbb{E}[\mathbf{p} \mid \mathbf{p} \geq P_i(\mathbf{q})], \end{aligned}$$

where $\mathbf{q} \sim U[0, 1]$ independent of \mathbf{p} and P_i is the inverse demand associated with Q_i . Since type i buys with probability $Q_i(p)$ at each price p , \underline{v}_i and \bar{v}_i are, respectively, the expected price conditional on the event that consumer of type i does or does not make a purchase. Accordingly, we refer to \underline{v}_i as the *buying price expectation* and to \bar{v}_i as the *non-buying price expectation*. Note that, since Q_i is downward sloping, $\underline{v}_i \leq \bar{v}_i$.

Lemma 3. *A value v_i is consistent with Q_i if and only if $\underline{v}_i \leq v_i \leq \bar{v}_i$.*

This result follows directly from (3) since

$$\mathbb{E}[\mathbf{p} \mid Q_i(\mathbf{p}) \leq \mathbf{q}] = \frac{\mathbb{E}[\mathbf{p}(1 - Q_i(\mathbf{p}))]}{\mathbb{E}[1 - Q_i(\mathbf{p})]}$$

and

$$\mathbb{E}[\mathbf{p} \mid Q_i(\mathbf{p}) \geq \mathbf{q}] = \frac{\mathbb{E}[\mathbf{p}Q_i(\mathbf{p})]}{\mathbb{E}[Q_i(\mathbf{p})]}.$$

We divide Theorem 1 into its sufficiency and necessity claims. To prove sufficiency, we first show by construction that each surplus between 0 and $\mathbb{E}[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]$ is consistent with the data. For the necessity claim, we prove that no other levels of surplus are consistent with the data. For the latter, we show that the stochastic value associated with the construction yielding the upper bound on surplus provides an upper bound with respect to the increasing convex order. (For necessity, it suffices to consider only the upper bound: since the lower bound on surplus is 0, it follows trivially from (1) that no lower surplus can be obtained.)

To prove the sufficiency claim, we construct for each $s \in [0, E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]]$ a model that generates surplus s . Let the type \mathbf{i} be uniformly distributed on $[0, 1]$ and let each realization i generate the demand function $Q_i(p) = \mathbb{1}_{p \leq P(i)}$. Thus type i always buys when the price is below $P(i)$ and never buys at prices above $P(i)$. Note that the average demand across all types is equal to the observed demand Q , as needed for the model to rationalize the data:

$$E[Q_{\mathbf{i}}(p)] = \Pr(P(\mathbf{i}) \geq p) = \Pr(\mathbf{i} \leq Q(p)) = Q(p).$$

By Lemma 3, a value v_i is consistent with demand Q_i if $\underline{v}_i \leq v_i \leq \bar{v}_i$. Due to the choice of Q_i , we have $\underline{v}_i = \underline{P}(i)$ and $\bar{v}_i = \bar{P}(i)$. Since type i buys if and only if $i \leq Q(p)$, taking $v_i = \bar{v}_i$ for all i gives ex ante surplus $E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]$; we refer to this model as the *upper threshold model*. At the other extreme, taking $v_i = \underline{v}_i$ for all i gives $E[\mathcal{CS}(Q, \underline{P}; \mathbf{p})] = 0$ since the value $\underline{P}(i)$ of each type i is equal to its expected expenditure. For any $s \in (0, E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})])$, taking $v_i = \lambda \bar{P}(i) + (1 - \lambda) \underline{P}(i)$ with $\lambda = s / E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]$ yields surplus s . This completes the proof that lying in the interval $[0, E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]]$ is sufficient for s to be consistent with the data.

We now shift our attention to the other direction, namely, that lying in the interval $[0, E[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]]$ is a necessary condition for consistency of surplus with the data. We say that a stochastic value \mathbf{v} is *consistent with data* (Q, F) if there exists a model satisfying $\mathbf{v} \stackrel{d}{=} v_{\mathbf{i}}$ that rationalizes (Q, F) . We provide an upper bound on stochastic values consistent with the data with respect to the increasing convex order. Let $\bar{\mathbf{v}} := \bar{P}(\mathbf{i})$ and $\underline{\mathbf{v}} := \underline{P}(\mathbf{i})$ for $\mathbf{i} \sim U[0, 1]$. Thus $\bar{\mathbf{v}}$ and $\underline{\mathbf{v}}$ are, respectively, the stochastic values associated with the upper threshold model and the corresponding model for the lower bound constructed above. For the proof of Theorem 1, we make use only of $\bar{\mathbf{v}}$; $\underline{\mathbf{v}}$ is needed in subsection 6.2.

The following lemma is the core technical insight underlying the necessity part of Theorem 1.

Lemma 4. *If a stochastic value \mathbf{v} is consistent with data (Q, F) , then $\bar{\mathbf{v}} \succeq_{\text{icx}} \mathbf{v}$.*

The proof of this lemma, which is in the appendix, starts by considering an arbitrary model rationalizing the data with values v_i and demands $Q_i(p)$ for each type i . We then amend the model in two steps such that (i) each step leads to an in-




Figure 2: Example illustrating the effect of splitting the demand $Q(p) = 1 - p$ into threshold demands for $\mathbf{p} \sim U[0, 1]$. The shaded triangle represents the region of non-buying prices. If the demand is that of a single type, the maximal value is the non-buying price expectation $E[\mathbf{p} \mid \mathbf{p} \geq P_i(\mathbf{q})]$, which is located at the height of the centroid of the shaded triangle. When the demand is split into threshold demands, each value of q is identified with a distinct type $j = q$ that has a threshold demand with threshold $1 - j$ (and thus has non-buying price expectation $E[\mathbf{p} \mid \mathbf{p} \geq P_i(j)]$). The expected maximal value then becomes $E[\bar{v}_j] = \int_0^1 E[\mathbf{p} \mid \mathbf{p} \geq P_i(j)] dj$.

crease in the stochastic value with respect to the increasing convex order and (ii) in combination, the two steps transform the original stochastic value $\mathbf{v} = v_i$ to $\bar{\mathbf{v}}$.

In the first step, we replace the value v_i of each type i with i 's non-buying price expectation \bar{v}_i (given the demand Q_i). Since, by Lemma 3, $\bar{v}_i \geq v_i$, this replacement leads to a first-order stochastic dominance increase in the stochastic value and hence also to an increase with respect to the increasing convex order.

In the second step, we decompose the demand of each type into demands of the form $Q_j(p) = \mathbb{1}_{p \leq \rho_j}$ for some ρ_j ; we refer to such functions as *threshold demands*. More specifically, we replace each type i with a stochastic type \mathbf{j} such that each realization j has a threshold demand and the average demand across \mathbf{j} is Q_i . (If Q_i is itself a threshold demand, then such a decomposition is trivial.) We assign to each j the value \bar{v}_j equal to its non-buying price expectation. Replacing each i with the corresponding \mathbf{j} clearly increases the spread in the values. For this change to be an increase with respect to the increasing convex order, it suffices to show that it also increases the means, i.e., that $E[\bar{v}_j] \geq \bar{v}_i$ for each i . To see why the last inequality holds, notice that, by the Law of Iterated Expectations, the expected non-buying price expectation *conditional on not buying* is unaffected by the decomposition of the demand $Q_i(p)$; that is, since $\mathbf{p} \geq P_j(\mathbf{q})$ is the condition under which type j does not

buy,

$$\bar{v}_i = \mathbb{E}[\bar{v}_{\mathbf{j}} \mid \mathbf{p} \geq P_{\mathbf{j}}(\mathbf{q})].$$

Since higher values of \bar{v}_j are associated with lower probabilities of not buying, when compared to the conditional expectation on the right-hand side of the last equation, the relative weight assigned to higher values \bar{v}_j in the unconditional expectation is larger, and therefore $\mathbb{E}[\bar{v}_{\mathbf{j}}] \geq \mathbb{E}[\bar{v}_{\mathbf{j}} \mid \mathbf{p} \geq P_{\mathbf{j}}(\mathbf{q})] = v_i$. See Figure 2 for an illustration.

Taken together, the two steps transform the original model into another one that rationalizes the data and in which all types have threshold demands and values equal to their non-buying price expectations. The associated stochastic value is $\bar{\mathbf{v}}$, as needed for the proof of Lemma 4.

We now establish an interim upper bound on consumer surplus that holds for each realization p of the random price \mathbf{p} . We say that a function $s(p)$ is an *interim consumer surplus consistent with the data* if there exists a model that rationalizes the data for which

$$s(p) = \mathbb{E}[(v_{\mathbf{i}} - p) Q_{\mathbf{i}}(p)].$$

Lemma 5. *If $s(p)$ is an interim consumer surplus consistent with the data, then $s(p) \leq \mathcal{CS}(Q, \bar{P}; p)$.*

Proof. First note that, for a given stochastic value \mathbf{v} , the consumer can suffer from two types of losses relative to optimal behavior: (i) the probability of purchase at a given price may not be optimal, i.e., $Q(p)$ may differ from $Q^s(p; \mathbf{v})$, and (ii) the set of types purchasing the good at a given price may not be those with the highest values. Starting from any model that rationalizes the data, reallocating demands across types to eliminate this latter loss (ignoring incentive compatibility) gives an upper bound $\mathcal{CS}(Q, \mathbf{v}; p)$ on the interim surplus at each p for models with stochastic value \mathbf{v} . Therefore, the surplus $s(p)$ generated by any such model satisfies

$$s(p) \leq \mathcal{CS}(Q, \mathbf{v}; p) \leq \mathcal{CS}(Q, \bar{\mathbf{v}}; p) = \mathcal{CS}(Q, \bar{P}; p),$$

where the middle inequality follows from Lemmas 2 and 4. □

Since the interim surplus is bounded from above by $\mathcal{CS}(Q, \bar{P}; p)$ for each price p , the ex ante surplus is bounded by $\mathbb{E}[\mathcal{CS}(Q, \bar{P}; \mathbf{p})]$, as needed. This concludes the proof of Theorem 1.

6 Multiple Datasets

The bounds on consumer surplus can be narrowed if the analyst observes the consumer's choices under varying market conditions, which we refer to as *regimes*. We assume that the consumer's types' values are fixed across regimes, but the regimes may differ in the distribution of prices or in the purchasing behavior of each type at any given price (or both). For example, one such regime may correspond to a publicly announced "sale" associated with a low distribution of prices while another corresponds to the same market in the absence of a sale; the sale announcement may affect the consumer's stochastic choice at each price through changes in attention or salience. Alternatively, the regimes may differ only in how prices are presented to consumers, as in the empirical studies of Chetty, Looney, and Kroft (2009) and Finkelstein (2009).

The analyst observes a *profile* of datasets (Q^k, F^k) , $k = 1, \dots, K$, where $Q^k(p)$ and $F^k(p)$ are, respectively, the probability that the consumer makes a purchase at each price p and the distribution of prices in regime k , and each (Q^k, F^k) satisfies the assumptions on data made in section 2. The consumer has a stochastic type \mathbf{i} , with each realization i specifying her value v_i for the good and her (nonincreasing) demand function $Q_i^k(p)$ in each regime. The distribution of types and the value of each type are the same across all regimes. A model for the analyst consists of a distribution of types together with a specification of $(v_i, Q_i^1, \dots, Q_i^K)$ for each type i .

We say that a model *rationalizes the profile of datasets* $(Q^k, F^k)_k$ if, for each regime k , it rationalizes dataset (Q^k, F^k) when each type i has demand Q_i^k . In particular, within each regime, we impose the basic rationality assumptions described by (1) and (2). A stochastic value \mathbf{v} is *consistent with the profile of datasets* $(Q^k, F^k)_k$ if there exists a model that rationalizes this profile and satisfies $\mathbf{v} \stackrel{d}{=} v_{\mathbf{i}}$.

Example 3. Consider two regimes. The data in regime 1 consist of the linear demand $Q^1(p) = 1 - p$ and uniform price distribution $\mathbf{p} \sim U[0, 1]$, as in Example 1. The data in regime 2 consist of the step-function demand $Q^2(p) = \mathbb{1}_{p \leq 2/3}$ and uniform price distribution $\mathbf{p} \sim U[2/3 - \varepsilon, 2/3 + \varepsilon]$, where $0 < \varepsilon \leq 1/3$. These two regimes are jointly rationalizable, for example by a single type with value $2/3$ and demand $Q^k(p)$ for $k = 1, 2$. \triangle

Given a model, the *interim consumer surplus in regime k at price p* is

$$s^k(p) = \mathbb{E} [Q_{\mathbf{i}}^k(p) (v_{\mathbf{i}} - p)] .$$

We write $s^k = \mathbb{E} [s^k(\mathbf{p})]$ for the ex ante surplus in regime k (with $\mathbf{p} \sim F^k$). Interim consumer surplus $s^k(p)$ in regime k is *consistent with the profile of datasets $(Q^k, F^k)_k$* if there exists a model that rationalizes this profile and generates surplus $s^k(p)$ in regime k (and analogously for the ex ante surplus).

6.1 Upper bound

The next result provides an upper bound on the surplus within each regime that generally improves upon the bounds that can be obtained for each regime separately. The basic idea is to derive the upper bounds on the stochastic value with respect to the increasing convex order when considering each regime separately and to combine them in such a way as to generate a tighter bound. The approach therefore requires combining bounds on random variables with respect to that stochastic order. To do so, building on ideas of Gentzkow and Kamenica (2016) and Kolotilin et al. (2017), we exploit the representation of a random variable in terms of a convex function described in section 4—in this case, the stochastic value in terms of the standard consumer surplus. According to Lemma 1, comparisons of stochastic values in the increasing convex order correspond to comparisons of the standard consumer surplus. Using this connection, we find the largest random variable that satisfies the bounds on the stochastic value across all of the regimes by finding the largest convex function lying below the corresponding bounds on the standard consumer surplus.

Let $\bar{\mathbf{v}}^k$ be the upper bound on stochastic values consistent with the data for regime k with respect to the increasing convex order, as in Lemma 4.¹⁶ For each k , the bound $\bar{\mathbf{v}}^k$ corresponds to the convex function $CS^s(p; \bar{\mathbf{v}}^k)$. The upper bound using data across all regimes therefore corresponds to the largest convex function that lies below each $CS^s(p; \bar{\mathbf{v}}^k)$. Accordingly, let $CS_*^s(p)$ denote the convex closure of the function $\min_k CS^s(p; \bar{\mathbf{v}}^k)$.¹⁷ We refer to CS_*^s as the *convexification* of $\min_k CS^s(p; \bar{\mathbf{v}}^k)$.

¹⁶That is, $\bar{\mathbf{v}}^k = \bar{P}^k(\mathbf{i})$ with $\mathbf{i} \sim U[0, 1]$, where the elevated demand for regime k is $\bar{P}^k(i) = \mathbb{E} [\mathbf{p} \mid \mathbf{p} \geq P^k(i)]$ with $\mathbf{p} \sim F^k$ and P^k is the inverse demand to Q^k .

¹⁷Recall that the convex closure of a function $g(p)$ is the function that maps each p to




Figure 3: Convexification for the two regimes described in Example 3 with $\varepsilon = 1/3$. Note that the graphs depict only prices $p \in [1/2, 1]$ since the convexification is trivial for $p < 1/2$. (a) Standard consumer-surplus functions: $CS^s(p; \bar{\mathbf{v}}^1)$ (dashed), $CS^s(p; \bar{\mathbf{v}}^2)$ (dotted), and the convexification $CS_*^s(p)$ (thick). (b) Standard demands associated with stochastic values: $Q^s(p; \bar{\mathbf{v}}^1)$ (dashed), $Q^s(p; \bar{\mathbf{v}}^2)$ (dotted), and $Q_*^s(p)$ (thick).

To map CS_*^s back to a stochastic value, recall from section 4 that $CS^s(p; \mathbf{v})$ is the integral of the upper tail of the standard demand $Q^s(p; \mathbf{v})$, which is the complementary distribution function of \mathbf{v} . Define the demand function $Q_*^s(p) = -\partial_- CS_*^s(p)$, where ∂_- denotes the left derivative. Note that $1 - Q_*^s$ is a distribution function and let $\bar{\mathbf{v}}_*$ be a stochastic value associated with this distribution.¹⁸ See Figure 3 for an illustration.

The following result is the main step underlying the upper bound for multiple regimes.

Lemma 6. *If a stochastic value \mathbf{v} is consistent with the profile of datasets, then $\bar{\mathbf{v}}_* \succeq_{\text{icx}} \mathbf{v}$.*

Proof. Follows from Theorem 3.2 of Müller and Scarsini (2006). \square

A direct argument is as follows. If \mathbf{v} is consistent with the profile of datasets, then it is consistent with each dataset separately; thus $\bar{\mathbf{v}}^k \succeq_{\text{icx}} \mathbf{v}$ for each regime k . By

$\inf \{s : (p, s) \in \text{co}(g)\}$, where $\text{co}(g)$ denotes the convex hull of the graph of the function g . In the terminology of convex analysis, CS_*^s is the biconjugate function to $\min_k CS^s(p; \bar{\mathbf{v}}^k)$.

¹⁸Since CS_*^s is convex, its left derivative exists, and Q_*^s is nonincreasing and left-continuous. Additionally, $CS_*^s(p) = 0$ for $p > \bar{p}$ and $CS_*^s(p)$ has slope -1 for $p < \underline{p}$; hence, $\lim_{p \rightarrow -\infty} Q_*^s(p) = 1$ and $\lim_{p \rightarrow +\infty} Q_*^s(p) = 0$. Thus, $1 - Q_*^s$ is a distribution function.

Lemma 1, $\min_k CS^s(p; \bar{\mathbf{v}}^k) \geq CS^s(p; \mathbf{v})$. Since $CS^s(p; \mathbf{v})$ is convex in p , $CS^s(p; \mathbf{v})$ is no greater than the convexification of $\min_k CS^s(p; \bar{\mathbf{v}}^k)$. Finally, again by Lemma 1, $\bar{\mathbf{v}}_* \succeq_{\text{icx}} \mathbf{v}$.

Combining Lemmas 2 and 6 leads to the following upper bound on the consumer surplus within each regime.

Theorem 2. *If the interim consumer surplus $s^k(p)$ in regime k is consistent with the profile of datasets, then $s^k(p) \leq \mathcal{CS}(Q^k, \bar{\mathbf{v}}_*; p)$.*

As an immediate consequence, the ex ante consumer surplus in regime k consistent with the profile of datasets is bounded from above by $\mathbb{E}[\mathcal{CS}(Q^k, \bar{\mathbf{v}}_*; \mathbf{p})]$, where the expectation is with respect to the distribution of prices in regime k .

Proof of Theorem 2. Given a stochastic value \mathbf{v} , the interim consumer surplus $s^k(p)$ in regime k is at most $\mathcal{CS}(Q^k, \mathbf{v}; p)$ because this is the surplus associated with having the measure $Q^k(p)$ of types with the highest values buy at p . By Lemma 2, $\mathcal{CS}(Q^k, \mathbf{v}; p)$ is nondecreasing in \mathbf{v} with respect to the increasing convex order. Finally, by Lemma 6, any stochastic value \mathbf{v} consistent with the profile of datasets is bounded by $\bar{\mathbf{v}}_*$ in the increasing convex order. \square

Lemma 6, when combined with Lemma 1, also provides an upper bound on the counterfactual consumer surplus that would arise if the consumer perfectly optimized.

Corollary 1. *Given a profile of datasets $(Q^k, F^k)_k$, the consumer surplus that would arise at price p if the consumer chose optimally is no greater than $CS_*^s(p)$.*

6.2 Lower bound

An analogous construction to that for the upper bound can be used to obtain a non-trivial lower bound on surplus using data from multiple regimes. Given a stochastic value \mathbf{v} and demand $Q(p)$, we can compute a lower bound on surplus by supposing that the measure $Q(p)$ of the *lowest* types purchase the good at each p (as opposed to the highest types we used for the upper bound). Under this assignment, roughly speaking, lower means and greater spreads of the stochastic value both reduce the lower bound on surplus. Consequently, the relevant ordering of stochastic values is \succeq_{sosd} (as opposed to \succeq_{icx} for the upper bound).

Just as $\bar{\mathbf{v}}$ is the highest and the most spread out stochastic value consistent with the data in a single regime, $\underline{\mathbf{v}}$ (as defined in section 5) is the lowest and the most spread out such stochastic value. More precisely, $\underline{\mathbf{v}}$ is a lower bound with respect to \succeq_{sosd} on all \mathbf{v} consistent with the data. While the central step of the proof of Lemma 4 was to show that a decomposition into threshold demands induces a mean-*increasing* spread of the *non-buying* price expectations, a symmetric argument implies that the same decomposition induces a mean-*decreasing* spread of the *buying* price expectations.

To represent the second-order stochastic dominance order, we define the *complementary standard consumer surplus*

$$\widehat{CS}^s(p; \mathbf{v}) := \int_{-\infty}^p (1 - Q^s(p'; \mathbf{v})) dp'$$

and note that it is nondecreasing and convex in p . By the well known characterization of Hadar and Russell (1969) and Rothschild and Stiglitz (1970), the ranking of stochastic values \mathbf{v} with respect to \succeq_{sosd} implies the opposite ranking of $\widehat{CS}^s(p; \mathbf{v})$, and the converse also obtains provided the latter ranking holds uniformly across all p .

With multiple regimes, following the analogous construction to that for the upper bound, $\min_k \widehat{CS}^s(p; \underline{\mathbf{v}}^k)$ is an upper bound on $\widehat{CS}^s(p; \mathbf{v})$, where \mathbf{v} is a stochastic value consistent with the data in each regime and, for each k , $\underline{\mathbf{v}}^k$ is the lower bound on stochastic values with respect to \succeq_{sosd} consistent with the dataset (Q^k, F^k) . Since $\min_k \widehat{CS}^s(p; \underline{\mathbf{v}}^k)$ is not generally convex, it may not correspond to any stochastic value; accordingly, let $\widehat{CS}_*^s(p)$ denote its convexification. Let $\underline{\mathbf{v}}_*$ be the stochastic value associated with \widehat{CS}_*^s .¹⁹ Along the same lines as in Lemma 6, $\underline{\mathbf{v}}_*$ is a lower bound with respect to \succeq_{sosd} on stochastic values \mathbf{v} consistent with the profile of datasets.

Let $\widehat{P}(q; \mathbf{v}) := P^s(1 - q; \mathbf{v})$ denote the q th lowest quantile of \mathbf{v} .

Theorem 3. *If the interim consumer surplus $s^k(p)$ in regime k is consistent with the profile of datasets, then $s^k(p) \geq CS(Q^k, \widehat{P}(\cdot, \underline{\mathbf{v}}_*); p)$.*

Once again, taking expectations with respect to the price in each regime gives a

¹⁹That is, let $1 - Q^s(\cdot; \underline{\mathbf{v}}_*)$ be the right derivative of \widehat{CS}_*^s , observe that it is a distribution function, and let $\underline{\mathbf{v}}_*$ be a random variable with this distribution.

lower bound on the ex ante surplus in that regime.

To understand this result, consider a consumer with stochastic value \mathbf{v} . According to the data for regime k , a measure $Q^k(p)$ of types buy at each price p . Selecting the types with the lowest values generates surplus $\mathcal{CS}\left(Q^k, \hat{P}(\cdot; \mathbf{v}); p\right)$ in regime k ; this lower bound is nondecreasing in \mathbf{v} with respect to second-order stochastic dominance. Finally, because stochastic values consistent with the profile of datasets are bounded from below with respect to \succeq_{sosd} by $\underline{\mathbf{v}}_*$, the bound on $s^k(p)$ from the theorem applies.

Example 4. To illustrate the lower bound, consider the regimes from Example 3 with $\varepsilon = 1/3$. In this case, $\underline{\mathbf{v}}_1$ is uniformly distributed on $[0, 1/2]$ and $\underline{\mathbf{v}}_2$ is almost surely equal to $1/2$. Thus $\underline{\mathbf{v}}_2$ second-order stochastically dominates $\underline{\mathbf{v}}_1$, making the convexification trivial with $\underline{\mathbf{v}}_* = \underline{\mathbf{v}}_2$. The lower bound from Theorem 3 on the ex ante consumer surplus in regime 1 is therefore $1/2 \cdot 1/2 - 1/6 = 1/12$ and the lower bound in regime 2 is 0. \triangle

6.3 Tightness of the bounds

Theorem 1 provides tight bounds on consumer surplus for data from a single market regime; for each value within the bounds, we have constructed a model for which the surplus is equal to that value. While the bounds on surplus in Theorems 2 and 3 are generally tighter within each regime than the bounds obtained from the data in that regime alone, they are not themselves tight bounds.

Example 5. To illustrate, consider the upper bound for the two regimes from Example 3 with $\varepsilon \leq 1/6$. In this case, one can show that $CS^s(p; \bar{\mathbf{v}}_2) \leq CS^s(p; \bar{\mathbf{v}}_1)$ for all p , making the convexification trivial: $CS_*^s(p) \equiv CS^s(p; \bar{\mathbf{v}}_2)$. Therefore, $\bar{\mathbf{v}}_* \stackrel{d}{=} \bar{\mathbf{v}}_2$ almost surely takes on the value $2/3 + \varepsilon/2$, which is the non-buying price expectation for regime 2. However, this value is not consistent with the data for regime 1 since the non-buying price expectation of at least some types must be no more than $2/3$ (the non-buying price expectation for demand $Q^1(p)$). Thus, the upper bound on consumer surplus in regime 1 constructed in Theorem 2 is not attainable in this case. \triangle

If the analyst observes only one market regime, then, to determine the range of values of consumer surplus, it suffices to consider types with simple threshold

demands. With multiple regimes, decompositions into threshold demands are not generally sufficient; it can happen that the regimes are not jointly rationalizable by any model with threshold demands (but can be rationalized by other models).²⁰ The constructions in Theorems 2 and 3 circumvent this complication by using a bound on surplus in each regime based on threshold demands. The upside of this approach is that the combined bound is simple. The downside is that the combined bound need not correspond to a model that rationalizes the profile of datasets, and hence the bound is not generally tight.

7 Bounds on Counterfactual Demand

Returning to the original model in which the analyst observes a single dataset (Q, F) that may result from imperfect optimization, we now consider the counterfactual demand that would arise if instead the consumer were to perfectly optimize and purchase precisely when her value v_i exceeds the price p . These bounds apply equally to a counterfactual market with a fixed, deterministic price where our minimal rationality conditions imply that the consumer would choose optimally.

As for consumer surplus, bounds on counterfactual demand correspond to bounds on the consumer's stochastic value, albeit with respect to a different stochastic order: while the increasing convex order and second-order stochastic dominance provide the relevant bounds for consumer surplus, the bounds for counterfactual demand correspond to first-order stochastic dominance.

To state these bounds, define the *doubly elevated* and *doubly lowered* inverse demands, respectively, by

$$\begin{aligned} \overline{\overline{P}}(q) &:= \mathbb{E}[\mathbf{p} \mid \mathbf{p} \geq P(\mathbf{q}), \mathbf{q} \leq q] \\ \text{and} \quad \underline{\underline{P}}(q) &:= \mathbb{E}[\mathbf{p} \mid \mathbf{p} \leq P(\mathbf{q}), \mathbf{q} \geq q], \end{aligned}$$

where $\mathbf{q} \sim U[0, 1]$ and $\mathbf{p} \sim F$. Both functions are non-increasing. Relative to the

²⁰Example 3 provides one such example when $\varepsilon < 1/6$. If the demand from regime 1 is decomposed into threshold demands, then a nonzero mass of types must have thresholds below $1/3 - 2\varepsilon$. The non-buying price expectation of such types is less than $2/3 - \varepsilon$. Therefore, these types would not buy at any price that occurs in regime 2, contradicting that $Q^2(p) = 1$ for $p \leq 2/3$.




Figure 4: Doubly elevated and doubly lowered inverse demands, $\overline{\overline{P}}$ (thick grey) and $\underline{\underline{P}}$ (thick dashed), for given inverse demand P (black). For each q , $\overline{\overline{P}}(q)$ is the expected price conditional on \mathbf{p} and \mathbf{q} lying in the upper-left grey area. Similarly, $\underline{\underline{P}}(q)$ is the expected price conditional on the lower-right grey area. For comparison, the thin grey and thin dashed curves depict \overline{P} and \underline{P} , respectively.

elevated and lowered inverse demands \overline{P} and \underline{P} , these inverse demands are further elevated and lowered, i.e., $\overline{\overline{P}}(q) \geq \overline{P}(q)$ and $\underline{\underline{P}}(q) \leq \underline{P}(q)$ for all q . To see this, observe that $\overline{\overline{P}}(q)$ is a convex combination of $\overline{P}(q')$ across $q' \in [0, q]$ and \overline{P} is non-increasing; a symmetric argument shows that $\underline{\underline{P}}(q) \leq \underline{P}(q)$. See Figure 4 for an illustration.

Unlike \overline{P} and \underline{P} , the doubly elevated and lowered inverse demands $\overline{\overline{P}}$ and $\underline{\underline{P}}$ could not arise under the counterfactual of perfect optimization: the corresponding stochastic values are inconsistent with the observed data. However, *for each* q , there exist models rationalizing the data for which the counterfactual inverse demands are $\overline{\overline{P}}(q)$ and $\underline{\underline{P}}(q)$, respectively.

Theorem 4. *For every stochastic value \mathbf{v} consistent with data (Q, F) , the standard inverse demand function satisfies*

$$\underline{\underline{P}}(q) \leq P^s(q; \mathbf{v}) \leq \overline{\overline{P}}(q)$$

for all q . These bounds are tight in the sense that for each q , there exists a stochastic value \mathbf{v} consistent with the data such that $P^s(q; \mathbf{v}) = \overline{\overline{P}}(q)$, and similarly for $\underline{\underline{P}}(q)$.

We sketch the argument for the upper bound; the argument for the lower bound is analogous. For each q , given any stochastic value \mathbf{v} , the standard inverse demand $P^s(q; \mathbf{v})$ is a particular quantile of \mathbf{v} (namely, the $(1 - q)$ th quantile). The model that maximizes the counterfactual inverse demand at q among those rationalizing

the data is therefore the one that maximizes this quantile. Accordingly, bounds on counterfactual demand correspond to bounds on stochastic values with respect to first-order stochastic dominance.

How can we maximize a given quantile of \mathbf{v} (among stochastic values consistent with the data)? Recall that the highest value compatible with a type's demand is its non-buying price expectation. It turns out that this non-buying price expectation is maximized when no other type has a higher value and the demand of this type is as large as possible. Accordingly, to maximize the value at the $(1 - q)$ th quantile, we use a model in which the type with the highest value has measure q and demand $\min\{Q(p)/q, 1\}$. By construction, the non-buying price expectation of this type is exactly $\bar{P}(q)$. To see that the bound is tight, note that such a type can be part of a model that rationalizes the data (in which the remaining measure $1 - q$ of types generate the residual demand).

As with consumer surplus, data from multiple market regimes can be used to tighten the bounds on counterfactual demand. Assuming, as in section 6, that preferences are stable across regimes, a tighter bound can be obtained by simply taking the minimum and maximum, respectively, of the upper and the lower bounds from Theorem 4 across all of the regimes.

8 Discussion

If the analyst does not know whether the consumer engages in optimal choice behavior, the consumer surplus cannot be point identified from price and demand data. Nonetheless, weak rationality assumptions impose significant restrictions on the levels of surplus consistent with the data. Identification of the consumer surplus can be further sharpened by combining data from market regimes with varying priors or consumer demands.

Two relevant questions related to this project remain open. First, the bounds we provide under multiple regimes are not tight; our bounds rely on separate rationalizations for each regime, whereas, in principle, identification of the surplus can be tightened by simultaneously rationalizing the profile of datasets. Second, in the interest of generality, we have imposed minimal structure on the relationship between

the consumer's value and her demand. Depending on the context, there may be additional structure that could be used to narrow the bounds on surplus or counterfactual demand.

A Proofs

Proof of Lemma 2. Note that

$$\mathcal{CS}(\tilde{Q}, \mathbf{v}; p) = \mathcal{CS}(\tilde{Q}, P^s(\cdot; \mathbf{v}); p) = \int_0^{\tilde{Q}(p)} P^s(q; \mathbf{v}) dq - \tilde{Q}(p)p.$$

Consider any \mathbf{v} and \mathbf{v}' such that $\mathbf{v}' \succeq_{\text{icx}} \mathbf{v}$. Since the expenditure $\tilde{Q}(p)p$ does not depend on the stochastic value, it suffices to prove that $\int_0^{q^*} P^s(q; \mathbf{v}') dq \geq \int_0^{q^*} P^s(q; \mathbf{v}) dq$ for each $q^* \in [0, 1]$. Fix q^* . For $p = P^s(q^*; \mathbf{v}')$,

$$\begin{aligned} \int_0^{q^*} P^s(q; \mathbf{v}') dq &= CS^s(p; \mathbf{v}') + pq^* \\ &\geq CS^s(p; \mathbf{v}) + pq^* \\ &\geq \int_0^{q^*} (P^s(q; \mathbf{v}) - p) dq + pq^* \\ &= \int_0^{q^*} P^s(q; \mathbf{v}) dq; \end{aligned}$$

the first inequality follows from Lemma 1 while the second follows from the observation that $CS^s(p; \mathbf{v}) = \max_{q'} \int_0^{q'} (P^s(q; \mathbf{v}) - p) dq$. \square

Proof of Lemma 4. Step 1: Consider a model such that each type i has value v_i and demand $Q_i(p)$. Let $\mathbf{v} = v_{\mathbf{i}}$ be the associated stochastic value. Let $\mathbf{v}' = \bar{v}_{\mathbf{i}}$, where $\bar{v}_i = E[\mathbf{p} \mid \mathbf{q} \geq Q_i(\mathbf{p})]$ for $\mathbf{q} \sim U[0, 1]$ denotes the non-buying price expectation associated with demand Q_i . By Lemma 3, $\bar{v}_i \geq v_i$ for each i . Thus, $\mathbf{v}' \succeq_{\text{icx}} \mathbf{v}$ (because \mathbf{v}' first-order stochastically dominates \mathbf{v}).

Step 2: For each type i , define a random variable $\bar{\mathbf{v}}_i$ as follows. Let $P_i(q)$ be the inverse demand to demand Q_i , let $\bar{P}_i(q) = E[\mathbf{p} \mid \mathbf{p} \geq P_i(q)]$ be the elevated demand of type i , and define the stochastic value $\bar{\mathbf{v}}_i = \bar{P}_i(\mathbf{q})$ for $\mathbf{q} \sim U[0, 1]$. Finally, let $\mathbf{v}'' = \bar{\mathbf{v}}_{\mathbf{i}}$; thus \mathbf{v}'' is a spread of \mathbf{v}' that replaces $v'_i = \bar{v}_i$ with $\bar{\mathbf{v}}_i$ for each i .

We will show that $\mathbf{v}'' \succeq_{\text{icx}} \mathbf{v}'$ (and hence $\mathbf{v}'' \succeq_{\text{icx}} \mathbf{v}$). It suffices to show that $\bar{v}_i \leq E[\bar{\mathbf{v}}_i]$ for each i . Indeed, for $\mathbf{q} \sim U[0, 1]$, by the Law of Iterated Expectations,

$$\begin{aligned} \bar{v}_i &= E[\mathbf{p} \mid \mathbf{q} \geq Q_i(\mathbf{p})] \\ &= E[\mathbf{p} \mid \mathbf{p} \geq P_i(\mathbf{q})] \\ &= E[E[\mathbf{p} \mid \mathbf{p} \geq P_i(\mathbf{q}), \mathbf{q}] \mid \mathbf{p} \geq P_i(\mathbf{q})] \\ &= E[\bar{P}_i(\mathbf{q}) \mid \mathbf{p} \geq P_i(\mathbf{q})] \\ &= E[\bar{P}_i(\mathbf{q}) \mid \mathbf{q} \geq Q_i(\mathbf{p})]. \end{aligned}$$

Since \mathbf{q} conditional on $\mathbf{q} \geq Q_i(\mathbf{p})$ first-order stochastically dominates \mathbf{q} itself and $\bar{P}_i(q)$ is nonincreasing, it follows that

$$\bar{v}_i \leq E[\bar{P}_i(\mathbf{q})] = E[\bar{\mathbf{v}}_i],$$

as needed.

Step 3: We conclude by proving that $\mathbf{v}'' \stackrel{d}{=} \bar{\mathbf{v}}$. Consider any p at which Q is continuous and let $\tilde{v} = E[\mathbf{p} \mid \mathbf{p} \geq p]$. For any $j \in [0, 1]$,

$$j = \Pr(\bar{\mathbf{v}} \geq \tilde{v}) \implies \bar{P}(j) = \tilde{v} \implies P(j) = p \implies j = Q(p).$$

Hence $\Pr(\bar{\mathbf{v}} \geq \tilde{v}) = Q(p)$. Likewise, $\Pr(\bar{\mathbf{v}}_i \geq \tilde{v}) = Q_i(p)$ for almost all i (i.e., for all i except those for which Q_i is discontinuous at p), and thus

$$\Pr(\mathbf{v}'' \geq \tilde{v}) = \Pr(\bar{\mathbf{v}}_i \geq \tilde{v}) = E[Q_i(p)] = Q(p) = \Pr(\bar{\mathbf{v}} \geq \tilde{v})$$

for all \tilde{v} from a dense subset of the support of $\bar{\mathbf{v}}$ and \mathbf{v}'' , as needed. \square

Proof of Theorem 3. Consider a model consistent with the profile of datasets and let \mathbf{v} be its associated stochastic value. Recall that $\hat{P}(q; \mathbf{v}) = P^s(1 - q; \mathbf{v})$ is the q th lowest quantile of \mathbf{v} . Note that

$$s^k(p) \geq \mathcal{CS}\left(Q^k, \hat{P}(\cdot; \mathbf{v}); p\right)$$

for each k since the right-hand side is the expected consumer surplus if the measure

$Q^k(p)$ of types with the lowest values buy at price p .

For any price p and any two stochastic values \mathbf{v} and \mathbf{v}' such that $\mathbf{v} \succeq_{\text{sosd}} \mathbf{v}'$ and any demand function \tilde{Q} , we claim that

$$\mathcal{CS}(\tilde{Q}, \hat{P}(\cdot; \mathbf{v}); p) \geq \mathcal{CS}(\tilde{Q}, \hat{P}(\cdot; \mathbf{v}'); p).$$

The proof of this claim is analogous to that of Lemma 2. In particular, we may disregard expenditures since they depend only on the first and the last arguments of \mathcal{CS} . It suffices to prove that if $\mathbf{v} \succeq_{\text{sosd}} \mathbf{v}'$, then $\int_0^{q^*} \hat{P}(q; \mathbf{v}) dq \geq \int_0^{q^*} \hat{P}(q; \mathbf{v}') dq$ for every q^* . Fixing q^* and letting $p = \hat{P}(q^*; \mathbf{v}')$, we have

$$\begin{aligned} \int_0^{q^*} \hat{P}(q; \mathbf{v}') dq &= q^* p - \widehat{CS}^s(p; \mathbf{v}') \\ &\leq q^* p - \widehat{CS}^s(p; \mathbf{v}) \\ &\leq q^* p - \int_0^{q^*} (p - \hat{P}(q; \mathbf{v})) dq \\ &= \int_0^{q^*} \hat{P}(q; \mathbf{v}) dq; \end{aligned}$$

the first inequality follows from the integral condition for $\mathbf{v} \succeq_{\text{sosd}} \mathbf{v}'$ and the second from the fact that $\widehat{CS}^s(p; \mathbf{v}) = \max_{\hat{q}} \int_0^{\hat{q}} (p - \hat{P}(q; \mathbf{v})) dq$.

Therefore, if a stochastic value \mathbf{v} is consistent with the profile of datasets, then

$$s^k(p) \geq \mathcal{CS}(Q^k, \hat{P}(\cdot; \mathbf{v}); p) \geq \mathcal{CS}(Q^k, \hat{P}(\cdot; \mathbf{v}_*); p)$$

since $\mathbf{v} \succeq_{\text{sosd}} \mathbf{v}_*$. □

Proof of Theorem 4. We prove only the upper bound; the argument for the lower bound is analogous.

For any $q \in (0, 1]$ consider a demand function \tilde{Q} that attains values in $[0, q]$, i.e., a nonincreasing function from $[p, \bar{p}]$ onto $[0, q]$. Let

$$\begin{aligned} \tilde{v}(\tilde{Q}; q) &:= \mathbb{E}[\mathbf{p} \mid \mathbf{q} \geq \tilde{Q}(\mathbf{p})] \\ \text{and} \quad w(\tilde{Q}; q) &:= q \Pr(\mathbf{q} \geq \tilde{Q}(\mathbf{p})) \end{aligned}$$

for $\mathbf{q} \sim U[0, q]$ and $\mathbf{p} \sim F$. To interpret these two functions, consider a model with type distribution M and a subset $I' \subseteq I$ of types such that $\Pr(\mathbf{i} \in I') = q$ and $\tilde{Q}(p) = \int_{i \in I'} Q_i(p) dM(i)$. Then $\tilde{v}(\tilde{Q}; q)$ is the expected price conditional on a type randomly drawn from I' not making a purchase and $w(\tilde{Q}; q)$ is the probability that a randomly drawn type lies in I' and does not buy.

Note the following identity. For any $q_a, q_b \in (0, q]$ such that $q_a + q_b = q$ and any two demands Q_a and Q_b that attain values in $[0, q_a]$ and $[0, q_b]$, respectively, such that $Q_a + Q_b = \tilde{Q}$,

$$\tilde{v}(\tilde{Q}; q) = \frac{w(Q_a; q_a) \tilde{v}(Q_a; q_a) + w(Q_b; q_b) \tilde{v}(Q_b; q_b)}{w(Q_a; q_a) + w(Q_b; q_b)}. \quad (6)$$

Given any model and a subset I' of types such that $\Pr(\mathbf{i} \in I') = q$, let $v_* := \inf_{i \in I'} v_i$. To establish the upper bound, it suffices to show for each q that the supremum of v_* across all models that rationalize the data and subsets I' such that $\Pr(\mathbf{i} \in I') = q$ is at most $\overline{P}(q)$.

Fix a model with type distribution M on I and types (v_i, Q_i) that rationalizes the data. Fix a set I' of types such that $\Pr(\mathbf{i} \in I') = q$. Let $\tilde{Q}(p) = \int_{i \in I'} Q_i(p) dM(i)$ be the demand generated by the types in I' . Note that

$$\inf_{i \in I'} v_i \leq \tilde{v}(\tilde{Q}; q)$$

since, by Lemma 3, $v_i \leq \bar{v}_i$ for each type i , where \bar{v}_i is the non-buying price expectation associated with the demand Q_i of type i and $\tilde{v}(\tilde{Q}; q)$ is a convex combination of \bar{v}_i across $i \in I'$.

Let $Q^*(p) := \min\{Q(p), q\}$ and observe that $\overline{P}(q) = \tilde{v}(Q^*; q)$. It suffices to show that

$$\tilde{v}(\tilde{Q}; q) \leq \tilde{v}(Q^*; q) \quad (7)$$

for all q and all demands \tilde{Q} that can be generated by a subset I' of types from a model that rationalizes the data and satisfies $\Pr(\mathbf{i} \in I') = q$. For all such demands \tilde{Q} , both \tilde{Q} and $Q(p) - \tilde{Q}(p)$ are nonnegative and nonincreasing because they are the demands induced by types in I' and $I \setminus I'$, respectively.

Let $\tilde{Q}(p)$ be any demand function attaining values in $[0, q]$ such that $Q(p) -$




Figure 5: Illustration of the definitions of Q^* , \tilde{Q} , and Q_0 .

$\tilde{Q}(p)$ is nonnegative and nonincreasing. Let $p^* := P(q)$ and $q^* := \tilde{Q}(p^*)$. Since $\tilde{Q}(p) \leq Q^*(p) \leq q$ for all p , we have that $q^* \leq q$. Define the demand function $Q_0(p) := \min\{\tilde{Q}(p), q^*\}$ that attains values in $[0, q^*]$ and let $Q_1(p) := Q^*(p) - Q_0(p)$ and $Q_2(p) := \tilde{Q}(p) - Q_0(p)$. See Figure 5 for an illustration.

Note that $Q_1(p)$ is nonincreasing: it is equal to $q - q^*$ for $p \leq p^*$ and to $Q(p) - \tilde{Q}(p)$ for $p \geq p^*$. The function $Q_2(p)$ is also nonincreasing since it is equal to $\tilde{Q}(p) - q^* \geq 0$ for $p \leq p^*$ and to 0 for $p > p^*$. Let P_0 , P_1 , and P_2 be the inverse demand functions associated with Q_0 , Q_1 , and Q_2 , respectively. Note that, on their respective domains, P_0 and P_1 only attain values above p^* , while P_2 only attains values below p^* .

Recall that $\tilde{v}(\tilde{Q}; q)$ can be written as $E[\mathbf{p} \mid \mathbf{p} \geq \tilde{P}(\mathbf{q})]$ for $\mathbf{q} \sim U[0, q]$, where \tilde{P} is the inverse demand to \tilde{Q} ; similarly, $w(\tilde{Q}; q)$ can be written as $q \Pr(\mathbf{p} \geq \tilde{P}(\mathbf{q}))$. It follows that $\tilde{v}(Q_1; q - q^*) \geq \tilde{v}(Q_2; q - q^*)$, $\tilde{v}(Q_0; q^*) \geq \tilde{v}(Q_2; q - q^*)$, and $w(Q_2; q - q^*) \geq w(Q_1; q - q^*)$. Finally, since $Q^* = Q_0 + Q_1$ and $\tilde{Q} = Q_0 + Q_2$, we have from (6) that

$$\tilde{v}(Q^*; q) = \frac{w(Q_0; q^*) \tilde{v}(Q_0; q^*) + w(Q_1; q - q^*) \tilde{v}(Q_1; q - q^*)}{w(Q_0; q^*) + w(Q_1; q - q^*)}$$

and

$$\tilde{v}(\tilde{Q}; q) = \frac{w(Q_0; q^*) \tilde{v}(Q_0; q^*) + w(Q_2; q - q^*) \tilde{v}(Q_2; q - q^*)}{w(Q_0; q^*) + w(Q_2; q - q^*)}.$$

Therefore,

$$\begin{aligned}
\tilde{v}(Q^*; q) &= \tilde{v}(Q_0; q^*) + \frac{w(Q_1; q - q^*)}{w(Q_0; q^*) + w(Q_1; q - q^*)} (\tilde{v}(Q_1; q - q^*) - \tilde{v}(Q_0; q^*)) \\
&\geq \tilde{v}(Q_0; q^*) + \frac{w(Q_1; q - q^*)}{w(Q_0; q^*) + w(Q_1; q - q^*)} (\tilde{v}(Q_2; q - q^*) - \tilde{v}(Q_0; q^*)) \\
&\geq \tilde{v}(Q_0; q^*) + \frac{w(Q_2; q - q^*)}{w(Q_0; q^*) + w(Q_2; q - q^*)} (\tilde{v}(Q_2; q - q^*) - \tilde{v}(Q_0; q^*)) \\
&= \tilde{v}(\tilde{Q}; q),
\end{aligned}$$

which establishes inequality (7), as needed. \square

References

- Allen, R. and J. Rehbeck (2021). Counterfactual and welfare analysis with an approximate model. Working paper.
- Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review* 105(3), 921–57.
- Bergemann, D., B. Brooks, and S. Morris (2017). First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica* 85(1), 107–143.
- Bergemann, D., B. Brooks, and S. Morris (2022). Counterfactuals with latent information. *American Economic Review* 112(1), 343–68.
- Bernheim, B. D. and A. Rangel (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *Quarterly Journal of Economics* 124(1), 51–104.
- Bernheim, B. D. and D. Taubinsky (2018). Behavioral public economics. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics: Applications and Foundations*, Volume 1, pp. 381–516. Elsevier.

- Bronnenberg, B. J., J.-P. Dubé, M. Gentzkow, and J. M. Shapiro (2015). Do pharmacists buy Bayer? Informed shoppers and the brand premium. *Quarterly Journal of Economics* 130(4), 1669–1726.
- Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review* 99(4), 1145–77.
- Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe (2016). Using big data to estimate consumer surplus: The case of Uber. NBER Working Paper 22627.
- Condorelli, D. and B. Szentes (2020). Information design in the holdup problem. *Journal of Political Economy* 128(2), 681–709.
- Condorelli, D. and B. Szentes (2022). Surplus sharing in Cournot oligopoly. *Theoretical Economics* 17(3), 955–975.
- Dickson, P. R. and A. G. Sawyer (1990). The price knowledge and search of supermarket shoppers. *Journal of Marketing* 54(3), 42–53.
- Feldman, N. E., P. Katuščák, and L. Kawano (2016). Taxpayer confusion: Evidence from the child tax credit. *American Economic Review* 106(3), 807–35.
- Finkelstein, A. (2009). E-ztax: Tax salience and tax rates. *Quarterly Journal of Economics* 124(3), 969–1010.
- Gentzkow, M. and E. Kamenica (2016). A Rothschild-Stiglitz approach to Bayesian persuasion. *American Economic Review* 106(5), 597–601.
- Goolsbee, A. and A. Petrin (2004). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica* 72(2), 351–381.
- Gruber, J. and B. Köszegi (2001). Is addiction “rational”? Theory and evidence. *Quarterly Journal of Economics* 116(4), 1261–1303.
- Hadar, J. and W. R. Russell (1969). Rules for ordering uncertain prospects. *American Economic Review* 59(1), 25–34.
- Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review* 104(2), 537–63.

- Kang, Z. Y. and S. Vasserman (2022). Robust bounds for welfare analysis. NBER Working Paper 29656.
- Kleiner, A., B. Moldovanu, and P. Strack (2021). Extreme points and majorization: Economic applications. *Econometrica* 89(4), 1557–1593.
- Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017). Persuasion of a privately informed receiver. *Econometrica* 85(6), 1949–1964.
- Manzini, P. and M. Mariotti (2014). Stochastic choice and consideration sets. *Econometrica* 82(3), 1153–1176.
- Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2012). Revealed attention. *American Economic Review* 102(5), 2183–2205.
- Müller, A. and M. Scarsini (2006). Stochastic order relations and lattices of probability measures. *SIAM Journal on Optimization* 16(4), 1024–1043.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110(4), 705–729.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. a definition. *Journal of Economic Theory* 2(3), 225–243.
- Sandomirskiy, F. and P. Ushchev (2022). The geometry of consumer preference aggregation. Working paper.
- Shaked, M. and J. G. Shanthikumar (2007). *Stochastic orders*. Springer.
- Taubinsky, D. and A. Rees-Jones (2018). Attention variation and welfare: theory and evidence from a tax salience experiment. *Review of Economic Studies* 85(4), 2462–2496.
- Tipoe, E. (2021). Price inattention: A revealed preference characterisation. *European Economic Review* 134, 103692.
- Varian, H. R. (1985). Price discrimination and social welfare. *American Economic Review* 75(4), 870–875.

Yang, K. H. and A. K. Zentefis (2023). Extreme points of first-order stochastic dominance intervals: Theory and applications. arXiv preprint arXiv:2302.03135.