Demand in the Dark*

Pavel Kocourek
CERGE-EI

Jakub Steiner
University of Zurich and CERGE-EI

Colin Stewart
University of Toronto

March 30, 2022

Abstract

A growing body of evidence suggests that consumers are not fully informed about prices, contrary to a critical assumption of classical consumer theory. We analyze a model in which consumer types can vary in both their preferences and their information about prices. Given data on demand and the distribution of prices, we identify the set of possible values of the consumer surplus. Each surplus in this set can be rationalized with simple information structures and preferences. We also show how to narrow down the set of values using richer datasets and provide bounds on counterfactual demands at perfectly observed prices.

1 Introduction

A key implicit assumption of the standard approach to consumer demand and welfare is that consumers perfectly observe and understand prices. 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, changing the way that prices are presented to consumers can have significant effects on demand (Chetty, Looney, and Kroft, 2009; Finkelstein, 2009).

We analyze consumer surplus and demand in the market for a single good without assuming complete information about prices. Unlike uncertainty about the value of the good, uncertainty about prices invalidates the classical approach to computing consumer surplus.[1] Under the standard assumptions of complete information and quasilinear preferences, a choice at a given price reveals that the marginal utility of the quantity chosen is equal to that price. Calculating total consumer surplus is then simply a matter of adding up these marginal utilities (net of expenditures). When the consumer is uncertain about prices, however, a choice instead reveals only that the marginal utility of the quantity chosen is equal to the *expected* price given the consumer's information. If this

*We have benefited from the comments of Tibor Heumann, Zi Yang Kang, Anton Kolotilin, Marek Pycia, Andy Zapechelnyuk, and various seminar audiences. Stewart is grateful to PSE for their hospitality. This work was supported by ERC grant 770652 and by the Social Sciences and Humanities Research Council of Canada.

[1]The key distinction is that the analyst observes the price but not the consumer's value. To accommodate incomplete information about the value on the part of the consumer, one can simply reinterpret the value as the expected value conditional on the consumer's information.

expectation is unobserved by the analyst, it is no longer possible to identify the marginal utilities directly.

Uncertainty about prices can result from several factors. The consumer may not be fully attentive to the price of a good they buy, either because the cost is small relative to their budget, because it is a habitual purchase, or because prices involve complexities that require nontrivial effort to understand. It may also happen that the consumer does not check the price of a good—or does so only if they notice that the good is on sale—under the belief that the price is likely to exceed their willingness to pay. Alternatively, uncertainty could reflect incomplete information about the value of money, for example because the consumer does not know the prices of other goods they might purchase.

In our main setting, an analyst observes data consisting of a distribution of prices for a particular good together with a downward-sloping demand function indicating the frequency with which the consumer buys the good at each price. The analyst rationalizes the data with a *model* describing a distribution of types of the consumer, with each type specifying the consumer's value for the good together with the structure of information this type receives about prices. A key feature of our approach is that we impose few restrictions on the consumer's information beyond the assumption that the consumer is Bayesian. In particular, the information structure can be type-dependent, reflecting the idea that consumers' attentiveness to prices may vary from day to day, and this variation may be correlated with their preferences.

We consider two questions. First, what can the analyst infer from the observed data about the surplus the consumer receives from participating in the market for this good? Second, what demand can the analyst forecast if the price were to become deterministic or the consumer perfectly informed?

One way the analyst can rationalize the data is to suppose the consumer perfectly observes prices and to attribute all stochasticity in behavior to changing preferences, as in a random utility model. This rationalization leads to the standard level of consumer surplus (averaged across prices). There are, however, other rationalizations leading to different levels of surplus; consequently, it is not possible to pin down the surplus exactly. For example, the analyst could go to the opposite extreme by supposing that the consumer has a fixed value for the good and attributing all stochasticity in behavior to the intrinsic randomness of information. The simplest such models involve binary signals indicating whether or not to buy the good (with the probabilities at each price chosen to match the observed demand). For such an information structure, there is generally a range of values that rationalize the data. At the low end, one obtains a surplus of 0 if the consumer's value is equal to the expected price conditional on receiving the signal to buy; in that case, the consumer receives a net expected benefit of 0 after each signal realization. At the high end, the consumer's value is equal to the expected price conditional on receiving the signal *not* to buy; if it were any higher, the consumer would prefer to switch to buying after both signals. The surplus associated with such a rationalization can be greater than that associated with complete information. More complex models that combine random utility with incomplete information can also rationalize the data, further enlarging the set of possible values of the surplus.

We fully characterize the set of levels of consumer surplus that are consistent with the data. This set consists of 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.

One model that achieves the upper bound involves particularly simple information structures: each type observes only whether the price is above or below a particular (type-dependent) threshold.[2] Each type has a corresponding simple threshold demand. Choosing the distribution of these thresholds appropriately ensures that the average demand agrees with the observed demand. Given the individual types' demands, the surplus is maximized when each type's value for the good is equal to the expected price conditional on not buying (as described above). We refer to this model as the *upper threshold model*.

The main object of our analysis is the random variable describing the consumer's value for the good, which we refer to as the *stochastic value*. 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), which can be viewed as the analogue of second-order stochastic dominance (SOSD) for a decision-maker who is risk-loving instead of risk averse. Thus where SOSD favors higher means and smaller spreads, ICX favors higher means and larger spreads.

The connection between the ICX order and consumer surplus is most direct under complete information: not only do higher values increase consumer surplus, so do mean-preserving spreads because they provide the consumer with an option value. Notice, however, that it is not necessary for this argument that the consumer fully optimize; the same connection holds as long as, at each price, the set of types that buy are those with the highest values (even if the cutoff type is not chosen optimally). Since the upper threshold model has exactly this feature and the association of highest values to purchases maximizes the consumer surplus for any given stochastic value, maximization in the ICX order gives the stochastic value corresponding to the upper bound on consumer surplus even under incomplete information.

We prove that the upper threshold model induces a stochastic value that dominates all other stochastic values consistent with the data with respect to the ICX order. Starting from an arbitrary stochastic value and corresponding model that rationalizes the data, we amend the model in two steps, each of which increases the stochastic value with respect to the ICX order (while continuing to rationalize the data). First, fixing the disaggregation of the observed demand into demands

[2] While it is important for the upper bound that the consumer receives no additional information when the price is above the threshold, the information at prices below the threshold can be arbitrarily precise. Thus, for example, it could be that a sale is announced if and only if the price drops below the threshold and the consumer checks the price only upon seeing a sale.

of individual types, we can maximize the value of each type by setting it equal to the expected price conditional on that type not buying—we call this value the "non-buying price expectation." This is the highest value for which there exists an information structure making the type's demand incentive compatible. Second, if the demand of any type is not a simple threshold demand then it can be further disaggregated into such demands. We show that this disaggregation results in a mean-increasing spread of the non-buying price expectations. Combining these two steps generates an increase in the stochastic value with respect to the ICX order and leads to the stochastic value associated with the upper threshold model.

Richer data can be used to tighten the bounds on consumer surplus. We extend our analysis to datasets comprising two or more regimes that may differ in the distribution of prices and/or the information obtained by the consumer. 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 considers all rationalizations of the datasets in which the value of each type is fixed across regimes (though each type's information may vary in a way that is unobservable to the analyst). We propose a simple procedure for tightening the bounds on surplus within each regime using the data from the other regimes. The key idea is to associate each stochastic value with the convex function describing the consumer surplus at each price under complete information. Increases in the stochastic value with respect to the ICX order correspond to increases in the associated consumer surplus function. Since the stochastic value is consistent across regimes, its associated consumer surplus function cannot, at any price, exceed that for the upper bound in any of the regimes. Therefore, the consumer surplus function can be no larger than the convex closure of the minimum of these functions for the upper bounds across the regimes. Mapping this convex closure back to a stochastic value yields a new upper bound with respect to the ICX order, and from there an upper bound on consumer surplus. An analogous construction with SOSD in place of the ICX order allows us to establish a nontrivial lower bound on the consumer surplus in each regime.

To predict counterfactual demand at a perfectly observed (or deterministic) price, the relevant comparison of stochastic values is with respect to first-order stochastic dominance (FOSD). This order captures the observation that, with complete information, the demand at a given price is equal to the probability that the consumer's value exceeds this price; thus the counterfactual demand increases with a FOSD increase in the stochastic value. We construct FOSD bounds on the stochastic values consistent with the data, which give rise to tight bounds on the counterfactual demand.

For simplicity of exposition, we focus on a single consumer with unit demand. In Appendix B, we show that our results extend to general quasilinear utilities: for any rationalization of the data with these utilities, there exists a rationalization with unit demands that generates the same surplus. Our model can also be interpreted as describing a population of consumers for which the analyst observes only the aggregate demand. Under this interpretation, each type corresponds to an individual consumer.

Our work is related to several streams of empirical and theoretical research. A number of studies in various contexts have found evidence that consumers do not perfectly observe or understand prices. For example, Chetty, Looney, and Kroft (2009) show that demand depends on whether sales taxes are included in the posted price for groceries and alcohol. Finkelstein (2009) finds that the switch to elecfronic collection of road tolls reduces drivers' responsiveness to changes in the cost associated with the toll.[3] Taubinsky and Rees-Jones (2018) and Tipoe (2021) find that there is significant heterogeneity in attention to prices.

Chetty (2012) considers identification of parameters in a large class of behavioral models including those featuring price misperceptions. Whereas his main identifying assumption uses bounds on the cost of deviations from optimal behavior, ours is based on Bayes rationality of the consumer.

The early literature on consumer surplus and price stability originating with Waugh (1944) (see also, e.g., Samuelson (1972) and Rogerson (1980)) shares with our model an important role for stochastic prices. The main question in that literature is whether the consumer benefits from stochasticity in prices under the implicit assumption that prices are perfectly observed. In our model, stochasticity allows for the consumer to be uncertain about prices. Perhaps surprisingly, the greatest surplus consistent with a given demand curve is generally higher when prices are uncertain than when they are perfectly observed: by allowing for uncertainty, the set of preferences consistent with the observed behavior is larger. (For any *given* preferences, of course, the consumer cannot benefit from incomplete information.)

Gul, Pesendorfer, and Strzalecki (2017) study general equilibrium when agents have coarse perceptions of the state of the world, which they interpret as capturing inattention to prices. Their main focus is on the resulting equilibrium prices.

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. We share with the information design literature the approach of imposing minimal assumptions on the information structure.[4] Bergemann, Brooks, and Morris (2015, 2017) identify the range of surplus that can be attained for given preferences as information varies in a monopolistic market or a first-price auction. Condorelli and Szentes (2020, 2021) characterize the range of surplus 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

[3]See also Ito (2014) for empirical evidence that consumers do not correctly account for marginal electricity pricing and Feldman, Katuščák, and Kawano (2016) for related evidence regarding marginal tax rates.

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

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 if the price were perfectly observed.

Perhaps the most closely related papers to ours in spirit are Varian (1985) and Kang and Vasserman (2022), which study a complementary problem of identifying bounds on consumer surplus. In their models, consumers perfectly observe prices but there are gaps in the demand observed by the analyst. Kang and Vasserman (2022) discuss how to interpret their problem in terms of concavification, along the lines of Kamenica and Gentzkow (2011); our problem can also be viewed as one of concavification—a connection we discuss in subsection 5.1—but in a different space and with a different constraint. In both cases, the space over which concavification occurs is very large, making standard techniques inapplicable.

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 $Q : [\underline{p}, \overline{p}] \longrightarrow [0, 1]$, 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 $Q(p)$, that is, $P(q) \coloneqq \inf\{p : Q(p) \leq q\}$.[5] Prices are distributed according to the continuous distribution $F(p)$ with support $[\underline{p}, \overline{p}]$, where $\underline{p} > 0$. As is standard when measuring consumer welfare, we assume that the analyst observes the choke price, i.e., $Q(\overline{p}) = 0$; similarly, we assume that $Q(\underline{p}) = 1$. We consider richer data involving demands under multiple market regimes in section 6 and non-unit demands in Appendix B.

The demand Q is the aggregation of many choices made by the consumer across which both her preferences and her information about prices 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 $I \subset \mathbb{R}$; the type is independent of the price. Each type i specifies the consumer's value v_i for the good and her information structure. The information structure consists of a distribution $\Phi_i(x \mid \mathbf{p} = p)$ of signals $x \in \mathbb{R}$ for each p. The information structure may be correlated with the value, as one might expect if the consumer has some discretion over her information acquisition. Let $\pi_i(x) \coloneqq \mathrm{E}_i[\mathbf{p} \mid \mathbf{x} = x]$ denote type i's posterior expected price upon observing signal realization x generated by Φ_i.

Each type has quasilinear preferences and unit demand. Given a realized signal x, a consumer of type i purchases the good with probability $q \in [0, 1]$ to maximize $q\,(v_i - \pi_i(x))$; we denote the chosen quantity by $q_i^*(x)$.[6] Each type i generates a demand function $Q_i(p) \coloneqq \mathrm{E}_i\left[q_i^*(\mathbf{x}) \mid \mathbf{p} = p\right]$,

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

[6] If $v_i = \pi_i(x)$, making all $q \in [0, 1]$ optimal, we include the chosen probability $q_i^*(x)$ in the description of the type

where the signal $\mathbf{x}$ is generated by Φ_i. We restrict $Q_i(p)$ to be non-increasing for each type i; this restriction is guaranteed to hold if, for example, each information structure satisfies the monotone likelihood ratio property.

The analyst looks to explain the data using a *model* that consists of a distribution M of types $i \in I$ together with a specification $(v_i, \Phi_i)_{i \in I}$ of values and information structures for each type. We say that a given model *rationalizes* data (Q, F) if $Q(p) = \mathrm{E}\left[Q_{\mathbf{i}}(p)\right]$ for all p. Given a model that rationalizes the data, the (ex ante) *consumer surplus*, s, is the consumer's expected utility; that is,

$$s \coloneqq \mathrm{E}\left[q^*_{\mathbf{i}}(\mathbf{x})\,(v_{\mathbf{i}} - \mathbf{p})\right] = \mathrm{E}\left[Q_{\mathbf{i}}(\mathbf{p})\,(v_{\mathbf{i}} - \mathbf{p})\right]. \tag{1}$$

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

Example 1. The analyst observes the linear demand function $Q(p) = 1 - p$ and prices uniformly distributed on $[0, 1]$. There are many possible rationalizations of this data. For instance, it could be that, as in the standard analysis, the consumer perfectly observes the price and her value, $v_{\mathbf{i}}$, is uniformly distributed on $[0, 1]$. 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 rationalization is therefore $\mathrm{E}\left[(1-\mathbf{p})^2/2\right] = 1/6$.

Alternatively, the data can be rationalized by a model in which the consumer has incomplete information about prices. Perhaps the simplest such rationalization features a consumer with a single type. For each price realization $p \in [0, 1]$, the consumer observes signal 1 with probability $Q(p)$ and signal 0 with the remaining probability $1 - Q(p)$. The consumer purchases if and only if she receives signal 1; that is, $q^*(x) = x$. This model trivially rationalizes the data as long as the strategy q^* is incentive compatible. To ensure that buying is optimal at $x = 1$, it must be that the consumer's value v is at least $\mathrm{E}[\mathbf{p} \mid \mathbf{x} = 1] = 1/3$. Similarly, to ensure that not buying is optimal at $x = 0$, it must be that $v \leq \mathrm{E}[\mathbf{p} \mid \mathbf{x} = 0] = 2/3$. Letting $\underline{v} = 1/3$ and $\overline{v} = 2/3$, incentive compatibility therefore holds if and only if $v \in [\underline{v}, \overline{v}]$. Taking $v = \overline{v}$ leads to a surplus of $(\overline{v} - \underline{v})\,\Pr(\mathbf{x} = 1) = 1/6$ (since $\underline{v}$ is the expected price conditional on buying). Taking $v = \underline{v}$ leads to a surplus of $(\underline{v} - \underline{v})\,\Pr(\mathbf{x} = 1) = 0$. Using values of v in between these two extremes, any surplus in $[0, 1/6]$ can be obtained.[7]

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

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

but omit this selection from the notation.

[7]That the upper bound of 1/6 is equal to the complete information 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 the consumer is indifferent whenever she buys the good.

Since $(Q_1+Q_2)/2=Q$, these demands are consistent with the observed data. Each type's demand Q_i can be rationalized in the same manner as in the single-type case by assuming that type i observes a binary signal $x\in\{0,1\}$ with $\Pr(\mathbf{x}=1\mid\mathbf{p}=p)=Q_i(p)$ and purchases if and only if $x=1$. Incentive compatibility requires that the value v_i of type i lie between $\underline{v}_i=\mathrm{E}_i[\mathbf{p}\mid\mathbf{x}=1]$ and $\overline{v}_i=\mathrm{E}_i[\mathbf{p}\mid\mathbf{x}=0]$. Using the maximal values gives surplus $\frac{1}{2}\sum_i\overline{v}_i\Pr_i(\mathbf{x}=1)-\mathrm{E}[\mathbf{p}Q(\mathbf{p})]=2/9$ (whereas the minimal values again give a surplus of 0).[8] By splitting the demand Q into demands Q_1 and Q_2, it is therefore possible to obtain higher values of the surplus than can be obtained with complete information or with only one type.

Since we have restricted attention to simple binary signals, one might wonder whether it is possible to enlarge the set of attainable surpluses by allowing for more complex information structures. The answer is no: any surplus that is consistent with the data can be rationalized by a model in which every type has a binary signal. The question remains, however, as to whether further splitting of the two demands Q_1 and Q_2—or some other splitting of Q—can expand the range of attainable surpluses. As we show in the next section, it turns out that a "maximal" splitting of the demand can rationalize higher values of the surplus (up to 1/4 in this case). $\triangle$

3 Main Result

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

$$\mathcal{CS}\left(\tilde{Q},\hat{P}\right)\coloneqq\mathrm{E}\left[\int_0^{\tilde{Q}(\mathbf{p})}\left(\hat{P}(q)-\mathbf{p}\right)dq\right],\tag{2}$$

where $\mathbf{p}\sim F$. When applied to the observed demand function Q and its inverse demand P, this functional returns the expected consumer surplus under complete information about the price. In that case, the inverse demand describes the consumer's marginal benefit of consumption at each q. Under incomplete information about the price, the marginal benefit is not generally equal to the inverse demand: the value of the marginal type is equal to her expectation of the price, not to the price itself. Nonetheless, if $\hat{P}(q)$ describes the marginal benefit, then $\mathcal{CS}(Q,\hat{P})$ is the consumer surplus.

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

$$\overline{P}(q)\coloneqq\mathrm{E}\left[\mathbf{p}\mid\mathbf{p}\geq P(q)\right]$$
$$\text{and}\qquad\underline{P}(q)\coloneqq\mathrm{E}\left[\mathbf{p}\mid\mathbf{p}\leq P(q)\right],$$

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

[8] Type 1 buys with probability 3/4 and has maximal value $\overline{v}_1=5/6$ and type 2 buys with probability 1/4 and has maximal value $\overline{v}_2=11/18$, giving a surplus of $1/2\cdot5/6\cdot3/4+1/2\cdot11/18\cdot1/4-1/6=2/9$.

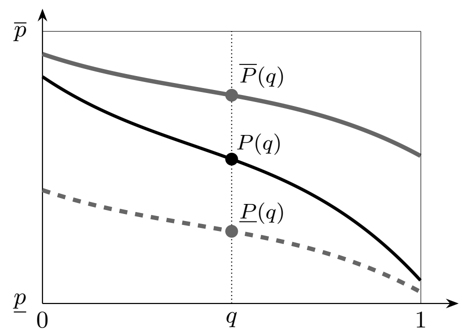

Figure 1: Elevated and lowered inverse demands for a particular inverse demand $P(q)$ with uniformly distributed prices.

Figure 1 for an illustration. Note that, letting $\mathbf{q} \sim U[0,1]$,

$$\begin{aligned}\mathcal{CS}\left(Q, \underline{P}\right) &= \mathrm{E}\left[\mathbb{1}_{\mathbf{q} \leq Q(\mathbf{p})}\left(\underline{P}(\mathbf{q}) - \mathbf{p}\right)\right] \\ &= \mathrm{E}\left[\mathbb{1}_{\mathbf{p} \leq P(\mathbf{q})}\left(\mathrm{E}[\mathbf{p} \mid \mathbf{p} \leq P(\mathbf{q})] - \mathbf{p}\right)\right] \\ &= \mathrm{E}\left[\mathbb{1}_{\mathbf{p} \leq P(\mathbf{q})}\left(\mathbf{p} - \mathbf{p}\right)\right] = 0.\end{aligned}$$

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

$$0 = \mathcal{CS}\left(Q, \underline{P}\right) \leq s \leq \mathcal{CS}\left(Q, \overline{P}\right).$$

The proof of this result is provided in subsection 5.1.

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 $\overline{P}(q) = (1 + P(q))/2 = 1 - q/2$, and $\mathcal{CS}\left(Q, \overline{P}\right) = 1/4$. Theorem 1 therefore indicates that any surplus in the interval $[0, 1/4]$ is consistent with the data. In this case, the consumer surplus can be up to $3/2$ times as large as the complete information surplus. △

In section 6, we extend the model to allow for the possibility that the analyst observes demand in two or more market regimes that may differ in the distribution of prices or in the consumer's information (or both). Assuming that preferences do not vary across regimes, combining the data allows for the construction of bounds on surplus within each regime that are generally narrower than the bounds in Theorem 1. In particular, data from multiple market regimes can lead to a non-trivial positive lower bound on the surplus.

Instead of welfare, the analyst may be interested in predicting the consumer's demand in some counterfactual market regime. In section 7, we derive tight bounds on the demand that would arise at a counterfactual deterministic price—or, equivalently, a perfectly observed realization of a stochastic price—for any given data.

4 Preliminaries

The distribution of the consumer's value of the good plays a central role in our analysis. As we explain in this section, when the value is viewed as a random variable, stochastic orders on the value are relevant for comparisons of demand and of consumer surplus.

Given a model, let $\mathbf{v} := v_{\mathbf{i}}$ be the consumer's *stochastic value* of the good. Thus $\mathbf{v}$ is a partial description of the model that disregards information structures. Let $Q^{\text{CI}}(p; \mathbf{v}) := \Pr(\mathbf{v} > p)$; the superscript CI indicates complete information. For any price p, $Q^{\text{CI}}(p; \mathbf{v})$ is the probability with which the consumer would buy the good if she perfectly observed the price p (except possibly at atoms of $\mathbf{v}$).[9] Therefore, we refer to $Q^{\text{CI}}(p; \mathbf{v})$ as the *complete information demand function* for $\mathbf{v}$. Note that the complete information demand function is the complementary distribution function of $\mathbf{v}$. Likewise, the complete information inverse demand function $P^{\text{CI}}(q; \mathbf{v})$—which is the inverse to the demand $Q^{\text{CI}}(p; \mathbf{v})$—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 complete information demands. Indeed, the following statements are equivalent: (i) $\mathbf{v}'$ first-order stochastically dominates $\mathbf{v}$; (ii) $Q^{\text{CI}}(p; \mathbf{v}') \geq Q^{\text{CI}}(p; \mathbf{v})$ for all p; and (iii) $P^{\text{CI}}(q; \mathbf{v}') \geq P^{\text{CI}}(q; \mathbf{v})$ for all q.

Gentzkow and Kamenica (2016) and Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) identify a mapping from random variables to convex functions that is useful for making comparisons with respect to the convex order. Define the function $CS^{\text{CI}}(\cdot; \mathbf{v}) : \mathbb{R} \longrightarrow \mathbb{R}$ by

$$CS^{\text{CI}}(p; \mathbf{v}) := \int_p^\infty Q^{\text{CI}}\left(p'; \mathbf{v}\right) dp'. \tag{3}$$

This mapping has a natural interpretation in our context: it is the area between the price line and the complete information demand function, i.e., the surplus of the consumer when she perfectly observes p.[10] Observe that $CS^{\text{CI}}(p; \mathbf{v})$ is convex in p because Q^{CI} is downward-sloping.

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

$$\mathbf{v}' \succeq_{\text{icx}} \mathbf{v} \text{ if and only if } CS^{\text{CI}}(p; \mathbf{v}') \geq CS^{\text{CI}}(p; \mathbf{v})$$

for every price p.

[9] If $\mathbf{v}$ has an atom at p, then the demand if p is perfectly observed lies in the closed interval between the left and right limits of $Q^{\text{CI}}(\cdot; \mathbf{v})$ at p.

[10] Gentzkow 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.

Proof. The result follows from Theorem 4.A.2 of (Shaked and Shanthikumar, 2007) given (3) together with the fact that $Q^{\text{CI}}(p; \mathbf{v})$ is the complementary distribution function of $\mathbf{v}$. □

The increasing convex order is closely related to second-order stochastic dominance, denoted here by $\succeq_{\text{sosd}}$.[11] (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 distribution functions.

The next result suggests how the increasing convex order can be useful even when there is incomplete information about prices. Consider a consumer with stochastic value $\mathbf{v}$ and demand $\tilde{Q}(p)$. The highest possible surplus for this consumer is obtained when, for each p, the highest types of measure $\tilde{Q}(p)$ are the ones that buy the good. In that case, the consumer surplus is $\mathcal{CS}\left(\tilde{Q}, P^{\text{CI}}(\cdot; \mathbf{v})\right)$.

Lemma 2. *For every demand function $\tilde{Q}$, $\mathcal{CS}\left(\tilde{Q}, P^{\text{CI}}(\cdot; \mathbf{v})\right)$ 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 $\mathcal{CS}\left(\tilde{Q}, P^{\text{CI}}(\cdot; \mathbf{v})\right)$. 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 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 Bounds on Consumer Surplus

5.1 Proof of Theorem 1

We first show by construction that each surplus between 0 and $\mathcal{CS}\left(Q, \overline{P}\right)$ is consistent with the data. To prove that no other levels of surplus are consistent with the data, we show that the stochastic value associated with the construction yielding surplus $\mathcal{CS}\left(Q, \overline{P}\right)$ provides an upper bound with respect to the increasing convex order. (For this step, it suffices to consider only the upper bound: since the lower bound on surplus is 0, it holds trivially that no lower surplus can be obtained by a rational consumer.)

Consider a decomposition of the observed demand Q into demands $(Q_i)_i$ of individual types i. We begin by identifying the range of stochastic values consistent with this decomposition and then consider varying the decomposition itself.

[11] Recall 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 $\mathbf{z}$.

We say that a value v_i is *consistent with demand* Q_i if v_i together with some information structure Φ_i generates demand Q_i. Given Q_i, let

$$\underline{v}_i \coloneqq \mathrm{E}\left[\mathbf{p} \mid Q_i(\mathbf{p}) \geq \mathbf{q}\right] = \mathrm{E}\left[\mathbf{p} \mid \mathbf{p} \leq P_i(\mathbf{q})\right],$$
$$\overline{v}_i \coloneqq \mathrm{E}\left[\mathbf{p} \mid Q_i(\mathbf{p}) \leq \mathbf{q}\right] = \mathrm{E}\left[\mathbf{p} \mid \mathbf{p} \geq P_i(\mathbf{q})\right],$$

where $\mathbf{q} \sim U[0,1]$ 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 $\overline{v}_i$ 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 $\overline{v}_i$ as the *non-buying price expectation*. Note that, since Q_i is downward sloping, $\underline{v}_i \leq \overline{v}_i$.

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

Proof. Since the average price across all signal realizations at which type i buys is $\underline{v}_i$, incentive compatibility requires that $v_i \geq \underline{v}_i$ for otherwise this type could improve its expected utility by never buying. Similarly, incentive compatibility at signals where this type does not buy requires that $v_i \leq \overline{v}_i$. This proves the "only if" statement.

For the "if" statement, suppose $\underline{v}_i \leq v_i \leq \overline{v}_i$. Consider a binary information structure Φ_i generating signals 0 and 1 with $\Pr(\mathbf{x} = 1 \mid \mathbf{p} = p) = Q_i(p)$. Let type i's strategy be $q_i^*(x) = x$; this generates demand $Q_i(p)$. Since the posterior price expectations satisfy $\pi_i(1) = \underline{v}_i \leq v_i$ and $\pi_i(0) = \overline{v}_i \geq v_i$, the strategy q_i^* is optimal as needed. □

For each $s \in \left[\mathcal{CS}\left(Q,\underline{P}\right), \mathcal{CS}\left(Q,\overline{P}\right)\right]$, we construct 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:

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

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

We now shift our attention to the other direction, namely, that lying in the given interval is a necessary condition for consistency with the data. We say that a stochastic value $\mathbf{v}$ is *consistent with data* (Q,F) if there exists a model satisfying $\mathbf{v} \overset{d}{=} v_{\mathbf{i}}$ that rationalizes (Q,F). We provide an upper bound on stochastic values consistent with the data with respect to the increasing convex

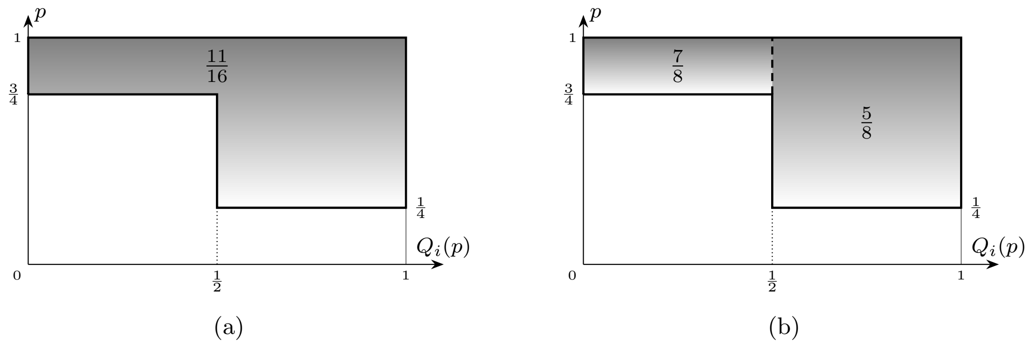

Figure 2: Example illustrating the effect of splitting demand into simple demands. (a) For $\mathbf{p} \sim U[0,1]$, the non-buying price expectation associated with demand Q_i is the expected price across the shaded region with both coordinates uniformly distributed, which is $\frac{11}{16}$. (b) Splitting type i into two (equally likely) types with simple demands gives non-buying price expectations of $\frac{7}{8}$ and $\frac{5}{8}$. The average of these two expectations conditional on not buying is $\frac{7}{8} \cdot \frac{1}{4} + \frac{5}{8} \cdot \frac{3}{4} = \frac{11}{16}$, whereas the unconditional average is $\frac{7}{8} \cdot \frac{1}{2} + \frac{5}{8} \cdot \frac{1}{2} = \frac{3}{4} > \frac{11}{16}$.

order. Let $\overline{\mathbf{v}} \coloneqq \overline{P}(\mathbf{i})$ and $\underline{\mathbf{v}} \coloneqq \underline{P}(\mathbf{i})$ for $\mathbf{i} \sim U[0,1]$. Thus $\overline{\mathbf{v}}$ 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 $\overline{\mathbf{v}}$; $\underline{\mathbf{v}}$ is needed in subsections 5.2 and 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* $\overline{\mathbf{v}} \succeq_{\text{icx}} \mathbf{v}$.

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

In the first step, we replace the value v_i of each type i with i's non-buying price expectation $\overline{v}_i$ (given the demand Q_i). Since, by Lemma 3, $\overline{v}_i \geq v_i$, this replacement leads to a first-order stochastic dominance increase in the stochastic value and hence also 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 $Q_j(p) = \mathbb{1}_{p \leq \rho}$ for some ρ; 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 $\overline{v}_j$ 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 $\mathrm{E}[\overline{v}_{\mathbf{j}}] \geq \overline{v}_i$ 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 $Q_i(p)$. Since higher values of $\overline{v}_j$ are associated with lower probabilities of not buying, when compared to the conditional expectation, the relative weight assigned to higher values $\overline{v}_j$ in the unconditional expectation is larger, as claimed. See Figure 2 for an illustration.

Taken together, the two steps transform the original model into one in which all types have threshold demands and values equal to their non-buying price expectations, and the average demand is Q; in other words, the result is the upper threshold model. The associated stochastic value is therefore $\overline{\mathbf{v}}$, as needed for the proof of Lemma 4.

We now use Lemma 4 to establish the upper bound on consumer surplus in Theorem 1. First note that, for a given stochastic value $\mathbf{v}$, incomplete information can lead to two types of loss for the consumer: (i) the probability of purchase at a given price may not be ex post optimal, i.e., $Q(p) \neq Q^{\mathrm{CI}}(p; \mathbf{v})$, 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 $\mathcal{CS}\left(Q, P^{\mathrm{CI}}(\cdot; \mathbf{v})\right)$ on the expected surplus for models with stochastic value $\mathbf{v}$. Therefore, the surplus s generated by any such model satisfies

$$s \leq \mathcal{CS}\left(Q, P^{\mathrm{CI}}(\cdot; \mathbf{v})\right) \leq \mathcal{CS}\left(Q, P^{\mathrm{CI}}\left(\cdot; \overline{\mathbf{v}}\right)\right) = \mathcal{CS}\left(Q, \overline{P}\right),$$

where the middle inequality follows from Lemmas 2 and 4. This concludes the proof of Theorem 1.

5.2 Related results

Theorem 1 has connections with several other questions.

First, Lemma 1 and Lemma 4 together have immediate implications for the counterfactual consumer surplus that would arise under complete information about the price.

Corollary 1. *Given data (Q, F), the consumer surplus that would arise under complete information about price p is no greater than $CS^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}\right)$.*

Second, as in Bayesian persuasion, the upper bound on consumer surplus can be viewed as the value of a concavification problem. (However, due to the high dimensionality of the problem, standard concavification techniques are not sufficient to identify a solution.) Just as Kamenica and Gentzkow (2011) split the prior belief into posterior beliefs under a Bayes-plausibility constraint, we split the aggregate demand $Q(p)$ into individual types' demands $Q_i(p)$ under the constraint $\mathrm{E}\left[Q_{\mathbf{i}}\right] = Q$. The objective in the persuasion problem is to maximize the expected value across posteriors. Likewise, our objective is to maximize the expected surplus $\mathrm{E}\left[s\left(Q_{\mathbf{i}}\right)\right]$, where $s\left(Q_i\right)$ is the highest surplus for type i consistent with demand Q_i. Threshold demands can be viewed as analogous to degenerate posteriors insofar as neither can be further split. Since the upper threshold model splits the original demand to threshold demands, it is analogous to full disclosure in the persuasion problem. The optimality of this model is nontrivial since the objective function $s(Q_i)$ is not convex.

Finally, all of the results stated so far have symmetric counterparts regarding a different welfare measure. Consumer surplus, as defined in (1), captures the consumer's benefit from freely choosing whether to buy relative to not having the option to buy the good. Define the *complementary consumer surplus*

$$\widehat{s} \coloneqq \mathrm{E}\left[\left(1 - q^*_{\mathbf{i}}(\mathbf{x})\right)\left(\mathbf{p} - v_{\mathbf{i}}\right)\right],$$

which captures the consumer's benefit from freely choosing relative to being forced to buy the good, i.e., not having the option *not* to buy the good. One can think of the complementary surplus as the gain relative to universal provision of the good financed by a tax equal to the average price.

Whereas consumer surplus is maximized when the consumer's value is high, complementary consumer surplus is maximized when the consumer's value is low. In both cases, however, greater spreads in values are associated with higher (complementary) surplus. Consequently, the relevant ranking of stochastic values for the complementary surplus is (the reverse of) second-order stochastic dominance: a lower bound with respect to $\succeq_{\text{sosd}}$ provides an upper bound on $\widehat{s}$.

Just as $\overline{\mathbf{v}}$ is the highest and the most spread out stochastic value consistent with the data, $\underline{\mathbf{v}}$ is the lowest and the most spread out such stochastic value. More precisely, $\underline{\mathbf{v}}$ is a lower bound with respect to $\succeq_{\text{sosd}}$ 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.

By analogy to the functional $\mathcal{CS}$, let

$$\widehat{\mathcal{CS}}\left(\tilde{Q}, \hat{P}\right) \coloneqq \mathrm{E}\left[\int_{\tilde{Q}(\mathbf{p})}^{1} \left(\mathbf{p} - \hat{P}(q)\right) dq\right].$$

Note that, just as $\mathcal{CS}\left(Q, \underline{P}\right) = 0$, $\widehat{\mathcal{CS}}\left(Q, \overline{P}\right) = 0$. Complementary consumer surplus $\widehat{s}$ is consistent with data (Q, F) if and only if

$$0 = \widehat{\mathcal{CS}}\left(Q, \overline{P}\right) \leq \widehat{s} \leq \widehat{\mathcal{CS}}\left(Q, \underline{P}\right);$$

this is the mirror image of Theorem 1. Since the proof of this result is analogous to that of Theorem 1, we omit the details.

Example 3. Consider the same data as in Example 1, namely, $Q(p) = 1 - p$ and $\mathbf{p} \sim U[0,1]$. In this case, the empirical inverse demand is $P(q) = 1 - q$, the lowered inverse demand is $\underline{P} = P(q)/2 = (1-q)/2$, giving an upper bound on complementary consumer surplus of $\widehat{\mathcal{CS}}(Q, \underline{P}) = 1/4$.

6 Multiple Datasets

The bounds on consumer surplus can be tightened if the analyst observes the consumer's choices under varying market conditions, which we refer to as *regimes*. We assume that the consumer's preferences are fixed across regimes, but the regimes may differ in the distribution of prices or in the information the consumer receives about prices (or both). For example, one such regime may correspond to a publicly announced "sale" associated with low distribution of prices while another corresponds to the same market in the absence of a sale. Alternatively, it could be the consumer's information that varies across regimes, for instance due to changes in how prices are presented as in the empirical studies of Chetty, Looney, and Kroft (2009) and Finkelstein (2009).

The analyst observes a *profile* of datasets $\left(Q^k, F^k\right)$, $k = 1, \ldots, K$, 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 $\mathbf{i}$, which specifies her value v_i for the good and her information structure Φ_i^k in each regime. The distribution of types and the value of each type are the same across all regimes. In each regime k, type i knows the prior distribution F^k and the information structure $\Phi_i^k(x \mid \mathbf{p} = p)$ that generates her signal. A model consists of a distribution of types together with a specification of $\left(v_i, \Phi_i^1, \ldots, \Phi_i^K\right)_i$ for each type.

We say that a model *rationalizes the profile of datasets* $\left(Q^k, F^k\right)_k$ if, for each regime k, it rationalizes dataset (Q^k, F^k) when each type i has information structure Φ_i^k. A stochastic value $\mathbf{v}$ is *consistent with the profile of datasets* $\left(Q^k, F^k\right)_k$ if there exists a model that rationalizes this profile and satisfies $\mathbf{v} \overset{d}{=} v_{\mathbf{i}}$.

Given a model, *consumer surplus in regime* k is

$$s^k = \mathrm{E}\left[q_{\mathbf{i}}^{k*}(\mathbf{x})\,(v_{\mathbf{i}} - \mathbf{p})\right],$$

where the signal $\mathbf{x}$ is generated according to Φ_i^k and q_i^{k*} is an optimal strategy for type i in regime k. Consumer surplus s^k 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 in regime k.

If, as in the previous sections, the analyst observes a single dataset, then, to determine rationalizable levels of consumer surplus, it suffices to consider types with threshold demands, corresponding to information structures that indicate only whether the price is above or below a particular (type-dependent) threshold. If, however, the analyst observes multiple regimes, some profiles of datasets can be rationalized only by types with more complex demands (even though each type's information structure, and therefore its demand, is allowed to depend on the regime). The following example illustrates this point.

Example 4. 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 by a single type with value $2/3$ who observes signals 0 or 1 with probability $\Phi^k(1 \mid \mathbf{p} = p) = Q^k(p)$ for $k = 1, 2$ and buys upon observing signal 1. Note that this type does not have a threshold demand in regime 1.

For $\varepsilon < 1/6$, this profile is not rationalizable with types having threshold demands (or, equivalently, threshold information structures). 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—which is their maximal incentive-compatible value—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 $p\leq 2/3$. △

Data from other regimes can help to pin down the surplus within a given regime. For instance, if ε is small in Example 4, then only values close to $2/3$ are consistent with the demand in regime 2. The surplus in regime 1 must therefore be approximately $2/3\times 1/2-1/6=1/6$.[12]

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 combine the upper bounds on the stochastic value with respect to the increasing convex order when considering each regime separately in such a way as to generate a tighter bound. The approach therefore requires combining bounds imposed on random variables with respect to a stochastic order. To do so, building on ideas of Gentzkow and Kamenica (2016) and Kolotilin et al. (2017), we exploit the connection described in section 4 between random variables and convex functions—in this case, the stochastic value and the complete information consumer surplus. According to Lemma 1, comparisons of stochastic values in the increasing convex order correspond to comparisons of complete information 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 complete information consumer surplus bounds. That this indeed provides an upper bound follows from an abstract result due to Müller and Scarsini (2006).

Let $\overline{\mathbf{v}}^k$ 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.[13] For each k, the bound $\overline{\mathbf{v}}^k$ corresponds to the convex function $CS^{\mathrm{CI}}\left(p;\overline{\mathbf{v}}^k\right)$. The upper bound using data across all regimes therefore corresponds to the largest convex function that lies below each $CS^{\mathrm{CI}}\left(p;\overline{\mathbf{v}}^k\right)$. Accordingly, let $CS_*^{\mathrm{CI}}(p)$ denote the convex closure of the function $\min_k CS^{\mathrm{CI}}\left(p;\overline{\mathbf{v}}^k\right)$.[14] By extension of the terminology of Kamenica and Gentzkow (2011), we refer to CS_*^{CI} as the *convexification* of $\min_k CS^{\mathrm{CI}}\left(p;\overline{\mathbf{v}}^k\right)$.

To map CS_*^{CI} back to a stochastic value, recall from section 4 that $CS^{\mathrm{CI}}(p;\mathbf{v})$ is the integral of the right tail of the complete information demand $Q^{\mathrm{CI}}(p;\mathbf{v})$, which is the complementary distribution function of $\mathbf{v}$. Define the demand function $Q_*^{\mathrm{CI}}(p)=-\partial_- CS_*^{\mathrm{CI}}(p)$, where ∂_- denotes the left derivative. Note that $1-Q_*^{\mathrm{CI}}$ is a distribution function and let $\overline{\mathbf{v}}_*$ be a stochastic value associated with this distribution.[15] Finally, let P_*^{CI} be the inverse demand to Q_*^{CI}. See Figure 3 for an illustration.

[12] Here, $1/2$ is the purchase probability and $1/6$ the expected expenditure in regime 1.

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

[14] Recall that the convex closure of a function $g(p)$ is the function that maps each p to $\inf\{s:(p,s)\in \mathrm{co}(g)\}$, where $\mathrm{co}(g)$ denotes the convex hull of the graph of the function g. In the terminology of convex analysis, CS_*^{CI} is the biconjugate function to $\min_k CS^{\mathrm{CI}}\left(p;\overline{\mathbf{v}}^k\right)$.

[15] Since CS_*^{CI} is convex, its left derivative exists, and Q_*^{CI} is nonincreasing and left-continuous. Additionally,

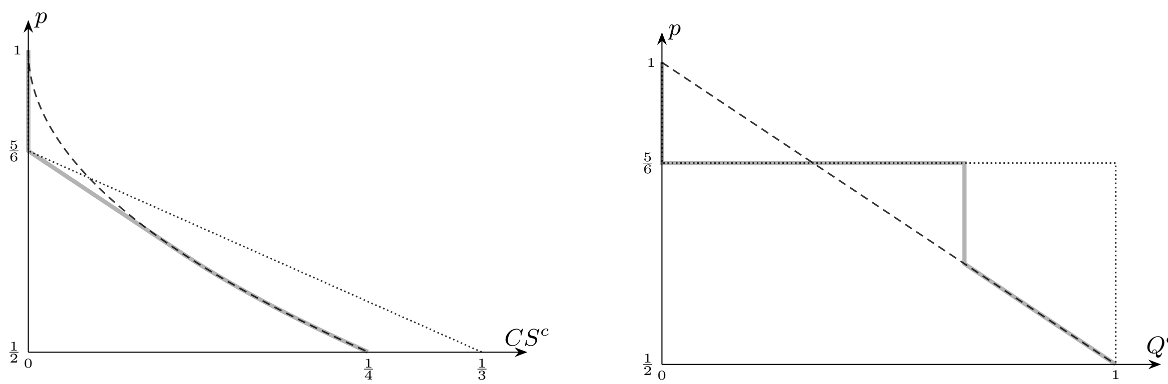

(a) (b)

Figure 3: Convexification for the two regimes described in Example 4 with $\varepsilon = 1/3$. The graphs depict prices $p \in [1/2, 1]$; for $p < 1/2$, the convexification is trivial. (a) Complete-information consumer-surplus functions: (dashed) $CS^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^1\right)$, (dotted) $CS^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^2\right)$, (thick) the convex closure $CS^{\mathrm{CI}}_*(p)$ (b) Complete-information demands associated with stochastic values: (dashed) $Q^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^1\right)$, (dotted) $Q^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^2\right)$, (thick) $Q^{\mathrm{CI}}_*(p)$.

	Bound from Theorem 1	Bound from Theorem 2
Regime 1	$\mathcal{CS}\left(Q^1, \overline{P}^1\right) = 0.25$	$\mathcal{CS}\left(Q^1, P^{\mathrm{CI}}_*\right) \approx 0.238$
Regime 2	$\mathcal{CS}\left(Q^2, \overline{P}^2\right) \approx 0.167$	$\mathcal{CS}\left(Q^2, P^{\mathrm{CI}}_*\right) \approx 0.125$

Table 1: Upper bounds on consumer surplus for the two regimes from Example 4 with the parameter $\varepsilon = 1/3$.

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

Lemma 5. *If a stochastic value* $\mathbf{v}$ *is consistent with the profile of datasets, then* $\overline{\mathbf{v}}_* \succeq_{\mathrm{icx}} \mathbf{v}$.

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

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 $\overline{\mathbf{v}}^k \succeq_{\mathrm{icx}} \mathbf{v}$ for each regime k. By Lemma 1, $\min_k CS^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^k\right) \geq CS^{\mathrm{CI}}(p; \mathbf{v})$. Since $CS^{\mathrm{CI}}(p; \mathbf{v})$ is convex in p, $CS^{\mathrm{CI}}(p; \mathbf{v})$ is no greater than the convexification of $\min_k CS^{\mathrm{CI}}\left(p; \overline{\mathbf{v}}^k\right)$. Finally, again by Lemma 1, $\overline{\mathbf{v}}_* \succeq_{\mathrm{icx}} \mathbf{v}$.

Combining Lemma 5 and Lemma 2 leads to the following upper bound on the consumer surplus within each regime. See Table 1 for an illustration.

Theorem 2. *If consumer surplus* s^k *in regime* k *is consistent with the profile of datasets, then* $s^k \leq \mathcal{CS}\left(Q^k, P^{\mathrm{CI}}_*\right)$.

Proof. Given a stochastic value $\mathbf{v}$, the consumer surplus in regime k is at most $\mathcal{CS}\left(Q^k, P^{\mathrm{CI}}(\cdot; \mathbf{v})\right)$ because this is the surplus associated with having the measure $Q^k(p)$ of types with the highest

$CS^{\mathrm{CI}}_*(p) = 0$ for $p > \overline{p}$ and $CS^{\mathrm{CI}}_*(p)$ has slope -1 for $p < \underline{p}$; hence, $\lim_{p \to -\infty} Q^{\mathrm{CI}}_*(p) = 1$ and $\lim_{p \to +\infty} Q^{\mathrm{CI}}_*(p) = 0$. Thus, $1 - Q^{\mathrm{CI}}_*$ is a distribution function.

values buy at each p. By Lemma 2, $\mathcal{CS}\left(Q^k, P^{\text{CI}}(\cdot;\mathbf{v})\right)$ is nondecreasing in $\mathbf{v}$ with respect to the increasing convex order. Finally, by Lemma 5, any stochastic value $\mathbf{v}$ consistent with the profile of datasets is bounded by $\overline{\mathbf{v}}_*$ in the increasing convex order. □

While the bound in this theorem is generally tighter than the bound from Theorem 1 obtained from the data in regime k alone, it is not itself a tight bound. For example, consider the two regimes from Example 4 with $\varepsilon \leq 1/6$. In this case, $CS^{\text{CI}}(p;\overline{\mathbf{v}}_2) \leq CS^{\text{CI}}(p;\overline{\mathbf{v}}_1)$ for all p, making the convexification trivial: $CS_*^{\text{CI}}(p) \equiv CS^{\text{CI}}(p;\overline{\mathbf{v}}_2)$. Therefore, $\overline{\mathbf{v}}_* \overset{d}{=} \overline{\mathbf{v}}_2$ 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 $Q^1(p)$).

Handling multiple regimes is complicated by the fact illustrated in Example 4 that decompositions into threshold demands are not generally sufficient to rationalize the data. The construction in Theorem 2 circumvents this complication by using a bound on surplus in each regime based on threshold demands. The downside of this approach 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.

Lemma 5, when combined with Lemma 1, also provides an upper bound on the counterfactual consumer surplus that would arise if prices were perfectly observed.

Corollary 2. *Given a profile of datasets $\left(Q^k, F^k\right)_k$, the consumer surplus that would arise under complete information about price p is no greater than $CS_*^{\text{CI}}(p)$.*

Like Theorem 2, this corollary improves on the bound for a single regime from Corollary 1 but is not generally tight.

6.2 Lower bound

An analogous construction to that for the upper bound can be used to obtain a nontrivial 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 $\succeq_{\text{sosd}}$ (as opposed to $\succeq_{\text{icx}}$ for the upper bound).

Define the *complete information complementary consumer surplus*

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

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 $\succeq_{\text{sosd}}$ implies the opposite ranking of $\widehat{CS}^{\text{CI}}(p;\mathbf{v})$, and the converse also holds provided the latter ranking is consistent across all p. Following the analogous construction to that for the upper bound, let $\widehat{CS}_*^{\text{CI}}(p)$ be the convexification of $\min_k \widehat{CS}^{\text{CI}}\left(p; \underline{\mathbf{v}}^k\right)$, where, for each k, $\underline{\mathbf{v}}^k$ is the lower bound on stochastic values with respect to $\succeq_{\text{sosd}}$ consistent with the dataset (Q^k, F^k) (see subsection 5.2). Let $\underline{\mathbf{v}}_*$ be the stochastic value associated with $\widehat{CS}_*^{\text{CI}}$.[16] Along the same lines as in Lemma 5, $\underline{\mathbf{v}}_*$ is a lower bound with respect to $\succeq_{\text{sosd}}$ on stochastic values $\mathbf{v}$ consistent with the profile of datasets.

Let $\widehat{P}(q; \mathbf{v}) \coloneqq P^{\text{CI}}(1-q; \mathbf{v})$ denote the qth lowest quantile of $\mathbf{v}$.

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

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 $\mathcal{CS}\left(Q^k, \widehat{P}(; \mathbf{v})\right)$; 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 $\succeq_{\text{sosd}}$ by $\underline{\mathbf{v}}_*$, the bound on s^k from the theorem applies.

To illustrate the lower bound, consider the regimes from Example 4 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 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.

7 Bounds on Counterfactual Demand

Returning to the original model in which the analyst observes a single dataset (Q, F), we now consider the counterfactual demand that would arise if the distribution of prices F were replaced by a deterministic price, or equivalently, the consumer were to perfectly observe the realized price. 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

$$\overline{\overline{P}}(q) \coloneqq \mathrm{E}\left[\mathbf{p} \mid \mathbf{p} \geq P(\mathbf{q}), \mathbf{q} \leq q\right]$$
$$\text{and} \qquad \underline{\underline{P}}(q) \coloneqq \mathrm{E}\left[\mathbf{p} \mid \mathbf{p} \leq P(\mathbf{q}), \mathbf{q} \geq q\right],$$

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

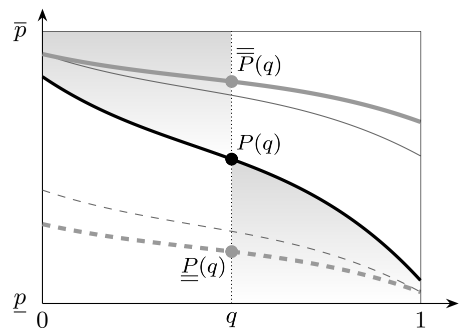

Figure 4: Doubly elevated and doubly lowered inverse demands, $\overline{\overline{P}}$ (thick grey) and $\underline{\underline{P}}$ (thick dashed), for given inverse demand P (black). For each q, $\overline{\overline{P}}(q)$ is the expected price conditional on $\mathbf{p}$ and $\mathbf{q}$ lying in 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.

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

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

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

for all q.

The bounds in this theorem are tight in the sense that for each q, there exists a stochastic value $\mathbf{v}$ consistent with the data such that $P^{\mathrm{CI}}(q;\mathbf{v}) = \overline{\overline{P}}(q)$, 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 complete information inverse demand $P^{\mathrm{CI}}(q;\mathbf{v})$ 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 incentive-compatible value given 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 $\overline{\overline{P}}(q)$. 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 perfectly observes and understands prices, the consumer surplus cannot be point identified from price and demand data. Nonetheless, Bayes rationality imposes 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 information.

Several 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 involving simple information structures (which are sufficient in the single-regime case), whereas, in principle, identification of the surplus can be improved by simultaneously rationalizing the profile of datasets using more complex information structures. Second, one could consider richer data than those studied here. For instance, outside of the unit-demand environment, our results apply if the analyst observes only the average demand of the consumer. If instead the dataset includes the full distribution of quantities chosen, our bounds continue to apply but are not generally tight if the data are inconsistent with simple information structures. Identifying tight bounds therefore again requires consideration of more complex information structures than those needed in the unit-demand or average-demand case. Finally, the tightness of our bounds depends on how much freedom the analyst has in the choice of information structures to rationalize the data. In some contexts, it may be plausible to assume that the consumer has at least some minimal information consisting of a particular signal. If this information is also observed by the analyst—as, for example, if the first digit of the price were perfectly observed by the consumer—then each realization of the signal can be treated as a different regime. In any case, since additional information restricts the set of models that can rationalize the data, it could potentially be used to narrow the bounds on surplus or counterfactual demand.

A Proofs

Proof of Lemma 2. Note that

$$\mathcal{CS}\left(\tilde{Q}, P^{\mathrm{CI}}(\cdot;\mathbf{v})\right) = \mathrm{E}\left[\int_0^{\tilde{Q}(\mathbf{p})} P^{\mathrm{CI}}(q;\mathbf{v})dq\right] - \mathrm{E}\left[\tilde{Q}(\mathbf{p})\mathbf{p}\right].$$

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

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

the first inequality follows from Lemma 1 and the second because $CS^{\text{CI}}(p;\mathbf{v}) = \max_{q'} \int_0^{q'} (P^{\text{CI}}(q;\mathbf{v}) - p)dq$. □

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

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

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

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

$$\overline{v}_i \leq \mathrm{E}\left[\overline{P}_i(\mathbf{q})\right] = \mathrm{E}[\overline{\mathbf{v}}_i],$$

as needed.

Step 3: We conclude by proving that $\mathbf{v}'' \overset{d}{=} \overline{\mathbf{v}}$. Consider any p at which Q is continuous and let

$\tilde{v} = \mathrm{E}[\mathbf{p} \mid \mathbf{p} \geq p]$. For any $j \in [0,1]$,

$$j = \Pr\left(\overline{\mathbf{v}} \geq \tilde{v}\right) \quad \Longrightarrow \quad \overline{P}(j) = \tilde{v} \quad \Longrightarrow \quad P(j) = p \quad \Longrightarrow \quad j = Q(p).$$

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

$$\Pr\left(\mathbf{v}'' \geq \tilde{v}\right) = \Pr\left(\overline{\mathbf{v}}_{\mathbf{i}} \geq \tilde{v}\right) = \mathrm{E}\left[Q_{\mathbf{i}}(p)\right] = Q(p) = \Pr\left(\overline{\mathbf{v}} \geq \tilde{v}\right)$$

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

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

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

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 each price p.

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

$$\mathcal{CS}\left(\tilde{Q}, \widehat{P}(\cdot;\mathbf{v})\right) \geq \mathcal{CS}\left(\tilde{Q}, \widehat{P}(\cdot;\mathbf{v}')\right).$$

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 argument of $\mathcal{CS}$. In fact, we prove a somewhat stronger statement: if $\mathbf{v} \succeq_{\mathrm{sosd}} \mathbf{v}'$, then $\int_0^{q^*} \widehat{P}\left(q;\mathbf{v}\right) dq \geq \int_0^{q^*} \widehat{P}\left(q;\mathbf{v}'\right) dq$ for every q^*. Fixing q^* and letting $p = \widehat{P}\left(q^*;\mathbf{v}'\right)$, we have

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

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

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

$$s^k \geq \mathcal{CS}\left(Q^k, \widehat{P}(\cdot;\mathbf{v})\right) \geq \mathcal{CS}\left(Q^k, \widehat{P}\left(\cdot;\underline{\mathbf{v}}_*\right)\right)$$

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

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 $[\underline{p}, \overline{p}]$ onto $[0,q]$. Let

$$\tilde{v}\left(\tilde{Q};q\right) := \mathrm{E}\left[\mathbf{p} \mid \mathbf{q} \geq \tilde{Q}(\mathbf{p})\right]$$
$$\text{and} \qquad w\left(\tilde{Q};q\right) := q\,\mathrm{Pr}\left(\mathbf{q} \geq \tilde{Q}(\mathbf{p})\right)$$

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 $\mathrm{Pr}\left(\mathbf{i} \in I'\right) = q$ and $\tilde{Q}(p) = \int_{i \in I'} Q_i(p) dM(i)$. Then $\tilde{v}\left(\tilde{Q};q\right)$ is the expected price conditional on a type randomly drawn from I' not making a purchase and $w\left(\tilde{Q};q\right)$ is the probability that a randomly drawn type lies in I' and does not buy.

Note the following recursion. 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 $Q_a + Q_b = \tilde{Q}$,

$$\tilde{v}\left(\tilde{Q};q\right) = \frac{w\left(Q_a;q_a\right)\tilde{v}\left(Q_a;q_a\right) + w\left(Q_b;q_b\right)\tilde{v}\left(Q_b;q_b\right)}{w\left(Q_a;q_a\right) + w\left(Q_b;q_b\right)}. \tag{4}$$

Given any model and a subset I' of types such that $\mathrm{Pr}\left(\mathbf{i} \in I'\right) = q$, 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 $\mathrm{Pr}\left(\mathbf{i} \in I'\right) = q$ is at most $\overline{\overline{P}}(q)$.

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

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

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

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

$$\tilde{v}\left(\tilde{Q};q\right) \leq \tilde{v}\left(Q^*;q\right) \tag{5}$$

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 $\mathrm{Pr}\left(\mathbf{i} \in I'\right) = q$. 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 $Q(p) - \tilde{Q}(p)$ is nonnegative and nonincreasing. Let $p^* := P\left(q\right)$ and $q^* := \tilde{Q}\left(p^*\right)$. Since $\tilde{Q}(p) \leq Q^*(p) \leq q$ for all p, we have that $q^* \leq q$. Define the demand function $Q_0(p) := \min\left\{\tilde{Q}(p), q^*\right\}$ 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.

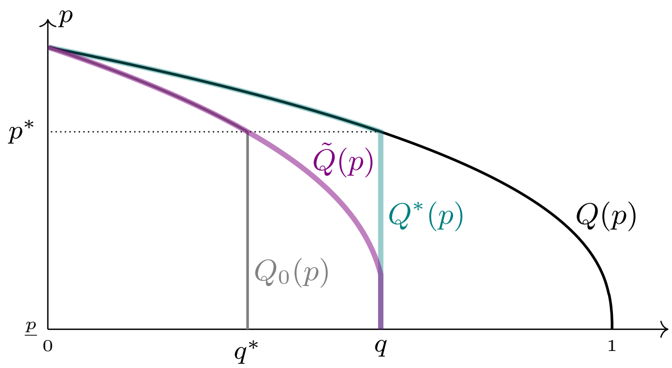

Figure 5: Illustration of the definitions of Q^*, $\tilde{Q}$, and Q_0.

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 P_0 and P_1 lie above P_2 in the strong sense that both 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 $\mathrm{E}\left[\mathbf{p} \mid \mathbf{p} \geq \tilde{P}(\mathbf{q})\right]$ 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 $q \Pr\left(\mathbf{p} \geq \tilde{P}(\mathbf{q})\right)$. It follows that $\tilde{v}\left(Q_1; q - q^*\right) \geq \tilde{v}\left(Q_2; q - q^*\right)$, $\tilde{v}\left(Q_0; q^*\right) \geq \tilde{v}\left(Q_2; q - q^*\right)$, and $w\left(Q_2; q - q^*\right) \geq w\left(Q_1; q - q^*\right)$. Finally, since $Q^* = Q_0 + Q_1$ and $\tilde{Q} = Q_0 + Q_2$, we have from (4) that

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

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

Therefore,

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

which establishes inequality (5), as needed. □

B General quasilinear utility

In this appendix, we show that the assumption of unit demands is without loss of generality within the class of quasilinear preferences. As in the main setup, the analyst observes data $(\hat{Q}, F)$, where F is the distribution of prices with support $\left[\underline{p}, \overline{p}\right]$. The *average demand* $\hat{Q}$ is a nonincreasing and nonnegative function with $\hat{Q}\left(\overline{p}\right) = 0$. While the demand $Q(p)$ in the main setup is the probability that the consumer purchases one unit of good, $\hat{Q}(p)$ denotes the average demand of a consumer who may purchase variable quantities. Equivalently, $\hat{Q}(p)$ can be viewed as the aggregate average demand of a population of consumers. Without loss of generality, we normalize $\hat{Q}(\underline{p})$ to 1.

The consumer's types are indexed by $i \in I \subset \mathbb{R}^K$ for some $K \geq 1$. A consumer of type i who chooses quantity $q \geq 0$ at price p receives quasilinear utility $u_i(q) - pq$, where the gross utility $u_i : [0, \infty) \to [0, \infty]$ is a nondecreasing concave function with $u_i(0)$ normalized to 0. Just as in the main setup, a consumer of type i receives a signal $x \in \mathbb{R}$ generated by an information structure $\Phi_i(x \mid \mathbf{p} = p)$ and we let $\pi_i(x) = \mathrm{E}[\mathbf{p} \mid \mathbf{x} = x]$. A consumer of type i chooses quantity $q_i^*(x) \in \arg\max_{q \geq 0} \{u_i(q) - \pi_i(x)q\}$.[17] Each type i generates a demand function $\hat{Q}_i(p) = \mathrm{E}\left[q_i^*\left(\mathbf{x}\right) \mid \mathbf{p} = p\right]$ that we assume is nondecreasing. There is no fixed upper bound on the values of q_i^* or $\hat{Q}_i$.

A *model* consists of a Borel probability measure μ on I together with a complete specification $(u_i, \Phi_i)_{i \in I}$ of preferences and information structures. We say that a given model *rationalizes* data $(\hat{Q}, F)$ if $\hat{Q}(p) = \mathrm{E}\left[\hat{Q}_{\mathbf{i}}(p)\right]$ for all prices p. Given a model that rationalizes the data, the consumer surplus is $s = \mathrm{E}\left[u_{\mathbf{i}}\left(q_{\mathbf{i}}^*(\mathbf{x})\right) - \mathbf{p}q_{\mathbf{i}}^*(\mathbf{x})\right]$.

For the purpose of analyzing the levels of consumer surplus consistent with the data, it turns out that it is without loss of generality to restrict attention to unit demand preferences. We say that type i has *unit demand* with value v_i if $u_i(q) = v_i \min\{q, 1\}$. Such a type purchases one unit of the good if $v_i > \pi_i(x)$ and zero units if $v_i < \pi_i(x)$.

Proposition 1. *Suppose that s is the consumer surplus in some model that rationalizes the data. Then there exists a model with unit demands that also rationalizes the data and generates consumer surplus s.*

The proposition implies that, given the data, it is impossible to distinguish between rationalizations in which types may have complicated quasilinear utility functions and those in which each type has unit demand. Moreover, rationalizations with unit demands can generate any consumer surplus that can be obtained by some other rationalization.

To illustrate the idea behind the proposition, consider a model consisting of a single type that receives marginal utility 2 from the first half-unit of the good, 1 from the second half-unit, and 0 from all subsequent units. To match the behavior in this model using only unit demands, we replace the single type with two equally likely types, one having value 2 and the other value 1. Each of the new types has the same information structure as the single type in the initial model. At each signal realization, both the total expected demand and the total expected utility of these

[17] As in the main setup, if there are multiple optimal choices, we include a selection among them in the description of the type but omit this selection from the notation.

two types are identical to those of the original type. Therefore, the aggregate demand and the consumer surplus at each price are identical to those in the original model.

While Proposition 1 implies that Theorem 1 extends immediately to general quasilinear demands, and the same is true of our upper bound with multiple regimes (Theorem 2), there is a subtle issue when applying it to the lower bound with multiple regimes (Theorem 3). For the construction of the lower bound, we associated purchases with the types having the lowest values. In the case of general quasilinear utilities, it could be that these lowest types do not purchase at all in the initial data, and that the total demand is explained by a subset of types that buy multiple units. Without additional assumptions, then, it is not possible to obtain a nontrivial lower bound.

It is important for Proposition 1 that the analyst observes only the average demand and not the consumer's full stochastic choice behavior, i.e., the full distribution of quantities chosen at each price. (While these two types of data are equivalent with unit demands, with general quasilinear utilities they are not.) Our bounds trivially apply to such richer data since the analyst could simply ignore the additional information and use only the average demand. However, the bounds in our main result are not generally tight in this case. The key issue is whether the observed data can be rationalized with unit demand preferences; if so, the bounds remain tight. However, there are many such datasets that cannot be rationalized in this way, such as those with deterministic demands that take on at least three distinct quantities. In that case, obtaining tight bounds appears to be much more difficult.

Proof of Proposition 1. For any type $i \in I$ in the original model and any quantity $a \in [0, \hat{Q}_i(\underline{p})]$, define a new type $j = (i, a) \in I \times (0, \infty]$ that has the original information structure $\Phi_j = \Phi_i$ and gross utility $u_j(q) = v_j \min\{q, 1\}$, where v_j is the left derivative of u_i at a.[18] Let the measure of types j be $\tilde{\mu} = \mu \times \lambda$, where λ is the Lebesgue measure. Note that, because of the normalization $\hat{Q}(\underline{p}) = 1$, $\tilde{\mu}$ is indeed a probability measure.

Given any $a \in [0, \hat{Q}_i(\underline{p})]$ and any signal x, if $q_i^*(x)$ is optimal for type i, then $q^*_{(i,a)}(x) = \mathbb{1}_{a \leq q_i^*(x)}$ is optimal for type (i, a). Indeed, the optimality condition for $q_i^*(x)$ is equivalent to $v_{(i,a)} \geq \pi(x)$ for all $a \in (0, q_i^*(x)]$ and $v_{(i,a)} \leq \pi(x)$ for all $a \in [q_i^*(x), \overline{q}]$. Thus the demand of type i at any signal realization x is equal to the aggregate demand of the types $(i, a)_a$ at x. Furthermore, the original gross utility function u_i can be written as an integral of its left derivative:

$$u_i(q) = \int_0^q v_{(i,a)} da = \int_0^{\hat{Q}_i(\underline{p})} u_{(i,a)} \left(\mathbb{1}_{a \leq q}\right) da.$$

[18]Since u_i is a concave function, the left derivative exists.

Since $\Phi_i = \Phi_{(i,a)}$ for all (i, a), the two models generate identical average demands:

$$\begin{aligned}\int_I \hat{Q}_i(p) d\mu(i) &= \int_I \mathrm{E}\left[q_i^*(\mathbf{x}) \mid \mathbf{p} = p\right] d\mu(i) \\ &= \int_I \mathrm{E}\left[\int_0^{\hat{Q}_i(p)} \mathbb{1}_{a \leq q_i^*(\mathbf{x})} da \mid \mathbf{p} = p\right] d\mu(i) \\ &= \int_J \mathrm{E}\left[q_j^*(\mathbf{x}) \mid \mathbf{p} = p\right] d\tilde{\mu}(j) \\ &= \int_J \hat{Q}_j(p) d\tilde{\mu}(j),\end{aligned}$$

where $J \coloneqq I \times [0, \infty)$.

Likewise, the two models agree on the gross consumer surplus at each price:

$$\begin{aligned}\int_I \mathrm{E}\left[u\left(q_i^*(\mathbf{x})\right) \mid \mathbf{p} = p\right] d\mu(i) &= \int_I \mathrm{E}\left[\int_0^{\hat{Q}_i(p)} u_{(i,a)}\left(\mathbb{1}_{a \leq q_i^*(\mathbf{x})}\right) da \mid \mathbf{p} = p\right] d\mu(i) \\ &= \int_J \mathrm{E}\left[u_j\left(q_j^*(\mathbf{x})\right) \mid \mathbf{p} = p\right] d\tilde{\mu}(j).\end{aligned}$$

Since the two models generate the same average demand, they have identical total expenditure at each price, and therefore the same average net consumer surplus. □

References

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.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica 80*(3), 969–1018.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review 99*(4), 1145–77.

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 (2021). Surplus sharing in Cournot oligopoly. Forthcoming in *Theoretical Economics*.

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.

Gul, F., W. Pesendorfer, and T. Strzalecki (2017). Coarse competitive equilibrium and extreme prices. *American Economic Review 107*(1), 109–37.

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.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kang, Z. Y. and S. Vasserman (2022). Robust bounds for welfare analysis. NBER Working Paper 29656.

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.

Rogerson, W. P. (1980). Aggregate expected consumer surplus as a welfare index with an application to price stabilization. *Econometrica 48*(2), 423–436.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. a definition. *Journal of Economic Theory 2*(3), 225–243.

Samuelson, P. A. (1972). The consumer does benefit from feasible price stability. *Quarterly Journal of Economics 86*(3), 476–493.

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.

Waugh, F. V. (1944). Does the consumer benefit from price instability? *Quarterly Journal of Economics 58*(4), 602–614.