NBER WORKING PAPER SERIES

STRESS TESTING STRUCTURAL MODELS OF UNOBSERVED HETEROGENEITY: ROBUST INFERENCE ON OPTIMAL NONLINEAR PRICING

Aaron L. Bodoh-Creed
Brent R. Hickman
John A. List
Ian Muir
Gregory K. Sun

Working Paper 31647
http://www.nber.org/papers/w31647

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023, Revised November 2023

John List was Chief Economist at Lyft when this research was carried out. He is now Chief Economist at Walmart. We would like to thank the following individuals for useful conversations in developing this work: Barton Hamilton, Stephen Ryan, Rob Clark, Dan Sacks, Jean-Francois Houde, Barry Nalebuff, Jakub Kastl, Ismael Mourifie, Phil Haile, Joachim Freyberger, and seminar participants at WashU-Olin, Queen's University, Prince- ton, and participants of the 2022 IIOC meetings, and 2023 NASMES meetings. Note: Original Version: February 2022. Current Version: October 2023. Previous versions of this manuscript circulated under the title “Stress Testing a Structural Model of Nonlinear Pricing: Robust Inference on Intensive-Margin Consumer Demand.” The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w31647

Stress Testing Structural Models of Unobserved Heterogeneity: Robust Inference on Optimal Nonlinear Pricing
Aaron L. Bodoh-Creed, Brent R. Hickman, John A. List, Ian Muir, and Gregory K. Sun
NBER Working Paper No. 31647
August 2023, Revised November 2023
JEL No. B4,C14,C51,C52,C93,D04,J02,L1,L5

ABSTRACT

We propose a suite of tools for empirical market design in adverse-selection settings where point identification based on exogenous price variation is hampered by multi-dimensional unobserved heterogeneity. Despite significant data limitations, we are able to derive informative bounds on counterfactual consumer demand under out-of-sample price changes. These bounds arise because empirically plausible DGPs must respect the Law of Demand and the observed shift(s) in aggregate demand resulting from a known exogenous price change(s). The bounds facilitate robust policy prescriptions using rich, internal data sources similar to those available in many real- world settings, including our empirical application to rideshare demand. Our partial identification approach enables viable, welfare-improving, nonlinear pricing design while achieving robustness against worst-case deviations from baseline model assumptions. As a side benefit, our framework also provides novel insights into optimal experimental design for pricing RCTs.

Aaron L. Bodoh-Creed
Amazon, Inc
abodoh.creed@gmail.com

Brent R. Hickman
Olin Business School
Washington University in Saint Louis
One Brookings Drive, Campus Box #1133
Saint Louis, MO 63130
hickmanbr@wustl.edu

John A. List
Department of Economics
University of Chicago
1126 East 59th
Chicago, IL 60637
and Australian National University
and also NBER
jlist@uchicago.edu

Ian Muir
Walmart
muir.ian.m@gmail.com

Gregory K. Sun
Washington University in Saint Louis
greg.s@wustl.edu

A data appendix is available at
http://www.nber.org/data-appendix/w31647
Stress Testing Structural Models of Unobserved Heterogeneity:
Robust Inference on Optimal Nonlinear Pricing available here:
https://drive.google.com/file/d/1zDEYdcNs0q_8qx2Q9rPfWrsAaFsVpgmx/view?usp=sharing

1. Introduction

Demand analysis has been a staple of academic and policy-oriented research in industrial organization for several decades. Beginning with the seminal work of Berry (1994), Berry et al. (1995), and Nevo (2003), empirical discrete-choice demand systems became an especially prolific branch of the industrial organization literature. While powerful and useful given a wealth of rich market-level data sets in recent years (e.g., Nielsen's scanner data), this family of methodologies focuses primarily on extensive-margin consumer decisions—that is, which among many substitutable products to buy—and often places less emphasis on, or abstracts away from intensive-margin consumer decisions—i.e., *how much* of a product to buy from a particular producer/provider.

More recently, subscription-based consumption platforms have proliferated, both in brick-and-mortar applications (e.g., Costco; Sam's Club; Club Car Wash), in e-commerce (e.g., Instacart, Audible.com), in services (e.g., Uber Technologies, Lyft, Inc., and YouTubeTV, Tesla, and Chargepoint), and even in meal-kits (e.g., Hello Fresh, Home Chef, and Blue Apron). Each of these firms (and many others) share an interest in optimal nonlinear pricing; that is, they all offer some sort of subscription program (or a menu of subscription options) to consumers, which involves an up-front fee in exchange for a percentage-based volume discount. The existing body of empirical methodologies for demand analysis are not well-suited to these business models for several reasons. First, intensive-margin demand heterogeneity (in addition to cross-platform substitution) is a central concern for market design in these settings. Second, typical data used by discrete-choice demand systems are rich in their coverage of many products or firms within and across markets, but are typically less rich when it comes to consumer-level transaction data within a given firm. Third, in light of prevailing data limitations, discrete-choice demand approaches typically achieve tractability by restricting the distributions of consumer tastes to parsimonious parametric families.

In contrast, many firms have generated a wealth of internal datasets with the opposite strengths and weaknesses: they tend to be rich on transactions with the firm's consumer base, often even including randomized controlled trials (RCTs) on pricing, but are anemic or silent on prices and market shares of rival producers. Moreover, aside from facilitating a study of individual demand-intensity variation, rich consumer-level data may reduce our dependence on parametric assumptions. Our goal is to develop a new complementary approach with a focus on intensive-margin demand in settings where explicit randomized pricing and consumer-transaction-level data are plentiful, but market-level data are not. This capability, relevant to both policy-makers and business practitioners, will facilitate answers to new and interesting questions on optimal design of nonlinear pricing in settings characterized by unobserved heterogeneity in consumer demand intensity.

While our empirical case study in this paper is from an industrial organization context, our econometric methodology applies to a broad class of adverse-selection models with applications including procurement, regulation, taxation, labor-supply, and insurance/healthcare demand. To fix ideas, consider a setting where the principal incentivizes some desired activity q on the part of a continuum of agents who vary by their willingness to engage in the activity. Assuming some exogenous source of variation in payoffs offered to agents from say $P_0(q)$ to $P_1(q)$, the question we examine is, what can be learned from the resulting aggregate shift in the distribution of agents' choices from $G_0(q)$ to $G_1(q)$? We show that the answer to this question hinges crucially on whether agent unobserved heterogeneity is single- or multi-dimensional. If agents with the same observed choices q under default incentives $P_0(q)$ are homogeneously price responsive, then inference on the underlying (single-dimensional) type distribution is straightforward. However, an important challenge arises if agents with similar observed choices are heterogeneously responsive to price changes, in which case inference on the underlying (multi-dimensional) type distribution is problematic.

We begin with a parsimonious model of intensive-margin demand, and we discuss how recent results from econometric theory (Torgovitsky (2015); D'Haultfoeuille and Février (2015); D'Haultfœuille and Février (2020)) may suffice to identify model primitives (e.g., utility functions and demand-type distributions), using exogenous price variation. We discuss how this baseline result is due to a key condition which we refer to as *rank stability*—that is, when prices exogenously vary (in the cross-section), the mapping between observed demand quantiles and unobserved consumer type quantiles remains stable.[1] We then discuss various plausible phenomena that may violate the central rank-stability condition and complicate model identification.

We first show how rank-stability violations of known forms—e.g., consumer behavioral mistakes in subscription purchasing—can be explicitly modeled to restore identification using basic available observables. We then discuss a more difficult empirical problem: rank-stability violations of *unknown* forms, either because the underlying model of consumer behavior is incomplete, or because the econometrician lacks sufficient observables to explicitly control for them. These may stem from phenomena such as (but not limited to) brand loyalty heterogeneity and unobserved substitution between the firm's product/service and that of its rival. This is a particularly thorny issue because such problems cannot be modeled directly using rich, internal firm data, which lack prices and market shares of rival firms. We develop a robust empirical strategy by deriving identifiable bounds on counterfactual demand distributions under (out-of-sample) price changes, despite data limitations.

Our partial identification approach focuses on bounding the maximal and minimal price sensitivity of consumers most likely to purchase a subscription, subject to consistency with observables and the law of demand (LoD). It assumes availability of plausibly exogenous, cross-sectional, price variation, where a random set of agents are assigned to "control" status, making demand decisions

[1] Our identification approach in this paper hinges on *cross-sectional* price variation, so the term "robust" is taken to signify robustness against arbitrary, unobserved rank-stability violations in the latent demand system, holding prices/features of rival goods fixed. We leave questions about responses by rival platforms to future research.

under default price p_0, while others are assigned to "treatment" status with a discount of d_0 off the default price. We show that predictions made by the fully specified rank-stable model correspond to a sharp upper bound on price responsiveness. Thus, point estimates under the (potentially flawed) rank-stability assumption still provide a useful benchmark for empirical market design. We also derive a sharp lower bound on price responsiveness at various consumer quantiles, subject to consistency with observables and the LoD. Thus, we show that the partially specified model still places informative restrictions on the underlying data-generating process.

While derivations of the bounds are fairly technical, there is some simple intuition behind our method. Consider a firm that implements a nonlinear pricing scheme in the form of a subscription offering, where consumers may pay $\$S$ upfront for a discount of $d \in (0,1)$ over the ensuing period. Suppose that the firm assumes rank stability and forecasts counterfactual demand under arbitrary (S,d) pairs in order to optimize the subscription program. Suppose further that consumers are characterized by two dimensions of unobserved heterogeneity—namely, demand intensity θ_i, and brand loyalty to the firm α_i—but only the aggregate demand distributions under control (i.e., price p_0) and treatment (i.e., price $p_0(1-d_0)$) are observed, while substitution between the firm and its rivals is unobserved. The least loyal consumers will buy less from the firm at baseline prices, but will be *more* elastic than their loyal counterparts, as they shift consumption away from the outside option under discount pricing. Moreover, one can characterize a maximally "adversarial" consumer base which behaves in such a way as to minimize profits from the rank-stable optimum subscription offer (S^*, d^*). This worst-case scenario involves disloyal consumers accounting for a maximal fraction of the observed aggregate shift from the control demand CDF to the treatment demand CDF. This in turn implies that loyal consumers, who purchase most under default prices, are minimally price sensitive, meaning that they are more likely to buy a subscription but increase spending very little, so that (S^*, d^*) is a (nearly) zero-sum cash transfer from the firm.

As a side benefit, our analysis also produces various novel insights to guide effective RCT design. Our modelling approach (see Section 2.4.1) highlights how randomized pricing may be insufficient alone to deliver point identification in the presence of multi-dimensional agent heterogeneity. We show that this problem may even persist despite continuous price variation (see Section 4.3). Finally, our derivations of robust bounds to grapple with this issue (see Remark 4, Section 2.5, Section 4 introduction, and Section 4.3) produce novel, concrete guidance on optimal randomized incentives for a researcher wishing to maximize ex-post inferential power.

We implement our approach in an empirical case study with rideshare data. Using the above insights, we estimate the baseline, single-dimensional (rank-stable) model, and we perform a series of model specification tests to probe for evidence of rank-stability violations. Interestingly, we find that our first over-identifying test, using multiple arms of the same RCT, fails to reject the rank-stable model, while a more stringent test combining datasets from two separate RCTs does reject the rank-stable model. We argue that this is so because of how the two RCT designs induce differing selection patterns of unobserved substitution behavior. This finding sheds further light on subtle potential limitations of inference from RCT data. If the underlying model is fundamentally mis-specified (e.g., single- vs multi-dimensional heterogeneity), then not only may point identification be compromised, but the researcher's ability to detect a mis-specification may also be compromised given standard over-identifying restrictions from multiple arms of a single RCT.

Using baseline (rank-stable) point estimates of the structural model, we compute a menu of profit-maximizing subscription offerings, as well as an optimal single subscription plan, in order to establish best-case profits. We find that the optimal single offering does nearly as well as the continuum menu of offerings that achieves a fully separating equilibrium. This result helps explain why real-world firms like Costco and Lyft tend to prefer simple subscription programs with only one or a small number of offered options. We then estimate robust lower bounds to show how an optimal subscription offer should be adjusted under less favorable circumstances where the baseline model overstates price sensitivity of likely subscribers. Intuitively, these adjustments make the subscription program somewhat less generous (i.e., higher up-front fees and/or lower discounts) to hedge profitability against worst-case unobserved phenomena. We find that hedging the baseline optimal policy against the worst case requires only relatively small adjustments for robustness, up to a point where non-rank-stable consumers reach a critical mass and subscription plans cease to be an effective business strategy. We also show how pre-RCT data can be used to pin down the most empirically relevant worst-case scenario, by estimating the mass of non-rank-stable consumers. Our estimate of this mass (roughly 16%) suggests strongly that a meaningful range of profit-improving subscription offers does exist under data-generating processes that are consistent with observables. Moreover, the robust subscription offer is able to fully insure the principal against deviations from baseline model assumptions, while achieving roughly 95% of baseline projected subscription profits.

Our partial identification framework is applicable outside of just demand estimation in IO. It can be used to stress-test the policy implications of a variety of adverse-selection models under potential multi-dimensional heterogeneity, including several prominent examples. First, Laffont and Tirole (1986) model a procurer that wishes to design a contract to incentivize a monopolist to complete a project while exerting effort to reduce costs. Monopolists are unobservably heterogeneous along a single dimension of productivity, but one might expect that they also differ by costs of unobserved managerial effort. Second, in the regulation context, Kang and Silveira (2021) analyze a novel framework where firms have private information on costs of externality abatement per unit of output. If the regulator can only observe total pollution, then one can imagine related settings where firms vary both by their baseline emission level (i.e., production technology) as well as by their abatement technology. In such a world, inferences on both dimensions may be empirically relevant to optimally incentivizing abatement on a fixed regulatory budget. Third, in the optimal taxation context, Mirrlees (1971) modeled households as varying only by (unobserved) labor productivity, but one might expect them to vary by consumption preferences as well. Fourth, in the labor context, D'Haultfœuille and Février (2020) derive optimal empirical piece-rate contracts for census workers that exhibit heterogeneous labor-supply costs. One can imagine other labor settings where variation in the productivity of effort is also crucial to the design of optimal incentive schemes.[2] To the extent that higher dimensions of agent heterogeneity are empirically salient, our robust bounds approach provides a path forward for viable inference in the face of formidable data limitations, such as availability of only a single source of exogenous incentive variation.

1.1. **Related Literature.** Our paper contributes to several literatures. First, we build on the work of Maskin and Riley (1984), who provide a theoretical framework for how a monopolist should set prices under heterogeneous consumer demand intensity.[3] A related empirical paper is Luo et al. (2018), who establish conditions under which structural demand primitives are identified from observations of equilibrium prices and quantities. Our approach is complementary to theirs: they achieve identification using equilibrium conditions while we do so using exogenous cross-sectional price variation under minimal assumptions and market-level data limitations.

Second, we build on and contribute to a recent literature on identification of non-game-theoretic adverse-selection models (D'Haultfoeuille and Février (2011), D'Haultfoeuille and Février (2015), Torgovitsky (2015), Hedblom et al. (2022), Kang and Silveira (2021), and Cotton et al. (2023)).[4] Here, point identification of model primitives typically hinges on the crucial exclusion restriction of rank stability. We study the question of point and set identification when this exclusion restriction does not hold. In doing so, we provide a theoretical foundation for formalizing rank-stability violations: we show how they arise given a second dimension of unobserved agent heterogeneity which alters responsiveness to incentive changes, conditional on a fixed baseline demand.

Third, we contribute to the demand-estimation literature in IO, pioneered by McFadden (1974), Berry (1994), Berry et al. (1995), and Nevo (2003), and recently surveyed in (Berry and Haile, 2021, BH). Our work represents a significant departure from much of the literature surveyed by BH, which focuses on questions of how *market-level competition* affects consumer welfare, given estimated *extensive-margin* substitution patterns.[5] By contrast, in our paper we are primarily

[2]Hedblom et al. (2022) and Cotton et al. (2023) study labor-supply settings with multiple dimensions of worker heterogeneity, but they require a combination of panel data *and* incentive variation to achieve point identification. We focus on more common settings where only cross-sectional incentive variation is available to the econometrician. In the IO context, this may stem from explicit price randomization, while D'Haultfœuille and Février (2020) and Kang and Silveira (2021) provide examples of plausibly exogenous incentive variation in observational data.

[3]Multi-product monopoly settings (e.g., Armstrong (1996) and Luo et al. (2011)) are beyond the current scope.

[4]The ideas in this literature are also related to an identification strategy proposed by Guerre et al. (2009).

[5]A smaller strand of the discrete-choice demand literature allows for mixed discrete-continuous consumer decisions. See Dubin and McFadden (1984), Hendel (1999), Dubé (2004), McManus (2007), and Richards and Bonnet (2016).

interested in how a *single firm* should optimally set volumetric prices. This focus implies *intensive-margin demand* responses of the firm's consumer base as the primary concern. Despite these differences, our results provide a case study in thinking about key issues raised in the broader discussion of demand estimation. BH point out that randomized pricing may be neither necessary nor sufficient to identify the counterfactual market shares of interest, which motivate the need to estimate (extensive-margin) demand in the first place.[6] This paper explores similar questions within the world of intensive-margin demand: our model further illustrates why randomized price variation may not suffice when consumers make continuous choices either. However, our partial identification framework and empirical application demonstrate that random price variation can nonetheless facilitate informative counterfactual bounds and meaningfully guide policy decisions.

Fourth, our robust bounds approach is related to a previous literature that studies the (partial) identifying power of weak assumptions derived from economic theory, often involving incomplete or partially specified models of decision making. Examples include Haile and Tamer (2003) and Hortaçsu and McAdams (2010) in the context of auctions, Heckman et al. (1997) in the context of program evaluation, Freyberger and Larsen (2021) in the context of bargaining, and Kang and Vasserman (2022), in the context of consumer demand.

The rest of this paper is organized as follows. Section 2 lays out our basic model of intensive-margin demand and discusses how we achieve point identification or bound identification under various circumstances. Section 3 discusses data, an estimation strategy and tests of the basic modelling assumptions. Section 4 presents robust design of optimal nonlinear pricing.

2. Model and Identification

Our basic one-dimensional demand system is an adverse-selection framework that captures salient features of the producer-consumer relationship: unobservably heterogeneous customers (agents) make volumetric consumption choices, q, subject to prevailing prices, while the supplier (principal) wishes to maximize profits. For now, individuals are heterogeneous along a single dimension (we relax this assumption in Section 2.4.1), with $\theta_i > 0$ indexing consumer i's idiosyncratic demand intensity: higher θ individuals are willing to consume more q at any given price. We assume that demand types follow an absolutely continuous distribution $\Theta \sim F_\theta \in \mathcal{C}^2$ with density f_θ that is strictly positive on compact support $[\underline{\theta}, \bar{\theta}]$. For our baseline model specification we assume multiplicatively separable (MS) agent utility $U_i(q;\theta_i) = \theta_i u(q)$, where $u(q)$ satisfies standard regularity conditions, being strictly increasing $u'(q) > 0$, $\forall q \in \mathbb{R}_+$, and concave $u''(q) < 0$, with $\lim_{q\to\infty} u'(q) = 0$. For convenience, we impose scale and shift normalizations as well: $u(0) = 0$ and $u'(0) = 1$.

The firm's pricing schedule for q units of consumption is $P(q)$. Specializing to the case of linear pricing, we have that $P(q) = pq$, so each consumer chooses q by solving $\max_q \{\theta u(q) - pq\}$. The

[6]Chan and Hamilton (2006) make a similar point in the context of medical RCTs.

first-order condition (FOC) to this optimization problem is given by

$$\theta u'(q) = p \quad \Rightarrow \quad q^*(p;\theta) = (u')^{-1}\left(\frac{p}{\theta}\right). \tag{1}$$

Under our assumptions, a unique solution $q^*(p;\theta)$ exists, and is strictly decreasing in price p and strictly increasing in type θ.[7] Given pricing schedule P, we denote the distribution of consumer demand as $G(q|P)$, though we will generally drop the conditioning on the price schedule unless needed for clarity. Note that G may have a mass point at $q=0$, but above that mass point (if it exists) G is absolutely continuous with a well-defined density $g(q)$ on a compact support $[\underline{q},\overline{q}]$.

By way of characterizing optimal choices, if consumer i is observed to demand more than consumer j at price p, $q^*(p;\theta_i)>q^*(p;\theta_j)$, then at any other price p', $q^*(p';\theta_i)>q^*(p';\theta_j)$ will be true as well.[8] This follows because $\theta_i=\frac{p}{u'(q^*(p;\theta_i))} > \frac{p}{u'(q^*(p;\theta_j))}=\theta_j$ and $q^*(p;\theta)$ is increasing in θ, holding price fixed. This implies what we refer to as the *Rank Stability* condition (henceforth, RS):

Condition 1 (**RS**)**.** An individual whose demand is at the r^{th} quantile of G under price $p > 0$ will also have demand at the r^{th} quantile under any other price $p' > 0$. That is, for all $(\theta,p,p')\in[\underline{\theta},\bar{\theta}]\times\mathbb{R}^2_{++}$ we have $q^*(p;\theta)=G^{-1}(r|p) \Rightarrow q^*(p';\theta)=G^{-1}(r|p')$ for each $r\in[0,1]$.

This condition will be a central focus of discussion throughout the paper, but first we consider other basic implications of the model. Another restriction that separable utility (or any quasilinear utility model) places on the data is the "*Law of Demand*" (henceforth, LoD):

Assumption 1. Each individual's optimal choice $q^*(p;\theta)$ is non-increasing in price, or $p<p' \Rightarrow q^*(p;\theta)\leq q^*(p';\theta)$ for every $\theta\in[\underline{\theta},\bar{\theta}]$.

An empirically testable implication of this basic assumption built into the model is first-order stochastic dominance of the conditional demand distributions: $p<p' \implies G(q|p)\leq G(q|p')$.

D'Haultfoeuille and Février (2015), Torgovitsky (2015), and D'Haultfœuille and Février (2020) (henceforth, DF/T) provide a thorough treatment of point identification within this baseline setup. Briefly though, with an exogenous price change we can identify the CDFs $G(q|p)$ and $G(q|p')$, and therefore the corresponding quantile functions as well. Rank stability then implies that the r^{th} quantile treatment effect is the individual treatment effect for the consumer whose type θ is at the r^{th} quantile of F_θ. The separable utility model also provides a within-consumer mapping from consumption level $q^*(p;\theta)$ under price p to counterfactual consumption $q^*(p';\theta)$ under price p':

$$q^*(p;\theta_i) = (u')^{-1}\left(\frac{p}{p'}u'(q^*(p';\theta_i))\right) \quad \Leftrightarrow \quad \frac{u'(G^{-1}(r|p))}{u'(G^{-1}(r|p'))} = \frac{p}{p'}, \;\; r\in[0,1]. \tag{2}$$

[7]More precisely, demand is strictly decreasing in p (increasing in θ) in the sense that for any p such that if $q^*(p;\theta)>0$, then $q^*(p';\theta)<q^*(p;\theta)$ whenever $p<p'$ ($q^*(p;\theta') < q^*(p;\theta)$ whenever $\theta' < \theta$).

[8]Since marginal utility is bounded, each type θ has a finite choke price $\overline{p}(\theta) = \theta u'(0)$, where they choose $q^*\left(\overline{\theta};\theta\right)=0$. Given our setup, this is the sole way for strict monotonicity to be violated, but for expositional simplicity we abstract from this detail until we prove our main results in Appendices A and B.

A remarkable fact implied by MS utility is that under relatively weak conditions, the RS condition plus equation (2) can pin down the values of the utility function $u(q)$ uniquely up to affine transformations for $q \in \mathcal{O}$ for some identified set of points $\mathcal{O}$. A cumbersome technicality is that $u(q)$ may only be partially identified for $q \notin \mathcal{O}$ (see DF/T for details). The sharp identified sets for $u(q)$ are quite informative though. Thus, for expositional simplicity we assume that the utility function u is known by the econometrician to belong to a set $\mathcal{U}$ satisfying the following:

Assumption 2. Fix $q_0 \geq 0$. Define a sequence of points $\mathcal{O} = \{q_k\}_{k=-\infty}^{\infty}$ recursively via the identity $q_{k+1} = G_{d'}^{-}(G_d(q_k))$.[9] The family of admissible utility functions $\mathcal{U}$ is such that if $u, v \in \mathcal{U}$ and $u(q_k) = v(q_k)$ for all $q_k \in \mathcal{O}$, then $u(q) = v(q)$ for all $q \in \mathbb{R}_+$.[10]

Given an identified $u(q)$, the consumer-level θ's are also identified (and hence, F_θ as well), since they can be obtained by inverting the first order condition (1).[11] Identification of structural counterfactuals in the basic model with an exogenous price change arises from three main restrictions on the data-generating process (DGP): (*i*) the RS condition, (*ii*) the LoD, and (*iii*) (via MS utility) the model provides a means of extrapolating from demand under observed prices to counterfactual demand under out-of-sample prices. The RS and LoD conditions buy the econometrician a lot of inferential power: if a firm can exogenously vary its price, it can nonparametrically identify all relevant parameters necessary for finding the optimal nonlinear price schedule. The LoD is testable, and with more than two prices, the extrapolation quality of MS utility may even be testable as well. However, the RS condition is arguably the most stringent and least testable model implication. Moving forward we will refer to the basic setup as the "rank-stable model."

This leaves an open question: how robust are model-based policy prescriptions to deviations from perfect rank stability within the underlying DGP? We begin our analysis by exploring RS violations generated by behavioral mistakes in subscription choices. We show that the researcher can restore point identification by explicitly modeling mistakes, and we propose a direct estimator of behavioral parameters, which is of independent interest. We then proceed to our main methodological contribution by exploring RS violations of *unknown form*; i.e., when the phenomena producing the violation are not well understood, or when requisite data for point identification are unavailable. We identify sharp bounds on the set of counterfactual demand CDFs consistent with observables and the LoD. We show that these bounds can be used to derive useful pricing prescriptions that are robust to the most extreme unobserved consumer behaviors not ruled out by the data and LoD.

[9] Here, $G_{d'}^{-}(r)$ is the *pseudo-inverse* of $G_{d'}$, that is $G_{d'}^{-}(r) = \inf\{q : G_{d'}(q) \geq r\}$.

[10] Since the set $\mathcal{O}$ typically includes values across the support of q, the family of utility functions $\mathcal{U}$ can be made quite flexible while still satisfying Assumption 2. E.g., $\mathcal{U}$ can be the set of all smoothing splines constrained to take on values $u(q_k) = u_k$ for $q_k \in \mathcal{O}$. Moreover, Assumption 2 can be relaxed under either of the following mild data augmentations: (*i*) $\frac{\partial G(q|p)}{\partial p}$ is identified or (*ii*) $G(q|p'')$ is identified for one additional price level $p'' \notin \{p, p'\}$.

[11] $F_\theta(\theta)$ is only bound identified up to the largest θ type consuming nothing under the lowest price.

2.1. **Structural Identification in An Explicit Model of RS Violations.** Experimental incentive variation has become a common tool among practitioners and researchers, but often, an impediment to useful inference from controlled randomization is apparent deviations from fully rational decision making on the part of experimental participants. Such deviations can include both under-reaction (e.g., Chetty et al. (2009)) and over-reaction (e.g., Arieley et al. (2003)). As motivation for this concern, we note that our raw data from a rideshare RCT—where some consumers were randomly selected to receive an offer to purchase a subscription (see Section 3)—show direct evidence of both over-reaction and under-reaction within the treatment group.

Now consider a randomized experiment with two treatment arms. A control group must pay the original price, p_0, for each unit of consumption while a treatment group has the option to buy a subscription contract where they pay $\$S$ upfront for a discount rate of d_0, so that their new per-unit price will be $p_0(1-d_0)$. Therefore, treated consumers must first *choose* whether to buy the subscription. A perfectly rational consumer of type θ should make this decision in the following way. Utility with a subscription is $\theta u(q^*(p_0(1-d_0);\theta)) - p_0(1-d_0)q^*(p_0(1-d_0);\theta) - S$, while without a subscription it is $\theta u(q^*(p_0;\theta)) - p_0 q^*(p_0;\theta)$. Thus, consumer type θ should subscribe if and only if $\theta u(q^*(p_0(1-d_0);\theta)) - p_0(1-d_0)q^*(p_0(1-d_0);\theta) - S \geq \theta u(q^*(p_0;\theta)) - p_0 q^*(p_0;\theta)$.

Rational consumers should buy the subscription only if surplus is greater than or equal to the upfront fee S. Surplus from discount d_0 is at least $(p_0 d_0) \times q^*(p_0;\theta)$ (i.e., the price drop times demand under default price) and no larger than $(p_0 d_0) \times q^*(p_0(1-d_0);\theta)$ (i.e., the price drop times counterfactual discounted demand), from which follows two testable predictions: rational individuals should subscribe if $q^*(p_0;\theta) > \frac{S}{p_0 \times d_0}$, and should not subscribe if $q^*(p_0(1-d_0);\theta) < \frac{S}{p_0 \times d_0}$. Thus, we should never see fully rational consumers decline a subscription and choose $q > \frac{S}{p_0 \times d_0}$, nor should we see them subscribe and choose $q \leq \frac{S}{p_0 \times d_0}$. In contrast, among treated consumers who did not subscribe, 29% would have unambiguously saved money, while 6% of subscribers had low enough demand in the subsequent month that they did not recover the subscription fee.

We now augment the model to allow for mistakes in the discrete choice of subscribing, while maintaining our assumption that consumers are able to choose the optimal quantity $q^*(p;\theta)$ without error.[12] We introduce a *salience* parameter, $\rho \in (0,1]$, representing the probability that a given (treated) consumer receives the relevant messaging and is cognizant of the subscription offer. We also allow for some fraction $\delta \in [0,1)$ of consumers to be *eager* and (conditional on salience) always purchase the subscription even without weighing costs and benefits. Finally, we allow for

[12] Our behavioral model assumes that i makes mistakes *only at the beginning* of the period, when i is unsure whether a subsequent demand shock ε_i will justify subscribing. Afterward, i's demand shock is revealed in a timely enough manner that i neither under-consumes nor over-consumes. This is reasonable if a period (e.g., a month) is composed of K sub-periods (e.g., days) within which i knows the sub-period shock ε_{ik} before choosing sub-period demand, but i cannot fully anticipate future sub-period shocks.

consumers to be imperfect at forecasting their own demand intensity over the period following their subscription decision. Thus, rather than evaluating their uptake decision based on their true type θ, they instead base subscription choice on a noisy estimate, $\hat{\theta}$, of their demand type. Moving forward, it will be convenient to re-parameterize unobserved types as an individual's demand under the baseline price, letting $\eta(\theta) \equiv q^*(p_0;\theta)$ denote their type. Because consumption is strictly increasing in θ (absent forecast errors), this is a one-to-one mapping, which allows us to measure errors in the same (directly observable) units as consumption. We assume that consumers mis-estimate their type η as $\hat{\eta} = \eta + \varepsilon$ where $\varepsilon \sim H_\varepsilon(\varepsilon)$ represents forecast error and satisfies the following:[13]

Assumption 3. H_ε is an absolutely continuous, unimodal distribution, with well-defined density $h_\varepsilon(\varepsilon)$ that is strictly positive on a connected, compact support.

To see why this behavioral framework produces RS violations, consider two consumers with types $\theta_i < \theta_j$, such that $q^*(p_0;\theta_i) < q^*(p_0;\theta_j) < q^*(p_0(1-d_0);\theta_i) < S/(p_0 \times d_0)$. If both consumers were fully rational, neither would buy the subscription. However, consumer i could mistakenly purchase the subscription while j does not. In absence of a subscription offer, θ_i would consume less than θ_j, but due to behavioral mistakes θ_i will now consume more than θ_j, which violates the RS condition.

From above we can see that the *ex-post* discounted surplus change from subscribing is increasing in θ, and there exists some cutoff θ^* where the change exactly offsets the fee S. Let $q_s^* = q^*(p_0;\theta^*)$ denote the analogous cutoff in consumption space. We now define an uptake indicator, $\upsilon_i \equiv \mathbb{1}[i \textit{ subscribes}]$, and an *uptake function* as $\Upsilon(q) \equiv \Pr[\textit{subscribe}|q^*(p_0;\theta) = q] = \rho\delta + \rho(1-\delta)H_\varepsilon(q - q_s^*)$.

We now show how to identify the utility function $u(q)$ (and hence q_s^*) and the uptake function $\Upsilon(q)$, by applying some basic ideas from the literature on the LATE interpretation of instrumental variables (e.g., Imbens and Angrist (1994), Imbens and Rubin (1997)). Recall that identification hinges on knowing the distributions of demand with and without an exogenous discount. While the researcher can observe a demand CDF for the sub-population within the treatment group that receives the discount, $G(q|p_0(1-d_0), \upsilon = 1)$, the complementary demand distribution, $G(q|p_0, \upsilon = 1)$, is not directly observable because the set of would-be uptakers in the control group is not known.

However, if we think of treatment status—i.e., whether a consumer is offered a subscription plan—as an instrumental variable for uptaker status (and hence, who gets a discount), we can identify the baseline demand distribution $G(q|p_0, \upsilon = 1)$ and the uptake function $\Upsilon(q)$. Let $G_c(q|p_0)$ denote the demand CDF for consumers in control, and let $G_t(q)$ be the unconditional demand CDF for all consumers in treatment, regardless of their subscription choice, and note that each of these is directly observable to the researcher. Letting τ denote the proportion of uptakers in the treatment

[13] Note that our model represents a static, one-time decision process of whether to purchase a subscription. In a dynamic model where the consumer makes this decision month after month, one could interpret $\hat{\eta}$ as fixed, long-run average demand intensity, with $\eta_t = \hat{\eta} - \varepsilon_t$ representing transitory demand intensity for month t.

group, we have $G_t(q) = \tau G(q|p_0(1-d_0), \upsilon=1) + (1-\tau)G(q|p_0, \upsilon=0)$. Similarly, we can decompose the control CDF as $G_c(q|p_0) = \tau G(q|p_0, \upsilon=1) + (1-\tau)G(q|p_0, \upsilon=0)$. Combining these two identifies and rearranging allows us to express $G(q|p_0, \upsilon=1)$ in terms of observable quantities

$$G(q|p_0, \upsilon=1) = G(q|p_0(1-d_0), \upsilon=1) - \frac{G_t(q) - G_c(q|p_0)}{\tau}. \tag{3}$$

This relationship tells us that for uptakers (referred to as "compliers" in the usual LATE parlance), we can identify both of the counterfactual CDFs, $G(q|p_0(1-d_0), \upsilon=1)$ and $G(q|p_0, \upsilon=1)$. The term $\frac{G_t(q)-G_c(q|p_0)}{\tau}$ in equation (3) characterizes quantile-specific consumption shifts among would-be uptakes within the control-group, had they received discount d_0. With $G(q|p_0, \upsilon=1)$ and $G(q|p_0(1-d_0), \upsilon=1)$ known, we have all requisite information to apply the DF/T identification approach to pin down the utility function $u(q)$ and demand types θ.

Finally (ignoring mass points at $q=0$), note that $G_c(q|p_0)$ and $G(q|p_0, \upsilon=1)$ have densities $g_c(q|p_0)$ and $g(q|p_0, \upsilon=1)$, and therefore $\Upsilon(q)$ is nonparametrically identified by Bayes' rule:

$$\Upsilon(q) = \frac{g(q|p_0, \upsilon=1)\tau}{g_c(q|p_0)}. \tag{4}$$

This pins down the behavioral parameters ρ, δ, and H_ε. First, note that $\lim_{q\to\infty}\Upsilon(q) = \rho$ and $\lim_{q\to 0}\Upsilon(q) = \rho\delta$.[14] Finally, with ρ and δ known, we can use the definition of $\Upsilon(q)$ above to nonparametrically identify $H_\varepsilon(q)$ using the relationship $H_\varepsilon(q) = \frac{\Upsilon(q+q_s^*)-\rho\delta}{\rho-\rho\delta}$.

2.2. **Robust Inference Under RS Violations of Unspecified Form.** We now note that explicitly controlling for RS violations is not always possible. For example, suppose that *Ian* typically purchases more volume than *John* from *Firm A*, but in a month where *A* offered a discount, *John* made more purchases than *Ian* from *A*. This could arise if, for example, *John* has low brand loyalty toward *Firm A* and thus views *Firm B*'s product as a closer substitute than *Ian*, who is more loyal to *Firm A*. As a result, sale prices at *Firm A* are more effective at switching *John*'s purchasing behavior away from its rival, *Firm B*. The inferential problem stems from the fact that *Firm A* knows more about its own internal prices and sales than it will ever know about its rival, making it unclear from *A*'s perspective how to interpret *John*'s apparent change in consumption.

Note that RS requires the same ordering of purchasing behavior by *Ian* and *John* in both the default (p_0) and discount ($p_0(1$–$d_0)$) states. However, the presence of a competitor and idiosyncratic brand loyalty may drive heterogeneous price sensitivity. Moreover, *Firm A* would not have the requisite internal observables to explicitly model cross-firm substitution by its consumers. More broadly, *any* phenomenon causing idiosyncratic price sensitivity by consumers with similar demands under default pricing—e.g., heterogeneity in budget constraints, income effects, or complementary with other goods—could render point identification impossible given available data.

[14]In general this may lead us to over-estimate δ if $H_\varepsilon(-q_s^*) > 0$. However, one can test for this problem: unimodality of H_ε implies that if $\Upsilon(q)$ is flat within a neighborhood of $q=0$ then significant upward bias in δ is unlikely.

This discussion highlights the fact that the RS condition required by the basic identification strategy indeed rules out some economically plausible behavior, which, notably, is not even precluded by controlled, experimental randomization in pricing. Rather than attempting to formalize all possible RS violations, we adopt the approach of deriving robust bounds on counterfactual demand (and in turn, on firm profits) projected by the model in the presence of RS violations of unknown form. The sale-price example described above is an intuitive way of fixing ideas, and is inspired by the prior demand estimation literature which focuses on consumer substitution patterns; however, our robust bounds approach does not hinge on this particular interpretation of RS violations.

As a precursor, a comment on the focus of our robust inference method will be helpful. Typical structural approaches to partial identification (e.g., Haile and Tamer (2003), Hortaçsu and McAdams (2010), and Freyberger and Larsen (2021)) often focus on bounding structural primitives like θ and F_θ. However, in our case a prominent source of the partial identification problem is the fact that the effective consumer base may shift between the control (p_0) and treatment ($p_0(1-d_0)$) states due to unobserved substitution between the firm and its rivals. If this is so, then the notion of pinning down a type distribution F_θ for a single, stable consumer base becomes problematic. For that reason, our approach focuses instead on bounding counterfactual demand distributions under alternate pricing levels.

To begin, we focus on the case of a single price change, though in Section 2.5 we show how richer price variation can be used for bound refinements. For each consumer, there are quantities (Q_{ci}, Q_{di}) such that consumer i would choose Q_{ci} if given the control price p_0 and Q_{di} if given some discounted price $p_0(1-d)$. We use the common convention of denoting random variables by upper-case letters, while realizations of random variables are denoted by lower-case. For the present purpose, we abstract away from the consumer's choice of whether or not to purchase a subscription plan, and we simply assume that a subset of consumers are exposed to an exogenous price drop from p_0 to $p_0(1-d_0)$. Let $G_{cd}(q_c, q_d)$ denote the joint distribution function of (Q_c, Q_d), with marginal distributions $G_c(q) = G(q|p_0)$ and $G_d(q) = G(q|p_0(1-d))$.

In this section we obtain bounds on conditional probabilities of the form $\Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$. These bounds are important because profitability of a subscription offering will be largely determined by the sub-population of most-likely subscribers for whom $Q_c \geq \frac{S}{p_0 \times d}$, which we refer to as *strong uptakers*. Formally, *strong uptakers* are the set $SU(p_0, S, d) = \left\{\text{consumer } i : q_{ci} \geq \frac{S}{p_0 \times d}\right\}$. Under *any* model of the underlying preferences, consumers will find it advantageous to subscribe if $Q_c \geq \frac{S}{p_0 \times d}$, while consumers for whom $Q_d < \frac{S}{p_0 \times d}$ will find subscription unambiguously disadvantageous. The question of whether or not *intermediate consumers* for whom $Q_c < \frac{S}{p_0 \times d} \leq Q_d$ should subscribe still depends on the particular model of preferences. In the interest of robustness we therefore focus on strong uptakers, since the baseline predictions of the RS model—which determine whether intermediate consumers subscribe—may not hold.

To facilitate robust inference on counterfactual profits, our empirical objects of study are *strong uptaker distributions* (SUDs), defined as $\Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$. These represent the conditional counterfactual demand CDFs of most-likely subscribers for a given (S,d) pair, and our goal is to construct a set of sharp SUD bounds, $\overline{\mathcal{B}}_{d_0}(q;S,d) \leq \Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right] \leq \underline{\mathcal{B}}_{d_0}(q;S,d)$. Since SUDs are CDFs, the object $\overline{\mathcal{B}}_{d_0}$ (also a CDF) is point-wise below the SUD, but is an *upper bound* in the sense that it stochastically dominates the true SUD. Similarly, $\underline{\mathcal{B}}_{d_0}$ (also a CDF) is point-wise above the true SUD, meaning the former is dominated by the latter.

We first construct a rank-stable mapping $\overline{\mathcal{Q}}_{d_0}$ from a given consumer's baseline consumption level into the space of counterfactual consumption. We show that this mapping defines a limiting DGP whose SUDs constitute the sharp upper bound $\overline{\mathcal{B}}_{d_0}(q;S,d)$ on SUDs for other DGPs which do not necessarily adhere to rank stability but are consistent with the observable demand distributions under control and treatment. We then construct a similar mapping $\underline{\mathcal{Q}}_{d_0}$ that characterizes maximal RS violations that still respect the data and the LoD. We show that this mapping defines a DGP whose SUDs constitute the sharp lower bound $\underline{\mathcal{B}}_{d_0}(q;S,d)$. The bound subscripts denote their dependence on observed demand CDFs (G_c, G_{d_0}) under prices p_0 and $p_0(1-d_0)$, respectively.

As before, let the CDF of Q_c be G_c and the CDF of Q_d be G_d, while G_{d_0} represents the observed demand CDF under the actual discount d_0 defining treatment within the RCT. Henceforth, we refer to this as the *in-sample* discount, to distinguish it from other *out-of-sample* (i.e., experimentally untested) discounts $d \neq d_0$. Given our current focus on identification rather than on estimation, this is somewhat of an abuse of terminology, but it is useful in distinguishing inferences directly based on d_0, from those that can be made somewhat less directly based on alternate $d \neq d_0$.[15] Additionally, we define the quantile functions $G_c^{-1}(r) = \inf\{q : G_c(q) \geq r\}$ and $G_d^{-1}(r) = \inf\{q : G_d(q) \geq r\}$, and note that these may represent the inverse CDFs, if they exist, or the quasi-inverses otherwise.

2.3. **Construction of the Upper Bound.** While $\overline{\mathcal{B}}_{d_0}(q;S,d)$ lives in probability space, the functional value of $\overline{\mathcal{Q}}_{d_0}$ represents an upper bound on consumption under counterfactual price $p_0(1-d)$, if baseline consumption (under p_0) is q. More formally, we define a (stochastic) mapping

$$\overline{\mathcal{Q}}_{d_0}(q;d,v) \equiv G_d^{-1}\left(a(q) + b(q)v\right), \quad a(q) \equiv \lim_{q' \to q^-} G_c(q'), \quad \text{and} \quad b(q) \equiv G_c(q) - \lim_{q' \to q^-} G_c(q'), \tag{5}$$

where v is a realization of a random variable $V \sim Uniform(0,1)$ that is independent of (Q_c, Q_d). In equation (5), $a(q)$ and $b(q)$ are to deal with possible mass points in the control CDF, G_c; V plays the role of a "tie-breaking" rule; $a(q)$ is the mass of all consumers with baseline demand below q; and $b(q)$ is the size of the probability mass at q. Given this definition, we can now also define

$$\overline{\mathcal{B}}_{d_0}(q;S,d) \equiv \Pr\left[\overline{\mathcal{Q}}_{d_0}(Q_c;d,v) \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]. \tag{6}$$

[15]This abuse is less awkward if one considers the word "sample" as denoting a dataset with infinite observations.

The formal proof is fairly technical, but we show that the mapping $\overline{\mathcal{Q}}_{d_0}(Q_c; d, v) \sim G_d(q)$ (see Appendix A.1), meaning that it constitutes a data-generating process that must be consistent with observables. Moreover, we show that $\overline{\mathcal{B}}_{d_0}$ is an upper bound on the SUDs $\Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$ for each (S, d) pair. This, combined with the previous fact, implies that it is the *sharp upper bound*, since $\overline{\mathcal{Q}}_{d_0}(q; d, v)$ is an admissible DGP that attains the upper bound $\overline{\mathcal{B}}_{d_0}(q; S, d)$.

For some intuition, consider the hypothetical dataset depicted in Panel (A) of Figure 1. Characterizing the SUD upper bound reduces to the question of, what is the maximal fraction of the aggregate shift from demand CDF G_c to discounted demand CDF G_{d_0} that was accounted for by high-demand consumers responding to the discount by increasing their purchase volume? For the special case where G_c has no mass points, we get the simpler expression $\overline{\mathcal{Q}}_{d_0}(q; d, v) = \overline{\mathcal{Q}}_{d_0}(q; d) = G_d^{-1}(G_c(q))$, whose geometric interpretation provides an answer: at most, *all* of the aggregate shift represents rank-stable price responses by high-demand consumers. The distribution of $\overline{\mathcal{Q}}_{d_0}(Q_c; d)$ is simply the distribution of Q_d, since $\overline{\mathcal{Q}}_{d_0}(Q_c; d)$ is chosen to match the quantiles of G_c to the corresponding quantiles of G_d. Intuitively then, the upper-bound DGP $\overline{\mathcal{Q}}_{d_0}$ is one where strong uptakers with high baseline demand are *maximally* price responsive, *none* of the aggregate demand shift is due to unobserved substitution, and therefore no "rank-jumping" happens at all. Interestingly, this first result, stated formally in Proposition 1, indicates that the naive and potentially mis-specified RS model still serves as a relevant empirical benchmark.

Before stating our first result we adopt an assumption on the underlying model of aggregate demand. In what follows there will be a distinction between *in-sample bounds*—e.g., the SUD upper bound $\overline{\mathcal{B}}_{d_0}(q; S, d_0)$ given the in-sample discount d_0—and *out-of-sample bounds*—e.g., the SUD upper bound $\overline{\mathcal{B}}_{d_0}(q; S, d)$ for an experimentally untested discount level $d \neq d_0$. The assumption has no impact on in-sample bounds, but provides structure for deriving out-of-sample bounds.

Assumption 4. The set of aggregate demand CDFs G_d arising from out-of-sample discounts d is such that if the reduced-form aggregate distributions of demand (G_c, G_{d_0}) are known for in-sample prices, $(p_0, p_0 \times (1 - d_0))$, then aggregate demand G_d is also known for out-of-sample discounts $d \neq d_0$ and is given by a (known) function $G_d(q) = G_d^{oos}(q; G_c, G_{d_0})$.

Remark 1. Assumption 4 is stated in this way in order to highlight the flexibility of our partial identification approach. It covers various methods for extrapolation of aggregate counterfactual demand, ranging from the basic theoretic demand system above, to more general demand models, and also allows for atheoretical, reduced-form approaches to extrapolation. Options include,

(1) The observed and counterfactual demand CDFs are consistent with the MS utility model: Assumption 2 is satisfied and there exists a (u, F_θ) pair, such that $G_d(q) = \int_\Theta \mathbb{1}[q^*(p_0(1 - d); \theta) \leq q] dF_\theta(\theta)$ for all $(d, q) \in (-\infty, 1) \times \mathbb{R}_+$, where $q^*(\cdot; \theta)$ is defined in equation (1).

(2) Observed and counterfactual demand CDFs are consistent with a φ-separable utility model:

$$U(q;\theta,\varphi)=\begin{cases}\int_0^q (u(t)+\theta)^{\frac{1}{1-\varphi}}\,\mathrm{d}t, & \varphi<1\\ \int_0^q \exp(u(t)+\theta)\,\mathrm{d}t, & \varphi=1.\end{cases} \tag{7}$$

Moreover, Assumption 2 holds and there exists a (φ,u,F_θ) triple, with known φ, such that $G_d(q)=\int_\Theta \mathbb{1}[q^*(p_0(1-d);\theta,\varphi)\le q]dF_\theta(\theta)$ for all $(d,q)\in(-\infty,1)\times\mathbb{R}_+$.[16]

(3) (Chernozhukov et al. (2010)) Aggregate quantile shifts are linear in the discount d, up to rearrangement: $G_d(q)=\Pr\left[Q_c+\frac{d\times(G_{d_0}^{-1}(G_c(Q_c))-Q_c)}{d_0}\le q\right]$.

In our main empirical application we use extrapolation option (1), but in Appendix A.3, we discuss the more general family of utility functions in option (2), and in Appendix F, we show that our counterfactuals are insensitive to the exact form of extrapolation. Option (2) illustrates how the econometrician may decide that some model other than the MS paradigm may be more appropriate for projecting counterfactual demand shifts. For example, consider the additively-separable (AS) utility form, $U(q;\theta)=u(q)+\theta q$ (e.g., see Maskin and Riley (1984) and Laffont and Tirole (1986)). The main difference between MS utility and AS utility is their implications for counterfactual demand extrapolation. MS implies that the price *elasticity* of demand, $\frac{\partial q^*(p;\theta)}{\partial p}\frac{p}{q^*(p;\theta)}$ depends on p and θ only through their implied level of demand $q^*(p;\theta)$, while additive separability implies the same property for the price *derivative* of demand, $\frac{\partial q^*(p;\theta)}{\partial p}$. The φ-separable utility family nests both MS ($\varphi=1$) and AS ($\varphi=0$) as special cases.

In options (1) and (2), counterfactual extrapolations stem from an explicit, rank-stable, structural model of demand. At first glance this may seem philosophically awkward, given that the purpose of our paper is to provide inferential tools in scenarios where validity of the RS assumption is in question. As we have seen above, $\overline{\mathcal{Q}}_{d_0}$ happens to be precisely the rank-stable DGP, so for the purposes of bounding counterfactual demand from above (Proposition 1), options (1) and (2) are less problematic as aids for out-of-sample inference. Options (1) and (2) of Remark 1 should not be interpreted as assuming directly that individual-level demand within the DGP is rank stable; rather, they simply state that the researcher is confident in a particular utility specification for producing sensible counterfactual projections for *aggregate* demand shifts.

A researcher who is skeptical of (1) and (2) may opt for a more agnostic method of extrapolation, with (3) being an example of one such approach. In Remark 3 below we discuss how all three options augment identifying power specifically by imposing local smoothness conditions on aggregate demand shifts under a price change from p_0 to $p_0(1-d_0)$. The intuition behind the third option, a special case of methods proposed by Chernozhukov et al. (2010), is that we can use the two available pieces of information, $G_c^{-1}(r)$ and $G_{d_0}^{-1}(r)$, $r\in[0,1]$, to approximate quantile treatment

[16]See Sun (2023b) and Appendix F for a complete discussion on identification under the φ-separable utility model. If the data include multiple exogenous price shifts, then φ need not be known ex-ante to satisfy Assumption 4.

effects of a discount as being locally linear in the discount d.[17] However, the linear approximation may lead to technical problems where, for discounts outside the range $(0, d_0)$, quantile projections cross. Chernozhukov et al. (2010) solve this problem by a method they refer to as *rearrangement*, which guarantees that quantile projections are monotone in r. We will return to the discussion of different approaches to out-of-sample extrapolation after we derive the lower-bound DGP $\underline{\mathcal{Q}}_{d_0}$.

Assumption 4 and Remark 1 are relatively mild for two reasons. First, additional RCT variation can be used to probe its validity: given multiple experimental discounts $(d_0, d_1, \ldots)$ one may test whether utility specifications like (1) or (2) produce realistic out-of-sample projections. Second, it is important to keep in mind that Assumption 4 deals with extrapolation of *reduced-form, aggregate demand*—i.e., both subscribers and non-subscribers—rather than with the primary structural primitive, counterfactual demand among subscribers only. Put another way, Assumption 4 directly concerns the *marginal distributions* of demand under alternate prices p_0 and $p_0(1-d)$, while our primary objects of study are bounds on the *copula* between consumer i's demand Q_{ci} and Q_{di} under those two prices. We now state our first result, relegating a formal proof to Appendix A.1. □

Proposition 1. *Under Assumptions 1 and 4, if $G_c(q)$ and $G_{d_0}(q)$ are known and are discontinuous at countably many mass points, then $\overline{\mathcal{B}}_{d_0}(q; S, d)$ constitutes an identified, sharp upper bound on the strong uptaker distributions. That is, for any (potentially out-of-sample) discount $d \in (0,1)$, and subscription fee $S \geq 0$, we have $\overline{\mathcal{B}}_{d_0}(q; S, d) \equiv \Pr\left[\overline{\mathcal{Q}}_{d_0}(Q_c; d, v) \leq q | Q_c \geq \frac{S}{p_0 \times d}\right] \leq \Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$, where the function $\overline{\mathcal{Q}}_{d_0}(q; d, v)$ is defined in (5).*

Remark 2. ***Out-of-Sample Inference*** The in-sample discount d_0 together with observed (G_c, G_{d_0}) and equations (5) and (6) are sufficient to pin down $\overline{\mathcal{Q}}_{d_0}(q; d_0, v)$ and $\overline{\mathcal{B}}_{d_0}(q; S, d_0)$. Then, Assumption 4 allows us to project aggregate demand $G_d^{oos}(q; G_c, G_{d_0})$ under arbitrary discount $d \neq d_0$—e.g., via options (1), (2), or (3) of Remark 1—where the parameter inputs denote dependence of this projection on the observables (G_c, G_{d_0}). The counterfactual CDFs G_c and G_d^{oos} can then be plugged back into equations (5) and (6) to get $\overline{\mathcal{Q}}_{d_0}(q; d, v)$ and $\overline{\mathcal{B}}_{d_0}(q; S, d)$, for arbitrary (S, d) pairs. □

2.4. **Construction of the Lower Bound.** We now construct an analogous mapping $\underline{\mathcal{Q}}_{d_0}(Q_c; d, v)$ that represents a lower-bound DGP for counterfactual consumption levels of likely subscribers consistent with the LoD and the data (G_c, G_{d_0}). To fix ideas on the sorts of phenomena that may produce empirically relevant violations of rank stability, we begin by generalizing the basic model from Section 2. This provides one especially salient (though not comprehensive) interpretation of RS violations arising from multi-dimensional agent heterogeneity within demand estimation.

2.4.1. *An Explicit Model of Unobserved Rank Stability Violations.* A central motivation behind our robustness exercise is the lack of information on consumer substitution patterns within a typical

[17]For values of $d \in (0, d_0)$, option (3) is exactly equivalent to linear quantile shifts. Moreover, with an additional discount, say d_0 and d_0', one could use Chernozhukov et al. (2010) to define a locally quadratic variant of (3) instead.

firm's internal data. The extended model we present here highlights how unobserved heterogeneity in brand loyalty may lead to apparent RS violations being more prominent among consumers with low (internal) demand at baseline pricing.

We generalize utility to be a function of total consumption across the default firm, L, and its competitor, firm C.[18] Each consumer i combines consumption from both sources into a composite good, "transportation," according to a constant elasticity of substitution "production" function $T_i = T_i(q_L, q_C) = \left[\alpha_i^{\frac{1}{\eta}} q_L^{\frac{\eta-1}{\eta}} + (1-\alpha_i)^{\frac{1}{\eta}} q_C^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}$. Here, α_i indexes consumer i's brand preference for the default firm L. A value of $\alpha_i = 0.5$ means i is perfectly indifferent between interacting with firm L versus firm C (when prices for both are the same), while $\alpha_i = 0$ ($\alpha_i = 1$) means that i would be unwilling to purchase from firm L (firm C) at any price. The parameter η is an elasticity of substitution between the default firm's services and services of its competitor which, for expositional simplicity, we take to be fixed in the population. As $\eta \to \infty$ the two services produced by L and C become perfect substitutes (holding brand loyalty fixed at $\alpha_i = 0.5$). Thus, α represents intrinsic utility from doing business with firm L specifically, while η determines how similar are the goods/services produced by each firm when divorced from their respective brand names.

Utility from total consumption is multiplicatively separable $U(T_i;\theta_i) = \theta_i u(T_i(q_L, q_C))$, where θ_i still indexes i's idiosyncratic demand intensity. Faced with prices (p_L, p_C), i chooses (q_L, q_C) to solve $\max_{(q_L,q_C)\in\mathbb{R}^2_+} \theta_i u(T_i(q_L, q_C)) - p_L q_L - p_C q_C$. Note that this more general formulation nests the basic model from Section 2 as a special case (when $\alpha_i = 1\ \forall i$).

To characterize this demand system, we take advantage of the fact that T_i is homogeneous of degree 1, meaning that if we wish to scale up composite consumption by some factor $\zeta \times T_i(q_L, q_C)$, we can accomplish this simply by scaling up the two inputs by the same factor $\zeta \times (q_L, q_C)$. As a result, we can solve the consumer's optimization problem in two steps. First, she solves an expenditure minimization to determine the optimal shares of the default firm and the competitor, per unit of composite consumption $T_i(q_L, q_C)$. Next, the consumer solves an outer utility maximization problem to determine the level of total transportation consumption.

$$\max_{t\in\mathbb{R}+} \theta_i u(t) - p_T t, \quad \textit{subject to} \quad p_T = \min_{(q_L,q_C)\in\mathbb{R}^2_+} p_L q_L + p_C q_C, \quad \textit{subject to} \quad T_i(q_L, q_C) = 1.$$

Then, letting q_L^h, q_C^h be the solutions to the cost minimization problem above (Hicksian demand), we have that individual i chooses $q_L = q_L^h t^*$ and $q_C = q_C^h t^*$ where t^* solves the outer utility maximization problem. Standard results on constant elasticity of substitution functions imply that

$$q_L^h = \alpha_i \times \left[\alpha_i + (1-\alpha_i)\left(\frac{p_L}{p_C}\right)^{\eta-1}\right]^{\frac{\eta}{1-\eta}}, \quad q_C^h = (1-\alpha_i) \times \left[(1-\alpha_i) + \alpha_i\left(\frac{p_C}{p_L}\right)^{\eta-1}\right]^{\frac{\eta}{1-\eta}}, \text{ and } p_T = \left[\alpha_i p_L^{1-\eta} + (1-\alpha_i) p_C^{1-\eta}\right]^{\frac{1}{1-\eta}}. \quad (8)$$

[18]The competitor firm C can be thought of as encompassing all of the consumer's outside options for substitutable goods/services, from competing private rideshare firms, to public transit, to walking instead.

For simplicity of discussion, we normalize prices $p_L = p_C = 1$, which in turn implies $p_T = 1$. If the price ratio is one this implies $q_L^h = \alpha_i$ in (8), meaning α_i is simply individual i's share of consumption supplied by the default firm. For any level of default firm consumption q_L, there is a locus of (θ_i, α_i) pairs that rationalize it. To see why, fix q_L and note that for each $\alpha_i \in (0,1)$ there is some θ_i such that, given parameters (θ_i, α_i), i consumes exactly q_L units from the default firm. This requirement is defined by combining the identity $\alpha_i t_i^* = q_L$ with the FOC of the outer utility maximization, $t_i^* = (u')^{-1}(1/\theta_i)$, to get $\alpha_i \times (u')^{-1}\left(\frac{1}{\theta_i}\right) = q_L$. Since u' is strictly increasing, this defines a curve in (θ, α)-space of types consistent with a fixed optimal choice q_L.

We now derive some comparative statics around prices $p_L = p_C = 1$ for a fixed type (θ_i, α_i). Using the chain rule and product rule, we have $\frac{\partial q_L}{\partial p_L} = \frac{\partial q_L^h}{\partial p_L} t^* + \frac{\partial t^*}{\partial p_T}\frac{\partial p_T}{\partial p_L} q_L^h$, which simplifies to $\frac{\partial q_L}{\partial p_L} = -\eta q_L \frac{1-\alpha_i}{\alpha_i^2} + \alpha_i^2 / u''\left(\frac{q_L}{\alpha_i}\right)$ since $q_L^h = \alpha_i$. This expression characterizes the responsiveness of observed demand for the default firm's service as its own price p_L changes. Importantly, note that the derivative of demand with respect to p_L changes as q_L is kept constant but α_i varies. I.e., if consumrs i and j choose the same quantity q_L under default pricing, but $\alpha_i < \alpha_j$, then their responsiveness to a price change will be different, thus creating an apparent violation of rank stability within the internal data available to default firm L.

Moreover, the model also implies that when services of firms L and C are sufficiently substitutable, a consumer with *lower* brand loyalty to L will be *more* sensitive to changes in p_L. Once again, suppose i and j consume the same firm-L quantity q_L, but consumer i has less brand loyalty, or $\alpha_i < \alpha_j$. Subtracting i's response to an infinitesimal price change from j's response gives $-\eta q_L\left(\frac{1-\alpha_j}{\alpha_j^2} - \frac{1-\alpha_i}{\alpha_i^2}\right) + \left(\frac{\alpha_j^2}{u''(q_L/\alpha_j)} - \frac{\alpha_i^2}{u''(q_L/\alpha_i)}\right)$. The first term is strictly positive because $\alpha_i < \alpha_j$ while the second term has a generally ambiguous sign. However, provided that consumption choices are bounded from above, so $q_L < M < \infty$, and $\alpha > 0$ for all consumers, the first term will dominate as η gets large. Intuitively, this implies that as the elasticity of substitution, η, gets sufficiently large, *individuals with lower brand loyalty will tend to be more responsive to changes in p_L in terms of their purchases from firm L* (i.e., a larger negative own-price elasticity).[19]

While the brand-loyalty interpretation is motivated by the consumer demand context, the same basic principle applies in other adverse-selection/principal-agent models as well. In Section 5.1 we briefly discuss five related settings—procurement, regulation of externalities, optimal taxation,

[19]In the limiting case $\eta \to \infty$, where the two goods become perfect substitutes (modulo brand loyalty), consumers solve $\max_{(q_L, q_C)\in\mathbb{R}_+^2} \theta_i u(\alpha_i q_L + (1-\alpha_i) q_C) - p_L q_L - p_C q_C$. Now, the expenditure minimization problem is simply to choose $q_L^h = \frac{1}{\alpha_i}$ if $\frac{p_L}{p_C} < \frac{\alpha_i}{1-\alpha_i}$ and $q_L^h = 0$ if the inequality is strict in the opposite direction. Consider now a price change from p_L to $p_L - \varepsilon$ for some small $\varepsilon > 0$. Individuals for whom $q_L^h = \frac{1}{\alpha_i}$ even before the price change will only slightly change their consumption levels from $q_L = \frac{1}{\alpha_i}(u')^{-1}\left(\frac{p_L}{\alpha_i\theta_i}\right)$ to $q_L = \frac{1}{\alpha_i}(u')^{-1}\left(\frac{p_L - \varepsilon}{\alpha_i\theta_i}\right)$. On the other hand, for some consumers we will have $\frac{p_L - \varepsilon}{p_C} < \frac{\alpha_i}{1-\alpha_i}$ but $\frac{p_L}{p_C} > \frac{\alpha_i}{1-\alpha_i}$. For them, q_L^h will change from 0 (under the original p_L) to $\frac{1}{\alpha_i}$ after the discount, and their consumption from firm L will "rank jump" from from $q_L = 0$ to $q_L = \frac{1}{\alpha_i}(u')^{-1}\left(\frac{p_L - \varepsilon}{\alpha_i\theta_i}\right)$.

labor contracts, and insurance/healthcare demand—including relevant applications of nonlinear pricing, and how exogenous price variation no longer suffices for point identification under multi-dimensional agent heterogeneity. These examples highlight how ideas analogous to those presented here can be applied to a wide variety of policy-relevant contexts in empirical market design.

2.4.2. *Formal Derivation of the SUD Lower Bound.* Having formalized a theoretical foundation for unobserved RS violations, we return our focus to derivation of the lower-bound DGP $\underline{\mathcal{Q}}_{d_0}$. That is, despite firm L lacking internal data to empirically model failures of RS, we derive a sharp lower bound $\underline{\mathcal{B}}_{d_0}(q;S,d) = \Pr\left[\underline{\mathcal{Q}}_{d_0}(Q_c;d) \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$ on counterfactual demand CDFs (SUDs) for strong uptakers. While the example of RS violations driven by unobserved substitution patterns and brand-loyalty heterogeneity is salient and empirically relevant, the lower bound we derive does not hinge on any specific underlying model of unobserved RS violations. Other plausible phenomena producing RS violations could include heterogeneous income effects and/or budget-constraint heterogeneity. Our method also allows for the fully-specified underlying model of RS violations to include multiple channels driving unpredictable heterogeneity in price sensitivity.

Our purpose here is to derive a sharp lower bound on the set of latent DGPs that are consistent with the LoD and the observables. To do so, we must characterize *maximal* RS violations that respect this *a priori* information. The LoD and the data (G_c, G_{d_0}) impose considerable discipline on the maximal mass of rank-jumpers and on plausible magnitudes of their rank-jumping behaviors. Recall that the rank-stable DGP is the least upper bound on counterfactual strong-uptaker demand under discount d. Suppose that a market designer naively optimizes profits, $\pi(S,d)$, from a single offer (S^*, d^*) under the RS assumption. This optimum fee structure must balance three things for each subscriber θ type in order for it to be profit-improving: it must (i) offer consumers a viable path to savings, or $(p_0 d^*) q^*(p_0(1-d^*);\theta) - S^* \geq 0$, and while (ii) discounted revenues are lower for each subscriber by $(p_0 d^*) q^*(p_0;\theta)$, their (iii) projected increase in demand volume $q^*(p_0(1-d^*);\theta) - q^*(p_0;\theta)$ is large enough so that the change in total revenues to the firm is positive:

$$S^* - (p_0 d^*) q^*(p_0;\theta) + p_0(1-d^*)\left[q^*(p_0(1-d^*);\theta) - q^*(p_0;\theta)\right] > 0. \tag{9}$$

Ensuring this is true requires an accurate forecast of demand responsiveness to discount d^*.

With that in mind, one can think of the lower-bound DGP $\underline{\mathcal{Q}}_{d_0}$ as being *maximally adversarial* in the sense of minimizing the naive market-designer's profits, $\pi(S^*, d^*)$, from subscription offer (S^*, d^*). For the hypothetical dataset depicted in Panel (A) of Figure 1, the naively presumed DGP $\overline{\mathcal{Q}}_{d_0}$ holds that the entirety of the shift from G_c (solid line) to G_d (dashed line) represents a rank-stable demand increase, which implies maximal demand responses $[q^*(p_0(1-d^*);\theta) - q^*(p_0;\theta)] = G_d^{-1}(F_\theta(\theta)) - G_c^{-1}(F_\theta(\theta))$ by high-θ subscribers (i.e., $q^*(p_0;\theta) \geq \frac{S^*}{p_0 \times d^*}$). In contrast, the lower-bound $\underline{\mathcal{Q}}_{d_0}$ asks, what is the *smallest* price response by uptaker θ types that cannot be ruled out by the LoD and data? This is equivalent to *minimizing* the firm's total revenue change from its subscription program (left-hand side of (9)). In the most extreme case, strong uptakers subscribe but then do not increase purchase volume at all, and the left-hand side of (9) becomes $S^* - (p_0 d^*) q^*(p_0;\theta) \leq 0$, denoting a zero-sum transfer from firm to consumer. This intuitive adversarial property of $\underline{\mathcal{Q}}_{d_0}$ will be discussed at length below, but first we formalize our primary objects of interest, $\underline{\mathcal{Q}}_{d_0}$ and $\underline{\mathcal{B}}_{d_0}$.

For simplicity, we temporarily assume that G_c and G_d are absolutely continuous and the difference $G_c(q) - G_d(q)$ is unimodal (i.e., quasi-concave). We relax both of these assumptions in our proofs in Online Appendix A.2, but to avoid tedious complications of exposition we limit discussion here to the simpler case.[20] The function $\underline{\mathcal{Q}}_{d_0}(q; d, v) : Q_c \to Q_d$ maps baseline consumption levels under p_0 into minimal plausible consumption levels under discount d, given the LoD and data (G_c, G_{d_0}). Since $G_c(q) - G_d(q)$ is unimodal, it is weakly increasing below its smallest maximizer, q^*_{min}, and weakly decreasing above its largest maximizer, q^*_{max}. Let q_{max} be the largest value for which G_c and G_d disagree and define $\bar{q}_{d_0}(q) = \inf\{q' \in [q^*_{max}, q_{max}] : G_c(q') - G_d(q') = G_c(q) - G_d(q)\}$.[21] In other words, $\bar{q}_{d_0}$ maps relatively low baseline consumption levels $q \leq q^*_{min}$ into discounted consumption levels $q' \geq q^*_{max}$ such that the condition described above is satisfied. Because G_c and G_d are continuous, the inf is attained, so $G_c(q) - G_d(q) = G_c(\bar{q}_{d_0}(q)) - G_d(\bar{q}_{d_0}(q))$. Letting V denote a $Uniform(0,1)$ random variable that is independent of (Q_c, Q_d), $\underline{\mathcal{B}}_{d_0}$ and $\underline{\mathcal{Q}}_{d_0}$ are defined by

$$\underline{\mathcal{B}}_{d_0}(q; S, d) \equiv \Pr\left[\underline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q \;\middle|\; Q_c \geq \tfrac{S}{p_0 \times d}\right]; \text{ and } \quad \underline{\mathcal{Q}}_{d_0}(q; d, v) = \begin{cases} \bar{q}_{d_0}(q) & q \leq q^*_{min}, v \leq \frac{g_c(q) - g_d(q)}{g_c(q)} \\ q & \text{otherwise.} \end{cases} \tag{10}$$

Graphical intuition for how $\bar{q}_{d_0}$ and $\underline{q}_{d_0}$ are constructed (and in turn, $\underline{\mathcal{Q}}_{d_0}$ as well) can be found in Panels (B) and (C) of Figure 1. Recall that $\underline{\mathcal{Q}}_{d_0}$ represents a maximally adversarial DGP from the perspective of a market designer who optimized a subscription offer (S^*, d^*) assuming the rank-stable DGP $\overline{\mathcal{Q}}_{d_0}$. Intuitively, to achieve the maximally adversarial property we start at the low end of the demand spectrum (i.e., low values of baseline demand q' under default price p_0) and we assume the largest possible mass of those consumers are rank-jumpers, which is depicted in Panel (B) of the figure. Moreover, we also assume that these low-demand rank-jumpers do so in the worst way from the naive market-designer's perspective, meaning that they rank jump by the largest possible margin that would not violate the shape of the treatment demand CDF G_{d_0}, which can be seen in Panel (C) of the figure. The CDF difference determines the maximal rank-jumping margin because, for large values of discounted demand q, it represents the excess mass of consumers who purchased at least q under discounted price $p_0 \times (1 - d_0)$, relative to the mass who purchased at least q under default price p_0. Then we continuously apply this adversarial re-allocation of low-demand

[20]The case of a unimodal CDF difference appears to be the most empirically relevant case, both in our setting and in a number of similar settings such as D'Haultfœuille and Février (2020) and Sun (2023a).

[21]If G_c, G_d have unbounded support and $G_c(q) > G_d(q)$ for all q, let $q_{max} = \infty$.

FIGURE 1. Proof Intuition For Case 1

(A) Demand CDFs (B) Demand PDFs

(C) Demand CDF Differences

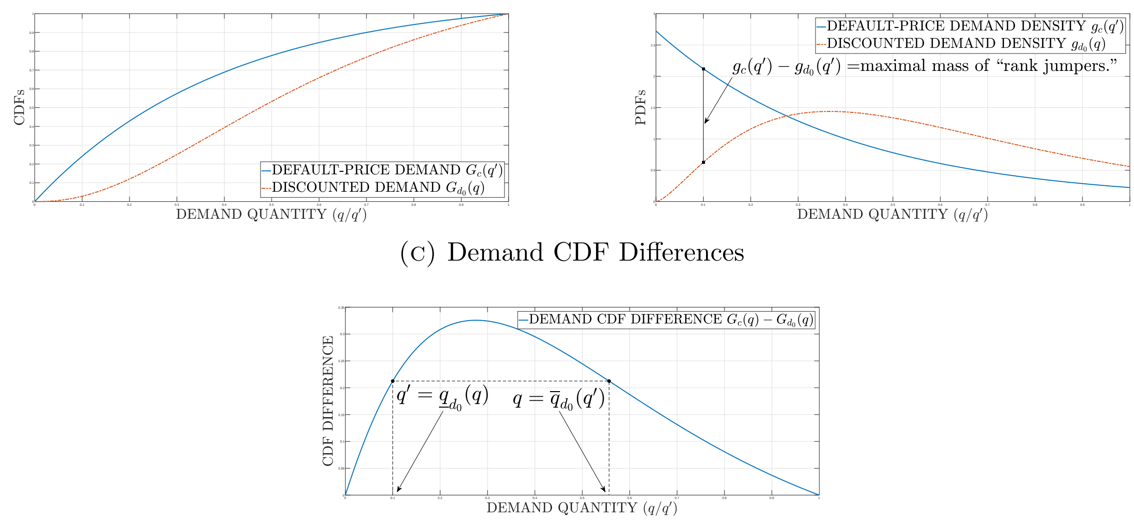

Notes: Panel (A) plots hypothetical demand CDFs G_c and G_{d_0}. Panel (B) plots the corresponding demand PDFs. Panel (C) plots the difference in the demand CDFs, $G_c - G_{d_0}$.

consumers for all values of q' between 0 and q^*_{min}, the minimal argmax of the CDF difference.[22] This ensures that the upper tail of the treatment demand CDF G_{d_0} is maximally populated by non-subscribers, and therefore minimally populated by subscribers. Conversely, this procedure also ensures that demand responses by strong uptakers collectively accounted for a minimal fraction of the upper-tail shift from G_c to G_{d_0}, which maximally exposes the naive market-designer to merely transferring money to them, with little or no compensating improvement in sales volume.

A remaining question concerns whether $\underline{Q}_{d_0}$ can be fully adversarial without accounting for the behavior of intermediate consumers, some of whom would subscribe. For another way of understanding the adversarial bound, note that the following must be true of any admissible DGP, whenever $S > 0$ and the discount is less than markup (i.e., $p_0(1-d) - c > 0$):

Fact 1. *Anybody who consumes $Q_c \geq \frac{S}{p_0 \times d}$ will always be a subscriber.*

Fact 2. *Anybody who consumes $Q_c < \frac{S}{p_0 \times d}$ will always bring profits to the firm if they subscribe, relative to the counterfactual of not subscribing.*

Fact 1 is true under the LoD because such an individual would surely save money by subscribing. To see why Fact 2 is true, note that the change in profits to the firm when a consumer buys the subscription is $\Delta\pi \equiv Q_d(p_0(1-d)-c) + S - Q_c(p_0 - c) = (Q_d - Q_c)(p_0(1-d)-c) + S - Q_c \times p_0 \times d$. Since the discount is less than the markup, $\frac{p_0 - c}{p_0}$, and $Q_d \geq Q_c$ by the LoD, the first term on the right-hand side is non-negative. Therefore, $S - Q_c \times p_0 \times d > 0$ implies $\Delta\pi > 0$, which is true whenever

[22] Note that in the example depicted in Figure 1, the CDF difference has a unique maximum corresponding to the single crossing point of the PDFs, so $q^*_{min} = q^*_{max}$. Otherwise, $g_c(q') = g_{d_0}(q')$ for $q' \in [q^*_{min}, q^*_{max}]$ would be true.

$Q_c < \frac{S}{p_0 \times d}$. From these two facts it follows that the firm could do no worse than if (i) everyone with $Q_c \geq \frac{S}{p_0 \times d}$ under default pricing (i.e., *all* strong uptakers) buys the subscription and then consumes minimal additional volume, while (ii) nobody with $Q_c < \frac{S}{p_0 \times d}$ under default pricing ever subscribes.

For example, this extreme scenario would occur if each consumer's demand curve had a discontinuity at price $p_0(1-d)$ and took the form $Q(p; Q_c, Q_d) = Q_c \mathbb{1}\{p > p_0(1-d)\} + Q_d \mathbb{1}\{p \leq p_0(1-d)\}$, where the joint distribution of (Q_c, Q_d) follows the DGP described by $\underline{\mathcal{Q}}_{d_0}$. Such an individual would not subscribe even though they are highly sensitive to the price change, because their consumer surplus from getting the discount is 0. Extreme demand patterns like this—where aggregate demand is virtually constant for prices $p \in (p_0(1-d), p_0]$, and increases precipitously at price $p_0(1-d)$—cannot be ruled out based solely on the information in (G_c, G_d).

Proposition 2. *Under Assumptions 1 and 4, if $G_c(q)$ and $G_{d_0}(q)$ are known and are discontinuous at countably many mass points, then $\underline{\mathcal{B}}_{d_0}(q; S, d)$ constitutes an identified upper bound on the strong uptaker distributions. That is, for any (potentially out-of-sample) discount $d \in (0,1)$, and subscription fee $S \geq 0$, we have $\Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right] \leq \underline{\mathcal{B}}_{d_0}(q; S, d)$, where the function $\underline{\mathcal{Q}}_{d_0}(q; d, v)$ is defined in (10) when (G_c, G_d) are absolutely continuous and $G_c(q) - G_d(q)$ is unimodal, and defined in (25) in Appendix B otherwise. Moreover, $\underline{\mathcal{B}}_{d_0}$ is sharp in the following ways:*

- (i) *The in-sample bound $\underline{\mathcal{B}}_{d_0}(q; S, d_0)$ (which does not depend on Assumption 4) is uniformly sharp with respect to q in the sense that $\forall q$ the bounding DGP is consistent with the LoD (i.e., $\underline{\mathcal{Q}}_{d_0}(q; d_0, V) \geq q$) and the shapes of (G_c, G_{d_0}) (i.e., $\Pr[\underline{\mathcal{Q}}_{d_0}(Q_c; d_0, V) < q] = G_{d_0}(q)$).*
- (ii) *$\underline{\mathcal{B}}_{d_0}(q; S, d)$ is uniformly sharp with respect to q and extrapolated aggregate demand in the sense that, $\forall q$ the bounding GDP is consistent with (G_c, G_d^{oos}) (i.e., $\Pr[\underline{\mathcal{Q}}_{d_0}(Q_c; d, V) < q] = G_d^{oos}(q)$) and satisfies the LoD with respect to prices p_0 and $p_0(1-d)$ (i.e., $\underline{\mathcal{Q}}_{d_0}(q; d, V) \geq q$).*

We relegate a technical proof of Proposition 2 to Appendix A.2. Our usage of the term *uniformly sharp* follows Firpo and Ridder (2019) and Molinari (2020). The notion of sharpness we have emphasized—uniform sharpness in q space, holding discount d fixed—is motivated by the interests of an empirical market designer using our methodology for robust nonlinear pricing: in Section 4.3 of our empirical application, a fixed discount level d^* (the optimum implied by the rank-stable upper bound) serves as a focal point for lower-bound computation. The following corollary provides simpler intuition behind the SUD lower bound for subscription offers that are not overly generous:[23]

Corollary 1. *Let "small" fee-to-discount ratios $\frac{S}{p_0 \times d}$ be those that are strictly less than the infimum of the right-most modal region of $G_c(q)$–$G_d(q)$ (e.g., if $G_c(q)$–$G_d(q)$ is unimodal then "small" means $\frac{S}{p_0 \times d} < q^*_{min}$). Then, under the assumptions of Proposition 2, for a subscription offer (S, d) that is not overly generous (i.e., where the fee-to-discount ratio is not small), the sharp SUD lower*

[23]As we will show in the empirical application, robust policies tend to not be *overly generous*.

bound $\underline{\mathcal{B}}_{d_0}$ is the same as the conditional control CDF of demand, given strong uptaker status, or $\underline{\mathcal{B}}_{d_0}(q;S,d)=\Pr\left[Q_c\leq q\middle|Q_c\geq \frac{S}{p_0\times d}\right]=\frac{G_c(q)-G_c(S/(p_0\times d))}{G_c(S/(p_0\times d))}=G_c(q|Q_c\geq S/(p_0\times d))$.

This corollary implies that for a wide range of potential (S,d) pairs, the SUD lower bound derived in Proposition 2 depends only on the in-sample demand distribution G_c and is therefore insensitive to the form of extrapolation used to calculate G_d, as we discuss further in the following remark.

Remark 3. ***Out-of-Sample Inference*** First note that the in-sample discount d_0 together with observables (G_c,G_{d_0}) and equation (10) (or equation (25)) suffice to pin down in-sample bounds, $\underline{\mathcal{Q}}_{d_0}(q;d_0,v)$ and $\underline{\mathcal{B}}_{d_0}(q;S,d_0)$, under d_0 coupled with arbitrary subscription fee S. A researcher reticent to impose *any* additional model structure (e.g., Assumption 4 and Remark 1) aside from the LoD, can apply Corollary 1 and Proposition 4 (Appendix A.3), giving a less informative out-of-sample lower bound which we label as $\underline{\mathcal{B}}^{LoD}_{d_0}(q;S,d)$. Little inferential power beyond the in-sample bounds is possible in that case, due to lack of a way to project out-of-sample reduced-form aggregate demand G_d. To see why, consider first a discount $d<d_0$ that is less generous than the in-sample discount. Here, without imposing any restrictions on the smoothness of aggregate demand shifts between prices p_0 and $p_0(1-d_0)$, we cannot rule out DGPs where demand is locally satiated, being arbitrarily close to G_c for any price $p\in(p_0,p_0(1-d_0))$. Thus, $\underline{\mathcal{B}}^{LoD}_{d_0}(q;S,d)=G_c(q|Q_c\geq\frac{S}{p_0\times d})$. Similarly, without imposing any local smoothness conditions on aggregate demand shifts, for $d>d_0$ that is more generous we cannot rule out DGPs where an arbitrarily large fraction of consumers have virtually satiated demand under price $p_0(1-d_0)$. Thus, $\underline{\mathcal{B}}^{LoD}_{d_0}(q;S,d)=\underline{\mathcal{B}}_{d_0}(q;S,d_0)$.

On the other hand, Assumption 4 and Remark 1 provide the researcher with concrete means of imposing local smoothness conditions on aggregate demand shifts between prices p_0 and $p_0(1-d)$. A researcher comfortable leaning on such restrictions may still harbor concerns about whether options (1) and (2) of Remark 1 induce excessive mis-specification bias in aggregate demand projections G^{oos}_d. In that case, the reduced-form linear quantile shifts option (3), or Corollary 1 are still useful. The former provides a relatively model-agnostic first-order approximation to aggregate demand quantiles for general (S,d) pairs, while the latter provides a precise characterization of structural counterfactual demand quantile bounds for (S,d) pairs that are not overly generous.

In either case, the lower bound under out-of-sample discounts $d\neq d_0$ are as in Remark 2. First, by Assumption 4, we can use in-sample demand distributions (G_c,G_{d_0}) to produce reduced-form aggregate demand projections $G^{oos}_d(q;G_c,G_{d_0})$ by some preferred method. The CDFs G_c and G^{oos}_d can then be plugged back into equations (10) to get $\underline{\mathcal{Q}}_{d_0}(q;d,v)$ and $\underline{\mathcal{B}}_{d_0}(q;S,d)$. In our empirical application, we explore inferences based on all three cases of Remark 1 and Corollary 1, and find that our main empirical conclusions are insensitive to one's choice among these options. □

Remark 4. ***Optimal Experimental Design*** The first paragraph of Remark 3 also suggests a key insight about optimal experimental randomization to maximize ex-post inferential power. If the experimentalist's prior indicates that the optimal discount lay in some neighborhood $d^*\in(\underline{d},\bar{d})$, then inferential power is maximized by choosing a *less generous* experimental discount $d_0 \le \underline{d}$. An alternative view of this insight is that an experimentalist who is comfortable with extrapolation but wishes to lean more on data than on model structure should also choose $d_0 \le \underline{d}$. □

Remark 5. ***Theory-Free Bound*** Without imposing the LoD, one can derive an alternate, theory-free lower bound $\widetilde{\underline{\mathcal{B}}}_{d_0}$ for the SUDs. The analogous DGP, $\widetilde{\underline{\mathcal{Q}}}_{d_0}$ would be one that swapped ranks: the individual whose consumption Q_c under default pricing was in the r^{th} quantile of G_c also has discounted consumption Q_d in the $(1-r)^{th}$ quantile of G_d. This DGP $\widetilde{\underline{\mathcal{Q}}}_{d_0}$ in general need not respect the LoD. In contrast, we have constructed our main lower bound $\underline{\mathcal{Q}}_{d_0}$ under the principle that we wish to maximizethe degree of RS violations *subject to* the LoD. The comparison between the resulting demand CDF bounds, $\widetilde{\underline{\mathcal{B}}}_{d_0}$ and $\underline{\mathcal{B}}_{d_0}$, helps to illuminate how structure from basic economic theory delivers useful inference. We explore this idea empirically in Section 4.3. □

2.5. **Lower-Bound Refinement Using Richer Experimental Variation.** We now explore how multiple discount treatment arms can enhance inferential power. While richer price variation can be used in various ways—e.g., testing and improving the extrapolation method—we focus here only on improving the informativeness of the lower bound, holding some extrapolation method fixed.[24] Consider now a set of additional discounts $\mathcal{D}_d = \{d' < d'' < \cdots < d^{(K-1)} < d\}$ that are *less generous* than d—i.e., $d^{(i)} \in (0, d)$, $i = 1, \ldots, K-1$—and the corresponding aggregate demand data $\mathcal{G}_{\mathcal{D}_d} = \{G_c, G_{d'}, \ldots, G_{d^{(K-1)}}, G_d\}$. To simplify discussion, we will consider the focal discount d^*, on which bounding inference is to be done, as belonging to the set $\mathcal{D}_d$, or in other words, $d^* = d$.[25]

Recall from our discussion in Section 2.4.2 that the DGP $\underline{\mathcal{Q}}_{d_0}$ satisfies the adversarial property, provided that all individuals consuming $Q_c < \frac{S}{p_0 \times d}$ under $\underline{\mathcal{Q}}_{d_0}$ *do not* buy the subscription. However, some adversarial behavior could be ruled out if we were able to observe the distribution of demand at a price discount between 0 and d, thus effecting a shift in the strong uptaker cutoff. The information in $\mathcal{G}_{\mathcal{D}_d}$ allows us to derive a set of *refined strong uptakers* by including some individuals who were previously considered part of the set of intermediate consumers (see Section 2.2 and Fact 2).

Consider a hypothetical scenario where we observe counterfactual demand $Q_{d'}$ at a single intermediate discount level $d > d' > 0$. When $Q_{d'} = q_{d'}$, then the demand curve passing through the price-quantity pair $(p_0(1-d'), q_{d'})$ which delivers the least consumer surplus relative to d is the demand curve $Q(p) = q_{d'} \mathbb{1}\{p \le p_0(1-d')\}$; in other words, where demand is constant for prices

[24]A complete treatment of extrapolation methods and their refinements is beyond the scope of this paper, but the interested reader is directed to Sun (2023b) for an in-depth analysis on this topic.

[25]To fix ideas, suppose the econometrician has a dataset $\{G_c, G_{0.10}, G_{0.15}\}$ with demand under default pricing and two treatment arms, but wishes to derive a refined lower bound relative to discount $d = 0.25$. Then our discussion assumes the econometrician will use some extrapolation method to first project aggregate demand $G^{oos}_{0.25}$, and then derive a refined bound based on $\mathcal{D} = \{0.10, 0.15, 0.25\}$ and $\mathcal{G}_{\mathcal{D}} = \{G_c, G_{0.10}, G_{0.15}, G^{oos}_{0.25}\}$. However, a researcher unwilling to lean on model structure that facilitates extrapolation would simply take $G^{oos}_{0.25}(q) = G_{0.15}(q)$, $\forall q$ instead.

below $p_0(1{-}d')$, and where demand jumps from 0 to $q_{d'}$ precisely at price $p_0(1{-}d')$. In this case, the change in consumer surplus relative to the base discount of d is $p_0(d{-}d')q_{d'}$. Thus, a consumer with counterfactual demand $Q_{d'} \geq \frac{S}{p_0(d-d')}$ at the intermediate price point must obtain surplus at least as large as $p_0(d-d') \times \frac{S}{p_0(d-d')} = S$, and thus, these consumers will unambiguously wish to subscribe to (S, d) under any DGP consistent with the LoD. Thus, we can define the refined strong uptaker set given $\mathcal{D}_d = \{d', d\}$ as $RSU(p_0, S, \mathcal{D}_d) = SU(p_0, S, d) \cup \left\{\text{consumer } n : q_{d'n} \geq \frac{S}{p_0(d-d')}\right\}$.

In reality, we do not directly observe an individual's complete $(q_c, q_{d'}, q_d)$ triple. However, recall that the adversarial DGP for an (S, d) pair is one where refined strong uptakers consume minimal incremental q after subscribing, subject to consistency with observables $\mathcal{G}_{\mathcal{D}_d} = \{G_c, G_{d'}, G_d\}$ and the LoD. Characterizing this scenario is equivalent to maximizing a RSU consumer's Q_c while minimizing that same consumer's Q_d. First, if we re-define $\widetilde{p}_0 = p_0(1-d')$ as the default price, and $\widetilde{d} = \frac{d-d'}{1-d'}$, then Corollary 1 applies to the price change from $\widetilde{p}_0$ to $\widetilde{p}_0(1-\widetilde{d})$—or equivalently, the price change from $p_0(1-d')$ to $p_0(1-d)$—which indicates that $Q_d = Q_{d'}$ is the DGP with minimal price responsiveness for RSUs. On the other hand, for the price change from p_0 to the intermediate discounted price $p_0(1-d')$, we can apply the logic of Proposition 1 and conclude that the rank-stable DGP where $Q_c = G_c^{-1}\left[G_{d'}(Q_{d'})\right]$ is true maximizes Q_c, subject to consistency with observables and the LoD. Finally, note that whenever $G_{d'}(q) < G_c(q)\ \forall q$ we have $Q_d = Q_{d'} > Q_c$.

This last finding indicates that we have achieved a tightening of the lower-bound DGP $\underline{\mathcal{Q}}_{d_0}(q; d, v)$, relative to the case with only a single price change, where we could not rule out $Q_d = Q_c$. More formally, we can define our *refined lower-bound DGP* as[26]

$$\underline{\mathcal{Q}}^R_{\mathcal{D}_d}(q; S, d, v) \equiv \underline{\mathcal{Q}}_{d_0}(q; d, v)\mathbb{1}\left\{G_{d'}^{-1}[G_c(q)] < \frac{S}{p_0(d-d')}\right\} + G_{d'}^{-1}[G_c(q)]\mathbb{1}\left\{G_{d'}^{-1}[G_c(q)] \geq \frac{S}{p_0(d-d')}\right\}. \tag{11}$$

The first term encompasses previous inferences about consumer types when only a single price change d was available: for demand levels $Q_{d'}$ between 0 and $G_{d'}^{-1}\left[G_c(q)\right]$ the adversarial demand projection does not change. The second term characterizes how the additional intermediate demand distribution $G_{d'}$ allows us to update the worst-case scenario. Specifically, it eliminates some degree of weak price responsiveness by requiring that RSUs—i.e., consumers with $q_{d'} \geq \frac{S}{p_0(d-d')}$ under discount d'—if they were given the more generous discount d, would have to increase counterfactual demand by at least a margin of $G_{d'}^{-1}\left[G_c(q_{d'})\right] - q_{d'}$ in order to be consistent with observables.

We can apply identical logic to any other price triple $(p_0, p_0(1-d^{(i)}), p_0(1-d))$ for $d^{(i)} \in \mathcal{D}_d = \{d', \ldots, d^{(K-1)}, d\}$, and show that the worst-case DGP consistent with $\{G_c, G_{d^{(i)}}, G_d\}$, and the LoD implies that all individuals with demand quantile rank $r > G_{d^{(i)}}\left(\frac{S}{p_0(d-d^{(i)})}\right)$ under discount $d^{(i)}$ must behave consistently with rank-stability when prices shift from p_0 to $p_0(1-d^{(i)})$. We can

[26]For simplicity, equation (11) assumes that G_c is absolutely continuous. Otherwise one can replace the "G_c" terms with $a(q) + b(q)v$, where a and b are defined in (5).

aggregate these individual worst-case DGPs into a single refined DGP, $\underline{\mathcal{Q}}^R_{\mathcal{D}_d}(q;S,d)$, consistent with *all* observables, $\mathcal{G}_{\mathcal{D}_d}=\{G_c,G_{d'},\ldots,G_{d^{(K-1)}},G_d\}$, and the LoD.

Whenever $d^{(i)}<d^{(j)}$, if consumers with baseline demand $Q_c=G_c^{-1}(r)$ behave in a rank-stable manner when prices shift from p_0 to $p_0(1-d^{(j)})$ in the worst-case DGP consistent with $\{G_c,G_{d^{(j)}},G_d\}$, then they must also behave in rank-stable fashion when prices shift from p_0 to $p_0(1-d^{(i)})$ in the worst-case DGP consistent with $\{G_c,G_{d^{(j)}},G_d\}\cup\{G_{d^{(i)}}\}$. Formally extending this argument, let $i^*(r)\equiv\max\left\{i\in\{1,\ldots,K\}:r>G_{d^{(i)}}\left(\frac{S}{p_0(d-d^{(i)})}\right)\right\}$, where, by convention, $\max\emptyset\equiv 0$. Thus, the refined lower-bound DGP and refined SUD bound are

$$\underline{\mathcal{Q}}^R_{\mathcal{D}_d}(q;S,d,v)=\begin{cases}\underline{\mathcal{Q}}_{d_0}(q;d,v) & i^*(r)=0\\ G^{-1}_{d^{(i^*)}}\left[G_c(q)\right] & i^*(r)\neq 0\end{cases}\quad\text{and}\quad\underline{\mathcal{B}}^R_{\mathcal{D}_d}(q;S,d)\equiv\Pr\left[\underline{\mathcal{Q}}^R_{\mathcal{D}_d}(Q_c;S,d,V)\leq q|RSU(p_0,S,\mathcal{D}_d)\right],\quad(12)$$

where $RSU(p_0,S,\mathcal{D}_d)=SU(p_0,S,d)\cup\bigcup_{i=1}^{K-1}\left\{\text{consumer }n:q_{d^{(i)}n}\geq\frac{S}{p_0(d-d^{(i)})}\right\}$. Intuitively, each of the intermediate price shifts, with their corresponding demand CDFs $G_{d^{(i)}}(q)$, $i=1,\ldots,K-1$, imposes some lower bound on the minimal price responsiveness of consumers at the r^{th} quantile under the control demand distribution $G_c(q)$. Equation (12) aggregates this information as the upper envelope of minimal price responsiveness, for each $r\in[0,1]$.

Proposition 3. *For $\mathcal{D}_d=\{d',\ldots,d^{(K-1)},d\}$, if $\mathcal{G}_{\mathcal{D}_d}=\{G_c,G_{d'},\ldots,G_{d^{(K-1)}},G_d\}$ are known and G_c is discontinuous at countably many points, then $\underline{\mathcal{B}}^R_{\mathcal{D}_d}(q;S,d)$ defined in (12) is an identified refined lower bound on RSU demand* $\Pr\left[Q_d\leq q|RSU(p_0,S,\mathcal{D}_d)\right]\leq\underline{\mathcal{B}}^R_{\mathcal{D}_d}(q;S,d)$.

In general, it cannot be established that $\underline{\mathcal{B}}^R_{\mathcal{D}_d}$ must always be below $\underline{\mathcal{B}}_{d_0}$ for all q because the former bounds worst-case behavior of a different set of consumers than the latter—namely, Refined Strong Uptakers rather than Strong Uptakers—making it difficult to rule out crossings of the two for low values of q. However, the above discussion does establish a general, concrete sense in which $\underline{\mathcal{B}}^R_{\mathcal{D}_d}$ refines or tightens the range of plausible inferences, relative to $\underline{\mathcal{B}}_{d_0}$: if $\pi^A(S,\mathcal{D}_d)$ denotes worst-case, adversarial profits from subscription offer (S,d) consistent with experimental discounts in the set $\mathcal{D}_d$, then refined lower-bound profits are weakly greater, or $\pi^A(S,d)\leq\pi^A(S,\mathcal{D}_d)$.

Remark 6. ***Refined Out-of-Sample Inference: a Single Observed Price Change*** Thus far, the discussion in this section has considered the additional discounts in $\mathcal{D}_d\backslash\{d_0\}$ and the additional CDFs in $\mathcal{G}_{\mathcal{D}_d}\backslash\{G_c,G_{d_0}\}$ as primarily coming from raw data, being the product of richer experimental variation. However, introducing Assumption 4 affords an alternative interpretation of Proposition 3: a researcher who has only $\{G_c,G_{d_0}\}$ in raw data, but is comfortable with imposing smoothness conditions on aggregate demand shifts, may choose each new discount $d^{(i)}\in\mathcal{D}_d\setminus\{d_0\}$ from the interval $(0,d)$—where d is the focal discount of interest (e.g., the naive optimal discount projection from the rank-stable upper bound)—and then use out-of-sample projections (e.g., via options (1), (2), or (3) of Remark 1) to populate the set $\mathcal{G}_{\mathcal{D}_d}=\{G^{oos}_{d'}(\cdot;G_c,G_{d_0}),G^{oos}_{d^{(K)}}(\cdot;G_c,G_{d_0})\}\cup\{G_c,G_{d_0}\}$.

Such an approach would provide a means of refining worst-case profit bounds based on more limited data and the researcher's *a priori* beliefs about smoothness of aggregate demand shifts, by ruling out locally satiated demand scenarios similar to those discussed in Remark 3. Combining this insight with the discussion in Remarks 3 and 5, our proposed method allows for different levels of reliance upon raw data and model structure to facilitate various levels of increasing inferential lower-bound precision, summarized as follows:

(*i*) ***(Theory-Free)*** $\widetilde{\underline{\mathcal{B}}}_{d_0}$ using only raw data $\{G_c, G_{d_0}\}$
(*ii*) ***(No-Extrapolation)*** $\underline{\mathcal{B}}^{LoD}_{d_0}$ using $\{G_c, G_{d_0}\}$, the LoD (Assumption 1), and Proposition 4
(*iii*) ***(Extrapolation-Light)*** $\underline{\mathcal{B}}_{d_0}$ using $\{G_c, G_{d_0}\}$, Assumptions 1 & 4, and Proposition 2
(*iv*) ***(Extrapolation-Full)*** $\underline{\mathcal{B}}^{R}_{d_0}$ using $\{G_c, G_{d_0}\}$, Assumptions 1 & 4, and Proposition 3. □

3. Data, Estimation Strategy, and Results

We execute an empirical case study using a rich internal dataset from Lyft, a popular rideshare platform in the United States. Like many firms that offer subscription programs—e.g., Club Car Wash, Hello Fresh, Chargepoint, Audible.com—rideshare platforms have ongoing relationships with customers whose demand fluctuates over time, and they collect a wealth of internal transaction data, often including some platform-imposed exogenous price variation. They also typically have little or no information on their customers' demand intensity for goods/services of rival firms.

We begin by analyzing data from a *subscriptions RCT* of the form described in Section 2.1: in each of two treatment arms, a random set of consumers were offered the option to buy a monthly subscription with a discount of $d_a \times 100\%$, $a = 1, 2$, for a month, in exchange for an upfront subscription fee of $\$S$.[27] The control group consisted of all other consumers within the same sample population who did not receive a subscription offer, and thus all made demand decisions over the ensuing month under default pricing. We use our explicit model of behavioral mistakes (Section 2.1) to allow for deviations from full rationality in uptake decisions, while estimating parameters of the baseline multiplicatively separable utility model.

We complement this analysis with data from a second pricing RCT that precludes considerations of behavioral mistakes in subscription decisions. A random subset of consumers (treatment group) received a fixed discount off all rides over a two-week period. The control group consisted of all other rideshare consumers who operated under default pricing. We refer to this as the *uniform-discount RCT*, since d was automatically applied to all treated individuals, thus eliminating concerns over salience, eagerness, or forecast error. We can fit the point-identified RS structural model to data from each of the three experiments, as a test of basic model assumptions like rank stability. Structural estimates here also set the stage for our robust market-design exercise in Section 4.

[27]The value of S was the same for both treatment arms, but to protect Lyft's internal data security we do not report the amount of S in this paper. See Section 3.1 for full discussion on reported consumption units.

3.1. **Reported Consumption Units.** Lyft rides are heterogeneous, so we need a measure of aggregated consumption that can be interpreted as q within the model. The most straightforward way is to define q as the cost of the ride in absence of discounts. For example, an individual in control who takes two rides at \$20 and \$15 is recorded as having $q=35$, while another who took the exact same rides, but had a 10% discount would pay $\$31.5=(1-0.1)\cdot(\$20+\$15)$ but would also be recorded as having *non-discount equivalent* (NDE) consumption $q=35$. This convention of using NDE as our measure of q is convenient for two reasons. First, it allows us to normalize the baseline market price to $p_0=1$. Second, two different origin-destination pairs at the same point in time may differ in miles travelled, while a single origin-destination pair may be considered two very different goods at different points in time, and hence be priced differently. Measuring q as total NDE allows for a convenient comparison across these scenarios. This approach is akin to assuming hedonic valuation of ride attributes (Rosen (1974)), and is similar in spirit to "bid homogenization" in the auctions literature (e.g., Haile et al. (2006), Athey and Haile (2007)).[28]

When we report our results in tables and/or figures, we add one more normalization for q to maintain Lyft's data confidentiality. Rather than reporting units of q directly, we divide by the 98$^{\text{th}}$ percentile (monthly) q observed in our data, denoted $\overline{q}$. Thus, all plots/tables involving q represent various levels of consumption as fractions of $\overline{q}$. Note that regardless of how q is normalized, percentage discounts are still represented by multiplying prices by d. Since this is a reporting issue only, for simplicity we maintain the convention that default price for 1 unit of consumption is $p_0=1$.

3.2. **Experiment Overview and Descriptive Statistics.** We first analyze the subscriptions RCT data (from early 2019), where Lyft offered a random subset (treatment group) of its users an opportunity to buy a discount lasting one month for an upfront fee while baseline rideshare pricing remained unchanged for a control group. In both treatment arms, the upfront fee was roughly 3% of $\overline{q}$. The offered discount among treated individuals was also randomized to be either $d_1=15\%$ off or $d_2=25\%$ off. For this experiment we have a dataset containing an indicator vector $\boldsymbol{t}_n$ for which treatment consumer n was assigned to (t_{0n}=1 for Control, t_{1n}=1 for 15% off, t_{2n}=1 for 25% off), an indicator v_n for n's subscription choice, and consumption level, q_n.

We complement our analyses of the subscriptions experiments with data from a second RCT in 2019 where Lyft gave automatic discounts of $d_0=10\%$ (for a duration of 2 weeks) to a randomly chosen set of consumers. Since this treatment gave a default, uniform discount of 10% to each consumer in treatment, the lack of up-front subscription decisions eliminates previous concerns over behavioral phenomena such as offer salience and/or eagerness. For each individual n in our

[28] In empirical auctions, bids are often regressed on auction covariates to "homogenize" them. The residual is interpreted as bidder-specific demand intensity, and the regression terms are hedonic utilities of auction covariates. In our case, ride covariates (proxied for by prices) implicitly play a similar role, while θ represents demand intensity.

Table 1. Summary Statistics for Subscription Uptake Behavior

	Treatment Arm: 15% Off			Treatment Arm: 25% Off		
Variable	Estimate	Std.Err.	95%CI	Estimate	Std.Err.	95%CI
Proportion Strong Uptakers	0.1901	*(0.0007)*	[0.1887,0.1915]	0.2974	*(0.0008)*	[0.2958,0.2990]
Revenue Share of Strong Uptakers	0.6986	*(0.0012)*	[0.6962,0.7009]	0.8433	*(0.0007)*	[0.8420,0.8447]
Proportion Uptakers	0.0091	*(0.0002)*	[0.0088,0.0094]	0.0156	*(0.0002)*	[0.0152,0.0160]
Proportion Saved Money | Uptaker	0.8278	*(0.0007)*	[0.8265,0.8291]	0.9297	*(0.0005)*	[0.9288,0.9306]

Notes: The third row shows the proportion of the treatment group who actually were uptakers. The fourth row shows the proportion of actual uptakers who saved money as a result of buying the subscription. The sample sizes were, respectively: Control: $N_c = 318,949$, 15% Off: $N_{15} = 318,755$, 25% Off: $N_{25} = 319,547$.

uniform-discount data, we record an indicator t_n^{ud} for treatment status (i.e., whether or not they got the 10% discount), and consumption q_n over the ensuing two-week period.

Table 1 summarizes uptake decisions in the subscriptions RCT. As is typical of within-firm datasets, our overall sample size was fairly large ($N^s = 957,251$), so all estimates reported in the table are highly significant and statistically different across the two columns, despite low uptake rates. Strong uptaker status can only be directly confirmed within the control group, where baseline demand Q_c is observed for the entire subsample. We compute statistics about uptakers (i.e., those who were offered a subscription *and* purchased it) separately within the two treatment subsamples. The proportion of strong uptakers naturally increases with the offered discount d, as does the share of firm revenues which are derived from strong uptakers. Even though almost 30% of the population are strong uptakers, under 2% of the treatment group actually subscribed, which is evidence of salience and/or forecast error. This is not surprising, as the experiment was an initial randomized trial, and hence was not accompanied by the sorts of marketing efforts associated with launches of mature product lines. Moreover, while a large fraction of subscribers saved money, this fraction is less than one, providing evidence of eagerness and/or forecast error.

The primary data inputs to our structural analysis are various CDFs of consumption across different subgroups within the experiments, plotted in Figure 2. In Panel (A), we display the demand CDFs in the subscription experiments, comparing the control group (sample size $N_0^s = 318,949$) to the two treatment groups (sample sizes $N_1^s = 318,755$ and $N_2^s = 319,547$, respectively). In Panel (B), we plot demand CDFs of subscription uptakers in the two treatment arms. These naturally differ much more from the control demand CDF, due to a combination of selection effects (uptake decisions) and treatment effects (increased demand under discounted pricing). We can reject equality of each treated distribution and the control distribution, as well as equality of the two treatment distributions, at the 5% level or less in Panels (A) and (B), using a standard Kolmogorov-Smirnov test: the maximum p-value for all six possible pairwise comparisons in the two figures is 0.028. In Panel (C) we plot the demand CDFs in the uniform-discount RCT for the control group ($N_c^{ud} = 500,645$) and the treatment group ($N_d^{ud} = 450,634$). We firmly reject equality of these two CDFs using a Kolmogorov-Smirnov test (p-value $\leq 10^{-16}$). Finally, Table 2 contains a comparison of demand moments across the subscriptions and uniform-discount experiments.

Figure 2. Raw CDFs From Experiemnts

(A) Subscription Aggregate CDFs (B) Subscription Uptaker CDFs (C) Uniform-Discount CDFs

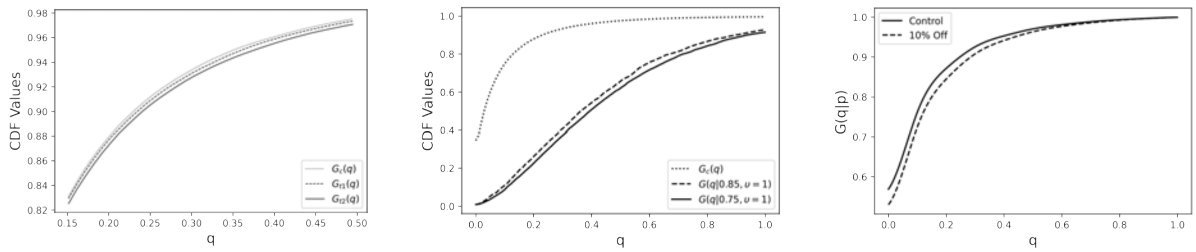

Notes: Panel (A) compares the demand CDF within the control group to the demand CDF in the treatment group. Panel (B) compares the demand CDF of the control group with the demand CDFs for uptakers within the two treatment arms. Panel (C) compares demand CDFs of control and treatment for the uniform-discount RCT.

Several facts are evident from the summary statistics and CDF plots. First, the LoD is empirically upheld in all cases, as lower prices drive stochastic dominance shifts in demand. Second, aggregate effects of the subscriptions treatment were small because the treatment condition here was an *offer to purchase a discount*, and only a small fraction of treated consumers did so. Third, uptakers within the subscription RCT are systematically different from the rest of the sample population: their demand CDFs differ dramatically from the control CDF. These differences encapsulate both selection and treatment effects, which our identification strategy is designed to tease apart. Fourth, there are some differences between demand distributions in the subscriptions and uniform-discount experiments: the mass point at zero is somewhat larger in the latter. As we discuss in Appendix D, these differences disappear after controlling for observable characteristics of the respective populations, and do not drive our main results or our model test results.

Table 2. Summary Statistics for q

	Min	1^{st} Quartile	Median	3^{rd} Quartile	Max	Mean	Std. Dev.	N
Subscription	0.000	0.000	0.029	0.100	1.000	0.081	0.131	961,003
Uniform Discount	0.000	0.000	0.000	0.102	1.000	0.081	0.148	946,681
Pre/Post Difference (UD)	-1.000	-0.084	0.000	0.013	1.000	-0.027	0.160	

Notes: Rows one and two present information about the distribution of q in the Subscription RCT and uniform-discount RCT, respectively. The third row contains information about the distribution of the diferences between the value of q when comparing the pre-experiment period to the post-experiment period in the uniform-discount RCT.

3.3. **Estimation: Subscriptions experiment.** In our main empirical application, we compute all estimates and counterfactual projections under the multiplicatively separable utility model. In Appendix F, we probe for robustness to mis-specification error by re-computing estimates and counterfactuals for the polar opposite extreme of the φ-separable family: additively separable utility (i.e., $\varphi = 0$). Our main market-design conclusions remain largely unchanged. Thus, the counterfactual extrapolation encoded in our particular application of Assumption 4 is not a key driver of our empirical results. Empirically, the conditions required by Corollary 1 turn out to be satisfied for the optimal contract (S^{*}, d^{*}) considered in Section 4 below, which implies that empirical market-design prescriptions should be largely invariant to one's choice of extrapolation method.

In Section 2.1, we established a constructive identification argument for two demand CDFs that satisfy the RS condition 1: $G(q|(1-d), \upsilon=1)$, for (observed) demand under discount for uptakers, and $G(q|1, \upsilon=1)$ for (counterfactual) demand under default pricing for uptakers. The former is known directly from raw data, while the latter is pinned down by equation (3) and the objects τ (uptake rate), $G_c(q)$ (control demand CDF), and $G_{t_a}(q)$ (treatment demand CDF in treatment arm $a=1,2$), which are known from raw data. Within the multiplicatively separable utility model, this in turn allowed us to identify the uptake function $\Upsilon(q)$, the behavioral parameters ρ (salience), δ (eagerness), $H_\varepsilon(\varepsilon)$ (forecast-error distribution); and the common utility function $u(q)$.

3.3.1. *Stage 1.* We begin by parameterizing the two rank-stable demand CDFs as cubic B-splines: $\widehat{G}(q|1-d_a, \upsilon=1; \boldsymbol{\omega}_a) \equiv \sum_{k=1}^{K_a+3} \omega_{ak}\mathcal{B}_{ak}(q)$, for $a=1,2$, and $\widehat{G}(q|1, \upsilon=1; \boldsymbol{\omega}_0) \equiv \sum_{k=1}^{K_0+3} \omega_k \mathcal{B}_{0k}(q)$, where the basis functions $\boldsymbol{\mathcal{B}}_a/\boldsymbol{\mathcal{B}}_0$ are uniquely determined by knot vectors $\boldsymbol{\kappa}_a = \{\kappa_{a1} < \kappa_{a2} < \cdots < \kappa_{a,K_a+1}\}$, $a=1,2$ and $\boldsymbol{\kappa}_0 = \{\kappa_{01} < \kappa_{02} < \cdots < \kappa_{0,K+1}\}$ (see de Boor (2001)), which are pre-specified by the econometrician, span the relevant support, and partition it into K_a and K_0 sub-intervals, respectively.[29] For each $a=1,2$ the B-spline forms facilitate a straightforward GMM estimator,

$$
\begin{aligned}
(\widehat{\boldsymbol{\omega}}_a, \widehat{\boldsymbol{\omega}}_0^a) =& \arg\min_{(\boldsymbol{\omega}_a, \boldsymbol{\omega}_0)} \Bigg\{ \sum_{n=1}^{N_{ua}} \Big(\widehat{G}(q_n|1-d_a, \upsilon=1; \boldsymbol{\omega}_a) - \tilde{G}(q_n|1-d_a, \upsilon=1)\Big)^2 \\
&+ \sum_{n=1}^{N_{ua}} \Big(\widehat{G}(q_n|1, \upsilon=1; \boldsymbol{\omega}_0) - \widetilde{G}(q_n|1-d_a, \upsilon=1) + \tfrac{\tilde{G}_{t_a}(q_n) - \tilde{G}_c(q_n)}{\tilde{\tau}_a}\Big)^2 \Bigg\}, \qquad (13) \\
& s.t.\ \ \omega_{ak} \le \omega_{a,k+1},\ k=1,\ldots,K_a+2; \quad \omega_{0k} \le \omega_{k+1},\ k=1,\ldots,K+2; \\
& \omega_{0,1} \ge 0,\ \ \omega_{a,1} \ge 0\ \ \omega_{0,K_0+3} = 1,\ \ \omega_{a,K_a+3} = 1; \text{ and } \widehat{G}(q_n|1, \upsilon=1; \boldsymbol{\omega}_0) \le \widehat{G}(q_n|1-d_a, \upsilon=1; \boldsymbol{\omega}_a)\ \ \forall n,
\end{aligned}
$$

where objects with tildes are empirical analogs of terms on the right-hand side of equation (3): $\tilde{G}(q_n|1-d_a, \upsilon_n=1) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \le q_n \cap t_{an'}=1 \cap \upsilon_{n'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{an'}=1 \cap \upsilon_{n'}=1)}$, $\tilde{G}_{t_a}(q_n) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \le q_n \cap t_{an'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{an'}=1)}$, $\tilde{G}_c(q_n) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \le q_n \cap t_{0n'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{0n'}=1)}$, and $\tilde{\tau} = \frac{\sum_{n=1}^{N^s} \mathbb{1}(t_{an}=1 \cap \upsilon_n=1)}{\sum_{n=1}^{N^s} \mathbb{1}(t_{an}=1)}$. The last line of the constraints represents terminal conditions CDFs must satisfy, and a stochastic dominance condition implied by the LoD. Resulting GMM estimates of the two main CDFs are plotted in Figure 12 in Online Appendix C.

3.3.2. *Stage 2.* After computing these estimates (still holding $a=1,2$ fixed), it is straightforward to plug them directly into Equation (4) to estimate the uptake function and behavioral parameters: $\widehat{\Upsilon}_a(q) = \frac{\widehat{g}(q|1, \upsilon=1; \widehat{\boldsymbol{\omega}}_0^a)\tilde{\tau}_a}{\widehat{g}_c(q; \widehat{\boldsymbol{\omega}}_c)}$, $\widehat{\rho}_a = \lim_{q\to\infty} \widehat{\Upsilon}_a(q)$, $\widehat{\delta}_a = \frac{\lim_{q\to 0} \widehat{\Upsilon}_a(q)}{\widehat{\rho}_a}$, and $\widehat{H}_\varepsilon^a(q) = \frac{\widehat{\Upsilon}_a(q+S/d_a) - \widehat{\rho}_a\widehat{\delta}_a}{\widehat{\rho}_a - \widehat{\rho}_a\widehat{\delta}_a}$, where $\widehat{g}_c(q; \widehat{\boldsymbol{\omega}}_c)$ is the derivative of a B-spline estimate of the control demand CDF.

[29] A B-spline representation is useful for its differentiability and ease of imposing shape constraints directly as linear restrictions on the parameters (Hickman et al. (2017)). This is especially important because the unsmoothed empirical analog of equation (3) for $G(q|1, \upsilon=1)$ is not guaranteed to be monotone in finite samples.

For utility estimation, we specify a knot vector, $\boldsymbol{\kappa}_u = \{\kappa_{u1} < \kappa_{u2} < \cdots < \kappa_{u,K_u+1}\}$, and parameterize the utility function as a flexible quartic B-spline $\widehat{u}(q;\boldsymbol{\omega}_u) = \sum_{k=1}^{K_u+4} \omega_{uk}\mathcal{B}_{uk}(q)$. Recall that the CDF $G(q|1,\upsilon=1)$ is a selection-corrected analog of the (observed) treatment CDF $G(q|1{-}d_a,\upsilon{=}1)$, for an identical population of control consumers who would have purchased a subscription, had they received the offer. Within this hypothetical population, the term $\frac{\tilde{G}_{t_a}(q_n)-\tilde{G}_c(q_n)}{\tilde{\tau}_a}$ represents quantile-specific demand shifts under discount d_a. For treatment arm a define $\widetilde{\mathcal{T}}_a(q_d) \equiv \widehat{G}^{-1}\left[\widehat{G}(q_d|1{-}d,\upsilon{=}1;\widehat{\boldsymbol{\omega}}_a)\middle|1,\upsilon{=}1;\widehat{\boldsymbol{\omega}}_0^a\right]$. This mapping is the prediction any RS model would make about how much an individual consuming q_d at discounted price $(1-d)$ would consume when the price is $p_0{=}1$ instead. For a given guess of the utility parameters $\boldsymbol{\omega}_u$, based on equation (2) we can define $\widehat{\mathcal{T}}(q_d;\boldsymbol{\omega}_u) \equiv (u')^{-1}\left[\frac{1}{1-d}u'\Big(q^*\left[1-d;\theta^*(1-d,q_d;\boldsymbol{\omega}_u);\boldsymbol{\omega}_u\right];\boldsymbol{\omega}_u\Big);\boldsymbol{\omega}_u\right] = (u')^{-1}\left(\frac{1}{1-d}u'(q_d;\boldsymbol{\omega}_u);\boldsymbol{\omega}_u\right)$, where $\theta^*(1-d,q_d;\boldsymbol{\omega}_u)$ is the consumer type that chooses q_d under discount pricing. This represents the model-derived prediction for how an individual consuming q_d at discounted price $(1-d)$ would consume under price $p_0{=}1$. Therefore, we can pin down $\boldsymbol{\omega}_u$ by minimizing the l_2 distance between $\widetilde{\mathcal{T}}$ and $\widehat{\mathcal{T}}$. Letting $0{=}q_1{<}q_2{<}\ldots{<}q_m{=}q_{max}$ be an evenly spaced grid of points, we have:

$$\begin{aligned}\widehat{\boldsymbol{\omega}}_u^a =& \arg\min_{\boldsymbol{\omega}_u} \textstyle\sum_{j=1}^m \left(\widehat{\mathcal{T}}(q_j;\boldsymbol{\omega}_u) - \widetilde{\mathcal{T}}_a(q_j)\right)^2 \\ & s.t.\ \ \omega_{u1}=0,\ \ \omega_{u1}=\tfrac{\kappa_{u5}-\kappa_{u2}}{3},\ \ \omega_{uk}\le\omega_{u,k+1}-\epsilon,\ k=1,\ldots,K_u+3,\ \epsilon>0,\ \text{and} \\ & \tfrac{\omega_{uk}-\omega_{u,k-1}}{\kappa_{u,k+3}-\kappa_{uk}} \le \tfrac{\omega_{u,k+1}-\omega_{uk}}{\kappa_{u,k+4}-\kappa_{u,k+1}} - \epsilon,\ k=2\ldots,K_u+3,\ \epsilon>0,\end{aligned} \tag{14}$$

where the first constraint is a boundary condition $u(0;\boldsymbol{\omega}_u)=0$, the second constraint is a scale normalization $u'(0;\boldsymbol{\omega}_u){=}1$, and the third and fourth enforce monotonicity and concavity.[30,31]

3.3.3. *Empirical Results.* Figure 3 plots the uptake functions (Panel (C)) and Table 3 reports structural behavioral parameters implied by the uptake functions. While the comparison is noisy, differences between the uptake functions for the 15% and the 25% treatment arms are as expected. Given a more attractive subscription offer, we see suggestive evidence that consumers are more willing to subscribe at every q; this is essentially another manifestation of the LoD.

Uptake parameter estimates are suggestive of three behavioral tendencies. First, we find a large degree of inattention: uptake is low, even among very high-consumption individuals. This is not necessarily surprising, given that the data came from a brand new product offering by Lyft. Second, even low-demand individuals had some positive probability of buying a subscription. Our

[30]The boundary derivative condition is equivalent to normalizing the demand type of the marginal consumer under p_0 to one, or $\theta^*(0,1){=}1$. Thus, all estimated demand types are relative to this marginal reference consumer.

[31]Knots are chosen so that $\boldsymbol{\kappa}_c=\boldsymbol{\kappa}_0=\boldsymbol{\kappa}_1=\boldsymbol{\kappa}_2=\boldsymbol{\kappa}_u$ in order to facilitate comparisons, with sizes $K_c=K_0=K_1=K_2=K_u=9$, as these afforded high flexibility and additional knots made little difference. For efficiency in smaller samples, one would choose knots uniformly in quantile-ranks and discipline choice of knot-vector size via cross-validation, or likelihood approaches (e.g., Bayesian/Akaike information criteria). The number of objective function (14) evaluations should be $m{\ge}K_u$; we chose $m{=}50$. For the tolerance on the shape constraints we chose $\epsilon{=}10^{-6}$.

Table 3. Subscription Uptake Parameters

Parameter	15% Off	95% CI	25% Off	95% CI	Joint	95% CI
Salience ρ	0.079	$[0.039, 0.140]$	0.093	$[0.060, 0.164]$	0.086	$[0.055, 0.137]$
Eagerness δ	0.019	$[0.004, 0.137]$	0.107	$[0.023, 0.236]$	0.046	$[0.066, 0.134]$
Forecast Err. Mean μ_ε	-0.315	$[-0.127, -0.539]$	-0.357	$[-0.201, -0.632]$	-0.334	$[-0.180, -0.548]$
Forecast Err. St. Dev. σ_ε	0.194	$[0.084, 0.305]$	0.234	$[0.118, 0.391]$	0.206	$[0.102, 0.317]$

Notes: This table reports point estimates and bootstrapped 95% confidence intervals (using 2,000 bootstrap samples) for the parameters summarizing mistakes consumers make when deciding whether or not to buy a subscription.

estimates are consistent with 2%-11% of the consumer base being over-eager, conditional on paying attention. Third, estimates suggest that consumers are fairly inaccurate at forecasting their own demand over a 1-month horizon. Not only is there considerable month-to-month demand variation, but perhaps most striking, there is non-trivial bias as well. Forecast mean bias is roughly a third of $\overline{q}$, meaning consumers on average act as if they require a very high degree of confidence that they will break even before purchasing a subscription. Our view on the behavioral parameters (ρ, δ, and H_ε) is that they represent short-run messaging/information problems that are solvable by targeted interventions and consumer learning over time. They suggest an important role for the marketing wing of the firm in the roll-out of a mature subscription plan offering. Evaluation of this viewpoint is left to future research, but the results illustrate why many real-world subscription programs include efforts to help consumers understand when it is worthwhile to subscribe.[32]

In Figure 3, we plot estimates of the utility function (Panel (A)) and elasticity function (Panel (B)) from separate estimation of the utility and uptake functions across the two treatment arms. Estimates for the 15% (25%) discount group are depicted by dashed (solid) lines, with 95% confidence bands depicted by thin lines. The utility function for the 15% arm is estimated less precisely than for the 25% arm. This is analogous to the fact that in regressions, standard errors tend to decrease when the regressors have higher variance. The 25% discount is a larger deviation from default pricing, and thus gives more information about the average consumer's responsiveness. We cannot reject the null hypothesis $H_0 : \widehat{u}_u(q; \widehat{\boldsymbol{\omega}}_u^1) = \widehat{u}_u(q; \widehat{\boldsymbol{\omega}}_u^2)$ that the two estimated utility functions are the same, since $\widehat{u}_u(q; \widehat{\boldsymbol{\omega}}_u^2)$ (and its confidence bounds) lay entirely within the confidence bounds of $\widehat{u}_u(q; \widehat{\boldsymbol{\omega}}_u^1)$. This serves as an over-identification test, and the results thus far suggest that the rank-stable, 1-dimensional multiplicatively separable model is not inconsistent with the data.

We also estimate the model with a single utility function fitted to data pooled across both treatment arms: $\widehat{\boldsymbol{\omega}}_u = \arg\min_{\boldsymbol{\omega}_u} \{w_1 \sum_{j=1}^m (\widehat{\mathcal{T}}(q_j; \boldsymbol{\omega}_u) - \widetilde{\mathcal{T}}_1(q_j))^2 + w_2 \sum_{j=1}^m (\widehat{\mathcal{T}}(q_j; \boldsymbol{\omega}_u) - \widetilde{\mathcal{T}}_2(q_j))^2\}$, where w_1 and w_2 are weights chosen so that their ratio equals the ratio of the within-treatment-arm mean pointwise variances of the underlying GMM estimators $\widehat{u}(q_j; \widehat{\boldsymbol{\omega}}_u^a)$ (across $\{q_1, \ldots, q_m\}$). The results are plotted in Figure 4, with jointly estimated behavioral parameters in Table 3.

[32]E.g., Costco nudges consumers in real-time at the checkout line when they would strictly benefit by increasing their membership to another level (with a higher up-front fee) in order to take advantage of a higher discount rate.

Figure 3. Stage-II Estimation: Utility and Uptake (separate estimation)

(A) Utility Function (B) Elasticity Function (C) Uptake Function

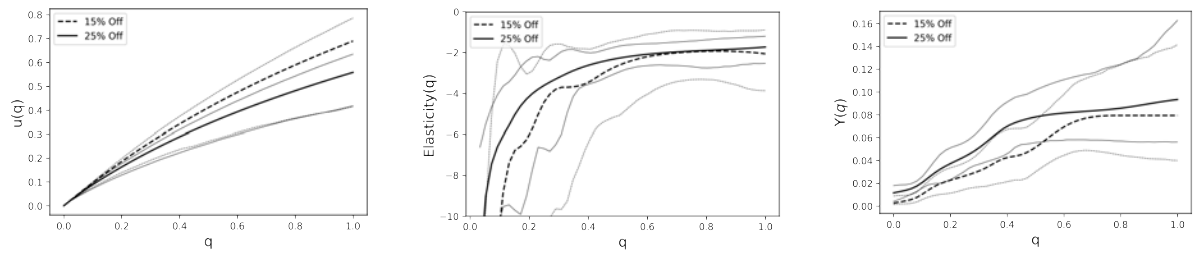

Notes: Thick lines are point estimates; Bootstrapped 95% confidence bands (using 2,000 bootstrap samples) are thin lines.

Figure 4. Stage-II Estimation: Utility (joint estimation)

(A) Utility Function (B) Elasticity Function

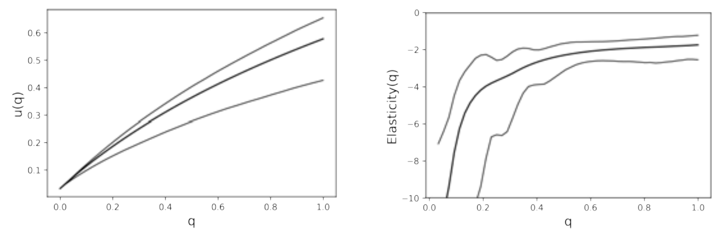

Notes: Thick lines are point estimates and thin lines are 95% confidence bands.

3.4. **Uniform-discount experiment.** Estimation on the uniform discount data is similar to the estimator in Section 3.3, but simpler because all treated consumers get the same 10% discount. Thus, the estimator is a special case of the previous one, where $\Upsilon(q)=1$ is trivially satisfied for all consumers in the treatment group. In the analogous stage 1 from Section 3.3, the relevant CDFs are directly known from raw data: $G(q|1-d, \upsilon=1)=G_t(q)$ and $G(q|1,\upsilon=1)=G_c(q)$. Thus, we first smooth the empirical CDFs with B-splines similarly as before:

$$(\widehat{\boldsymbol{\omega}}_d,\widehat{\boldsymbol{\omega}}_0)=\arg\min_{(\boldsymbol{\omega}_d,\boldsymbol{\omega}_0)}\left\{\sum_{n=1}^{N_d^{ud}}\left(\widehat{G}(q_n|1-d;\boldsymbol{\omega}_d)-\frac{\sum_{n'=1}^{N_d^{ud}}\mathbb{1}(q'_n\leq q_n)}{N_d^{ud}}\right)^2+\sum_{n=1}^{N_c^{ud}}\left(\widehat{G}(q_n|1;\boldsymbol{\omega}_0)-\frac{\sum_{n'=1}^{N_c^{ud}}\mathbb{1}(q'_n\leq q_n)}{N_c^{ud}}\right)^2\right\}$$

$$s.t.\;\; \omega_{dk}\leq\omega_{d,k+1},\; k=1,\ldots,K_a+2;\;\; \omega_{0k}\leq\omega_{0,k+1},\; k=1,\ldots,K_0+2,$$

$$\omega_{0,1}\geq 0,\;\; \omega_{d,1}\geq 0\;\; \omega_{0,K_0+3}=1,\;\; \omega_{d,K_d+3}=1;\;\text{ and }\;\widehat{G}(q_n|1;\boldsymbol{\omega}_0)\leq\widehat{G}(q_n|1-d;\boldsymbol{\omega}_d)\forall n=1,\ldots,N_c^{ud}.$$

After specifying a flexible B-spline utility function $\widehat{u}(q;\boldsymbol{\omega}_u)$, Stage-2 estimation follows Section 3.3.2 using the GMM estimator (14).[33] Finally, for each n we can estimate $\hat{\theta}_n$ (and hence, $\hat{F}_\theta$) within the multiplicatively separable model by inverting the consumer's FOC (1) for each q_n.[34]

[33] We again choose knots uniformly on the support $[0,\overline{q}]$, so we are only left to pick the number of subintervals. We must also specify the number of objective grid points, m. In practice, we chose $K_d=K_0=10$, $K_u=8$, and $m=50$.

[34] θ is only bound identified when $q_c=0$, so we back out F_θ by looking at the distribution of consumption in the treatment group and code θ to be the maximum type consistent with no consumption for all consumers with $q=0$. Given our normalization $u'(0)=1$, this amounts to setting $\theta=0.9$ for individuals in treatment with $q=0$.

3.4.1. *Results.* Utility function estimates and (bootstrapped) confidence bounds are in Figure 5.

Remark 7. ***Experimental Design*** The tight confidence bands (relative to subscription RCT estimates) are due to (*i*) lack of uptake failures arising from behavioral mistakes by consumers, and (*ii*) a uniform discount applied to both high- and low-demand patrons, which is less common when an up-front fee inhibits low-demand patrons from purchasing the discount. Thus, a uniform-discount RCT naturally treats a wider swathe of the population, which increases statistical power. This implies a novel methodological insight for optimal experiment design: if the goal is to optimize a nonlinear pricing scheme, then the best initial RCT to learn about latent agent heterogeneity is a uniform price shift, rather than a randomized screening mechanism. □

In Panel (C) of Figure 5, we plot the elasticity functions from the uniform-discount experiment and subscription experiment together. We find statistically and economically significant differences in elasticity estimates across the two experimental settings.[35] These differences reflect a more stringent test of the basic RS model of consumer demand: while the previous test of over-identifying restrictions (Section 3.3.3), based on comparisons of estimates from different arms of the *same* RCT, were unable to reject the model, a comparison of estimates from *two distinct* pricing RCT designs does reject it. Taking cues from the 2-dimensional model in Section 2.4.1, if unobserved substitution and brand-loyalty heterogeneity are present in the DGP, one would expect the two experimental conditions to induce different unseen selection patterns (via consumer switching behaviors), which could in principle account for the substantial differences in elasticity estimates. To be useful as a market-design tool these benchmark estimates must be evaluated for robustness to RS violations in the data. If unobserved consumer substitution is common and accounts for a large fraction of the aggregate shift from G_c to G_d, then the fee-discount offer (S, d) should be adjusted in order to preserve profitability, but if unobserved consumer substitution is less prevalent, then the baseline model will prescribe a subscription (S, d) that is closer to the true optimum. Using current structural estimates as a baseline reference point, we now implement a robust market-design exercise via our bounds approach from Sections 2.2–2.4.

4. Counterfactuals and Robust Policy Inference

We now use model estimates from the uniform-discount experiment to design an optimal subscription program; this choice is motivated by two ideas.[36] First, the uniform-discount RCT gives a more precise view of demand responses to discounts, as it is not complicated by behavioral mistakes. Second, since consumer mistakes can be mitigated in the long run through experience and

[35] In Appendix D, we show that systematic differences between the sample populations cannot account for the differences in Panel (C) of Figure 5, so one can rule out sampling differences as a viable explanation.

[36] Note that Lyft's implemented subscription plan pre-dates our analyses (e.g., Figures 6 and 16), so one cannot reverse-engineer Lyft's raw consumption quantiles (e.g., the normalizing constant $\overline{q}$) from our results.

FIGURE 5. Stage-II Estimation: Utility (uniform-discount RCT)

(A) Utility Function (B) Elasticity Function (C) Cross-RCT Model Test

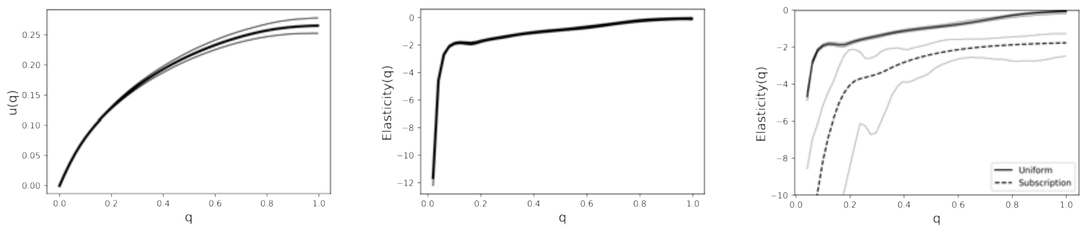

Notes: Thick lines are point estimates and bootstrapped 95% confidence bands (2,000 bootstrap samples) are depicted by thin lines.

firm marketing/information interventions,[37] and since firms have a hard time changing subscription programs once details are made public, they should base market design on their best approximation to long-run counterfactual demand shifts. We begin by characterizing optimal nonlinear pricing under baseline model estimates which assume rank stability. Although RS may be violated within the latent DGP, recall from Proposition 1 that it characterizes maximal price responsiveness by subscribers, and therefore still serves as a useful empirical reference point. After deriving a baseline (rank-stable) optimal policy, we employ our bounding approach to study how the model recommendation should shift under alternate, plausible DGPs. We also propose a simple data-driven method to estimate the degree of RS violations within the firm's latent DGP.

4.0.1. *Marginal Cost Imputation.* Henceforth we assume a constant marginal cost, c. Once again, in order to protect Lyft's internal data confidentiality, we do not incorporate raw information on its internal cost structure into our empirical analyses. Rather, we follow an imputation approach that is common to various strains of the industrial organization literature, including markup estimation in demand analysis (e.g., Ackerberg et al. (2007), MacKay and Miller (2021)). See Online Appendix G.1 for complete details on marginal cost imputation. Methodologically, this exercise will be of utility to researchers who lack access to internal cost data. Note also that it will become transparent below how changes in marginal cost c affect our derivation of the optimal subscriptions menu.

4.1. **Optimal Menu of Subscriptions.** Because our pricing problem is essentially a special case of the more general nonlinear pricing framework of Maskin and Riley (1984), we only sketch the key points here; see Online Appendix G.2 for additional technical details. We first derive a profit-maximizing continuous menu of subscription offers to produce a fully separating equilibrium by consumer types. The basic idea is that a firm's choice of discount as a function of θ is pinned down by an analog of the inverse-elasticity markup rule for monopoly pricing. Specifically, let $p(\theta)$ denote discounted price paid by subscribers of type θ (within the optimal menu) and let

[37] Our findings about the mistakes parameters, for example, spurred internal discussions within Lyft about ways the firm could help its consumers evaluate whether subscriptions made sense for them or not.

$\varepsilon(\theta) = -\frac{\theta f_\theta(\theta)}{1-F_\theta(\theta)}$ be the elasticity of the survivor function $1 - F_\theta(\theta)$ (interpretable as a demand curve). Then the firm's first order condition for $p(\theta)$ takes the form $\frac{p(\theta)-c}{p(\theta)} = -\frac{1}{\varepsilon(\theta)}$ which implies that $p(\theta) = \frac{c}{1+1/\varepsilon(\theta)}$. Having solved for $p(\theta)$, the optimal discount to give to each type θ is simply $d(\theta) = 1 - \frac{p(\theta)}{p_0}$. In turn, the optimal upfront cost schedule $S(\theta)$ is pinned down by a combination of the participation and incentive compatibility constraints. Participation constraints imply a boundary condition $S(\theta) = 0$ whenever $d(\theta) = 0$, while incentive compatibility implies the ordinary differential equation $S'(\theta) = -p'(\theta) q^*(p(\theta), \theta)$. Solving this ODE gives the optimal fee schedule, $S(\theta)$, and the locus $(S(\theta), d(\theta))$ constitutes the optimal continuum of subscription offerings. We plot the results in Figure 6. For Lyft's internal data confidentiality, we report d as a fraction of the (imputed) markup under default pricing, and S as a fraction of the maximum upfront fee from the optimal menu. Our optimal menu exhibits the "no-distortion-at-the-top" property familiar to mechanism design: the highest type, $\overline{\theta}$, buys a subscription where marginal price equals marginal cost: $p(\overline{\theta}) = c$.

4.2. **Optimal Single Subscription.** While a continuum menu is interesting to market-design researchers, real-world practitioners (e.g., Costco, Lyft, Charge Point, etc.) typically offer one or a small number of subscription plans. Firms may prefer simplicity for ease of implementation, or because customers find it difficult to select their optimal choice from a continuum. This begs the question, how much profits are left on the table in foregoing the fully separating equilibrium in favor of a simplistic subscription program? To answer this question we construct a grid of discounts d from 0 to marginal cost pricing and another grid of upfront fees S from 0 to 10% larger than the maximum upfront fee from the optimal menu. For each candidate subscription offer (S, d), we compute the (fully rational) model-implied subscription and consumption choices of each consumer type. Integrating over the type distribution gives us an estimate of total profits under each (S, d).

Results of this exercise are also shown in Figure 6, which depicts the optimal single contract, denoted by (S^*, d^*), with a dot, and a heatmap corresponding to the profitability of various (S, d) pairs. Lighter shades denote higher profitability. The heatmap, along with the superimposed optimal continuum menu shows an interesting pattern: There is a large region of high-profit contracts, with intermediate values of d (i.e., not too close to 0 or 1), where any single contract near the optimal menu performs fairly well.[38] The lower-right corner is shaded much more darkly than the upper-left corner, indicating that the main threat to profitability is offering a subscription plan that is overly generous (i.e., S being too low and/or d being too high). This insight will play a key role in our robust subscription offer derivation in the following section.

We find that the best single subscription offers a discount roughly half of the markup and captures 90% of profit gains from an optimal continuum menu. This suggests an answer to the previous question regarding the profit-simplicity tradeoff: the full continuum menu is only marginally better than

[38] There is reason to believe that the high-profit region in the heatmap of Figure 6 should generally be centered around the intermediate range of the optimal menu in other settings as well. See Online Appendix E for discussion.

FIGURE 6. Profitability of Subscription Offers

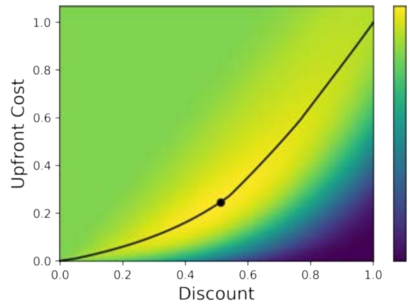

Notes: Discount d is expressed as a fraction of the markup under default pricing. Upfront fee S is expressed as a fraction of the maximal upfront fee in the optimal menu. Lighter shades in the heatmap denote higher profitability. We also plot the optimal menu of subscription prices, and a point representing the optimal single subscription.

a well-chosen simple menu. Thus, only mild or moderate concerns about implementing complex menus could rationalize the fact that most real-world subscription programs are low-dimensional.

4.3. **Robust SUD Bound Estimates.** Having optimized nonlinear pricing under the ideal but potentially faulty RS assumption, we turn to our robust bounds analysis, motivated by possible unobserved substitution between Lyft and rival services (Section 2.4.1), among other plausible RS failures (Section 4.3). Given the aggregate demand CDFs $(\widehat{G}_c, \widehat{G}_{d_0})$ from raw data, Proposition 1 enables us to construct upper bounds on counterfactual subscriber demand under (S^*, d^*), while Remark 6 provides a suite of options for constructing lower bounds, varying by reliance on data vs model structure. We compare these bounds in Figure 7. The thick solid line in Panel (A) is the estimated SUD upper bound, while the thick dashed line is the extrapolation-light SUD lower bound $\underline{\mathcal{B}}_{d_0}$ (option (iii) of Remark 6). We also plot the theory-free lower bound $\widetilde{\underline{\mathcal{B}}}_{d_0}$ (thick dotted line, option (i) of Remark 6) which does not impose the LoD. The theory-free bound is totally uninformative because it places too few restrictions on the latent DGP, while $\underline{\mathcal{B}}_{d_0}$, which does respect the LoD, is quite informative by comparison.[39] This finding demonstrates the inferential power to be had from the most basic behavioral assumptions within an incomplete structural model of demand. It mirrors similar findings on partial identification in the IO literature, including Haile and Tamer (2003), Hortaçsu and McAdams (2010), and Freyberger and Larsen (2021), who derived remarkably tight empirical bounds on private valuations within incomplete bidding/bargaining models, by assuming only that consumer behavior adheres to basic rationality constraints.

Panels (B) and (C) of Figure 7 also plot extrapolation-full bounds $\underline{\mathcal{B}}^R_{\mathcal{D}_d}$ (thin dash-dot line, option (iv) of Remark 6) based on a continuum of price changes, $\mathcal{D}_{d^*} = (0, d^*]$, and their projected (out-of-sample) aggregate demand CDFs, $G^{oos}_d(q; G_c, G_{d_0})$, obtained by imposing the local smoothness

[39] Bound tightness depends both on features of the dataset and the (S, d) pair under consideration. As $S/(p_0 \times d)$ approaches 0, strong uptakers encompass the entire population, and $\underline{\mathcal{B}}(q; 0, d) = \overline{\mathcal{B}}(q; 0, d) = G_d(q)$. Similarly, as $S/(p_0 \times d)$ approaches ∞ the set of strong uptakers vanishes, so once again, the gap between the bounds collapses.

FIGURE 7. Robust Bounds

(A) SUD Bounds

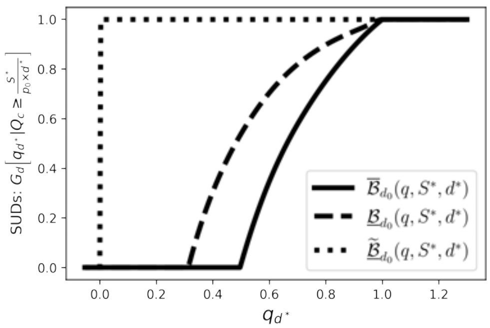

(B) Refined SUD Bounds

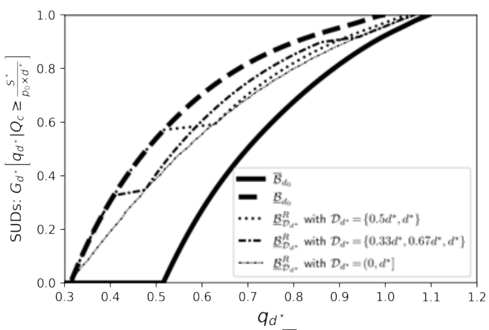

(C) SUD Elasticity Bounds

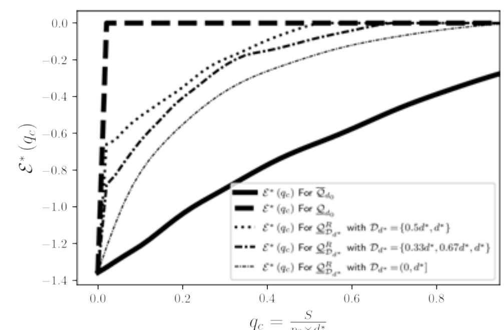

Notes: Panel (A) plots empirical bounds, $\overline{\mathcal{B}}_{d_0}(q; S^*, d^*)$ and $\underline{\mathcal{B}}_{d_0}(q; S^*, d^*)$, and the theory-free lower bound, $\underline{\widetilde{\mathcal{B}}}_{d_0}(q; S^*, d^*)$. Panels (B) and (C) plot refined lower bounds $\underline{\mathcal{B}}^R_{\mathcal{D}_{d^*}}(q; S^*, d^*)$ under hypothetical intermediate experimental prices $\mathcal{D}_{d^*}$, using simulated data to illustrate how additional exogenous price variation enables further inferential power. Panel (C) plots implied bounds on demand elasticities, holding discount d^* fixed, and varying upfront fee S; i.e., $\mathcal{E}^*(q_c) = \frac{\mathrm{E}[\varepsilon(Q_C)Q_C|Q_C \geq q_C]}{\mathrm{E}[Q_C|Q_C \geq q_C]}$, where $q_c = \frac{S}{p_0 \times d^*}$.

conditions in Remark 1 (option (1)).[40] For illustrative purposes we also depict hypothetical refined bounds from two intermediate discount sets: $\mathcal{D}^1_{d^*} = \{\frac{d^*}{2}, d^*\}$ and $\mathcal{D}^2_{d^*} = \{\frac{d^*}{3}, \frac{2d^*}{3}, d^*\}$ (medium dotted and dash-dot lines) as well. Recall from Section 2.5 that the thin and medium lines are also interpretable as bound refinements that could alternatively be obtained in fully nonparametric fashion from additional price variation within the RCT. The figure illustrates how adding new information—either through imposing smoothness conditions or by incorporating more data—leads to tighter bounds on counterfactual subscriber demand by ruling out various latent DGPs where subscribers are relatively price-insensitive. Panel (C) compares analogous demand elasticities of strong uptakers: the thick solid line is a lower-bound from the rank-stable DGP, while the other lines are upper bounds from various combinations of data and model structure.[41]

[40]In Figure 13 (Appendix C) we show that in the current context using rideshare data, out-of-sample CDF projections using options (2) or (3) of Remark 1 are remarkably similar. This suggests that imposing any one of the three varieties of smoothness conditions would have lead to nearly identical lower-bound refinements.

[41]A subtle insight on inference from RCT data arises from Panels (B) and (C): note that the thin dash-dot line assumes availability of fully continuous discount variation between 0 and the focal discount d^*, but multi-dimensional unobserved agent heterogeneity may still prevent point identification (i.e. upper and lower bounds coinciding).

While structural methods are often critiqued for opaqueness on the relation between empirical moments and model primitives, Figure 7 transparently illustrates how theory and raw data combine to deliver identifying power for a broad class of adverse selection models. The gap between the thick dotted line and the thick dashed line in Panel (A) represents inference derived from the observables using only basic structure from the LoD. The gaps between the thick dashed lines and the thin/medium lines in Panels (B) and (C) are interpretable as additional inference derived either from richer data, holding model structure fixed, or from imposition of additional structure (smoothness conditions), holding data fixed (see Remark 6). Finally, gaps between the thin dash-dot lines and the thick solid lines in panels (B) and (C) are inference derived from the observables by layering the full, rank-stable, multiplicatively separable utility model on top of the LoD.

4.3.1. *Robust Optimal Subscription Plans.* We now consider adjustments to the RS-optimal policy (S^*, d^*) based on bounds from Figure 7 in order to add robustness against unobserved RS violations. Recall that the adversarial DGP is one where strong uptakers minimally increase consumption after subscribing, subject to consistency with observables and the LoD. We can directly compute this worst case by simulating profits under the adversarial DGP, denoted by $\pi^A(S,d)$ for each (S,d), and we can compare these to rank-stable profits, denoted $\pi^{RS}(S,d)$, the opposite-extreme latent DGP not ruled out by data. Henceforth, we adopt the refined extrapolation-full lower bound $\underline{\mathcal{B}}^R_{\mathcal{D}_d}$ (Remark 6, option (iv)) as our main specification, because we view Assumption 4 and Remark 1, which require that aggregate demand evolves smoothly with price changes, as being reasonable. For comparison we also compute profits based on the extrapolation-light bound $\underline{\mathcal{B}}_{d_0}$ (Remark 6, option (iii)). The additional model structure that separates the former from the latter eliminates from consideration various pathological worst-case DGPs, where aggregate demand expands sporadically with price drops, and the experimentalist just happened to pick price levels p_0 and $p_0(1-d_0)$, between which *all* consumers are locally sated. We differentiate the corresponding adversarial profits under each extrapolation mode as $\pi^A_e(S,d),\ e \in \{l, f\}$. Recall from Sections 2.3–2.5 that these quantities bound true profits: $\pi^A_l(S,d) \leq \pi^A_f(S,d) \leq \pi(S,d) \leq \pi^{RS}(S,d)$.

In order to further refine market-design decisions, we can interpolate between the upper- and lower-bound DGPs by considering λ-*adversarial profits* $\pi^\lambda_e(S,d) \equiv \lambda\pi^A_e(S,d) + (1-\lambda)\pi^{RS}(S,d)$ if fraction λ of consumers behave according to the adversarial DGP, while $(1-\lambda)$ behave according to the rank-stable DGP. An intuitive interpretation would be that $(1-\lambda)$ of Lyft's users are loyal, having only a Lyft account or routinely checking only the Lyft app. The remaining λ fraction exhibit low loyalty, *and* are also assumed to substitute adversarially across the two platforms in response to Lyft price changes. Optimizing with respect to π^λ_e instead of π^A_e allows for improved robust decisions if the market designer believes that the pure adversarial DGP is overly pessimistic.[42] Our

[42]A common alternative approach would have been to compute the set of all subscription offers (S^*, d^*) that are optimal relative to *some* counterfactual strong-uptaker demand CDF consistent with the bounds in Figure 7. We

Figure 8. Path of Optimal Single Subscriptions

(A) Before Discontinuity (B) After Discontinuity

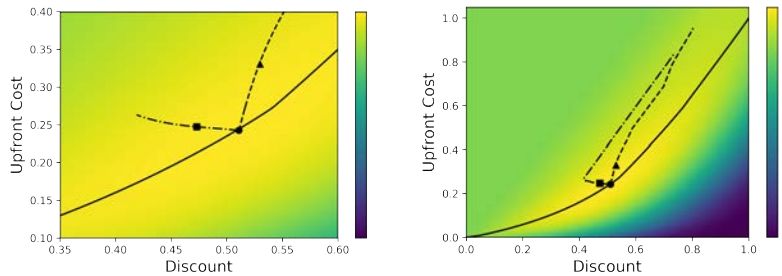

Notes: This figure plots evolution of robust optimal subscriptions as λ varies. The solid line and heat map are as in Figure 6, being from the (potentially mis-specified) baseline RS model. The circle is the RS-optimal subscription offer, $(S^*(0), d^*(0))$. The dashed line shows the path of robust adjustment using the extrapolation-light lower bound (Remark 6), and the dash-dot line is similar, but uses the extrapolation-full lower bound (Remark 6). The triangle and square are $(S^*(0.16), d^*(0.16))$ for the estimated value of $\widehat{\lambda} = 0.16$.

Figure 9. Robustness Tests

(A) Relative Adversarial Profits (B) Robust Profits Under RS

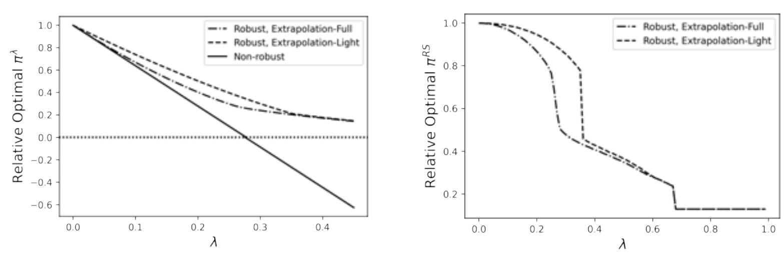

Notes: Panel (A) shows how optimal robust profits vary by λ, relative to RS profits—solid line, i.e., $\frac{\pi^\lambda(S^*(\lambda),d^*(\lambda))-\pi^\lambda(0,0)}{\pi^{RS}(S^*(0),d^*(0))-\pi^{RS}(0,0)}$—and how naive relative profits vary by λ—dashed line, i.e., $\frac{\pi^\lambda(S^*(0),d^*(0))-\pi^\lambda(0,0)}{\pi^{RS}(S^*(0),d^*(0))-\pi^{RS}(0,0)}$. Panel (B) depicts the cost of adopting a robust policy $(S^*(\lambda), d^*(\lambda))$ when the DGP is actually rank stable; i.e., $\frac{\pi^{RS}(S^*(\lambda),d^*(\lambda))-\pi^{RS}(0,0)}{\pi^{RS}(S^*(0),d^*(0))-\pi^{RS}(0,0)}$.

approach is conceptually related to Hansen and Sargent (2008), who consider a "structured" model as a benchmark and choose a policy to maximize the worst-case outcome in a family of unstructured models sufficiently "close" to the structured model.

For each $\lambda \in (0,1)$ we find the robust optimum, $(S_e^*(\lambda), d_e^*(\lambda))$, via grid search. Figure 8 plots the $(S_e^*(\lambda), d_e^*(\lambda))$ locus for our main specification (dash-dot line) along with the extrapolation-light analog (dashed line). Panel (A) zooms in on values of $\lambda \in [0, 0.35]$ for the main specification, while Panel (B) includes $\lambda > 0.35$ also. Figure 9 translates the robust (S,d) adjustments for various levels of λ into profit implications. Panel (A) characterizes a non-RS world, where λ is the actual degree of RS violations in the DGP. The dash-dot and dashed lines are robust excess profits above linear pricing, $\pi_e^\lambda(S_e^*(\lambda), d_e^*(\lambda)) - \pi^\lambda(0,0)$, relative to baseline excess profits $\pi^{RS}(S_e^*(0), d_e^*(0)) - \pi^{RS}(0,0)$, given extrapolation mode $e \in \{l, f\}$. The solid line depicts how naive excess profits,

do not pursue this approach partly because of the computational complexity involved in traversing the full set of CDFs between our bounds. More fundamentally though, as Aryal and Kim (2013) point out in a related context of setting auction reserve prices, the identified set of optimal policies may be of limited practical value: it includes many elements that may only yield small profit gains in the best case while causing large profit losses in the worst case.

TABLE 4. Fraction Subscriber Savings Retained: $(S_f^*(\lambda), d_f^*(\lambda))$ vs $(S^*(0), d^*(0))$

λ	Strong Uptaker Percentiles	0.1	0.25	0.5	0.75	0.9	Total
$\lambda = 0.35$	$SU(p_0, S_f^*(\lambda), d_f^*(\lambda))$	0.157	0.324	0.504	0.615	0.665	0.542
$\lambda = 0.35$	$SU(p_0, S^*(0), d^*(0))$	0	0	0.287	0.558	0.645	0.497
$\lambda = 0.16$	$SU(p_0, S_f^*(\lambda), d_f^*(\lambda))$	0.380	0.607	0.760	0.832	0.858	0.791
$\lambda = 0.16$	$SU(p_0, S^*(0), d^*(0))$	0	0.179	0.724	0.823	0.856	0.783

Notes: This table reports the retained savings ratio $\frac{q(r)d_f^*(\lambda)-S_f^*(\lambda)}{q(r)d^*(0)-S^*(0)}$, where $q(r)$ is the r^{th} quantile of Q_c among strong uptakers, for $r \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$. The final column is aggregate retained savings, or $\frac{\int_0^1 q(r)d_f^*(\lambda)-S_f^*(\lambda)dr}{\int_0^1 q(r)d^*(0)-S^*(0)dr}$.

$\pi_f^\lambda(S_f^*(0), d_f^*(0)) - \pi^\lambda(0,0)$, vary with λ, relative to baseline excess profits. The three lines are mechanically close to each other on the left where λ is near zero. Eventually, the solid line goes negative, whereas the robust adjustments prevent the other two lines from doing so. Panel (B) considers the implications of fixing a non-existent problem in a RS world, where an over-cautious market designer incorrectly chooses a robust offer $(S_e^*(\lambda), d_e^*(\lambda))$ anyway. There, we plot "paranoid" excess profits, $\pi^{RS}(S_e^*(\lambda), d_e^*(\lambda)) - \pi^{RS}(0,0)$, relative to true excess profits, as a measure of the insurance premium required to guard profits against the worst possible unobserved contingency.

For some intuition behind the profit discontinuity at $\lambda = 0.35$, first note that there is a set of contracts (S, d) that are good for profits near the RS DGP (λ near zero), but non-robust and very bad for profits near the adversarial DGP (λ near one). There is another disjoint set of contracts—with higher up-front fees S—that are very robust and profit-optimal near the adversarial DGP, but generally rendering low excess profits everywhere. In our main specification, for $\lambda < 0.35$ the market designer prioritizes profitability, and for $\lambda > 0.35$ she prioritizes robustness, with $\lambda = 0.35$ as the indifference point. Panel (A) traces out the locus of robust offers prior to the phase change: increasing S and reducing d (i.e., subscription generosity) both help hedge against profit shocks from unseen adverse consumer behavior. In Panel (B), we zoom out and show how the path jumps to a new region of (S, d) space with low profits when λ crosses the threshold.

The lesson from Figures 9 and 8 is that nonlinear pricing via subscription offers is only guaranteed to be a viable strategy when the adversarial fraction of consumers is below roughly one third under the main specification, and below one quarter under the unrefined extrapolation-light specification. Finally, Panel (B) of Figure 9 shows that the cost of achieving robustness is fairly low. If the true DGP really is RS but the market designer assumes $0 \le \lambda \le 0.35$, then robust profits $\pi^\lambda(S^*(\lambda), d^*(\lambda))$ still capture 80% or more of true excess profits. Thus, if we view robust policy design as insurance against making large errors, then fully hedging against that risk comes relatively cheaply for relevant values of λ where subscriptions are a viable business strategy at all.

Prior to the phase change, Figure 8 shows that policy prescriptions from the RS model are fairly robust to moderate perturbations in the underlying DGP. In our main specification we find that $S_f^*(0.35)$ is about 9% higher than $S^*(0)$, and $d_f^*(0.35)$ is about 18% lower than $d^*(0)$. Table 4 depicts implications of moving from the naive optimum to $(S_f^*(0.35), d_f^*(0.35))$ for consumer surplus. Among the set of strong uptakers relative to the robust optimum, $SU(p_0, S_f^*(\lambda), d_f^*(\lambda))$, the median ($90^{\text{th}}$ percentile) consumer retains half (two thirds) of would-be savings from the more generous but non-robust contract. Among the set of strong uptakers relative to the naive optimum, $SU(p_0, S^*(0), d^*(0))$, the lower quartile consumers retain none of their previous savings, since many of them transition from uptakers to non-uptakers when the contract becomes less generous.

4.3.2. *Estimating* λ. This discussion begs the question, what are relevant values of λ to focus on? To answer this question, it turns out that commonly available auxiliary data—individual-level consumption data prior to the RCT sample period—will suffice. Let Q_n^{pre} denote volume demanded in the two weeks prior to the start of the sampling period for the uniform-discount RCT. Note that the sample $\{q_n^{pre}\}_{n=1}^{N_c^{ud}+N_d^{ud}}$ is realized under default price p_0. Now, recall that the adversarial DGP $\underline{\mathcal{Q}}_{d_0}(q)$ maximally violates rank stability, and consider a comparison of the rank correlations between Q^{pre} and Q_{d_0} in the treatment group, and between Q^{pre} and Q_c in control. Under RS, we would expect these rank correlations to be identical, but if consumer behavior followed the adversarial DGP the pre-/post-RCT rank correlation should be lower within the treatment group.

This suggests a way to quantify the degree to which the data favor the rank-stable DGP $\overline{\mathcal{Q}}_{d_0}$ over the adversarial DGP $\underline{\mathcal{Q}}_{d_0}$. Let $\mathcal{S}_t$ denote Spearman's rank correlation between the pre- and post-RCT samples within the treatment group, or $\{q_n^{pre}, q_{dn}\}_{n=1}^{N_d^{ud}}$. Similarly, let $\mathcal{S}_c$ denote the rank correlation between the pre- and post-RCT samples in the control group, or $\{q_n^{pre}, q_{cn}\}_{n=1}^{N_c^{ud}}$. Finally, let $\mathcal{S}_a$ denote the rank correlation between the pre-RCT sample and the adversarial projection for the control group, or $\left\{q_n^{pre}, \underline{\mathcal{Q}}_{d_0}(q_{cn}, v)\right\}_{n=1}^{N_c^{ud}}$. We can then define $\hat{\lambda} \equiv 1-(\hat{\mathcal{S}}_t - \hat{\mathcal{S}}_a)/(\hat{\mathcal{S}}_c - \hat{\mathcal{S}}_a)$. Since $\mathcal{S}_c$ may be less than one due to within-consumer time-varying demand, we do not directly construct a ratio of $\mathcal{S}_t$ to $\mathcal{S}_c$; rather, we compare the differences $(\mathcal{S}_t - \mathcal{S}_a)$ and $(\mathcal{S}_c - \mathcal{S}_a)$ instead. Intuitively, the pre/post control rank correlation $\mathcal{S}_c$ is generally above the treatment rank correlation $\mathcal{S}_t$, which in turn is above the adversarial rank correlation $\mathcal{S}_a$.[43] Thus, $0 \leq \hat{\lambda} \leq 1$ should generally be true.

Table 6 in the Online Appendix reports raw rank correlations with 95% confidence intervals. Our point estimate is $\hat{\lambda} = 0.160$, with a 95% confidence upper bound of 0.281, well below the critical cutoff of $\lambda = 0.35$ where subscriptions cease to be an effective business strategy. The robust optimal subscription selects $S_f^*(0.16)$ at 2.0% above the naive optimum fee, and $d_f^*(0.16)$ at 8.0% below the naive optimum discount. This contract implies an implicit "insurance premium" of only 5% of baseline subscription profits (Figure 9, Panel (B)), to hedge against the worst likely case, while consumers retain nearly 80% of surplus gains, relative to the naive optimum. Thus, our empirical case study demonstrates that one may use partial identification to facilitate effective empirical

[43]If one assumes a well-behaved model of time-varying demand where persistence arises solely from a consumer's stable type θ, and period-t demand shocks ε_t similar to those described in Section 2.1 are independent across time, then the inequalities $\mathcal{S}_c \geq \mathcal{S}_t \geq \mathcal{S}_a$ follow as a direct consequence of Propositions 1 and 2.

market design, despite crucial limitations like unseen confounding choices by agents. We find that nonlinear pricing policies exist that are both surplus improving for agents, profit-improving for the principal, *and* robust against worst-case, unobserved contingencies.

5. Discussion and Conclusion

5.1. **Potential Applications in Broader Adverse-Selection Settings.** Our proposed methodology may facilitate new empirical analyses outside of consumer demand in IO. Here we give a brief overview of a few examples of related settings where non-linear pricing is of interest and multi-dimensional agent heterogeneity is a salient challenge to inference. Recall that these inferential challenges arise when agents with the same observed action are heterogeneously price sensitive.

5.1.1. *Regulation.* In the environmental regulation model of Kang and Silveira (2021), firms are heterogeneous with respect to θ_i, which parameterizes private benefits from negligently ignoring pollution regulations. A type-θ_i firm who chooses negligence level a receives gross utility $\theta_i b(a)$. Trading off these benefits, more negligent firms are more likely to get caught: the number of infractions observable to the regulator, K, is distributed $Poisson(a)$. The regulation punishes polluting firms by setting a penalty schedule mapping the number of violations k to a fee $\epsilon(k)$. The firm then chooses negligence level a to solve $\max_a \theta_i b(a) - \sum_{k=0}^{\infty} \epsilon(k) \frac{a_k}{k!}$. Now suppose that long-run adjustments to new, heavier regulations requires investment, firms start from various baseline levels of negligence, $\bar{a}_i$, and the cost of compliance is $\theta_i b(a) - \frac{C}{2}(a - \bar{a}_i)^2$. Under this richer model, if two firms make the same observed abatement decision a, then the one with larger $\bar{a}_i$ must have a correspondingly lower θ_i and hence will be less responsive to changes in the fine schedule.

5.1.2. *Income Taxation and Labor Supply.* In the optimal taxation model of Mirrlees (1971), workers are heterogeneous with respect to productivity W_i but have one unit of time to allocate, and homogeneous preferences over consumption and leisure $u(c, l)$. Given an income tax schedule $T((1-l)W_i)$, a type-W_i worker chooses leisure l to solve $\max_l u(c, l)$ s.t. $c \leq (1-l)W_i - T((1-l)W_i)$. In another labor-supply setting where census workers perform a low-skill task, D'Haultfœuille and Février (2020) adopt a model of labor supply where workers are heterogeneous with respect to their cost of effort, θ_i: they supply y units of effort and get paid piece-rate wage w_0, but incur a utility cost of this effort, $\theta_i c(y)$, with net utility given by $yw_0 - \theta_i c(y)$. In more general settings, it may be empirically relevant to allow for households or workers to differ both by (opportunity) cost of effort *and* by productivity of effort. In such settings, workers (households) with a fixed observed effort y (labor-supply $(1-l)$) will be heterogeneously responsive to changes in wage (tax) incentives.

5.1.3. *Insurance Demand.* Consumers are often thought of as being heterogeneous in their propensity for health spending λ_i. For example, consider a simple model where individuals who buy a given insurance policy pay $T(s)$ for s units of healthcare services. Individuals are homogeneous with respect to their utility of health $u(h)$, but differ in baseline illness rates λ_i. A type-λ_i individual who consumes s units of service has health level $h = s - \lambda_i$. Given a fixed insurance plan, this individual solves $\max_s u(s–\lambda_i)–T(s)$. In reality, some individuals are more at risk for diseases which are cheaper to treat while others are more at risk for diseases which can only be treated at a high price. This can be modeled as a second dimension of unobserved heterogeneity in the "marginal rate of transformation" from healthcare services to health, ω_i, where a type (ω_i, λ_i) individual who consumes s units of healthcare services has health $h = \omega_i s - \lambda_i$. This gives rise to the "selection on moral hazard" effect studied in Einav et al. (2013). Under their quadratic utility specification, if two individuals, i and j, choose the same consumption s, but $\omega_i > \omega_j$, then individual i's consumption will be more elastic to changes in $T(s)$ than individual j. Einav et al. (2013) are able to identify a multi-dimensional model of heterogeneity using rich menu variation in plan offerings, but such a strategy would be infeasible in many other settings such as the RAND health insurance experiment (Aron-Dine et al. (2013)), where plans are randomly assigned to consumers.

5.2. **Conclusion.** We propose a suite of tools that allow a market designer to flexibly estimate pricing counterfactuals. We clarify key conditions that underlay identification of the canonical adverse-selection model and highlight one key assumption, rank stability, as being especially problematic in the presence of multiple dimensions of unobserved heterogeneity. Despite significant data limitations, one can derive informative bounds on counterfactual demand under (out-of sample) price changes. These bounds arise because empirically plausible DGPs must respect the LoD and the observed shift(s) in aggregate demand resulting from a known experimental price change(s).

In the demand context, a (naive but) fully-specified rank-stable DGP corresponds to a sharp upper bound on consumer price responsiveness, and therefore still serves as a useful empirical benchmark for market design. The sharp lower bound on price responsiveness corresponds to a worst-case scenario (for profits) where the firm's loyal customer base is least price sensitive, and less loyal customers' unobserved substitution patterns account for a maximal fraction of the observed shift in aggregate demand. We also relax rank stability in a second way that can be explicitly modelled using rich internal data; namely, when customers fail to optimize subscriber decisions due to salience issues, over-eagerness, or an inability to perfectly forecast future demand.

Our estimated demand CDF bounds within the rideshare data turn out to be informative, despite lack of information that would facilitate structural identification of a more complete model of multi-dimensional agent heterogeneity. The bounds facilitate robust policy prescriptions using rich, internal data sources similar to those available in many other real-world applications. Our partial identification approach: (i) enables profitable nonlinear pricing design while achieving robustness against worst-case deviations from model assumptions, (ii) applies to a wide class of adverse-selection models, and (iii) serves as a novel guide for more effective experimental design.

References

Ackerberg, D., Benkard, C. L., Berry, S., and Pakes, A. (2007). Econometric tools for analyzing market outcomes. *Handbook of econometrics*, 6:4171–4276.

Arieley, D., Lowerstein, G., and Prelec, D. (2003). "coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, 118(1):73–106.

Armstrong, M. (1996). Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75.

Aron-Dine, A., Einav, L., and Finkelstein, A. (2013). The rand health insurance experiment, three decades later. *Journal of Economic Perspectives*, 27(1):197–222.

Aryal, G. and Kim, D.-H. (2013). A point decision for partially identified auction models. *Journal of Business & Economic Statistics*, 31(4):384–397.

Athey, S. and Haile, P. (2007). Nonparametric approaches to auctions. *Handbook of econometrics*, 6(ch. 60):3847–3965.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 7(1):242–262.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Berry, S. T. and Haile, P. A. (2021). Foundations of demand estimation. In *Handbook of Industrial Organization*, volume 4, pages 1–62. Elsevier.

Chan, T. Y. and Hamilton, B. H. (2006). Learning, private information, and the economic evaluation of randomized experiments. *Journal of Political Economy*, 114(6):997–1040.

Chernozhukov, V., Fernández-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78(3):1093–1125.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177.

Cotton, C., Hickman, B., List, J., Price, J., and Roy, S. (2023). Disentangling motivation and study productivity as drivers of adolescent human capital formation: Evidence from a field experiment and structural analysis. *Working paper, University of Chicago Economics Department.*

de Boor, C. (2001). *A Practical Guide to Splines Revised Edition.* New York: Springer-Verlag.

D'Haultfoeuille, X. and Février, P. (2011). Identification of a class of adverse selection models with contracts variation. *mimeo, CREST*, 2011-29.

D'Haultfoeuille, X. and Février, P. (2015). Identification of nonseparable triangular models with discrete instruments. *Econometrica*, 83(3):1199–1210.

D'Haultfœuille, X. and Février, P. (2020). The provision of wage incentives: A structural estimation using contracts variation. *Quantitative Economics*, 11(1):349–397.

Dubé, J.-P. (2004). Multiple discreteness and product differentiation: Demand for carbonated soft drinks. *Marketing Science*, 23(1):66–81.

Dubin, J. A. and McFadden, D. L. (1984). An econometric analysis of residential electric appliance holdings and consumption. *Econometrica: Journal of the Econometric Society*, pages 345–362.

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., and Cullen, M. R. (2013). Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219.

Firpo, S. and Ridder, G. (2019). Partial identification of the treatment effect distribution and its functionals. *Journal of Econometrics*, 213:210–234.

Freyberger, J. and Larsen, B. J. (2021). How well does bargaining work in consumer markets? a robust bounds approach. *Working Paper*.

Guerre, E., Perrigne, I., and Vuong, Q. (2009). Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4):1193–1227.

Haile, P., Hong, H., and Shum, M. (2006). Nonparametric tests for common values in first-price sealed-bid auctions. *Working Paper, Yale University*.

Haile, P. and Tamer, E. (2003). Inference with an incomplete model of english auctions. *The Journal of Political Economy*, 111(1):1–51.

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press.

Heckman, J. J., Smith, J., and Clements, N. (1997). Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies*, 64(4):487–535.

Hedblom, D., Hickman, B., and List, J. (2022). Toward an understanding of corporate social responsibility: Theory and field experimental evidence. *Working paper, Washington University in St Louis, Olin Business School*.

Hendel, I. (1999). Estimating multiple-discrete choice models: An application to computerization returns. *Review of Economic Studies*, 66(2):423–436.

Hickman, B. R., Hubbard, T. P., and Paarsch, H. J. (2017). Identification and estimation of a bidding model for electronic auctions. *Quantitative Economics*, 8(2):505–551.

Hortaçsu, A. and McAdams, D. (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy*, 118(5):833–865.

Imbens, G. and Angrist, J. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 52(2):467–475.

Imbens, G. and Rubin, D. (1997). Estimating outcome distributions for compliers in instrumental variables models. *Review of Economic Studies*, 64(4):555–574.

Kang, K. and Silveira, B. (2021). Understanding disparities in punishment: Regulator preferences and expertise. *Journal of Political Economy*, 129(10:2947–2992.

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

Laffont, J.-J. and Tirole, J. (1986). Using cost observation to regulate firms. *The Journal of Political Economy*, 94(3):614–641.

Lariviere, M. A. (2006). A note on probability distributions with increasing generalized failure rates. *Operations Research*, 54(3):602–604.

Luo, Y., Perrigne, I., and Vuong, Q. (2011). Multiproduct nonlinear pricing: Mobile voice service and sms. *Working Paper*.

Luo, Y., Perrigne, I., and Vuong, Q. (2018). Structural analysis of nonlinear pricing. *The Journal of Political Economy*, 126(6):2523–2568.

MacKay, A. and Miller, N. (2021). Estimating models of supply and demand: Instruments and covariance restrictions. *working paper, Harvard Business School*.

Maskin, E. and Riley, J. (1984). Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics*, 3(4):303–328.

McManus, B. (2007). Nonlinear pricing in an oligopoly market: The case of specialty coffee. *Rand Journal of Economics*, 38(2):513–533.

Mirrlees, J. A. (1971). An exploration in the theory of optimal taxation. *The Review of Economic Studies*, 38(2):175–208.

Molinari, F. (2020). Microeconometrics with partial identification. *Handbook of econometrics*, 7:355–486.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18:301–317.

Nevo, A. (2003). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Richards, T. J. and Bonnet, C. (2016). Models of consumer demand for differentiated products. *Working Paper # 16-741, Toulouse School of Economics*.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.

Sun, G. (2023a). Minimum cost minimum income guarantees. *Working Paper*.

Sun, G. (2023b). Separability, identification, and extrapolation in a family of adverse selection models. *Working Paper*.

Torgovitsky, A. (2015). Identification of nonseparable models using instruments with small support. *Econometrica*, 83(3):1185–1197.