

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

AARON BODOH-CREED*, BRENT HICKMAN**, JOHN A. LIST†, IAN MUIR††, AND GREGORY SUN‡

ABSTRACT. In this paper, we provide a suite of tools for empirical market design, including optimal nonlinear pricing in intensive-margin consumer demand, as well as a broad class of related adverse-selection models. Despite significant data limitations, we are able to derive informative bounds on demand under counterfactual 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). These bounds facilitate robust policy prescriptions using rich, internal data sources similar to those available in many real-world applications. Our partial identification approach enables viable nonlinear pricing design while achieving robustness against worst-case deviations from baseline model assumptions. As a side benefit, our identification results also provide useful, novel insights into optimal experimental design for pricing RCTs.

Keywords: Subscriptions, demand, consumer heterogeneity, behavioral mistakes, robust inference

JEL Codes: B4, C14, C51, C52, C93, D4, J2, L1, L5

Acknowledgements: 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, and seminar participants at WashU-Olin, Queen’s University, Princeton, and participants of the 2022 IIOC meetings.

Note: previous versions of this manuscript circulated under the title “Stress Testing a Structural Model of Nonlinear Pricing: Robust Inference on Intensive-Margin Consumer Demand”

Contact Information:

*Aaron Bodoh-Creed (Amazon Inc.): abodoh.creed@gmail.com

**Brent Hickman (*corresponding author*, WashU St. Louis, Olin Business School): hickmanbr@gmail.com

†John List (University of Chicago Economics Dept.): jlist@uchicago.edu

††Ian Muir (Walmart Inc.): muir.ian.m@gmail.com

‡Gregory Sun (WashU St. Louis, Olin Business School): 14sungre@gmail.com

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 component of optimal 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 of typical demand analysis data: 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 across rival producers. Moreover, aside from just facilitating a study of individual-level demand intensity variation, rich consumer-level data may allow for reduced dependence on parametric restrictions.

Of course, the auctions literature represents another instance of demand analysis methodologies which are typically more flexible on *a priori* assumptions about the distribution of consumer tastes, and which are well-suited for market design. However, empirical auction methods typically rely on game-theoretic interactions among consumers to identify demand intensities, making them ill-suited to most settings where consumers are price-takers. Our goal in this paper is to develop a

new complementary approach with a focus on intensive-margin demand in settings where consumer-transaction-level data are plentiful, but market-level data are not. This new capability will facilitate analyses on new and interesting questions relevant to policy-makers and business practitioners. In particular, we will use our approach to inform design of optimal nonlinear pricing in the form of a subscription plan offered to consumers.

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 (described briefly below) 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 the principal can exogenously vary payoffs offered to agents from say $P_0(q)$ to $P_1(q)$, the question we seek to answer 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 the unobserved heterogeneity among agents is single- or multi-dimensional. For example, 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 by developing a parsimonious model of intensive-margin demand, and we discuss how recent results from econometric theory (Torgovitsky (2015); D'Haultfoeulle and Février (2015); D'Haultfoeulle and Février (2020)) can be applied to achieve nonparametric identification of model primitives, including a flexible consumer utility function and a distribution of idiosyncratic demand intensities. 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. We then discuss various plausible phenomena—e.g., behavioral mistakes among consumers, and unobserved, extensive-margin substitution decisions—that can lead to failures of the central rank-stability condition needed for point identification.

We first show how rank-stability violations of known forms—e.g., consumer mistakes in subscription purchasing evident within raw data—can be explicitly added to the empirical model to restore identification, using basic available observables. We then discuss a more difficult empirical problem: violations of rank stability that are of *unknown* forms, either because the underlying model of consumer behavior is incomplete, or because the econometrician lacks sufficient observables to empirically model them. These may stem from phenomena such as (but not limited to) unobserved substitution between the firm's product/service and that of its rival, in the presence of unobserved

heterogeneity in consumer brand loyalty. 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 the firm’s consumers are assigned to “control” status, making demand decisions under default price p_0 , while others are assigned to “treatment” status with a discount of d_0 off the default price.¹ 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 by characterizing the best-case scenario in terms of profits derived from nonlinear pricing. 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.

As a side benefit, our analysis produces various novel insights to guide effective field experimental design as well. Our central theme and modelling approach (see Section 2.4.1) highlights how individually randomized price variation may be insufficient alone to deliver full econometric exogeneity and point identification in the presence of multiple dimensions of test-subject heterogeneity. We also show that this problem may even persist despite incredibly rich price variation (see Section 4.3). Finally, our derivations of robust bounds to grapple with this issue (see Section 2.5, Section 4 introduction, and Section 4.3) provide specific, concrete guidance on optimal pricing RCT design for a researcher wishing to maximize ex-post inferential power.

While derivations of the bounds are fairly technical, there is some simple intuition behind our method. Suppose the firm wishes to optimize a nonlinear pricing scheme in the form of a subscription offering where consumers pay $\$S$ upfront for a discount of $d \in (0, 1)$ over the ensuing period. Therefore, the firm has a need to forecast counterfactual demand under arbitrary (S, d) pairs in order to optimize the subscriptions 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 the outside option is unobserved. In such a world, 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

¹In the IO context, such variation commonly arises by firms’ use of internal pricing RCTs.

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 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 under discount d_0 . This 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 (nearly) a zero-sum transfer from the firm. Conversely, the best-case scenario is one where this problem is non-existent: all consumers are loyal, and the aggregate shift between control and treatment demand CDFs is all due to predictable, rank-stable demand responses.

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 agent heterogeneity), then not only may point identification be compromised, but also, the researcher’s ability to detect the mis-specification problem may also be compromised given standard over-identifying restrictions (e.g., availability of 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 to 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 rank-stability violations. 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 hedge against worst-case unobserved consumer behavior, while achieving roughly 95% of baseline projected subscription profits under the rank-stability assumption.

More broadly, our proposed robust bounds approach is applicable outside the narrow field of demand estimation in industrial organization. It can be used to stress-test the policy implications of a wide variety of adverse selection models, including a few prominent examples we discuss here. First, in the context of procurement, Laffont and Tirole (1986) consider a model where a government wishes to complete some project and must design a contract to incentivize a monopolist to exert effort to reduce costs. They assume that monopolists are unobservably heterogeneous along a single dimension of productivity, but one might expect that monopolists also differ by costs of unobserved managerial effort as well. Second, in a related context of environmental regulation, 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 is only able to observe total pollution, then one might imagine an alternate model 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 as the regulator seeks to optimally incentivize abatement on a fixed regulatory budget. Third, in the context of optimal income taxation considered by Mirrlees (1971), households are assumed to vary only by their (unobserved) labor productivity, but one might expect households to vary by consumption preferences as well. Fourth, in the context of labor contracts (e.g., D’Haultfoeulle and Février (2020)), workers may be unobservably heterogeneous by both productivity and leisure preferences which impact labor-supply costs, with both dimensions being crucial to the design of optimal nonlinear incentive schemes.²

Given exogenous incentive variation, adverse-selection models like those discussed above are typically non-parametrically identified, if agent heterogeneity can be plausibly restricted to one dimension. Conversely, these models typically are not identified, even with exogenous incentive variation, given multiple dimensions of heterogeneity. Our approach offers a new suite of tools for partial identification under high-dimensional unobserved heterogeneity. We show how to use a robust bounds approach to “stress-test” and adjust policy prescriptions derived from a potentially mis-specified one-dimensional structural model.

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

²Hedblom 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.

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 show how a monopolist’s first-order conditions can be used to infer demand parameters from equilibrium outcomes, and our approach uses exogenous (cross-sectional) price variation to infer parameters for optimal pricing 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, where the mapping between quantiles of the unobserved characteristic and quantiles of the observed action remain stable across price changes. We study the question of point and set identification when the central assumption of rank stability is violated. 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, both in terms of the counterfactual questions of interest and in the nature of the data required. The prior literature 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 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 about demand estimation. In particular, 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. This paper explores similar questions within the world of intensive-margin demand: our explicit model of rank-stability violations further illustrates why randomized price variation may not suffice when

³We focus on single-product, continuous, intensive-margin demand, so we do not consider complications arising in the case of a multi-product monopolist (e.g., Armstrong (1996) and Luo et al. (2011)). Our model also differs from settings where demand is discrete (often unitary), with pricing nonlinearity along the quality dimension. In practice, such applications tend to focus on discrete price/quality pairs (e.g., Leslie (2004) and McManus (2007)).

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

⁵A smaller strand of the discrete-choice demand literature allows for mixed discrete-continuous consumer decisions, including Hendel (1999) and Dubé (2004), and surveyed in Richards and Bonnet (2016).

consumers make continuous choices either. However, our partial identification framework and empirical application demonstrate that random price variation can nonetheless be used to construct informative counterfactual bounds that can still 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) wish to consume as much of the service as possible but are idiosyncratically deterred by high prices, while the supplier (principal) wishes to maximize profits. We denote by q the quantity of a homogeneous, finely divisible good or service demanded by a consumer.

We assume multiplicatively separable (MS) utility $U_i(q; \theta_i) = \theta_i u(q)$, parameterized by idiosyncratic type $\theta > 0$. Here, $u(q)$ satisfies standard regularity conditions, being strictly increasing $u'(q) > 0$, $\forall q \in \mathbb{R}_+$, and concave $u''(q) < 0$, with $\lim_{q \rightarrow \infty} u'(q) = 0$. For convenience, we also impose scale and shift normalizations: $u'(0) = 1$ and $u(0) = 0$. For now, individuals are heterogeneous along a single dimension (we relax this assumption in Section 2.4.1), with θ indexing each consumer'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 $[\theta, \bar{\theta}]$.

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 $\max_q \{\theta u(q) - pq\}$.⁶ The first-order condition (FOC) to this optimization problem is given by

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

⁶The pricing schedule can be nonlinear; a sufficient condition is that $P(q)$ is not too convex, or $P''(q) < |u''(q)|$ $\forall q \in \mathbb{R}_+$, in which case a well-defined, strictly monotone solution exists. Empirically, this is an innocuous restriction, given that volume *discounts* (i.e., *concave* pricing) is most common in real-world applications.

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 $[\underline{q}, \bar{q}]$.

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: $p < p' \Rightarrow G(q|p) \leq G(q|p')$.

D’Haultfoeuille and Février (2015), Torgovitsky (2015), and D’Haultfoeuille 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' :

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

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

⁷More 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 $\theta' < \theta$).

⁸Since marginal utility is bounded, each type θ has a finite choke price $\bar{p}(\theta) = \theta u'(0)$, where they choose $q^*(\bar{p}(\theta); \theta) = 0$. 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 B and C.

transformations for $q \in \mathcal{O}$ for some identified set of points $\mathcal{O}$. An 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

Remark 1. 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 $\square$

Structural identification 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 surprisingly high degree of identifying power: if a firm has capacity to exogenously vary prices for its product/service it can nonparametrically identify all relevant parameters necessary for finding the optimal (nonlinear) price schedule. The LoD is always testable, and with more than two prices, the extrapolation quality of MS utility may even be testable as well. On the other hand, the RS condition is arguably the most stringent and least testable model implication. Moving forward we will sometimes 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 a scenario where RS violations are generated by departures from full rationality in consumers' subscription choices. In this case, we briefly show that by explicitly modeling behavioral mistakes, the researcher can restore point identification using available data. Moreover, we propose a means of directly estimating behavioral parameters (e.g., offer salience and forecast errors), which is of independent interest to market designers engaged in nonlinear pricing via optimal subscription programs. We then proceed to our main methodological contribution by exploring RS violations of *unknown form*; i.e., when the precise 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

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

¹⁰Since 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) $\frac{\partial G(q|p_0)}{\partial p}$ is identified or (ii) $G(q|p_0(1-d_1))$ is identified for one additional discount level $d_1 \neq d_0$.

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

counterfactual demand CDFs consistent with observables and the LoD. We show that these bounds can be used to derive optimal pricing prescriptions that are robust to the most extreme profit-minimizing unobserved consumer behaviors not ruled out by the data and LoD.

2.1. Structural Identification in An Explicit Model of RS Violations. In recent decades, experimental incentive variation has become an increasingly common tool among firms, governments, and other organizations. A common impediment to useful inference from controlled randomization is the presence of apparent deviations from fully rational decision making on the part of experimental participants. Such deviations can include both under-reaction (e.g., salience deficiencies as in Chetty et al. (2009)) and over-reaction (e.g., when irrelevant information influences behavior as in Ariely et al. (2003)). As motivation for this concern in the consumer demand context, Lyft ran a RCT in early 2019 where it offered a random subset of its customers the opportunity to buy a subscription lasting for a month. The resulting data show direct evidence of both over-reaction and under-reaction within the treatment group offered a subscription plan.

We consider a setting where Lyft runs 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 $\theta u(q^*(p_0; \theta)) - p_0q^*(p_0; \theta)$. Thus, consumer type θ should subscribe if and only if $\theta u(q^*(p_0(1-d_0); \theta)) - p_0(1-d_0)q^*(p_0(1-d_0); \theta) - S \geq \theta u(q^*(p_0; \theta)) - p_0q^*(p_0; \theta)$.

Intuitively, consumers should buy the subscription only if the change in surplus is greater than or equal to the upfront fee S . Consumer surplus from discount d_0 is at least as large as $p_0 \times d_0 \times q^*(p_0; \theta)$ and no larger than $p_0 \times d_0 \times q^*(p_0(1-d_0); \theta)$, from which follows a set of testable predictions that do not require assuming MS utility or RS: rational individuals for whom $q^*(p_0; \theta) > \frac{S}{p_0 \times d_0}$ should subscribe, while individuals for whom $q^*(p_0(1-d_0); \theta) < \frac{S}{p_0 \times d_0}$ should not. Therefore, 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 if they had, while 6% of subscribers had low enough ride-share 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

¹²Our 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

(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, since subscription benefits last for a month, we allow for consumers to be imperfect at forecasting their own demand intensity over the ensuing month. 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 and no information is lost by this re-parameterization, 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.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 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 can now define an *uptake function* as $\Upsilon(q) \equiv \Pr[\text{subscribe} | q^*(p_0; \theta) = q] = \rho\delta + \rho(1-\delta)H_\varepsilon(q - q_s^*)$. Moving forward, it will be convenient to define a short-hand indicator function, $v_i \equiv \mathbb{1}[i \text{ subscribes}]$.

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, $G(q|p_0(1-d_0), v=1)$, the complementary demand distribution, $G(q|p_0, v=1)$, is not directly observable because the set of would-be uptakers in the control group is not known.

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.

¹³Note 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 .

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

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

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

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

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

This pins down the behavioral parameters ρ , δ , and H_ε . First, note that $\lim_{q \rightarrow \infty} \Upsilon(q) = \rho$ and $\lim_{q \rightarrow 0} \Upsilon(q) = \rho\delta$.¹⁴ Finally, with ρ and δ known, we can use the definition of $\Upsilon(q)$ above to identify $H_\varepsilon(q)$ using the relationship $H_\varepsilon(q) = \frac{\Upsilon(q+q_s^*) - \rho\delta}{\rho - \rho\delta}$. Summarizing, we have shown how deviations from fully rational subscription decisions may generate RS violations, but ones that can be fully addressed to restore point identification using available data.

2.2. Robust Inference Under RS Violations of Unspecified Form. We now shift focus to more econometrically challenging scenarios where explicitly controlling for RS violations is not possible. For example, firms usually do not have good information on consumers’ substitution between their products/services and those of their nearest competitor. Suppose that household 1 usually purchases more volume than household 2 from *Firm A*, but in a month where *Firm A* had a large sales promotion, household 2 made more purchases than household 1. This is potential evidence that RS is violated and could arise if, for example, household 2 has low brand loyalty

¹⁴In 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.

toward *Firm A* and thus views *Firm B* as a closer substitute than household 1, which has more brand loyalty toward *Firm A*. As a result, sale prices at *Firm A* are more effective at (unobservably) switching household 2's purchasing behavior away from its rival, *Firm B*.

Note that RS requires the same ordering of purchasing behavior by households 1 and 2 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., heterogeneous budget constraints or income effects from a given price change—could render point identification impossible given available data.

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.

Before moving on, 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 estimation of bounds on the structural primitives θ and F_θ . However, in our case a prominent source of the partial identification problem is the fact that the consumer base may shift between the control (p_0) and treatment ($p_0(1 - d_0)$) states due to unobserved substitution between the firm's product and that of its competitors. If this is so, then the notion of pinning down a type distribution F_θ for a single, stable consumer base becomes ill-defined. For that reason, our approach focuses instead on derivation of bounds on counterfactual demand behavior under alternate pricing levels. Importantly, *this approach still allows for out-of-sample inferences*, under discount levels d not observed in the data.

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 to refine our profit bounds. In particular, we suppose that corresponding to each consumer, there are quantities (Q_c, Q_d) such that the individual would consume Q_c if given the control price p_0 and would consume Q_d 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 (or fixed numbers) 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) . The observed marginal distributions are as defined earlier, $G_c(q) = G(q|p_0)$ and $G_{d_0}(q) = G(q|p_0(1-d_0))$.

Given known CDFs G_c and G_{d_0} , and assuming the LoD, 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 for whom $Q_c \geq \frac{S}{p_0 \times d}$, (i.e., baseline consumption is relatively high) which we refer to as *strong uptakers*. Formally, *strong uptakers* are the set $SU(p_0, S, d) = \left\{ \text{consumer } n : q_{cn} \geq \frac{S}{p_0 \times d} \right\}$. Given the option to purchase a subscription offering discount d for $\$S$ upfront, under *any* model of the underlying preferences consumers for whom $Q_c \geq \frac{S}{p_0 \times d}$ will find it advantageous to subscribe, while (fully rational) 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 seek here to impose only minimal assumptions, meaning that the baseline predictions of the RS model may not hold. Our approach is based on the idea that subscription profitability depends primarily on how strong uptakers increase consumption in response to the discount.¹⁵

To facilitate robust inference on counterfactual profits, we will be interested in evaluating conditional probabilities of the form $\Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right]$, or in other words, the conditional counterfactual demand CDFs, for various values of (S, d) . We will often refer to these focal conditional probabilities as *strong uptaker distributions* (or SUDs) for short. Our goal is to construct a set of sharp bounds for the SUDs: $\bar{\mathcal{B}}_{d_0}(q; S, d) \leq \Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right] \leq \underline{\mathcal{B}}_{d_0}(q; S, d)$. Since SUDs represent a conditional counterfactual demand CDF, the *upper bound* $\bar{\mathcal{B}}_{d_0}$ (also a CDF) is pointwise *below* the SUD; i.e., $\bar{\mathcal{B}}_{d_0}(q; S, d)$ stochastically dominates $\Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right]$. Similar logic applies to the lower bound $\underline{\mathcal{B}}_{d_0}$ (also a CDF): fixing the last two arguments (S, d) , $\underline{\mathcal{B}}_{d_0}(q; S, d)$ is pointwise *above* $\Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right]$, meaning the former is first-order dominated by the latter.

We first construct a rank-stable mapping $\bar{\mathcal{Q}}_{d_0}$ 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 $\bar{\mathcal{B}}_{d_0}(q; S, d)$ 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; note that the final argument of $\bar{\mathcal{B}}_{d_0}/\underline{\mathcal{B}}_{d_0}$ may represent out-of-sample discounts $d \neq d_0$.

¹⁵Our robust lower-bound on demand is characterized by worst-case, profit-minimizing, unobserved behavior by consumers. Note that if any of the relatively low intermediate consumers subscribe, it can only be unambiguously better for the firm, relative to the worst-case scenario.

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 that our discussion in this section focuses on identification rather than on estimation, this is somewhat of an abuse of conventional terminology, but it nevertheless is useful in distinguishing inferences that can be made directly based on d_0 , and those that can be made somewhat less directly based on alternate $d \neq d_0$.¹⁶ 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 either the inverses of CDFs, if they exist, or the quasi-inverses otherwise.

2.3. Construction of the Upper Bound. While $\bar{B}_{d_0}(q; S, d)$ lives in probability space, the functional value of $\bar{Q}_{d_0}$ 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

$$\bar{Q}_{d_0}(q; d, v) \equiv G_d^{-1}(a(q) + b(q)v), \quad a(q) \equiv \lim_{q' \rightarrow q^-} G_c(q'), \quad \text{and} \quad b(q) \equiv G_c(q) - \lim_{q' \rightarrow q^-} G_c(q'), \quad (5)$$

where v is a realization of a random variable $V \sim \text{Uniform}(0, 1)$ 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

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

The formal proof is fairly technical, but we show that the mapping $\bar{Q}_{d_0}(Q_c; d, v) \sim G_d$ (see Appendix B.1), meaning that it constitutes a data-generating process that must be consistent with observables. Moreover, we also show that $\bar{B}_{d_0}$ as defined in (6) is an upper bound on the SUDs $\Pr \left[Q_d \leq q \mid Q_c \geq \frac{S}{p_0 \times d} \right]$ for each (S, d) pair, which, combined with the previous fact implies that it is the *sharp upper bound*, because $\bar{Q}_{d_0}(q; d, v)$ is an instance within the class of admissible DGPs that attains the upper bound $\bar{B}_{d_0}(q; S, d)$.

For some brief intuition, consider the hypothetical dataset depicted in Panel (A) of Figure 1. When faced with potential RS violations due to unobserved substitution or some other phenomenon, characterizing the SUD upper bound reduces to the question of, what fraction of the aggregate shift from demand CDF G_c (solid line) to discounted demand CDF G_{d_0} (dashed line) at most could have been accounted for by high-volume 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 $\bar{Q}_{d_0}(q; d) = G_d^{-1}(G_c(q))$, whose geometric interpretation provides an answer to the question: at most, *all* of the aggregate shift represents rank-stable price responses by high-volume consumers.

¹⁶This abuse is less awkward if one considers the word “sample” as denoting a dataset with infinite observations.

Note that by construction here, the distribution of $\overline{Q}_{d_0}(Q_c; d)$ is simply the distribution of Q_d , since $\overline{Q}_{d_0}(Q_c; d)$ is chosen to match the quantiles of G_c to the corresponding quantiles of G_d . Intuitively then, the upper-bound DGP $\overline{Q}_{d_0}$ 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 below, indicates that the naive and potentially mis-specified RS model still serves as a relevant empirical benchmark for partial identification.

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 $\overline{B}_{d_0}(q; S, d_0)$ given the in-sample discount d_0 —and *out-of-sample bounds*—e.g., the SUD upper bound $\overline{B}_{d_0}(q; S, d)$ 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 2. Assumption 4 is stated in this way in order to highlight the flexibility of our partial identification approach. It covers various methods for counterfactual demand extrapolation, ranging from the basic theoretic demand system from Section 2 above, to more general demand models, and also allows for more atheoretical, reduced-form approaches to extrapolation. Examples of extrapolation methods satisfying Assumption 4 include the following:

- (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).¹⁷
- (2) Observed and counterfactual demand CDFs are consistent with a φ -separable utility model:

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

Moreover, Assumption 2 holds and there exists a (φ, u, F_θ) triple, with known φ , such that $G_d(q) = \int_{\Theta} \mathbb{1}[q^*(p_0(1 - d); \theta, \varphi) \leq q] dF_\theta(\theta)$ for all $(d, q) \in (-\infty, 1) \times \mathbb{R}_+$.¹⁸

¹⁷In our main empirical application we use extrapolation option (1) whenever extrapolations are necessary. In Appendix B.3, we discuss the more general family of utility functions in option (2), and in Appendix G, we show that our counterfactuals are insensitive to the exact form of extrapolation.

¹⁸See Sun (2023b) and Appendix G 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.

- (3) Quantile treatment effects of discounts on consumer demand are linear in the discount amount d : $G_d^{-1}(G_c(q)) = q + \frac{d}{d_0}[G_{d_0}^{-1}(G_c(q)) - q]$.

Option (2) above 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, $U(q; \theta) = u(q) + \theta q$, which can be found in much of the early theoretical principal-agent literature (e.g., see Maskin and Riley (1984) and Laffont and Tirole (1986)). The main difference between multiplicative and additive separability 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 the MS ($\varphi=1$) and AS ($\varphi=0$) models as special cases.

In options (1) and (2), the form of counterfactual extrapolation is derived from an explicit, rank-stable, structural model of demand. The reader may find this manner of extrapolation 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 (see Proposition 1), options (1) and (2) are less problematic as aids for out-of-sample inference. Assumption 4 and options (1) and (2) of Remark 2 should not be interpreted as stating directly that the underlying DGP is rank stable; rather, they simply state that the researcher is confident in a particular utility specification for producing sensible counterfactual projections. Option (2) allows one to base this confidence on a utility model with greater degrees of freedom. However, 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. We will return to the discussion of different approaches to out-of-sample extrapolation after we derive the lower-bound DGP $\underline{\mathcal{Q}}_{d_0}$ below.

We view Assumption 4 as relatively mild for two reasons. First, additional RCT variation can be used to probe its validity: given multiple experimental discounts $(d_0, d_1, \dots)$ one may test whether utility specifications like (1) or (2) produce realistic out-of-sample projections. One could even run an experiment where discounts are individually drawn from a continuum distribution in order to identify a richer extrapolation model. Second, it is important to keep in mind that Assumption 4 deals with extrapolation of *reduced-form, aggregate demand*, including both subscribers and non-subscribers, rather than with our primary object of interest, being counterfactual demand among subscribers only. Another way of articulating this distinction is that Assumption 4 directly concerns the *marginal distributions* of demand under alternate prices p_0 and $p_0(1-d)$, while our primary object of study is bounds on the *copula* between consumer i 's demand Q_{ci} and Q_{di} under those two prices. We now state our first main result, but we relegate a formal proof to Appendix B.1. $\square$

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 $\bar{B}_{d_0}(q; S, d)$ 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 $\bar{B}_{d_0}(q; S, d) \equiv \Pr \left[\bar{Q}_{d_0}(Q_c; d, v) \leq q | Q_c \geq \frac{S}{p_0 \times d} \right] \leq \Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right]$, where the function $\bar{Q}_{d_0}(q; d, v)$ is defined in (5).*

Remark 3. Out-of-Sample Inference First, the in-sample discount d_0 together with observables (G_c, G_{d_0}) and equations (5) and (6) are sufficient to pin down $\bar{Q}_{d_0}(q; d_0)$ and $\bar{B}_{d_0}(q; S, d_0)$. Then, Assumption 4 allows us to project counterfactual demand $G_d^{oos}(q; G_c, G_{d_0})$ under arbitrary discount $d \neq d_0$, 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 $\bar{Q}_{d_0}(q; d)$ and $\bar{B}_{d_0}(q; S, d)$, for arbitrary (S, d) pairs. $\square$

2.4. Construction of the Lower Bound. We now construct an analogous mapping $\underline{Q}_{d_0}(Q_c; d, v)$ 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 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 .¹⁹ 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 $\eta \rightarrow \infty$ the two services produced by L and C become perfect substitutes (holding brand loyalty fixed at $\alpha_i = 0.5$). Thus, α represents

¹⁹The 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.

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.

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

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

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

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 consumers 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).²⁰

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 Appendix A we briefly discuss five related settings—procurement, regulation of externalities, optimal taxation, 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 $\underline{Q}_{d_0}$. That is, despite firm L lacking internal data to empirically model failures of RS, we derive a sharp lower bound $\underline{B}_{d_0}(q; S, d) = \Pr \left[\underline{Q}_{d_0}(Q_c; d) \leq q | Q_c \geq \frac{S}{p_0 \times d} \right]$ 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 pre-suppose or hinge on a specific underlying model of unobserved RS violations.²¹

Our purpose here is to derive a sharp lower bound on the set of all DGPs involving RS violations consistent with the LoD and with the shapes of the observed CDFs G_c and G_{d_0} . To do so, we will need to characterize the sorts of *maximal* RS violations that respect this *a priori* information. As

²⁰In the limiting case $\eta \rightarrow \infty$, 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 $q_L = \frac{1}{\alpha_i} (u')^{-1} \left(\frac{p_L - \varepsilon}{\alpha_i \theta_i} \right)$.

²¹Other plausible phenomena producing RS violations include heterogeneous income effects and/or budget-constraint heterogeneity. Our proposed partial identification method allows for the fully-specified underlying model of RS violations to include multiple channels driving unpredictable heterogeneity in price sensitivity.

before let d_0 denote the in-sample discount, and let G_c and G_d denote the aggregate demand CDFs under prices p_0 and $p_0(1-d)$. Given subscription offer (S, d) we fix the usual strong uptaker cutoff $Q_c \geq \frac{S}{p_0 \times d}$. The LoD and the data (G_c, G_{d_0}) impose considerable discipline on the maximal masses of rank-jumpers and on the maximal plausible magnitudes of their rank-jumping behaviors. Once again, for some intuition it is informative to consider the hypothetical dataset depicted in Panel (A) of Figure 1. We learned in the previous section that the rank-stable DGP is the least upper bound on counterfactual strong-uptaker demand under discount d . Suppose now 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 two things for each strong-uptaker θ type in order for it to be profit-improving: while per-unit revenues are lower for each subscriber (due to discount d^*), their projected increase in demand volume $q^*(p_0(1-d^*); \theta)$ is large enough so that the change in total revenues to the firm is positive, or $S^* + p_0 [(1-d^*)q^*(p_0(1-d^*); \theta) - q^*(p_0; \theta)] > 0$. Ensuring the second part is true requires an accurate forecast of demand responsiveness to discount d^* .


With that in mind, one can think of the lower-bound DGP $\underline{Q}_{d_0}$ as being *maximally “adversarial”* in the sense of minimizing the naive market-designer’s profits, $\pi(S^*, d^*)$, from subscription offer (S^*, d^*) . Recall that the naively presumed DGP $\overline{Q}_{d_0}$ holds that the entirety of the shift from G_c (Panel (A) of Figure 1, solid line) to G_d (Panel (A), dashed line) represents a rank-stable demand increase, which implies maximal demand responses $(1-d^*)q^*(p_0(1-d^*); \theta) - q^*(p_0; \theta)$ by high-baseline-demand θ types who subscribe (i.e., $Q_c = q^*(p_0; \theta) \geq \frac{S^*}{p_0 \times d^*}$). In contrast, the lower-bound DGP, $\underline{Q}_{d_0}$, asks, what is the *smallest* price response $(1-d^*)q^*(p_0(1-d^*); \theta) - q^*(p_0; \theta)$ by uptaker θ types that cannot be ruled out by the LoD and data? This in turn is equivalent to *minimizing* the firm’s total revenue change from its subscription program. In the most extreme case, if strong uptakers subscribe to receive the discount, but then defy the model projection by not increasing their purchase volume at all, then the incorrectly calibrated fee schedule (S^*, d^*) merely becomes a zero-sum transfer from the firm to the consumer. This intuitive adversarial property of $\underline{Q}_{d_0}$ will be discussed at length below, but first we formalize our primary objects of interest, $\underline{Q}_{d_0}$ and $\underline{B}_{d_0}$.

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).²² We relax both of these assumptions in Online Appendix B.2, where we derive the lower bound in greater generality, but to avoid tedious complications of exposition we limit further discussion here to the simpler case. The function $\underline{Q}_{d_0}(q; d, v) : Q_c \rightarrow Q_d$ maps baseline consumption levels into minimal plausible counterfactual (discounted) consumption levels, given the LoD and data (G_c, G_{d_0}) . Since $\underline{Q}_{d_0}$ is a lower bound in the sense of minimizing price responsiveness by strong uptakers, it implies a lower-bound counterfactual demand CDF $\underline{B}_{d_0}$, or one that is weakly stochastically dominated by the true SUDs.


²²The 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’Haultfoeille and Février (2020) and Sun (2023a).

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, $G_c - G_{d_0}$.

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 $\bar{q}_{d_0}(q) = \inf\{q' \in [q_{max}^*, q_{max}]: G_c(q') - G_d(q') = G_c(q) - G_d(q)\}$.²³ 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 $V \sim Uniform(0, 1)$ denote a uniform random variable that is independent of (Q_c, Q_d) , $\underline{B}_{d_0}$ and $\underline{Q}_{d_0}$ are defined by

$$\underline{B}_{d_0}(q; S, d) \equiv \Pr \left[\underline{Q}_{d_0}(Q_c; d, V) \leq q \mid Q_c \geq \frac{S}{p_0 \times d} \right]; \text{ and } \underline{Q}_{d_0}(q; d, v) = \begin{cases} \bar{q}_{d_0}(q) & q \leq q_{min}^*, v \leq \frac{q_c(q) - g_d(q)}{g_c(q)} \\ q & \text{otherwise.} \end{cases} \quad (9)$$

Graphical intuition for how $\bar{q}_{d_0}$ and q_{d_0} are constructed (and in turn, $\underline{Q}_{d_0}$ as well) can be found in Panel (C) of Figure 1. Recall that $\underline{Q}_{d_0}$ represents a maximally adversarial DGP from the perspective of a market designer who optimized a subscription offer (S^*, d^*) assuming the rank-stable DGP $\bar{Q}_{d_0}$. Intuitively, to achieve the maximally adversarial property of the lower-bound DGP, we start at the

²³If G_c, G_d have unbounded support and $G_c(q) > G_d(q)$ for all q , let $q_{max} = \infty$.

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 consumers for all values of q' between 0 and q_{min}^* , the minimal argmax of the CDF difference.²⁴ 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} . In other words, under the DGP $\underline{Q}_{d_0}$, subscribers are minimally price responsive, which maximally exposes the naive market-designer to merely transferring money to them, with little or no compensating improvement in sales volume.

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 $Q_c < \frac{S}{p_0 \times d}$. From these two facts it follows that the firm could do no worse than if (i) everybody who consumes $Q_c \geq \frac{S}{p_0 \times d}$ under default pricing (i.e., *all* strong uptakers) buys the subscription and then consumes minimal additional volume, while (ii) nobody who consumes $Q_c < \frac{S}{p_0 \times d}$ ever buys the subscription. This constitutes the adversarial DGP, since it minimizes firm profitability for a given (S, d) pair, subject to consistency with observables and the LoD.

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 $Q(p; Q_c, Q_d) = Q_c \mathbb{1}\{p > p_0(1 - d)\} + Q_d \mathbb{1}\{p \leq p_0(1 - d)\}$,

²⁴Note 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.

where the joint distribution of (Q_c, Q_d) follows the DGP described by $\underline{Q}_{d_0}$. 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 constant for prices $p \in (p_0(1-d), p_0]$, and decreases 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 $\underline{B}_{d_0}(q; S, d)$ 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 $\Pr \left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d} \right] \leq \underline{B}_{d_0}(q; S, d)$, where the function $\underline{Q}_{d_0}(q; d, v)$ is defined in (9) when (G_c, G_d) are absolutely continuous and $G_c(q) - G_d(q)$ is unimodal, and defined in (24) in Appendix C otherwise. Moreover, $\underline{B}_{d_0}$ is sharp in the following ways:*

- (i) *The in-sample bound $\underline{B}_{d_0}(q; S, d_0)$ (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., $\underline{Q}_{d_0}(q; d_0, V) \geq q$) and the shapes of (G_c, G_{d_0}) (i.e., $\Pr[\underline{Q}_{d_0}(Q_c; d_0, V) < q] = G_{d_0}(q)$).*
- (ii) *$\underline{B}_{d_0}(q; S, d)$ 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., $\Pr[\underline{Q}_{d_0}(Q_c; d, V) < q] = G_d^{oos}(q)$) and satisfies the LoD with respect to prices p_0 and $p_0(1-d)$ (i.e., $\underline{Q}_{d_0}(q; d, V) \geq q$).*

We relegate a technical proof of Proposition 2 to Appendix B.2. In the proposition, our usage of the term *uniformly sharp* follows Firpo and Ridder (2019) and Molinari (2020).²⁵ 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:²⁶

Corollary 1. *For arbitrary (S, d) pairs, we define “small” fee-to-discount ratios $\frac{S}{p_0 \times d}$ as 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 an arbitrary subscription offer (S, d) where the fee-to-discount ratio is not small, the sharp SUD lower bound $\underline{B}_{d_0}$ is the same as the conditional control CDF of demand, given strong uptaker status, or $\underline{B}_{d_0}(q; S, d) = \Pr \left[Q_c \leq q | Q_c \geq \frac{S}{p_0 \times d} \right] = \frac{G_c(q) - G_c(S/(p_0 \times d))}{G_c(S/(p_0 \times d))} = G_c(q | Q_c \geq S/(p_0 \times d))$.*

²⁵Mourifie et al. (2020) use the term *functionally sharp* to refer to the same concept.

²⁶As we will show in the empirical application, robust policies tend to not be overly generous.

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 4. Out-of-Sample Inference First note that the in-sample discount d_0 together with observables (G_c, G_{d_0}) and equation (9) are sufficient to pin down in-sample bounds, $\underline{Q}_{d_0}(q; d_0, v)$ and $\underline{B}_{d_0}(q; S, d_0)$, under d_0 coupled with arbitrary subscription fee S . If the researcher is reticent to impose any additional model structure aside from the LoD—e.g., the extrapolation methods listed in Remark 2—then Corollary 1 and Proposition 4 (Appendix B.3) apply: in short, little inferential power beyond the in-sample bounds is possible in that case, due to lack of an obvious way to project reduced-form aggregate demand G_d . Consider first a discount $d > d_0$ that is more generous than the in-sample discount. Here we cannot rule out DGPs where an arbitrarily large fraction of consumers have virtually satiated demand under price $p_0(1 - d_0)$. Thus, $\underline{Q}_{d_0}(q; d, v) = \underline{Q}_{d_0}(q; d_0, v)$ and $\underline{B}_{d_0}(q; S, d) = \underline{B}_{d_0}(q; S, d_0)$ when $d > d_0$. If $d < d_0$ is less generous, then by similar logic 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, $\underline{B}_{d_0}(q; S, d) = G_c(q | Q_c \geq \frac{S}{p_0 \times d})$ whenever $d < d_0$.

On the other hand, a researcher comfortable leaning on some form of model structure for out-of-sample inference may still harbor concerns about whether options (1) and (2) of Remark 2 induce excessive mis-specification bias in aggregate demand projections G_d^{oos} . In that case, the reduced-form linear quantile shifts option (3), or Corollary 1 may still be of use. The former may provide a reasonable first-order approximation to reduced-form 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 $d \neq d_0$ follow similarly as in Remark 3. First, by Assumption 4, we can use in-sample demand distributions (G_c, G_{d_0}) to produce reduced-form aggregate demand projections $G_d^{oos}(q; G_c, G_{d_0})$. The counterfactual CDFs G_c and G_d^{oos} can then be plugged back into equations (9) to get $\underline{Q}_{d_0}(q; d, v)$ and $\underline{B}_{d_0}(q; S, d)$. In our empirical application, we explored inferences based on all three cases of Remark 2 and Corollary 1, and find that our main empirical conclusions are insensitive to one’s choice among these options. $\square$

Remark 5. Theory-Free Bound Without imposing the LoD, one can derive an alternate, theory-free lower bound $\tilde{\underline{B}}_{d_0}$ for the SUDs. The analogous DGP, $\tilde{\underline{Q}}_{d_0}$ 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 $\tilde{\underline{Q}}_{d_0}$ in general need not respect the LoD. In contrast, we have constructed our main bound $\underline{Q}_{d_0}(Q_c; d)$ under the heuristic that we wish to “maximize” the degree of RS violations *subject to* the LoD. The comparison between the resulting demand CDF bounds, $\tilde{\underline{B}}_{d_0}$ and $\underline{B}_{d_0}$, helps to illuminate how structure from basic economic theory delivers useful inference. We explore this idea empirically in Section 4.3. $\square$

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.²⁷ We first describe how bounds on SUDs can be translated into bounds on profits, $\pi(S, d)$, from a given subscription program. We then consider richer experimental variation including additional discounts $\mathcal{D}_d = \{d' < d'' < \dots < d^{(K-1)} < d\}$ that are *less generous* than d —i.e., $d^{(i)} \in (0, d)$, $i = 1, \dots, K-1$ —and the corresponding aggregate demand data $\mathcal{G}_{\mathcal{D}_d} = \{G_c, G_{d'}, \dots, G_{d^{(K-1)}}, G_d\}$. For the current exercise it will simplify discussion markedly to 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, $d^* = d$.²⁸

Based on our discussion in Section 2.4.2, under the assumptions of Proposition 2 the DGP $\underline{Q}_{d_0}$ satisfies the adversarial property, provided that all individuals consuming $Q_c < \frac{S}{p_0 \times d}$ under $\underline{Q}_{d_0}$ 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, 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), and obtain tighter bounds on aggregate worst-case post-subscription behavior.

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 $Q(p) = q_{d'} \mathbb{1}\{p \leq p_0(1-d')\}$; in other words, where demand is constant for prices 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. With that in mind, we can now define the 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

²⁷A 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.

²⁸To 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 $G_{0.25}^{oos}$, and then derive a refined bound based on $\mathcal{D} = \{0.10, 0.15, 0.25\}$ and $\mathcal{G}_{\mathcal{D}} = \{G_c, G_{0.10}, G_{0.15}, G_{0.25}^{oos}\}$. However, a researcher unwilling to lean on model structure that facilitates extrapolation would simply take $G_{0.25}^{oos}(q) = G_{0.15}(q)$, $\forall q$ instead.

the LoD. Characterizing this scenario is equivalent to maximizing an RSU consumer's Q_c while minimizing that same consumer's Q_d . First, if we re-define $\tilde{p}_0 = p_0(1-d')$ as the default price, and $\tilde{d} = \frac{d-d'}{1-d'}$, then Corollary 1 applies to the price change from $\tilde{p}_0$ to $\tilde{p}_0(1-\tilde{d})$ —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 $Q_c = G_c^{-1}[G_{d'}(Q_{d'})]$ 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 $\underline{Q}_{d_0}(q; d, v)$, 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* as²⁹

$$\underline{Q}_{\mathcal{D}_d}^R(q; S, d, v) \equiv \underline{Q}_{d_0}(q; d, v) \mathbb{1} \left\{ G_{d'}^{-1}[G_c(q)] < \frac{S}{p_0(d-d')} \right\} + G_{d'}^{-1}[G_c(q)] \mathbb{1} \left\{ G_{d'}^{-1}[G_c(q)] \geq \frac{S}{p_0(d-d')} \right\}. \quad (10)$$

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 $G_{d'}^{-1}[G_c(q)]$ 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 $G_{d'}^{-1}[G_c(q_{d'})] - q_{d'}$ 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 $d^{(i)} \in \mathcal{D}_d = \{d', \dots, d^{(K-1)}, d\}$, 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 aggregate these individual worst-case DGPs into a single refined DGP, $\underline{Q}_{\mathcal{D}_d}^R(q; S, d)$, consistent with *all* observables, $\mathcal{G}_{\mathcal{D}_d} = \{G_c, G_{d'}, \dots, G_{d^{(K-1)}}, G_d\}$, 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 $i^*(r) \equiv \max \left\{ i \in \{1, \dots, K\} : r > G_{d^{(i)}}\left(\frac{S}{p_0(d-d^{(i)})}\right) \right\}$, where, by convention, $\max \emptyset \equiv 0$. Thus, the refined lower-bound DGP and refined SUD bound are

$$\underline{Q}_{\mathcal{D}_d}^R(q; S, d, v) = \begin{cases} \underline{Q}_{d_0}(q; d, v) & i^*(r) = 0 \\ G_{d^{(i^*)}}^{-1}[G_c(q)] & i^*(r) \neq 0 \end{cases} \quad \text{and} \quad \underline{B}_{\mathcal{D}_d}^R(q; S, d) \equiv \Pr \left[\underline{Q}_{\mathcal{D}_d}^R(Q_c; S, d, V) \leq q | RSU(p_0, S, \mathcal{D}_d) \right], \quad (11)$$

²⁹For simplicity, equation (10) 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).

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 $G_{d^{(i)}}(q)$, $i = 1, \dots, K-1$, 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 (11) aggregates this information as the upper envelope of minimal price responsiveness, for each $r \in [0, 1]$.

Proposition 3. *For $\mathcal{D}_d = \{d', \dots, d^{(K-1)}, d\}$, if $\mathcal{G}_{\mathcal{D}_d} = \{G_c, G_{d'}, \dots, G_{d^{(K-1)}}, G_d\}$ are known and G_c is discontinuous at countably many mass points, then $\underline{\mathcal{B}}_{\mathcal{D}_d}^R(q; S, d)$ defined in (11) is identified and constitutes a lower bound on RSU demand $\Pr[Q_d \leq q | RSU(p_0, S, \mathcal{D}_d)] \leq \underline{\mathcal{B}}_{\mathcal{D}_d}^R(q; S, d)$.*

3. DATA, ESTIMATION STRATEGY, AND RESULTS

We now 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., Costco, 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 $\$S$.³⁰ 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 eliminates the need for considering behavioral mistakes in subscription purchases. In this dataset, 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 the discount 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. As a preview, the two arms of the subscriptions RCT produce structural estimates that are broadly consistent with each other. This preliminary test suggests that the baseline RS model appears consistent with the latent DGP. However, we find substantial

³⁰The 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.

differences between the shapes of the utility functions recovered from the subscriptions and uniform-discount RCTs. We interpret these differences as a rejection of rank stability in characterizing Lyft’s consumer base. In the consumer-demand context, a likely culprit is that the two RCT designs induce different unseen selection patterns as some consumers substitute consumption between Lyft and its competitors under discounted pricing. Structural estimates from this section set the stage for our robust market-design approach discussed in Section 4.

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)).³¹

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 98th percentile (monthly) q observed in our data, denoted $\bar{q}$. Thus, all plots/tables involving q represent various levels of consumption as fractions of $\bar{q}$. 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 $\bar{q}$. 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 $\mathbf{t}_n$ for which

³¹In 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

Variable	Treatment Arm: 15% Off			Treatment Arm: 25% Off		
	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: $N_{25} = 319,547$.

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 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 Experiments



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 E, 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 differences 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 G, 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: $G(q|(1-d), v=1)$, for (observed) demand under discount for uptakers, and $G(q|1, v=1)$ 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: $\hat{G}(q|1-d_a, v=1; \omega_a) \equiv \sum_{k=1}^{K_a+3} \omega_{ak} \mathcal{B}_{ak}(q)$, for $a=1, 2$, and $\hat{G}(q|1, v=1; \omega_0) \equiv \sum_{k=1}^{K_0+3} \omega_{0k} \mathcal{B}_{0k}(q)$, where the basis functions $\mathcal{B}_a/\mathcal{B}_0$ are uniquely determined by knot vectors $\kappa_a = \{\kappa_{a1} < \kappa_{a2} < \dots < \kappa_{a, K_a+1}\}$, $a=1, 2$ and $\kappa_0 = \{\kappa_{01} < \kappa_{02} < \dots < \kappa_{0, K_0+1}\}$ (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.³² For each $a=1, 2$ the B-spline forms facilitate a straightforward GMM estimator,

$$(\hat{\omega}_a, \hat{\omega}_0^a) = \arg \min_{(\omega_a, \omega_0)} \left\{ \sum_{n=1}^{N_{ua}} \left(\hat{G}(q_n|1-d_a, v=1; \omega_a) - \tilde{G}(q_n|1-d_a, v=1) \right)^2 + \sum_{n=1}^{N_{ua}} \left(\hat{G}(q_n|1, v=1; \omega_0) - \tilde{G}(q_n|1-d_a, v=1) + \frac{\tilde{G}_{t_a}(q_n) - \tilde{G}_c(q_n)}{\tilde{\tau}_a} \right)^2 \right\}, \quad (12)$$

$$s.t. \quad \omega_{ak} \leq \omega_{a, k+1}, \quad k=1, \dots, K_a+2; \quad \omega_{0k} \leq \omega_{k+1}, \quad k=1, \dots, K+2;$$

$$\omega_{0,1} \geq 0, \quad \omega_{a,1} \geq 0 \quad \omega_{0, K_0+3} = 1, \quad \omega_{a, K_a+3} = 1; \quad \text{and} \quad \hat{G}(q_n|1, v=1; \omega_0) \leq \hat{G}(q_n|1-d_a, v=1; \omega_a) \quad \forall n,$$

where objects with tildes are empirical analogs of terms on the right-hand side of equation (3): $\tilde{G}(q_n|1-d_a, v_n=1) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \leq q_n \cap t_{an'}=1 \cap v_{n'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{an'}=1 \cap v_{n'}=1)}$, $\tilde{G}_{t_a}(q_n) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \leq q_n \cap t_{an'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{an'}=1)}$, $\tilde{G}_c(q_n) = \frac{\sum_{n'=1}^{N^s} \mathbb{1}(q_{n'} \leq q_n \cap t_{0n'}=1)}{\sum_{n'=1}^{N^s} \mathbb{1}(t_{0n'}=1)}$, and $\tilde{\tau} = \frac{\sum_{n=1}^{N^s} \mathbb{1}(t_{an}=1 \cap v_n=1)}{\sum_{n=1}^{N^s} \mathbb{1}(t_{an}=1)}$. The penultimate line of the constraints represents terminal conditions

³²A 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 $G(q|1, v=1)$ is not guaranteed to be monotone in finite samples.

CDFs must satisfy, and the last line is a stochastic dominance condition implied by the LoD. Resulting GMM estimates of the two main CDFs are plotted in Figure 12 in Online Appendix D.

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: $\hat{\Upsilon}_a(q) = \frac{\hat{g}(q|1, v=1; \hat{\omega}_0^a) \hat{\tau}_a}{\hat{g}_c(q; \hat{\omega}_c)}$, $\hat{\rho}_a = \lim_{q \rightarrow \infty} \hat{\Upsilon}_a(q)$, $\hat{\delta}_a = \frac{\lim_{q \rightarrow 0} \hat{\Upsilon}_a(q)}{\hat{\rho}_a}$, and $\hat{H}_\varepsilon^a(q) = \frac{\hat{\Upsilon}_a(q + S/d_a) - \hat{\rho}_a \hat{\delta}_a}{\hat{\rho}_a - \hat{\rho}_a \hat{\delta}_a}$, where $\hat{g}_c(q; \hat{\omega}_c)$ is the derivative of a B-spline estimate of the control demand CDF.

For utility estimation, we specify a knot vector, $\kappa_u = \{\kappa_{u1} < \kappa_{u2} < \dots < \kappa_{u, K_u+1}\}$, and parameterize the utility function as a flexible quartic B-spline $\hat{u}(q; \omega_u) = \sum_{k=1}^{K_u+4} \omega_{uk} \mathcal{B}_{uk}(q)$. Recall that the CDF $G(q|1, v=1)$ is a selection-corrected analog of the (observed) treatment CDF $G(q|1-d_a, v=1)$, for an identical population of control consumers who would have purchased a subscription, had they received the offer. Within this hypothetical population, the term $\frac{\hat{G}_{ta}(q_n) - \hat{G}_c(q_n)}{\hat{\tau}_a}$ represents quantile-specific demand shifts under discount d_a . For treatment arm a define $\tilde{\mathcal{T}}_a(q_d) \equiv \hat{G}^{-1} \left[\hat{G}(q_d|1-d, v=1; \hat{\omega}_a) | 1, v=1; \hat{\omega}_0^a \right]$. 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 ω_u , based on equation (2) we can define $\hat{\mathcal{T}}(q_d; \omega_u) \equiv (u')^{-1} \left[\frac{1}{1-d} u' \left(q^* [1-d; \theta^*(1-d, q_d; \omega_u); \omega_u]; \omega_u \right) \right] = (u')^{-1} \left(\frac{1}{1-d} u'(q_d; \omega_u); \omega_u \right)$, where $\theta^*(1-d, q_d; \omega_u)$ 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 ω_u by minimizing the l_2 distance between $\tilde{\mathcal{T}}$ and $\hat{\mathcal{T}}$. Letting $0 = q_1 < q_2 < \dots < q_m = q_{max}$ be an evenly spaced grid of points, we have:

$$\begin{aligned} \hat{\omega}_u^a = \arg \min_{\omega_u} \sum_{j=1}^m \left(\hat{\mathcal{T}}(q_j; \omega_u) - \tilde{\mathcal{T}}_a(q_j) \right)^2 \\ \text{s.t. } \omega_{u1} = 0, \quad \omega_{u1} = \frac{\kappa_{u5} - \kappa_{u2}}{3}, \quad \omega_{uk} \leq \omega_{u, k+1} - \epsilon, \quad k = 1, \dots, K_u + 3, \quad \epsilon > 0, \quad \text{and} \\ \frac{\omega_{uk} - \omega_{u, k-1}}{\kappa_{u, k+3} - \kappa_{uk}} \leq \frac{\omega_{u, k+1} - \omega_{uk}}{\kappa_{u, k+4} - \kappa_{u, k+1}} - \epsilon, \quad k = 2, \dots, K_u + 3, \quad \epsilon > 0, \end{aligned} \quad (13)$$

where the first constraint is a boundary condition $u(0; \omega_u) = 0$, the second constraint is a scale normalization $u'(0; \omega_u) = 1$, and the third and fourth enforce monotonicity and concavity.^{33,34}


3.3.3. Empirical Results. In Figure 3, we plot results 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. We display

³³The 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.

³⁴Knots are chosen so that $\kappa_c = \kappa_0 = \kappa_1 = \kappa_2 = \kappa_u$ 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 (13) evaluations should be $m \geq K_u$; we chose $m=50$. For the tolerance on the shape constraints we chose $\epsilon = 10^{-6}$.

the utility function, the elasticity function, and the uptake function. 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 $H_0: \hat{u}_u(q; \hat{\omega}_u^1) = \hat{u}_u(q; \hat{\omega}_u^2)$ that the two estimated utility functions are the same, since $\hat{u}_u(q; \hat{\omega}_u^2)$ (and its confidence bounds) lay entirely within the confidence bounds of $\hat{u}_u(q; \hat{\omega}_u^1)$. Taken together, these two observations serve 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.

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



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

Moreover, while the comparison is noisy, differences between the uptake functions for the 15% and the 25% 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. Table 3 reports structural behavioral parameters implied by the uptake functions. We find that estimates from the two subscription-RCT treatment arms give similar results.

TABLE 3. Subscription Uptake Parameters

Parameter	15% Off	95% CI	25% Off	95% CI	Joint	95% CI
Saliency 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.

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

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 $\bar{q}$, 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.³⁵

We also estimate the model with a single utility function fitted to data pooled across both treatment arms: $\hat{\omega}_u = \arg \min_{\omega_u} \{w_1 \sum_{j=1}^m (\hat{T}(q_j; \omega_u) - \tilde{T}_1(q_j))^2 + w_2 \sum_{j=1}^m (\hat{T}(q_j; \omega_u) - \tilde{T}_2(q_j))^2\}$, 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 $\hat{u}(q_j; \hat{\omega}_u^a)$ (across $\{q_1, \dots, q_m\}$). The results are plotted in Figure 4. These look most similar to estimates from the 25% off data alone, but with tighter confidence bounds. We also jointly estimate behavioral parameters from the pooled data; results are displayed in Table 3.

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



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 the analysis is more straightforward because all treated consumers get the same uniform 10% discount. Thus, the uniform discount estimator is a special case of the previous one, where the RS condition is unencumbered by uptaker status; i.e., $\Upsilon(q) = 1$ 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: $G(q|1-d, v=1) = G_t(q)$ and $G(q|1, v=1) = G_c(q)$.

³⁵E.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.

Thus, we first smooth the empirical CDFs with B-splines similarly as before:

$$\begin{aligned}
(\hat{\omega}_d, \hat{\omega}_0) = & \arg \min_{(\omega_d, \omega_0)} \left\{ \sum_{n=1}^{N_d^{ud}} \left(\hat{G}(q_n|1-d; \omega_d) - \frac{\sum_{n'=1}^{N_d^{ud}} 1(q'_n \leq q_n)}{N_d^{ud}} \right)^2 + \sum_{n=1}^{N_c^{ud}} \left(\hat{G}(q_n|1; \omega_0) - \frac{\sum_{n'=1}^{N_c^{ud}} 1(q'_n \leq q_n)}{N_c^{ud}} \right)^2 \right\} \\
& s.t. \quad \omega_{dk} \leq \omega_{d,k+1}, \quad k = 1, \dots, K_d + 2; \quad \omega_{0k} \leq \omega_{0,k+1}, \quad k = 1, \dots, K_0 + 2, \\
& \omega_{0,1} \geq 0, \quad \omega_{d,1} \geq 0 \quad \omega_{0,K_0+3} = 1, \quad \omega_{d,K_d+3} = 1; \quad \text{and} \quad \hat{G}(q_n|1; \omega_0) \leq \hat{G}(q_n|1-d; \omega_d) \quad \forall n = 1, \dots, N_c^{ud}.
\end{aligned}$$

After specifying a flexible B-spline utility function $\hat{u}(q; \omega_u)$, Stage-2 estimation follows Section 3.3.2 using the GMM estimator (13).³⁶ 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 .³⁷

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

Remark 6. *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. $\square$


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.³⁸ Why? These differences reflect the results of another, more stringent test of the basic RS model of consumer demand. While the test of over-identifying restrictions in Section 3.3.3—based on comparisons across estimates from different arms of the *same* subscription experiment—were unable to reject the model, the results in this section—comparing estimates from *two distinct* pricing RCT designs—do reject it. Taking insights from the model we proposed in Section 2.4.1, if unobserved substitution and heterogeneity in brand loyalty are present in the DGP, then 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 that we see in the baseline elasticity estimates.

³⁶We again choose knots uniformly on the support $[0, \bar{q}]$, 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 $m = 50$.

³⁷ 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 $q = 0$.

³⁸In Appendix E, 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.

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



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

To be useful as a market-design tool these benchmark estimates must be adjusted according to the magnitude of RS violations in the data. Intuitively, if unobserved consumer substitution is very 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 maximize profitability. If unobserved substitution behavior is less prevalent, then the baseline model presented here will prescribe a subscription (S, d) that is closer to the true optimum. In datasets our methodology is designed for—rich internal firm data with no market-level demand shares—these concerns are impossible to control for directly. Therefore, using current structural estimates as a baseline reference point, we implement robust market-design 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 three ideas.³⁹ First, differences between the uniform-discount and subscription RCTs are largely driven by selection effects due to heterogeneity in mistake making. Second, we argue (Section 3.3.3) that it is feasible in the long run to mitigate consumer mistakes through experience and firm marketing/information interventions.⁴⁰ Third, since firms may have a hard time changing its subscription program 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,

³⁹Note 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 $\bar{q}$) from our results.

⁴⁰Our 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.

plausible DGPs. We also propose a simple data-driven method to estimate the degree of RS violations within the firm’s latent DGP. This enables us to produce a single model prescription that is robust to the worst-case unobserved behavior among the firm’s consumer base.


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., Akerberg et al. (2007), MacKay and Miller (2021)). See Online Appendix H.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. The interested reader can find additional technical details in Online Appendix H.2. 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 $\varepsilon(\theta) = -\frac{\theta f_\theta(\theta)}{1-F_\theta(\theta)}$ 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, henceforth 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, $\bar{\theta}$, buys a subscription where marginal price equals marginal cost: $p(\bar{\theta})=c$. This will be a key detail, moving forward.

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 either 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

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.

favor of a single optimal subscription offer? 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.⁴¹ The extreme lower-right corner is shaded much more darkly than the extreme 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 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 using our baseline estimates derived under the ideal but potentially faulty RS assumption, we now turn to our robust


⁴¹There 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 F for discussion.

FIGURE 7. Robust Bounds


(A) SUD Bounds



(B) Refined SUD Bounds



(C) SUD Elasticity Bounds



Notes: Panel (A) plots empirical bounds, $\bar{B}_{d_0}(q; S^*, d^*)$ and $\underline{B}_{d_0}(q; S^*, d^*)$, and the theory-free lower bound, $\tilde{B}_{d_0}(q; S^*, d^*)$. Panels (B) and (C) plot refined lower bounds $\underline{B}_{d^*}^R(q; S^*, d^*)$ 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., $\mathcal{E}^*(q_c) = \frac{\mathbb{E}[\varepsilon(Q_c)Q_c | Q_c \geq q_c]}{\mathbb{E}[Q_c | Q_c \geq q_c]}$, where $q_c = \frac{S}{p_0 \times d^*}$.

bounds analysis. While this approach admits a variety of interpretations, an especially salient motivating example (Section 2.4.1) is unobserved substitution between Lyft and competing transportation services. Recall from Sections 2.3 and 2.4 that characterizing sharp bounds on counterfactual demand under (potentially out-of-sample) discounts d hinges only on knowledge of the observable reduced-form demand CDFs ($\hat{G}_c(q), \hat{G}_{d_0}(q)$) which are estimable from raw data. We plot the empirical bounds in Figure 7— $\hat{\bar{B}}_{d_0}(q; S^*, d^*)$ as outlined in Remark 3 and $\hat{\underline{B}}_{d_0}(q; S^*, d^*)$ as outlined in Remark 4—which are interpretable as conditional demand CDF bounds for strong uptakers (i.e., consumers for whom $Q_c \geq \frac{S^*}{p_0 \times d^*}$) under the single RS-optimal subscription offer (S^*, d^*) . In our particular empirical application, the lower bound $\tilde{B}_{d_0}(q; S^*, d^*)$ is the same for all three forms of extrapolation considered in Remark 2. This is related to the fact that the conditions of Corollary 1 being satisfied, and thus the lower bound depends *only* on the in-sample CDF, G_c .

For comparison, we also plot what the theory-free SUD lower bound $\tilde{B}$ (dotted line) would have been had we not imposed our most basic assumption, the Law of Demand (LoD). This theory-free


bound is totally uninformative because it places too few restrictions on the latent DGP. On the other hand, our lower bound $\underline{\mathcal{B}}_{d_0}$, which does respect the LoD, is quite informative.⁴² This finding demonstrates the inferential power to be had from the most basic behavioral assumptions within an incomplete structural model of demand. In that sense, 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 consumers are not overly cavalier with their bidding, and they do not leave obvious gains from trade on the table. The main difference between our method and previous work is that we focus on directly bounding counterfactual demand, rather than bounding the structural primitive distribution F_θ .

Panels (B) and (C) of Figure 7 illustrate lower-bound refinements under hypothetically augmented experimental price variation (Section 2.5). We use the point-identified RS model to simulate additional sets of reduced-form demand CDFs under intermediate discount levels $d \in \mathcal{D}_{d^*}$, for three increasingly rich sets $\mathcal{D}_{d^*}^1 = \{\frac{d^*}{2}, d^*\}$, $\mathcal{D}_{d^*}^2 = \{\frac{d^*}{3}, \frac{2d^*}{3}, d^*\}$, and $\mathcal{D}_{d^*}^3 = (0, d^*]$. The final case characterizes (hypothetical) maximal continuous price variation between p_0 and $p_0(1 - d^*)$. We use the simulated datasets, $\mathcal{G}_{\mathcal{D}_{d^*}^1} = \{G_c, G_{\frac{d^*}{2}}, G_{d^*}\}$, $\mathcal{G}_{\mathcal{D}_{d^*}^2} = \{G_c, G_{\frac{d^*}{3}}, G_{\frac{2d^*}{3}}, G_{d^*}\}$, and $\mathcal{G}_{\mathcal{D}_{d^*}^3} = \{G_c; G_d(q|d), d \in (0, d^*)\}$, to compute bound refinements as an illustration of how additional exogenous price variation enhances identification by ruling out various worst-case (for profits) DGPs. Panel (B) compares bounds under the actual data, $\overline{\mathcal{B}}_{d_0}(q; S^*, d^*)$ and $\underline{\mathcal{B}}_{d_0}(q; S^*, d^*)$, to hypothetical bound refinements $\underline{\mathcal{B}}_{\mathcal{D}_{d^*}}^R(q; S^*, d^*)$ under the actual *and* simulated data. Panel (C) compares analogous bounds on demand elasticities. Although these bound refinements (thin lines) are merely based on simulated additional data—being therefore artifacts of our finite-sample CDF estimates $(\hat{G}_c, \hat{G}_{d_0})$ —they are useful for characterizing the partial identification frontier under unknown RS violations for a researcher who can design a pricing experiment with many treatment arms.

Figure 7 provides key insights on identification. Refined lower bounds under $\mathcal{D}_{d^*}^1$ and $\mathcal{D}_{d^*}^2$ cross because these sets are non-nested. Conversely, the limiting refined bound under $\mathcal{D}_{d^*}^3 = (0, d^*]$ is everywhere below the other two, because $\mathcal{D}_{d^*}^j \subset \mathcal{D}_{d^*}^3$, $j = 1, 2$. More broadly, while structural methods are often critiqued for opaqueness on the relation between empirical moments and model primitives, the Figure clearly illustrates how theory and raw data combine to deliver identifying power. 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 dash-dot lines in Panels (B) and (C) represent inference derived from richer data, holding model structure fixed. Finally, gaps between the thin dash-dot lines and


⁴²Bound 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 8. Robustness Tests



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., $\frac{\pi^{RS}(S^*(\lambda), d^*(\lambda)) - \pi^{RS}(0,0)}{\pi^{RS}(S^*(0), d^*(0)) - \pi^{RS}(0,0)}$.

FIGURE 9. Path of Optimal Single Subscriptions



Notes: Discounts are expressed as a fraction of the firm's markup under p_0 . Upfront fees S are expressed as a fraction of the maximal upfront fee in the RS-optimal continuum menu. This figure plots evolution of robust optimal subscriptions as λ varies. The black dot is the RS-optimal subscription offer, and the black triangle is $(S^*(0.16), d^*(0.16))$ for the calibrated value of $\lambda = 0.16$.

the thick solid lines in panels (B) and (C) represent 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 compute adjustments to the RS-optimal policy (S^*, d^*) from Figure 6 that add robustness against RS violations; we do so by computing profit bounds implied by the DGPs $\underline{\mathcal{B}}_{d_0}$ and $\overline{\mathcal{B}}_{d_0}$ depicted in Panel (A) of Figure 7. 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 $\pi^A(S, d)$ for each (S, d) pair, and we can compare this to profits implied by the RS model, denoted $\pi^{RS}(S, d)$. Recall from Sections 2.3 and 2.4 that, by construction, these two quantities bound true profits: $\pi^A(S, d) \leq \pi(S, d) \leq \pi^{RS}(S, d)$. π^{RS} corresponds to an ideal scenario where the basic model is not mis-specified, and $\pi^A(S, d) < \pi^{RS}(S, d)$ is a lower bound on profits when unseen RS violations wreak maximum havoc.

In order to further refine market-design decisions, we can interpolate between these extremes by considering λ -adversarial profits $\pi^\lambda(S, d) \equiv \lambda\pi^A(S, d) + (1-\lambda)\pi^{RS}(S, d)$ if fraction λ of consumers behaved according to the adversarial DGP, while $(1-\lambda)$ behaved 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, while λ fraction exhibit low loyalty, frequently comparing prices for Uber and Lyft, and then substituting adversarially across the two platforms. Optimizing with respect to π^λ instead of π^A allows for improved robust decisions if the market designer believes that the pure adversarial DGP is overly pessimistic. This is conceptually similar 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. Aryal and Kim (2013) apply a similar approach to partially identified auction models.

For each $\lambda \in (0, 1)$, we find the robust optimum $(S^*(\lambda), d^*(\lambda))$ via grid search. Figure 9 plots the $(S^*(\lambda), d^*(\lambda))$ locus; Panel (A) focuses on values of $\lambda \in [0, 0.35]$, while Panel (B) includes $\lambda > 0.35$ also. Figure 8 explores profit implications. Panel (A) characterizes a non-RS world, where the market designer is or isn’t correcting for RS violations. The solid line is robust excess profits above linear pricing, $\pi^\lambda(S^*(\lambda), d^*(\lambda)) - \pi^\lambda(0, 0)$, relative to baseline excess profits $\pi^{RS}(S^*(0), d^*(0)) - \pi^{RS}(0, 0)$. The dashed line depicts how naive excess profits under a fixed (S^*, d^*) offer, $\pi^\lambda(S^*(0), d^*(0)) - \pi^\lambda(0, 0)$, vary with λ , relative to baseline excess profits. The two lines are mechanically close to each other on the left, since λ controls the magnitude of the RS-violation problem, by construction. Eventually, the dashed line goes negative, whereas the solid line, by construction, does not for any value of λ . Panel (B) considers fixing a non-existent problem in a RS world where an over-cautious market designer chooses a robust offer $(S^*(\lambda), d^*(\lambda))$ anyway. There, we plot “paranoid” excess profits, $\pi^{RS}(S^*(\lambda), d^*(\lambda)) - \pi^{RS}(0, 0)$, relative to true excess profits.

First, note some subtle intuition behind the profit discontinuity at roughly $\lambda = 0.35$. On one hand, there is a set of contracts (S, d) that are good for profits near the RS DGP (i.e., for λ near zero), but non-robust and very bad for profits near the adversarial DGP (i.e., for λ near one). On the other hand, there is another set of contracts disjoint from the first one—having higher S and lower d —that are very robust and profit-optimal near the adversarial DGP, but generally rendering substantially lower profits. Because the corresponding optima are so different, for $\lambda < 0.35$ the market designer prefers to prioritize profitability, and for $\lambda > 0.35$ she prefers to prioritize 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., a less generous subscription plan) both help the firm hedge against profit shocks from unseen adverse consumer behavior. In Panel (B), we zoom out and show how the path jumps discontinuously to a new region of (S, d) space with low profits when λ crosses the threshold. The lesson from Figures 8 and 9 is that nonlinear pricing via subscription offers is only a viable strategy when the adversarial fraction of consumers is below roughly one third.

Prior to the phase change, however, Figure 9 suggests that policy prescriptions from the RS model are fairly robust to moderate perturbations in the underlying DGP. We find that $S^*(0.35)$ is about 9% higher than $S^*(0)$, and $d^*(0.35)$ is about 18% lower than $d^*(0)$. Table 4 depicts implications

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

λ	Strong Uptaker Percentiles	0.1	0.25	0.5	0.75	0.9	Total
$\lambda = 0.35$	$SU(p_0, S^*(\lambda), d^*(\lambda))$	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
$\lambda = 0.16$	$SU(p_0, S^*(\lambda), d^*(\lambda))$	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 $\frac{q(r)d^*(\lambda) - S^*(\lambda)}{q(r)d^*(0) - S^*(0)}$, 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 $\frac{\int_0^1 q(r)d^*(\lambda) - S^*(\lambda) dr}{\int_0^1 q(r)d^*(0) - S^*(0) dr}$.

of moving from the naive optimum to $(S^*(0.35), d^*(0.35))$ for consumer surplus (i.e., subscriber savings). Among the set of strong uptakers relative to the robust optimum, $SU(p_0, S^*(\lambda), d^*(\lambda))$, the median (90th percentile) consumer retains half (70%) 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. Finally, Panel (B) of Figure 8 shows that the cost of achieving robustness is fairly low. If the true DGP really is RS but the market designer assumes $0 \leq \lambda \leq 0.35$, then robust profits $\pi^\lambda(S^*(\lambda), d^*(\lambda))$ account for 80% or more of true optimal profits $\pi^{RS}(S^*(0), d^*(0))$. Thus, if we view robust policy design as insurance against making large errors, then hedging against that risk comes relatively cheaply for relevant values of λ where subscriptions are a viable business strategy at all.

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 $\underline{Q}_{d_0}(q)$ 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 $\bar{Q}_{d_0}$ over the adversarial DGP $\underline{Q}_{d_0}$. 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 $\{q_n^{pre}, \underline{Q}_{d_0}(q_{cn}, v)\}_{n=1}^{N_c^{ud}}$. 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 5 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 $S^*(0.16)$ at 2.0% above the naive optimum fee, and $d^*(0.16)$ at 8.0% below the naive optimum discount. This empirical case study demonstrates that our proposed methodology facilitates effective empirical market design, despite data limitations like unseen confounding choices by agents. We find that nonlinear pricing policies do indeed exist that are both profit-improving for the principal, while also being robust against worst-case, unobserved agent behavior.

5. CONCLUSION

In this paper, we provide 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 agent 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 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’ 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.

⁴³Note that 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.

REFERENCES

- Akerberg, 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.
- 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’Haultfoeuille, 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.
- 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.

- Fillmore, I. and Gallen, T. (2019). Heterogeneity in talent or in tastes? implications for redistributive taxation. *Working Paper*.
- 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.
- Leslie, P. (2004). Price discrimination in Broadway theater. *Rand Journal of Economics*, 35(3):520–541.
- 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.
- Mourifie, I., Henry, M., and Meango, R. (2020). Sharp bounds and testability of a roy model of stem major choices. *Journal of Political Economy*, 128(8):3220–3283.
- 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.

APPENDIX A. EXAMPLES OF POTENTIAL RS VIOLATIONS IN ADVERSE-SELECTION SETTINGS

A.1. Procurement. Laffont and Tirole (1986) provides a model of optimal procurement for a public project where the single dimension of heterogeneity is the efficiency β_i of the i^{th} firm while the cost of effort $\psi(e)$ is assumed to be homogeneous across firms. In such settings, common market design problems include derivation of an optimal menu of incentive contracts to induce firms (with privately known efficiency β_i) to minimize the cost of completing a project. Formally, a type- β_i firm chooses effort level e so that the cost of completing the project is $\beta_i - e$. The principal cannot directly observe effort e or type β_i , so she must set a reimbursement schedule $T(\beta_i - e)$ to induce effort on the part of the firm, who chooses effort to solve $\max_e T(\beta_i - e) - \psi(e)$.

Suppose, however, that effort costs are actually heterogeneous and given by $\psi(e; \alpha_i)$ with $\frac{\partial^2 \psi(e; \alpha_i)}{\partial e^2}$ increasing in α_i . Intuitively, β_i parameterizes vertical shifts in the firm's effort supply curve while α_i parameterizes rotations of the supply curve. Under this richer model, if two monopolists, given the same status quo incentive contract, produce at the same cost, the one with a higher α_i type will be less responsive to changes in incentives.

A.2. Regulation. In the environmental regulation model of Kang and Silveira (2021), firms are heterogeneous with respect θ_i , which parameterizes the degree to which they derive private benefits from negligently ignoring pollution regulations. A type- θ_i firm who chooses negligence level a receives gross utility $\theta b(a)$. 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 $\max_a \theta_i b(a) - \sum_{k=0}^{\infty} \epsilon(k) \frac{a^k}{k!}$.

A potential second source of unobserved heterogeneity in this context might arise if adjustment to new regulations requires significant investment from firms, and different firms start from different baseline levels of negligence. For instance, suppose that firms differ in the negligence of their current business practices $\bar{a}_i$ and the cost of complying with regulation is given by $\theta_i b(a) - \frac{C}{2}(a - \bar{a}_i)^2$. Under this richer model, suppose that at some baseline fine schedule $\epsilon(k)$, two firms are observed to make the same abatement decision a . Then the firm with the larger level of $\bar{a}_i$ must have a correspondingly lower θ_i and hence will be less responsive to changes in the fine schedule.

A.3. Income Taxation. In the optimal taxation model of Mirrlees (1971), workers are heterogeneous with respect to productivity W_i but have homogeneous preferences over consumption and labor $u(c, l)$. Given an income tax schedule $T(lW_i)$, a type- W_i worker chooses labor supply l to solve $\max_l u(c, l)$ s.t. $c \leq lW_i - T(lW_i)$. In reality, we might expect that different people, given the same wage, will nonetheless work different amounts, due to different preferences for leisure. Fillmore and Gallen (2019) parameterize heterogeneity in willingness to work by $u(c, l; \alpha_i) = v_1(c) - \alpha_i v_2(l)$. Higher values of α_i then correspond (all else equal) to higher marginal disutility of labor. Suppose

that under a status-quo tax policy, two workers have the same income, $W_i l$, but worker 1 has a higher value of α_i than worker 2. In this situation, worker 1 will exhibit a lower income elasticity of taxation relative to worker 2.

A.4. Labor Supply. D’Haultfœuille and Février (2020) introduces a model of labor supply where workers are heterogeneous with respect to their costs effort, θ_i . Specifically, a worker with type θ_i who supplies y units of effort given wages w_0 gets paid yw_0 , but incurs a utility cost of this effort equal to $\theta_i c(y)$, so her net utility is given by $yw_0 - \theta_i c(y)$. In this model, there is no scope for productivity differences, but Cotton et al. (2023) and Hedblom et al. (2022) consider settings where workers additionally vary in their productivity α_i . Thus, a worker of type (α_i, θ_i) may receive utility equal (on average) to $\alpha_i yw_0 - \theta_i c(y)$. This model at first glance *seems* isomorphic to the baseline one-dimensional model, since $\arg\max_y \alpha_i yw_0 - \theta_i c(y) = \arg\max_y yw_0 - \frac{\theta_i}{\alpha_i} c(y)$, and hence, type- (α_i, θ_i) individuals are observationally equivalent to type $(1, \theta_i/\alpha_i)$ individuals with respect to their choice of effort y . However, in many labor supply contexts, the econometrician may only realistically have access to data on total output, $\alpha_i y$ and can only indirectly observe effort y .⁴⁴ The partial identification approach in this paper may allow for the analysis of similar models using typical, observational data which does not contain sufficient observables to distinguish θ_i from α_i .

A.5. Insurance Demand. In the baseline adverse selection model of insurance demand, consumers are heterogeneous in their propensity for health spending λ_i . For example, consider a 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 $\max_s u(s - \lambda_i) - T(s)$.

The basic adverse selection model implicitly imposes homogeneity in the extent to which health-care consumption can improve an individual’s health. In reality, some individuals are more at risk for diseases which are cheaper to treat while other individuals are more at risk for diseases which can only be treated at a high price. This can be modeled as 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 . The authors are able to identify a multi-dimensional model of heterogeneity in their setting because individuals are offered multiple plans, and selection of one plan over another is informative about which type an individual could

⁴⁴Cotton et al. (2023) and Hedblom et al. (2022) address this issue by a field-experimental data collection procedure that is informative about both y and $\alpha_i y$.

be. Such a strategy requires fairly rich data, and might not be available if the analyst only has access to fully random assignment of insurance plans, such as in the RAND health insurance experiment (Aron-Dine et al. (2013)). Interestingly, in this setting, fully exogenous randomization may even be antithetical to the goal of identifying structural primitives of a model, due to rank-stability violations arising out of multi-dimensional heterogeneity.

APPENDIX B. PROOFS OF PROPOSITIONS 1 AND 2 (UNDER STANDARD ASSUMPTIONS)

In Appendix B.1, we complete the proof of Proposition 1 by modifying our construction of $\overline{Q}_{d_0}(q)$ to allow for the presence of mass points and show that our modified DGP attains the bound corresponding to $\overline{B}_{d_0}(q; S, d)$. In Appendix B.2, we complete the proof of Proposition 2 by constructing $\underline{Q}_{d_0}(q)$ and showing that it attains the bound corresponding to $\underline{B}_{d_0}(q; S, d)$ in the case where the CDF difference $G_{d_0}(q) - G_c(q)$ is quasi-concave, and both distributions are absolutely continuous. Proofs of the lower bound for the cases where the CDF difference is not quasi-concave and where mass points exist build on the basic ideas here, but involve tedious technical complications, so we defer them to Online Appendix C.

As in the main text, let the CDF of potential outcomes in control, Q_c be G_c and the CDF of potential outcomes under some discount d , Q_d , be denoted G_d . As in the body of the paper, we denote random variables by upper-case letters, while realizations of random variables (or fixed numbers) are denoted by lower-case. Additionally, 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 either the inverses of the CDFs, if they exist, or the quasi-inverses otherwise. Throughout this appendix, we maintain the assumption that the underlying data-generating process satisfies the Law of Demand (LoD), and that the econometrician has access to a dataset (G_c, G_{d_0}) , where $G_c(q)$ is observed demand under default price p_0 , and $G_{d_0}(q)$ is observed demand under a particular discounted price $p_0(1-d_0)$. The econometric challenge here essentially stems from the fact that the copula of the joint distribution of (Q_c, Q_{d_0}) is unknown. This is because for each consumer we only observe *either* Q_c *or* Q_{d_0} (but never both), and (G_c, G_{d_0}) were generated from two separate samples of consumers having similar distributions of unobserved taste characteristics.

B.1. Proof of Proposition 1 in the Presence of Mass Points in G_c . Recall from equation (5) that $\overline{Q}_{d_0}(q, v; d) \equiv G_d^{-1}(a(q) + b(q)v)$, where V is an independent uniform random variable, $a(q) \equiv \lim_{q' \rightarrow q^-} G_c(q')$ is the mass of consumers with baseline demand strictly below q , and $b(q) \equiv G_c(q) - \lim_{q' \rightarrow q^-} G_c(q')$ is the size of the mass point at $Q_c = q$. Intuitively, V is a device for “breaking ties” in rank that arise when a positive mass of consumers have the same baseline demand q . When G_c is left-continuous at a particular q then $b(q) = 0$, and the upper-bound DGP reduces to the simpler form $\overline{Q}_{d_0}(q, v; d) = G_d^{-1}(G_c(q))$. We break up our proof into two steps as follows.

Lemma 1. $\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0)$ as defined in Equation (5) is an admissible DGP that cannot be ruled out by the dataset (G_c, G_{d_0}) ; that is $\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0) \leq q] = G_{d_0}(q)$.

Proof. At any q where the quantile function G_c^{-1} is strictly increasing at $G_{d_0}(q)$, we have that $\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0) \leq q] = \Pr[G_d^{-1}(G_c(Q_c)) \leq q] = \Pr[G_c(Q_c) \leq G_d(q)] = G_d(q)$ where the last equality follows because $G_c(Q_c)$ is a *Uniform*(0, 1) random variable. Otherwise, the random variable Q_c has a mass point at $q_q^* \equiv G_c^{-1}(G_{d_0}(q))$. By definition of a , $a(q_q^*) = \Pr[Q_c < q_q^*]$, so

$$\begin{aligned} \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0) \leq q] &= \Pr\left[Q_c < q_q^* \text{ or } \left(Q_c = q_q^* \text{ and } V \leq \frac{G_{d_0}(q) - a(q_q^*)}{b(q_q^*)}\right)\right] \\ &= \Pr[Q_c < q_q^*] + \Pr\left[Q_c = q_q^* \text{ and } V \leq \frac{G_{d_0}(q) - a(q_q^*)}{b(q_q^*)}\right] \\ &= a(q_q^*) + \Pr[Q_c = q_q^*] \times \Pr\left[V \leq \frac{G_{d_0}(q) - a(q_q^*)}{b(q_q^*)}\right] \\ &= a(q_q^*) + b(q_q^*) \frac{G_{d_0}(q) - a(q_q^*)}{b(q_q^*)} = G_{d_0}(q). \quad \square \end{aligned} \tag{14}$$

Lemma 2. $\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q']$ constitutes an upper bound (in the first-order dominance sense) on Strong Uptaker Distributions; that is, $\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q'] \leq \Pr[Q_{d_0} \leq q | Q_c \geq q']$.

Proof. Consider the joint distribution of (Q_c, Q_d) with marginal distributions G_c and G_d . The upper bound property is equivalent to $\Pr\left[Q_d \leq q | Q_c \geq \frac{S}{p_0 \times d}\right] < \Pr\left[\overline{\mathcal{Q}}_{d_0}(Q_c; d, v) \leq q | Q_c \geq \frac{S}{p_0 \times d}\right]$ being impossible. Suppose for a contradiction that there exists a baseline consumption level $q' = \frac{S}{p_0 \times d}$ (under price p_0), and a counterfactual consumption level q (under price $p_0(1-d)$) satisfying this inequality. In that case,

$$\begin{aligned} &\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c \geq q'] \Pr[Q_c \geq q'] + \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c < q'] \Pr[Q_c < q'] \\ &= \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q] = \Pr[Q_d \leq q] \\ &= \Pr[Q_d \leq q | Q_c \geq q'] \Pr[Q_c \geq q'] + \Pr[Q_d \leq q | Q_c < q'] \Pr[Q_c < q'], \\ &\Rightarrow \Pr[Q_c < q'] (\Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c < q'] - \Pr[Q_d \leq q | Q_c < q']) \\ &\quad = \Pr[Q_c \geq q'] (\Pr[Q_d \leq q | Q_c \geq q'] - \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c \geq q']), \end{aligned} \tag{15}$$

where the first and third equalities follow from the law of total probability and the second follows from (14). Our supposition $\Pr[Q_d \leq q | Q_c \geq q'] < \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c \geq q']$ is thus equivalent to

$$\Pr[Q_d \leq q | Q_c < q'] > \Pr[\overline{\mathcal{Q}}_{d_0}(Q_c; d, V) \leq q | Q_c < q'], \tag{16}$$

since the last two lines of (15) have the same sign. By definition, $\overline{\mathcal{Q}}_{d_0}$ is non-decreasing; as a result, inequality (16) implies $\overline{\mathcal{Q}}_{d_0}(q'; d, v) = \overline{\mathcal{Q}}_{d_0}\left(\frac{S}{p_0 \times d}; d, v\right) > q$. To see why, note that if $\overline{\mathcal{Q}}_{d_0}(q'; d, v) \leq q$, then the conditioning event $Q_c \leq q'$ implies $\overline{\mathcal{Q}}_{d_0}(Q_c; d, v) \leq q$ as well, by monotonicity of $\overline{\mathcal{Q}}_{d_0}$. This in turn implies $\Pr\left[\overline{\mathcal{Q}}_{d_0}(Q_c; d, v) \leq q | Q_c \leq \frac{S}{p_0 \times d}\right] = 1$. But this would violate

(16) since $\Pr \left[Q_d \leq q | Q_c < \frac{S}{p_0 \times d} \right]$ cannot exceed 1. In other words, (16) requires that counterfactual consumption implied by the baseline consumption level $q' = \frac{S}{p_0 \times d}$ must weakly exceed the benchmark q . Furthermore,

$$\begin{aligned} G_d(q) &= \Pr [Q_d \leq q] = \Pr [Q_d \leq q | Q_c < q'] \Pr [Q_c < q'] + \Pr [Q_d \leq q, Q_c \geq q'] \\ &> \Pr [\bar{Q}_{d_0}(Q_c; d, V) \leq q | Q_c < q'] \Pr [Q_c < q'] + \Pr [Q_d \leq q, Q_c \geq q'], \end{aligned} \quad (17)$$

where the equality follows from the law of total probability and the inequality follows from (16). This last expression can be re-written as

$$\begin{aligned} &\Pr [\bar{Q}_{d_0}(Q_c; d, V) \leq q | Q_c < q'] \Pr [Q_c < q'] + \Pr [Q_d \leq q, Q_c \geq q'] \\ &= \Pr [\bar{Q}_{d_0}(Q_c; d, V) \leq q, Q_c < q'] + \Pr [Q_d \leq q, Q_c \geq q'] \\ &= \Pr [\bar{Q}_{d_0}(Q_c; d, V) \leq q] + \Pr [Q_d \leq q, Q_c \geq q'] \\ &= G_d(q) + \Pr [Q_d \leq q, Q_c \geq q'] \geq G_d(q). \end{aligned} \quad (18)$$

The second equality follows because if $\bar{Q}_{d_0}(Q_c; d, v) \leq q$ then the event $Q_c \geq q' = \frac{S}{p_0 \times d}$ is impossible since otherwise $\bar{Q}_{d_0}(Q_c; d, v) \geq \bar{Q}_{d_0}(q'; d, v) > q$. Note, however, that (17) and (18) imply that $G_d(q)$ is strictly greater than itself, a contradiction. $\square$

Taken together, Lemmas 1 and 2 imply that $\bar{Q}_{d_0}$ is a sharp upper bound on the range of DGPs consistent with the dataset (G_c, G_{d_0}) . Therefore $\bar{B}_{d_0}(q; S, d_0) \equiv \Pr [\bar{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q']$ in turn constitutes a sharp upper bound (in the first-order dominance sense) on strong uptaker distributions, or in other words, $\bar{B}_{d_0}(q; S, d_0) \leq \Pr [Q_{d_0} \leq q | Q_c \geq q']$.

More formally, the two lemmas show that the rank-stable DGP is in fact the (sharp) upper bound on the set of DGPs that cannot be ruled out by the dataset (G_c, G_{d_0}) . In light of Assumption 4, we can furthermore use the in sample distributions (G_c, G_{d_0}) to construct out-of-sample distribution $\bar{G}_d^{oos}(q; G_c, G_{d_0})$ for each value of d under consideration. Then, the logic of Lemmas 1 and 2 goes through exactly as before, but with $\bar{G}_d^{oos}$ replacing G_{d_0} in the definitions of $\bar{Q}_{d_0}(q, v; d)$ and $\bar{B}_{d_0}(q; S, d)$ and in equations (17) and (18). $\blacksquare$

B.2. Proof of Proposition 2. We split our proof into four cases. In *Case 1*, we complete the proof in the special case considered in the main text where $G_c - G_{d_0}$ is quasi-concave and G_c, G_{d_0} are absolutely continuous. In all empirical applications considered in the body of this paper, this case appears to be the most empirically relevant of the three. For generality, in Online Appendix C we also consider *Case 2* and *Case 3* as well, where we relax the quasi-concavity and absolute continuity requirements, respectively. Throughout this section, we will again be making extensive use of a tie-breaking random variable, V , which is independent of Q_c and distributed $\text{Unif}(0, 1)$. We also refer the reader to Figure 1 for intuition on our proof construction; Panel A plots a hypothetical dataset from a pricing RCT, including control CDF G_c and treatment CDF G_{d_0} .

B.2.1. *Case 1:* Since the CDF difference $G_c(q) - G_{d_0}(q)$ is unimodal, it is weakly increasing below its smallest maximizer, q_{min}^* , and weakly decreasing above its largest maximizer, q_{max}^* . Let q_{max} denote the largest value at which G_c and G_{d_0} disagree, and define $\bar{q}_{d_0}(q') = \inf \left\{ q \in [q_{max}^*, q_{max}] : G_c(q') - G_{d_0}(q') = G_c(q) - G_{d_0}(q) \right\}$. We also define a quasi-inverse of $\bar{q}_{d_0}(q')$, which we denote by $\underline{q}_{d_0}(q) = \sup \left\{ q' \in [0, q_{min}^*] : G_c(q') - G_{d_0}(q') = G_c(q) - G_{d_0}(q) \right\}$. When G_c and G_{d_0} are absolutely continuous, the infimum and supremum defined above are attained, so $G_c(q) - G_{d_0}(q) = G_c(\bar{q}_{d_0}(q)) - G_{d_0}(\bar{q}_{d_0}(q))$ and $G_c(q) - G_{d_0}(q) = G_c(\underline{q}_{d_0}(q)) - G_{d_0}(\underline{q}_{d_0}(q))$. We also note that because G_c and G_{d_0} have no mass points, $\bar{q}_{d_0}$ and $\underline{q}_{d_0}$ are strictly decreasing on their respective domains. Given these preliminaries, recall from equation (9) that the lower-bound DGP $\underline{Q}_{d_0}$ is defined by

$$\underline{Q}_{d_0}(q; d, v) = \begin{cases} \bar{q}_{d_0}(q) & \text{if } q \leq q_{min}^*, \ v \leq \frac{g_c(q) - g_{d_0}(q)}{g_c(q)}, \text{ and} \\ \underline{Q}_{d_0}(q; d, v) = q & \text{otherwise.} \end{cases} \quad (19)$$

This definition implies that $\underline{Q}_{d_0}(Q_c)$ respects the LoD, meaning $\underline{Q}_{d_0}(Q_c) \geq Q_c$. In what follows, we will often refer to the individuals for whom $Q_c \leq q_{min}^*$ and $V \leq \frac{g_c(Q_c) - g_{d_0}(Q_c)}{g_c(Q_c)}$ as “jumpers”. The maximal proportion of jumpers that could be consistent with the data (G_c, G_{d_0}) , conditional on Q_c , is given by the quantity $\frac{g_c(Q_c) - g_{d_0}(Q_c)}{g_c(Q_c)}$ and is visualized in Panel B of Figure 1.

Lemma 3. $\underline{Q}_{d_0}(Q_c, V; d_0)$ as defined in (19) is an admissible DGP that cannot be ruled out by the dataset, (G_c, G_{d_0}) ; that is $\Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q] = G_{d_0}(q)$.

Proof. By definition of $\underline{Q}_{d_0}$, we have the following:

$$\Pr \left[\underline{Q}_{d_0}(Q_c, V) \leq q \right] = \begin{cases} A(q) & q \leq q_{min}^*, \\ A(q_{min}^*) + B(q) & q_{min}^* < q < q_{max}^*, \text{ and} \\ A(q_{min}^*) + B(q_{max}^*) + C(q) & q \geq q_{max}^*, \end{cases} \quad (20)$$

where $A(q) = \int_0^q \left[1 - \frac{g_c(x) - g_{d_0}(x)}{g_c(x)} \right] g_c(x) dx$ covers the case where q is below the smallest maximizer q_{min}^* , $B(q) = \int_{q_{min}^*}^q g_c(x) dx$ covers the case where q is between the smallest and largest maximizers, and $C(q) = \int_{q_{max}^*}^q g_c(x) dx + \int_{\underline{q}_{d_0}(q)}^{q_{min}^*} \frac{g_c(x) - g_{d_0}(x)}{g_c(x)} g_c(x) dx$ covers the case where q is above the largest maximizer q_{max}^* . $A(q)$ corresponds to the probability that $\underline{Q}_{d_0}(Q_c, V) \leq q$ resulted because $Q_c \leq q_{min}^*$ and the values of (Q_c, V) do not imply a jumper. $B(q)$ and the first term of $C(q)$ together correspond to the case where $Q_c \in [q_{min}^*, q]$ since conditioning on $Q_c \in [q_{min}^*, q]$ implies $\underline{Q}_{d_0}(Q_c, V) = Q_c \leq q$ with probability 1. Finally, the second term of $C(q)$ corresponds to jumpers for whom $Q_c \in [\underline{q}_{d_0}(q), q_{min}^*]$, in which case, despite jumping, it is still true that $\underline{Q}_{d_0}(Q_c, V) \leq q$.

The expression for A can be simplified as $A(q) = \int_0^q g_{d_0}(x) dx = G_{d_0}(q)$. On the other hand, for $q_{min}^* < q < q_{max}^*$, $G_c(q) - G_{d_0}(q)$ is constant (by unimodality), which implies $g_c(q) - g_{d_0}(q) = 0$. Thus,

$B(q)$ can also be written $\int_{q_{min}^*}^q g_{d_0}(x) dx$, so $B(q) = G_{d_0}(q) - G_{d_0}(q_{min}^*)$. Finally, we also have

$$\begin{aligned} \int_{\underline{q}_{d_0}(q)}^{q_{min}^*} \frac{g_c(x) - g_{d_0}(x)}{g_c(x)} g_c(x) dx &= G_c(q_{min}^*) - G_{d_0}(q_{min}^*) - (G_c(\underline{q}_{d_0}(q)) - G_{d_0}(\underline{q}_{d_0}(q))) \\ &= G_c(q_{min}^*) - G_{d_0}(q_{min}^*) - (G_c(q) - G_{d_0}(q)), \end{aligned}$$

which implies $C(q) = G_{d_0}(q) - G_{d_0}(q_{max}^*)$. Plugging these identities into (20) shows that regardless of the value of q , $\Pr[\underline{Q}_{d_0}(Q_c) \leq q] = G_{d_0}(q)$. $\square$

Lemma 4. $\Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q']$ constitutes a lower bound (in the first-order dominance sense) on Strong Uptaker Distributions; that is, $\Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q'] \geq \Pr[Q_{d_0} \leq q | Q_c \geq q']$.

Proof. We must show that the inequality $\Pr[Q_{d_0} \leq q | q_c \geq q'] > \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q']$ is impossible for any (q, q') pair. Suppose then, for a contradiction, that it holds for some (q, q') pair. By similar logic as in equation (15), this is equivalent to $\Pr[Q_{d_0} \leq q | q_c < q'] < \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c < q']$. Since $\Pr[Q_{d_0} \leq q | q_c < q'] + \Pr[Q_{d_0} > q | q_c < q'] = 1 = \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c < q'] + \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) > q | Q_c < q']$, our supposition is further equivalent to

$$\Pr[Q_{d_0} > q | Q_c < q'] > \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) > q | Q_c < q'], \quad (21)$$

which we now show is impossible. To reduce notational clutter, we denote the RHS of (21) as $RHS(q, q') \equiv \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) > q | Q_c < q']$. We further split Lemma 4 into the following steps.

Step 1: If pair (q, q') satisfies (21) then $\underline{Q}_{d_0}(q', 0) < q$: We will construct a proof of Step 1 by contrapositive by splitting the analysis into two further sub-cases, but it will be useful to first note the following. When $q' < q_{min}^*$, $\underline{Q}_{d_0}(q', 0; d_0) = \bar{q}_{d_0}(q')$ can be interpreted as the minimum value of counterfactual consumption among jumpers for whom $Q_c < q'$, or $\underline{Q}_{d_0}(q', 0; d_0) = \inf\{\underline{Q}_{d_0}(q', v) : \underline{Q}_{d_0}(q', v; d_0) > q'\}$. On the other hand, if $q' \geq q_{min}^*$, then $\underline{Q}_{d_0}(q', 0; d_0) = q'$.

Case 1.1: $q \leq q'$: The Law of Demand implies that if $Q_{d_0} < q$, then $Q_c < q'$ also. This implies that

$$G_{d_0}(q) = \Pr[Q_{d_0} \leq q] = \Pr[Q_{d_0} \leq q, Q_c < q'] = \Pr[Q_{d_0} \leq q | Q_c < q'] \Pr[Q_c < q'].$$

Dividing the above equation by $\Pr[Q_c < q'] = G_c(q')$ shows that $\Pr[Q_{d_0} \leq q | Q_c < q'] = \frac{G_{d_0}(q)}{G_c(q')}$. Thus, $\Pr[Q_{d_0} > q | Q_c < q'] = 1 - \Pr[Q_{d_0} \leq q | Q_c < q'] = \frac{G_c(q') - G_{d_0}(q)}{G_c(q')}$. This line of reasoning relied *only* on the LoD, and applies to the pair of random variables $(Q_c, \underline{Q}_{d_0}(Q_c, V; d_0))$ as well, since $\underline{Q}_{d_0}$ was constructed to satisfy the LoD. Thus, $RHS(q, q') = \Pr[\underline{Q}_{d_0}(Q_c, V; d_0) > q | Q_c < q'] = \frac{G_c(q') - G_{d_0}(q)}{G_c(q')} = \Pr[Q_{d_0} > q | Q_c < q']$ which contradicts (21) for the case where $q \leq q'$.

Case 1.2: $q' < q \leq \underline{Q}_{d_0}(q', 0; d_0)$: In this case, $RHS(q, q')$ is constant in its first argument for all q on the closed interval $[q', \underline{Q}_{d_0}(q', 0; d_0)]$. To see why, note that conditional on the event $Q_c < q'$, we know by the definition of $\underline{Q}_{d_0}$ that either the value of V is high, so $\underline{Q}_{d_0}(Q_c, V; d_0) = \underline{Q}_{d_0}(Q_c, 1; d_0) = Q_c < q'$; or the value of V is low, in which case, $\underline{Q}_{d_0}(Q_c, V; d_0) = \underline{Q}_{d_0}(Q_c, 0; d_0) = \bar{q}_{d_0}(Q_c) > \bar{q}_{d_0}(q') =$

$\underline{Q}_{d_0}(q', 0; d_0)$, where the inequality follows from monotonicity of $\bar{q}_{d_0}$. In either case, $\underline{Q}_{d_0}(Q_c, V; d_0)$ lies outside of $[q', \underline{Q}_{d_0}(q', 0; d_0)]$ with certainty, so for any q on that interval, we have

$$\begin{aligned} 0 \leq RHS(q', q') - RHS(q, q') &= \Pr \left[\underline{Q}_{d_0}(Q_c, V; d_0) \in [q', q] | Q_c > q' \right] \\ &\leq \Pr \left[\underline{Q}_{d_0}(Q_c, V; d_0) \in [q', \underline{Q}_{d_0}(q', 0; d_0)] | Q_c > q' \right] = 0, \end{aligned}$$

where the first inequality and equality are by definition of RHS , and the second inequality is implied by the supposition of Case 1.2. Moreover, the logic employed in Case 1.1 above establishes that if we replace q with q' in the first argument of RHS , then we have $RHS(q', q') = \Pr [Q_{d_0} > q' | Q_c < q']$. Therefore, since $RHS(q, q')$ is constant in its first argument for all $q \in [q', \underline{Q}_{d_0}(q', 0; d_0)]$, we have $RHS(q, q') = RHS(q', q') = \Pr [Q_{d_0} > q' | Q_c < q'] \leq \Pr [Q_{d_0} > q | Q_c < q']$, which contradicts (21).

The contradictions in Cases 1.1 and 1.2 demonstrate that inequality (21) can only be satisfied when $q' \leq \underline{Q}_{d_0}(q', 0) < q$. The next step shows that this case leads to a contradiction as well.

Step 2: *Inequality (21) leads to a contradiction when $q' \leq \underline{Q}_{d_0}(q', 0) < q$:*

Using similar logic as in equations (17) and (18), we have

$$\begin{aligned} 1 - G_{d_0}(q) &= \Pr [Q_{d_0} > q] = \Pr [Q_{d_0} > q | Q_c < q'] \Pr [Q_c < q'] + \Pr [Q_{d_0} > q | Q_c \geq q'] \Pr [Q_c \geq q'] \\ &> \Pr [\underline{Q}_{d_0}(Q_c, V; d_0) > q | Q_c < q'] \Pr [Q_c < q'] + \Pr [Q_{d_0} > q, Q_c \geq q'] \quad (22) \\ &= \Pr [\underline{Q}_{d_0}(Q_c, V; d_0) > q, Q_c < q'] + \Pr [Q_{d_0} \geq q, Q_c \geq q'], \end{aligned}$$

where the strict inequality follows directly from (21). We now analyze each of the two terms on the RHS of (22) (last line) in turn. For the first term, note that the events $\underline{Q}_{d_0}(Q_c, V; d_0) > q$ and $Q_c < q'$ can simultaneously occur if and only if (Q_c, V) is a jumper *and* $\underline{Q}_{d_0}(Q_c, V; d_0) = \underline{Q}_{d_0}(Q_c, 0; d_0) = \bar{q}_{d_0}(Q_c) > q$.⁴⁵ But by monotonicity of $\underline{q}_{d_0}$, this is equivalent to $Q_c < \underline{q}_{d_0}(q)$, since $Q_c \leq \underline{q}_{d_0}(\bar{q}_{d_0}(Q_c))$, by definition—recall that $\underline{q}_{d_0}$ and $\bar{q}_{d_0}$ are quasi-inverses of each other—and since $\underline{q}_{d_0}(\bar{q}_{d_0}(Q_c)) < \underline{q}_{d_0}(q)$, which follows from $q_L(Q_c)$ being strictly greater than q . This further implies that the first term on the RHS of (22) satisfies

$$\begin{aligned} \Pr \left[\underline{Q}_{d_0}(Q_c, V; d_0) > q, Q_c < q' \right] &= \Pr \left[(Q_c, V) \text{ is a jumper} \cap Q_c < \underline{q}_{d_0}(q) \right] \\ &= \int_0^{\underline{q}_{d_0}(q)} \Pr \left[V \leq \frac{g_c(x) - g_{d_0}(x)}{g_c(x)} \right] g_c(x) \, dx \\ &= \int_0^{\underline{q}_{d_0}(q)} \frac{g_c(x) - g_{d_0}(x)}{g_c(x)} g_c(x) \, dx \\ &= G_c(\underline{q}_{d_0}(q)) - G_{d_0}(\underline{q}_{d_0}(q)) = G_c(q) - G_{d_0}(q). \end{aligned}$$

⁴⁵Note that $Q_c > q$ is inconsistent with the second event of the joint probability for the case considered in Step 2, where $Q_c < q' < q$.

Next, turning to the second joint probability on the RHS of (22), we have

$$\begin{aligned}\Pr[Q_{d_0} \geq q, Q_c \geq q'] &= \Pr[Q_{d_0} \geq q, Q_c \geq q] + \Pr[Q_{d_0} \geq q, q > Q_c \geq q'] \\ &= \Pr[Q_c \geq q] + \Pr[Q_{d_0} \geq q, q > Q_c \geq q'] \geq 1 - G_c(q)\end{aligned}$$

where the first equality follows from the law of total probability and the supposition of Step 2, and the second equality is true because $Q_c \geq q \Rightarrow Q_{d_0} \geq q$ by the LoD. As a result, the last line of (22) is greater than or equal to $G_c(q) - G_{d_0}(q) + (1 - G_c(q)) = 1 - G_{d_0}(q)$, which combined with the rest of inequality (22), leads to the contradiction that $1 - G_{d_0}(q) > 1 - G_{d_0}(q)$. $\square$

Together, Lemmas 3 and 4 imply that $\underline{Q}_{d_0}$ is a sharp upper bound on the range of DGPs that cannot be ruled out by the dataset (G_c, G_{d_0}) . Therefore, $\underline{B}_{d_0}(q; S, d_0) \equiv \Pr \left[\underline{Q}_{d_0}(Q_c, V; d_0) \leq q | Q_c \geq q' \right]$ constitutes a sharp upper bound (in the first-order dominance sense) on strong uptaker distributions, or in other words, $\underline{B}_{d_0}(q; S, d_0) \geq \Pr[Q_{d_0} \leq q | Q_c \geq q']$. $\blacksquare$

B.3. Out-of-Sample Inference Using Only the Law of Demand. In this Appendix, we further discuss the extent to which out-of-sample inference is possible without Assumption 4 but still main. If one assumes only the LoD, but is uncomfortable adding additional structure such as that discussed in Remark 2, then for out-of-sample discounts, counterfactual CDFs can be sharply bounded only as follows:

Proposition 4. *Under Assumption 1 and for arbitrary (S, d) pairs, if $G_c(q)$ and $G_{d_0}(q)$ are known and are discontinuous at countably many mass points, then the following constitute (identified) sharp bounds on SUDs:*

$$\bar{B}_{d_0}^{LoD}(q; S, d) \equiv \begin{cases} \bar{B}_{d_0}\left(q; \frac{Sd_0}{d}, d_0\right), & \text{if } d < d_0, \\ \bar{B}_{d_0}(q; S, d_0), & \text{if } d = d_0, \\ 0, & \text{if } d > d_0; \end{cases} \quad \underline{B}_{d_0}^{LoD}(q; S, d) \equiv \begin{cases} G_c\left(q | Q_c \geq \frac{S}{p_0 \times d}\right), & \text{if } d < d_0, \\ \underline{B}_{d_0}(q; S, d_0), & \text{if } d = d_0, \\ \underline{B}_{d_0}\left(q; \frac{Sd_0}{d}, d_0\right), & \text{if } d > d_0. \end{cases} \quad (23)$$

For some brief intuition on the bounds in Equation (23), note that when an out-of-sample discount is less generous than d_0 —i.e., $0 < d < d_0$ —the LoD implies that G_d lies somewhere between G_c and G_{d_0} but does not give any further information. The upper bound on DGPs that cannot be ruled out by existing data thus corresponds to the case where we make the most optimistic possible assumption about shifts in consumer demand, i.e., that $G_d = G_{d_0}$ for any $d \in (0, d_0)$. On the other hand, the lower bound corresponds to the case where we make the most pessimistic assumption that $G_d = G_c$ for any $d \in (0, d_0)$. Similarly, when an out-of-sample discount is more generous than d_0 —i.e., $d_0 < d$ —the lower bound corresponds to making the most pessimistic possible assumption that $G_d = G_{d_0}$, while the upper bound is completely uninformative, since we know only that G_d is located to the right of G_{d_0} , and demand is otherwise unconstrained from above by the LoD.