MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 1523-4614 | EISSN 1526-5498 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

The Power and Limits of Decision Making with Confounded Data: The Case of Pricing

Dimitris Bertsimas
Massachusetts Institute of Technology, dbertsim@mit.edu, http://web.mit.edu/dbertsim

Nathan Kallus
Cornell Tech and Cornell University, kallus@cornell.edu, http://www.nathankallus.com

While data-driven decision-making is transforming modern operations, most large-scale data is of an observational nature, such as transactional records. This poses unique challenges for managerial decision making problems where one must evaluate the effect of a decision on an uncertain reward based on historical data where observed decisions and outcomes may be confounded. The clearest example is pricing: we are interested in maximizing price times average demand at the price based on data where historical prices and demand are endogenous. Most often the observational data lacks the features, such as containing all confounders or a valid instrument, to correctly assess decisions' causal effects and/or the strong assumptions necessary are highly dubious. Nonetheless, the inevitable *statistical biases* in flawed causal inferences may in fact be *irrelevant* in as much as they may have limited impact on the downstream *optimization error* in any resulting decision. In this paper we seek to formalize this notion by considering optimization based on *predictive* rather than causal models and focusing on the case of pricing. We give bounds on the percent suboptimality of pricing using the *prediction* of historical demand given price as the demand model. That this can be perform well even when data is unconfounded demonstrates the power of decision making with confounded data. To study potential limits in real datasets, we develop a new hypothesis test for optimality. We apply the test to interest-rate-setting data to assess whether the performance of predictive approaches can even be distinguished from optimal to statistical significance. We empirically demonstrate that predictive approaches can generally be powerful in practice but with some limitations.

Key words: Mathematical Programming, Revenue Management and Pricing

History:

1. Introduction

Data-driven decision-making is transforming modern operations and is being rapidly adopted in practice (McAfee and Brynjolfsson 2012, Brynjolfsson and McElheran 2016). Such data-driven decisions crucially rely on the availability and reliability of very large-scale datasets. But in many decision-making instances, some data is missing and some decision alternatives are not fully characterized by the data. A primary example is in data-driven optimization under uncertainty when decisions affect the uncertain variable: when making a decision, one considers what would have performed well historically in the data but the data may only contains the realized outcomes of the historical decisions and never the counterfactual outcomes under potential alternative decisions, which must necessarily be modeled.

For example, in revenue management, a manager may wish to set prices based on a large-scale transactional dataset of prices and the demands that resulted from these prices by using demand modeling. Indeed, there are many examples of studies on pricing decisions based on large-scale historical datasets.[1] Besbes et al. (2010) study validating demand models and consider an example of setting loan interest rates based on a large historical dataset of loan offers, including customer acceptance and rejection of each these offers. Ferreira et al. (2016) consider pricing fashion items for an online retailer based on the company's extremely rich historical data on past sale events. And Cohen et al. (2014) study the problem of planning price discounts and consider the application to a large dataset of historical prices for goods in a supermarket and the historical sales generated. Price promotions may also be combined with marketing promotions and similar data-driven decision-making situations arise in these marketing settings. Similarly, in inventory management, there are examples where demand may be inventory-dependent (Lee et al. 2012).

[1] There is also an important stream of literature looking at pricing based on repeated experimentation (Bertsimas and Perakis 2006, Besbes and Zeevi 2009, Harrison et al. 2012, Nambiar et al. 2019), which does not consider pricing based on a *historical* dataset.

In all of these examples, the effect of a decision (e.g., price, order quantity, assortment, marketing campaign) on the outcome of an uncertain variable (e.g., demand, yield, returns) has to be modeled from historical, observational data. In general observational data, however, historical decisions and outcomes may be confounded, obscuring the isolated, causal effect of any one potential decision. For example, in econometric studies of marketplace supply-demand-price relationships, the analysis generally has to take such endogeneity into consideration (Berry et al. 1995, Bijmolt et al. 2005, Phillips et al. 2012). However, in data-driven decision-making, large-scale data is often purely transactional, as in some of the examples above, and may lack the features necessary to enable a precise and sound assessment of the causal effects of any one decision (e.g., by controlling for all common causes of decision and outcome or by using an instrumental variable), or the strong assumptions necessary for such an assessment may be unfounded. Instead of giving up completely, *predictive* – rather than *causal* – analyses may be done, modeling the outcome of a given decision as the best *prediction* of outcome given an observation of the decision in the historical data, namely, the conditional expectation. In general observational settings where data does not arise from experimental manipulations this identification of causal effects is *spurious*. As we study in this paper, notwithstanding the spurious identification of predictions as effects:

predictive models may give biased <u>causal estimates</u> but still lead to good <u>decisions</u>.

Indeed, all models are wrong but some are useful, and, in optimization settings, it is only the objective value of the final decision that matters, rather than the validity of any model used to arrive at it or the statistical error or bias of the model. Thus, while a predictive approach may fail to estimate the causal effect on the objective of the feasible decisions or may be on too shaky ground to provide reliable and valid causal estimates, it may still lead to good decisions down the line. In this paper, we seek to study both the power and limits of predictive approaches to observational-data-driven decision making by studying performance bounds and developing a hypothesis test for causal-effect objective optimality, focusing primarily on the case of *data-driven pricing*.

Let $z \in \mathcal{Z}$ denote the price decision variable and $r(z) = z - c$ the per-unit revenue where c is a procurement cost. Let $Y(z)$ represent the *potential* random demand we would observe *if* we set

Table 1 Example Demand Data for the Classic MIT Hoodie

Day (i)	Price (Z_i)	Observed Demand (Y_i^{obs})	Unobserved Demand (Y_i^{unobs})
1	20 ($)	1 (units)	0
2	28	0	1
3	28	1	1
4	20	2	0
⋮	⋮	⋮	⋮
$n-1$	20	1	0
n	28	0	1

price to z. The collection of variables $\{Y(z) : z \in \mathcal{Z}\}$ represents the random demand curve. Defining $y(z) := \mathbb{E}[Y(z)]$, the optimization problem of interest is then

$$z^* \in \arg\max_{z \in \mathcal{Z}} R(z), \qquad \text{where } R(z) := r(z)y(z). \tag{1}$$

While we focus on pricing, some insights extend to other problems. In an alternative example in inventory management with quantity-dependent demand (Lee et al. 2012), we may have that z is order quantity and $R(z) = \mathbb{E}[p \max\{z, Y(z)\} - cz]$, where p is a constant price. Another example is transshipment with pricing (Bertsimas and Kallus 2020, Section 3.2).

In *observational-data-driven decision-making*, we consider solving problem (1) based on observational data consisting of repeated observations of (Z, Y^{obs}) where $Y^{\text{obs}} = Y(Z)$. We call such data *observational* because, firstly, it does not contain counterfactual outcomes, $Y(z)$ for $z \neq Z$, and, secondly, the historical choice of Z may not be randomized and may have correlations with $Y(z)$ for any one z. Recall that in pricing, Z is the price and Y^{obs} is corresponding observed demand while $Y(z)$ is the demand curve for all potential prices z. Also note our convention where upper-cases (e.g., Z) denotes random variables and lower-cases (e.g., z) denote some fixed dummy values.

Note that in many cases, the decision is *customized* to a specific customer segment. For example, Besbes et al. (2010) consider separate interest rates customized for different customer classes defined by ranges of FICO scores, known generally as customized pricing (Phillips 2005). Therefore, we may simply think of the distribution in Problem (1) over which expectation is taken as specific to, or conditioned on, a particular targeted segment.

1.1. An Illustrative Example

For the purpose of illustrating the issues with observational data, we consider a very simple example. Table 1 displays the unit demand Y_i^{obs} observed at the MIT Coop on each day $i = 1, \ldots, n$ for the classic MIT hoodie and the price Z_i at which it was offered. Only two prices, \$20 and \$28, have been observed. For each day, there are the potential demands $Y_i(20)$ and $Y_i(28)$ that would have been observed if the price were set to \$20 or \$28, respectively – these two values together represent the full, unseen *demand curve*, $\{Y(z) : z \in \mathcal{Z}\}$, associated with that day. We only observe the realized demand $Y_i^{\text{obs}} = Y_i(Z_i)$. We do *not* observe the counterfactual demand $Y_i^{\text{unobs}} = Y_i(48 - Z_i)$. For example, on day 1, $Y_1^{\text{obs}} = Y_1(20) = 1$ and $Y_1^{\text{unobs}} = Y_1(28) = 0$. From the observed data only, we can compute

$$\tilde{y}_n(20) = \text{Average}\left(\left\{Y_i^{\text{obs}} \,:\, i = 1, \ldots, n,\, Z_i = 20\right\}\right),$$

$$\tilde{y}_n(28) = \text{Average}\left(\left\{Y_i^{\text{obs}} \,:\, i = 1, \ldots, n,\, Z_i = 28\right\}\right).$$

If the rows of Table 1 constitute independent and identically distributed (iid) data and $Z, Y^{\text{obs}}, Y(20), Y(28)$ represent a generic random draw, then $\tilde{y}_n(20)$ and $\tilde{y}_n(28)$ are our best guesses for the value of demand Y^{obs} in a new random instance where $Z = 20$ or $Z = 28$, respectively. In particular, $\tilde{y}_n(z) \to \tilde{y}(z) := \mathbb{E}[Y^{\text{obs}} \mid Z = z]$ almost surely as $n \to \infty$, and $\tilde{y}(z)$ is by definition the *best* predictor of demand given price (in squared error). Identifying $\tilde{y}(z)$ with the expected demand resulting from a pricing decision z leads to the optimization problem

$$\tilde{z} \in \operatorname*{arg\,max}_{z \in \mathcal{Z}} \tilde{R}(z), \qquad \text{where } \tilde{R}(z) := r(z)\tilde{y}(z). \tag{2}$$

In this particular example $\mathcal{Z} = \{20, 28\}$. When we substitute any estimate $\tilde{y}_n(z)$ of $\tilde{y}(z)$ (or any estimate $\tilde{R}_n(z)$ of $\tilde{R}(z)$) into (2), we call the result a *predictive approach* to pricing from data because it hinges on fitting a *predictive model* of the objective given decision. Given observational data, we can estimate $\tilde{y}(z)$ (or $\tilde{R}_n(z)$) using any regression method, fitting a linear, other parametric, or non-parametric regression of the response Y^{obs} (or, $R(Z)$) to regressor Z. For example, Besbes et al. (2010) use non-parametric Nadaraya-Watson kernel regression to construct an estimator $\tilde{R}_n(z)$ of $\tilde{R}(z)$ that is a universally consistent under mild conditions (Greblicki et al. 1984), that is $\tilde{R}_n(z) \to \tilde{R}(z)$ as $n \to \infty$ without specifying a model for $\tilde{R}(z)$.

On average, over each new day, $y(z) = \mathbb{E}[Y(z)]$ is the expected demand if price were set to z and the profit-optimizing price is given by Problem (1). In general, $y(z) \neq \tilde{y}(z)$ so this might be different from Problem (2). In particular, if the population distribution of the data is exactly the discrete distribution with weight 1/6 on each of the six displayed rows of Table 1, then

$$\tilde{y}(20) = 4/3 \neq 7/6 = y(20), \qquad \tilde{y}(28) = 1/3 \neq 1/6 = y(28).$$

In contrast, a corresponding sample-average estimate of $y(z)$ involves *unobserved* data:

$$y_n(z) = \frac{1}{n}\left(Y_1(z) + Y_2(z) + \cdots + Y_n(z)\right) = \text{Average}\left(\{Y_i(z) : i = 1, \ldots, n\}\right),$$

In particular, if population distribution is an equiprobable draw of the six displayed rows, we can imagine filling in the unseen values in Table 1 with anything, changing $y(z)$, and correspondingly Problem (1), but keeping the *observed* data completely unchanged. This is because the demand curve, which captures demand at any one price, can be distinct from the relationship of observed price and demand. The key distinction is between *any one price* and *the price observed*. This highlights that, in general, Problem (1) is not necessarily well-specified by the observed data in the most general settings.

However, in an optimization setting, this issue may potentially be moot. Regardless of how close $\tilde{y}(z)$ and $y(z)$ are, it may very well be the case that $\tilde{z}$ is close to z^* and, much more importantly, that $R(\tilde{z})$ is close to $R(z^*)$. For example, in the above example, if the data is distributed in the population like the six displayed rows and supposing the MIT Coop has procurement cost of \$19, we have $\tilde{z} = z^* = 28$ even though $y(z) \neq \tilde{y}(z)$ and $y(z)$ is not well-specified by the observed data, showing the potential power of predictive approaches in practice. But, we can also imagine a setting where we change the values of Y_i^{unobs} in Table 2 so that $\tilde{z} \neq z^*$ and, more importantly, $R(z^*)$ is much larger than $R(\tilde{z})$, showing the potential limits of predictive approaches in practice.

1.2. Overview of the Paper

In this paper, we provide a thorough exploration of this issue in data-driven pricing. We first provide a more thorough discussion of the limits of identifying z^* from observational data (Section 2). We then prove new performance guarantees on predictive approaches, some with potentially unrestricted confounding effects. Specifically, by leveraging the special structure of the optimization problem, these bound our profit-suboptimality even when the optimal price cannot be identified from the data (Section 3). This demonstrates the power of predictive approaches to observational-data-driven optimization. To study potential limits of predictive approaches in real datasets, we then develop a new hypothesis test to evaluate objective optimality based on observational data in problems that include data-driven pricing. The test allows us to determine whether the reward generated by any one decision-making algorithm, such as a predictive one, can be distinguished as suboptimal to a statistically significant degree based on purely observational data when such data does identify the demand curve. To develop the test, we show favorable asymptotics (consistency and asymptotic normality) for a non-parametric solution that works under certain identifiability conditions (Section 4.2) and use this to establish an asymptotic null distribution for a test statistic for the hypothesis of objective optimality (Section 4.3). We also present a parametric approach to solving Problem (1) under these conditions where the solution is identifiable (Section 5). Using our hypothesis test, we empirically study an interest-rate-setting problem with data from Columbia University Center for Pricing and Revenue Management (2012) (Section 6).

The message of the results, both theoretical and empirical, are twofold. First, for predictive approaches, the results show that even when the data does not identify the optimal decision, a good predictive approach can obtain good objective performance. We see this both in our theoretical bounds and in our empirical investigation. Beyond predictive approaches, in the rare cases where optimal decisions are identifiable by controlling for confounders, we add a dimension to the discussion on the sufficiency of parametric approaches by showing that a proposed parametric approach to observational-data-driven decision-making that accounts for confounders performs well and can hardly be distinguished from optimal by our objective-optimality test even though the model is certainly misspecified.

2. Identifiability and the Gap Between Problems (1) and (2)

As exemplified in the introduction, the mean-response function $y(z)$ that defines Problem (1) is not well-specified by observational data and is potentially distinct from the function $\tilde{y}(z)$, which is. In this section, we first consider a more detailed example and study its implications on the identifiability of z^* (we will define identifiability precisely below) and then we consider the gap between the functions $y(z)$ and $\tilde{y}(z)$ in preparation for bounding the revenue gap between their induced optima, $R(z^*)$ and $R(\tilde{z})$, in the next section.

Example 1 (Consulting for the MIT Coop). Alice and Bob are hired by the MIT Coop to help determine an optimal sale price for the classic MIT hoodie, which the MIT Coop procures at a unit price of $c = \$19$ ($r(z) = z - 19$). The MIT Coop is debating between a retail price of \$20 and a retail price of \$28. In any given day in the past, the MIT Coop has offered the hoodie at either of the two prices and observed either no units sold, one unit sold, or two units sold. The Coop has a great deal of observational data on joint observations of price Z and demand Y^{obs}.

Alice and Bob collate this data into a table that shows the frequency of each price-demand combination over history shown in Table 2(a). Due to the abundance of data, Alice and Bob are confident that this is a faithful representation of the joint distribution of price and demand, (Z, Y^{obs}). Naturally, the data only has the demand that was in fact observed on each day and the demand that would have been observed under any other price is missing. However, a full demand model must model the distribution of the demand curve, $(Y(20), Y(28))$, for each new sale event, which includes the unobserved demands. Together, a full model for all the data, observed and unobserved, would provide a joint probability distribution over $(Z, Y(20), Y(28))$.

Alice, in an attempt to produce such a model, regresses demand on price by computing a weighted average in each of the columns of Table 2(a) and finds that $\tilde{y}(20) = \mathbb{E}\left[Y^{\text{obs}} \middle| Z = 20\right] = 10/9$, $\tilde{y}(28) = \mathbb{E}\left[Y^{\text{obs}} \middle| Z = 28\right] = 1/9$. She then constructs a model, that is a joint distribution over $(Z, Y(20), Y(28))$, wherein $y(z) = \tilde{y}(z)$ would be satisfied, and she arrives at the one shown in Table 2(b). Alice verifies that her model fully agrees with the observed data: when she applies the

Table 2 Data for Example 1

Z	Y^{obs}	$\mathbb{P}$
20	0	0
20	1	8/18
20	2	1/18
28	0	8/18
28	1	1/18
28	2	0

(a) Joint distribution of historical price and demand

Z	$Y(20)$	$Y(28)$	$\mathbb{P}$
20	1	0	32/81
20	1	1	4/81
20	2	0	4/81
20	2	1	1/162
28	1	0	32/81
28	1	1	4/81
28	2	0	4/81
28	2	1	1/162

(b) Alice's demand model

Z	$Y(20)$	$Y(28)$	$\mathbb{P}$
20	1	0	40/99
20	1	1	4/99
20	2	0	0
20	2	1	1/18
28	1	0	4/9
28	1	1	2/45
28	2	0	0
28	2	1	1/90

(c) Bob's demand model

transformation $(Z, Y(20), Y(28)) \mapsto (Z, Y(Z))$, she recovers the frequencies in Table 2(a). She then computes $R(20) = (20-19) \times (\frac{72}{81} \times 1 + \frac{9}{81} \times 2) = 10/9$ and $R(28) = (28-19) \times (\frac{72}{81} \times 0 + \frac{9}{81} \times 1) = 1$, concluding that $z^\star = 20$ is the optimal price.

Bob, working from home that day and unaware of Alice's progress, has independently come up with another model, shown in Table 2(c), in order to explain the observed pricing data. Bob, too, verifies that his model completely agrees with the observed data (by applying the above transformation and checking that the frequencies match) and calculates $R(20) = (20-19) \times (\frac{14}{15} \times 1 + \frac{1}{15} \times 2) = 16/15$ and $R(28) = (28-19) \times (\frac{28}{33} \times 0 + \frac{5}{33} \times 1) = 15/11$, concluding, differently from Alice, that $z^\star = 28$ is in fact the optimal price.

Alice and Bob had both come up with demand models that fully concur with the observed data but recommended different prices as optimal. Both models support the data in Table 2(a) fully in the sense that they recover the frequencies in Table 2(a) under the transformation $(Z, Y(20), Y(28)) \mapsto (Z, Y(Z))$, which corresponds to the censoring of counterfactual demands. These frequencies can also be described by a homoskedastic linear model,

$$Y^{\text{obs}} = \frac{65}{18} - \frac{1}{8} Z + \omega, \qquad \omega = \begin{cases} -1/9, & \text{with prob. } 8/9, \\ 8/9, & \text{with prob. } 1/9, \end{cases} \qquad \omega \perp\!\!\!\perp Z,$$

where regressor Z is independent of zero-mean error ω. Note that the response variable is $Y^{\text{obs}} = Y(Z)$, and *not* $Y(z)$, and that the regressor is historically observed price Z, and not any one price decision $z \in \mathcal{Z}$. Since the two models agree on this form (since they agree on Table 2(a)) but recommended different prices, this also highlights that this is not an issue of misspecifying a regression function model.

The reason for this is simple: the data, the function $\tilde{y}(z)$, and $\tilde{z}$ all involve just the joint distribution of (Z, Y^{obs}), while the true demand model $y(z)$ and the optimal price z^* involve the joint distribution of $(Z, Y(20), Y(28))$, and the former joint distribution does not determine the latter joint. We make this formal in Section 2.1 below. More generally, as we explain in Section 2.2, the difference between $y(z)$ and $\tilde{y}(z)$ involves the correlation between $Y^{\text{obs}} - y(Z)$ and Z.

2.1. Identifiability

The issue we encountered above is one of identifiability and shows that z^* is non-identifiable in general.

Definition 1. Let $\Pi = \{\mathbb{P}_\theta : \theta \in \Theta\}$ be a model for the distribution of the observed data. We say that $\phi : \Theta \to \Phi$ is *identifiable* if for any $\mathbb{P}_{\theta_1}, \mathbb{P}_{\theta_2} \in \Pi$ such that $\mathbb{P}_{\theta_1} = \mathbb{P}_{\theta_2}$, we have $\phi(\theta_1) = \phi(\theta_2)$.

In the above definition, Θ and Φ may be arbitrary sets, that is, the model need not be parametric. Note that if any ϕ is not identifiable then any finer quantity, such as θ itself, is not identifiable.

To connect the above definition with our decision-making setting, we let θ denote the joint distribution of $(Z, \{Y(z) : z \in \mathcal{Z}\})$, we let $\mathbb{P}_\theta$ be the corresponding distribution of the data (Z, Y^{obs}) begotten via the transformation $(Z, \{Y(z) : z \in \mathcal{Z}\}) \mapsto (Z, Y(Z))$ applied to θ (i.e., this variable transformation describes the map $\theta \mapsto \mathbb{P}_\theta$), and we let ϕ map θ to the optimal decision $z^\star$ (or, set thereof) as described by Problem (1). Then, Example 1 above proves the following result by example:

Corollary 1. *The optimal decision $z^\star$ is not identifiable on the basis of observations of (Y^{obs}, Z).*

In fact, we have shown a stronger result of which Corollary 1 is a corollary:

Theorem 1. *The optimal decision $z^\star$ is not identifiable on the basis of observations of (Y^{obs}, Z) even under the Gauss-Markov assumptions:*

a. Linearity: there is a random variable ω such that $Y^{\text{obs}} = \beta_0 + \beta_1 Z + \omega$.

b. *Exogeneity of independent variables:* $\mathbb{E}\left[\omega \middle| Z\right] = 0$.

c. *Homoskedasticity:* $\mathrm{Var}\left(\omega \middle| Z\right) = \mathrm{Var}\left(\omega\right)$ *is constant.*

d. *No collinearity:* Z *is not constant.*

In Example 1, exogeneity and homoskedasticity are a consequence of $\mathbb{E}[\omega] = 0$ and $\omega \perp\!\!\!\perp Z$. Exogeneity implies $\mathrm{Cov}(\omega, Z) = 0$. Note that whenever the optimal price z^* is not identifiable, the mean-response function $y(z)$, a finer quantity (i.e., z^* is a function of $y(z)$), is not identifiable either.

When *is* z^* identifiable from observational data? Naturally, one case is when $y(z) = \tilde{y}(z)\,\forall z$, since $\tilde{y}(z)$ is always identifiable from observational data. Generally, however, the functions are distinct. Next, we study the gap between the two functions $y(z)$ and $\tilde{y}(z)$.

2.2. The Confounding Error Gap

Let $\epsilon(z) = Y(z) - y(z)$ be the deviation of the response curve from its mean at any point z. Let $\epsilon = \epsilon(Z) = Y^{\text{obs}} - y(Z)$, which is the deviation of the *observed* demand Y^{obs} under price Z from the mean demand that would be induced by the price Z. The degree to which ϵ is correlated with Z is known as *confounding* or *endogeneity*. In fact it can be directly related to the discrepancy between the functions $y(z)$ and $\tilde{y}(z)$:

$$E(z) := \tilde{y}(z) - y(z) = \mathbb{E}\left[Y^{\text{obs}} \mid Z = z\right] - y(z) = \mathbb{E}\left[Y^{\text{obs}} - y(Z) \mid Z = z\right] = \mathbb{E}\left[\epsilon \mid Z = z\right].$$

We call $E = E(Z) = \mathbb{E}[\epsilon \mid Z]$ the *confounding error*.

Thus, the error E is directly related to the association of historical decisions Z and the unique idiosyncrasies ϵ of historical outcome events. Independence of the two – i.e., that Z is independent of the particular event, as in a randomized controlled trial (RCT) – would immediately imply that the confounding error is exactly zero: $E(z) = \mathbb{E}[\epsilon \mid Z = z] = \mathbb{E}[\epsilon] = \mathbb{E}[Y(z)] - y(z) = 0$. Correspondingly, in this case, we *would* have $\tilde{y}(z) = y(z)$ for all z. Therefore, when the data is the result of experimental manipulations rather than observation, we must also have $\tilde{z} = z^*$.

Note the critical distinction between ϵ and the regression errors $\omega = Y^{\text{obs}} - \mathbb{E}[Y^{\text{obs}} \mid Z = z]$. Regression errors, by their very definition, will *always* have $\mathbb{E}[\omega \mid Z = z] = \mathbb{E}[Y^{\text{obs}} \mid Z = z] -$

$\mathbb{E}[Y^{\text{obs}} \mid Z = z] = 0$ even in the most general setting for observational data, but the same is generally only true of ϵ in experimental settings. Thus, in general settings, $\mathbb{E}[\epsilon \mid Z = z] \neq 0$ and, correspondingly, $\tilde{y}(z) \neq y(z)$, as functions.

However, even when the function $\tilde{y}(z)$ is very different from $y(z)$, what really matters is the performance of the resulting solution, that is, $R(\tilde{z})$. Figure 1 shows an example for the pricing problem with $r(z) = z$ (zero procurement cost) and $\mathcal{Z} = [0, 1]$. Figure 1(a) shows one true curve for $y(z)$ and three possible different curves for $\tilde{y}(z)$. The first two, $\tilde{y}_1(z)$ and $\tilde{y}_2(z)$, both differ from $y(z)$ by a constant amount 1/2 throughout the domain, while the second is also non-increasing. The third, $\tilde{y}_3(z)$, differs from $y(z)$ a lot in magnitude but not in its linear shape. Figure 1(b) shows the corresponding revenue curves. Optimizing each $r(z)\tilde{y}_k(z)$ for $k = 1, 2, 3$ to find $\tilde{z}_k$ and plugging each into $R(z)$, we see that $R(\tilde{z}_1)/R(z^\star) = 0.75$, $R(\tilde{z}_2)/R(z^\star) = 0.84$, and $R(\tilde{z}_3)/R(z^\star) = 0.96$. This shows a few examples of how discrepancies between $\tilde{y}(z)$ and $y(z)$ translate to performance, and potentially only very little revenue loss. Since in general settings $z^\star$ can never be known, to have performance that is guaranteed to be close to that of the unknown and optimal $z^\star$ would be very strong evidence for the strength of a method. Next, we address how to bound this optimality gap in general, without knowledge of $y(z)$, so that we can theoretically guarantee good performance even if the best decision cannot be pinned down.

3. The Power of Predictive Approaches: Suboptimality Bounds

In this section, we and demonstrate that even though the optimal price, z^* is not identifiable from data, the price generated by the predictive approach based on the data, $\tilde{z}$, can still yield good revenue. We consider $\mathcal{Z} = (c, \infty)$ throughout this section.

To show the power of predictive approaches, we establish bounds on how suboptimal a predictive approach may be in the pricing problem. Such bounds show that even if confounding is present and even if Problem (1) is not well-defined given the data, predictive approaches, which are implementable in practice, can still get us reasonably close to the optimum under certain conditions. Crucially, the bounds leverage both the special structure of the pricing optimization problem and

Figure 1 Predictive demand and revenue under confounding.

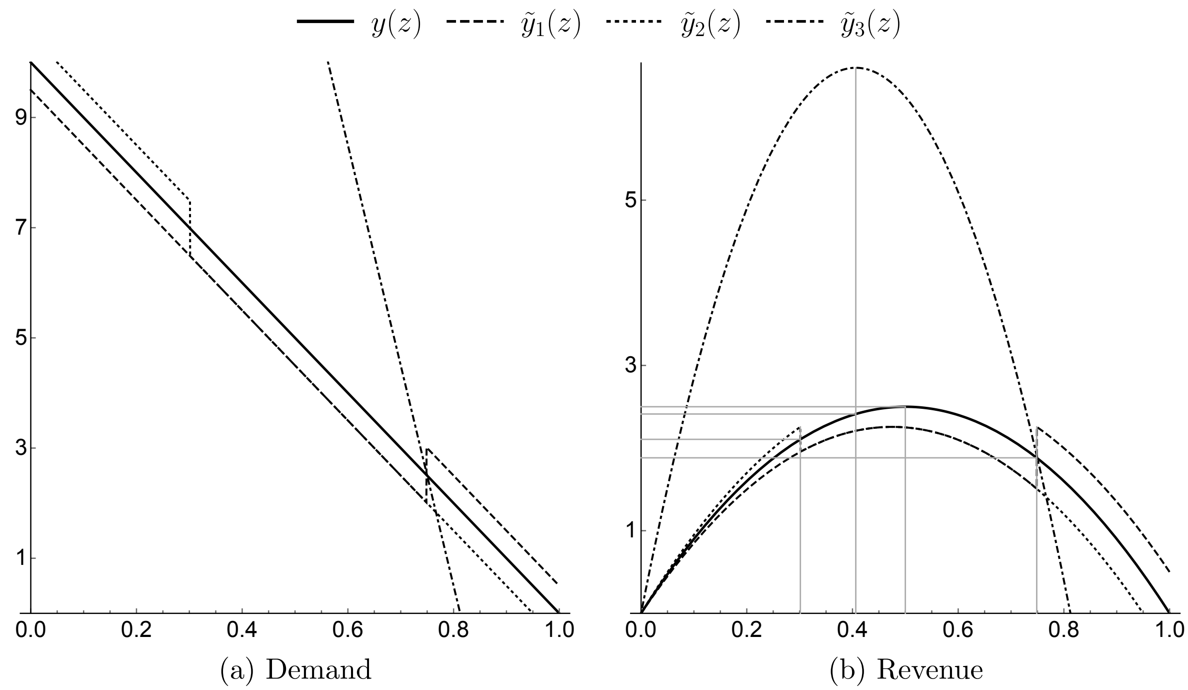

common features of price-demand relationships. The bounds address error in the metric of interest, which is true profit $R(z)$. The intent is to express this error in *relative* terms, using only *few* parameters, and using assumptions we can *reason* about. Beyond pricing, these bounds suggest a framework for establishing similar results in other observational-data-driven decision-making problems by appealing to the special structure of the problem and not just the statistical errors or bias in estimating the objective.

The bounds we present are expressed in terms relative to the *size of the market*:

$$y_0 = \sup_{z\in\mathcal{Z}} y(z).$$

The first bound is based on the *magnitude* of confounding error relative to the size of the market.

THEOREM 2. *In the pricing problem, if $y(z)$ is non-increasing and linear and $|E|/y_0 \le \gamma$ then*

$$1 - 4\gamma - 4\gamma^{3/2} - \gamma^2 \le \frac{R(\tilde{z})}{R(z^\star)} \le 1.$$

The proof is given in the appendix. Note that under the assumptions of the theorem $y_0 = y(c)$. The bound is non-negative (i.e., nontrivial) for values of γ up to $3-\sqrt{8}\approx 0.17$. We plot the bound

Figure 2 **The bounds of Theorems 2, 3, and 4 demonstrate that predictive approaches to pricing from observational data can actually perform well in practice, even when the true optimal price could never be identified from the data.**

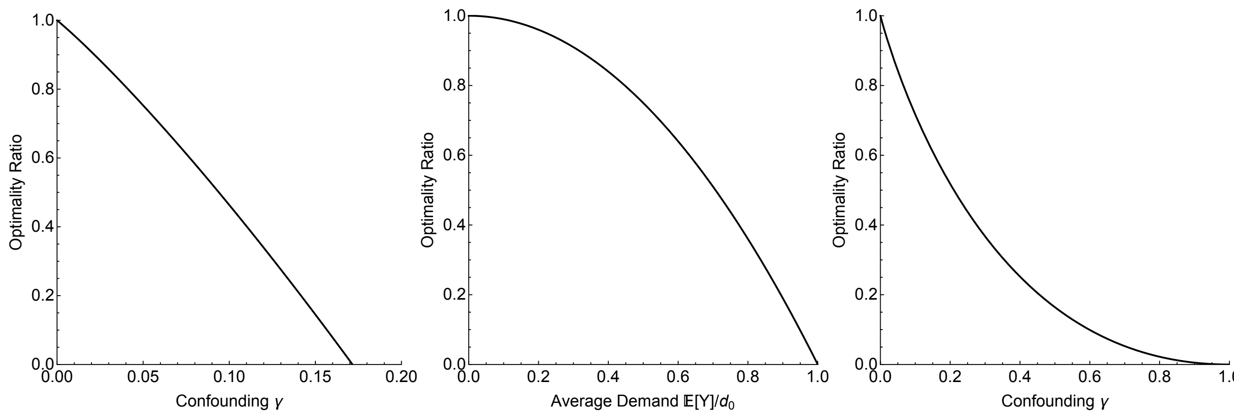

(a) Theorem 2 (b) Theorem 3 (c) Theorem 4

in Figure 2(a). Applying this bound to $\tilde{y}_1(z)$ of Figure 1, which has $\gamma = 0.05$, we get 0.71, which indeed bounds the observed ratio of 0.75.

In general, the magnitude of the confounding error is neither known nor estimable from data for the very same reasons that the function $y(z)$ is not. Our next bound seeks to leverage particular structure that is symptomatic of pricing data to express suboptimality in terms of average demand, which *can* be estimated from data.

THEOREM 3. *In the pricing problem, if $y(z)$ is non-increasing and linear and ϵ and Z are non-positively correlated and jointly normal then*

$$1 - \left(\frac{\mathbb{E}\left[Y^{\text{obs}}\right]}{y_0}\right)^2 \le \frac{R(\tilde{z})}{R(z^\star)} \le 1.$$

The proof is given in the appendix. Note that $\mathbb{E}\left[Y^{\text{obs}}\right]$ can be unbiasedly and consistently estimated by the sample average of observed demands in the data. The assumption of non-positive correlation between ϵ and Z corresponds to a common feature of pricing datasets. For example, promoting a product via advertising, which would increase its potential demand at any given price, would often coincide with promotion via price discounts and this would lead to such non-positive correlation. This assumption, which can be reasoned about in such a way, leads to a stronger bound that is completely independent of the size of confounding error. We plot the bound in Figure 2(b).

The curve $\tilde{y}_3(z)$ of Figure 1 can be seen to arise from a situation where $y(z)$ is as in Figure 1, Z and ϵ are jointly normal with a covariance that is -30 times the variance of Z, and $\mathbb{E}[Y^{\text{obs}}] = 2.5$. Applying Theorem 3, we get 0.94, which indeed bounds the observed ratio of 0.96. (In particular, we can achieve the bound by letting the covariance approach $-\infty$.)

If we consider a similar non-positive relationship between ϵ and Z in the setting of Theorem 2, we obtain an improved bound.

THEOREM 4. *In the pricing problem, if $y(z)$ is non-increasing and linear, $E(z)$ is non-increasing, and $|E|/y_0 \le \gamma \le 1$ then*

$$1 - 4\gamma + 4\gamma^{3/2} - \gamma^2 \le \frac{R(\tilde{z})}{R(z^\star)} \le 1.$$

The proof is given in the appendix. This result is similar to that in Theorem 2 but the additional assumption that $E(z)$ is decreasing allows us to achieve a strictly stronger bound that is non-negative for values of γ up to 1. That $E(z)$ is non-increasing captures the same common feature of pricing datasets in a more general way. We plot the bound in Figure 2(c). Applying this bound to $\tilde{y}_2(z)$ of Figure 1, which has $\gamma = 0.05$ and $E(z)$ decreasing, we get 0.79, which indeed bounds the observed ratio of 0.84.

4. Testing the Limits of Predictive Approaches: a Hypothesis Test for Causal-Effect Optimality

In the last sections we saw that predictive approaches, while in fact solving the *wrong* problem, can still often lead to good objective performance nonetheless. But, naturally, there are limits to this. To evaluate the performance of predictive approaches more generally and check whether they can actually be distinguished from optimal in practice, we next develop a hypothesis test for causal-effect optimality. We do this in the spirit of Besbes et al. (2010), who develop a test for objective optimality in *experimental* settings to test whether parametric models suffice to achieve good objective performance. Our test extends that work to the case of observational data and to causal effects. In this section, we treat a slightly more general problem – not just pricing – as we will not necessarily leverage special problem structure to obtain our results. Namely, $r(z)$ may be any function such that $\mathbb{E}[r(z)Y(z)]$ represents the objective value of a decision $z \in \mathcal{Z} \subseteq \mathbb{R}$.[2]

Suppose we wish to test the objective optimality in Problem (1) of a data-driven decision-making algorithm that prescribes the decision $\hat{z}_n$. The data in this case consists of n iid observations of (Z, Y^{obs}) and $\hat{z}_n$ is a map from the observed data to a decision in $\mathcal{Z}$. Later in this section we will also consider including additional variables in each observation. We formulate our test not in terms of whether the models used are consistent for $y(z)$, but rather in terms of whether $\hat{z}_n$ converges upon a price with revenue indistinguishable from optimal. Toward that end, we to define the limit of $\hat{z}_n$:

$$\hat{z} := \operatorname{plim}_{n\to\infty} \hat{z}_n,$$

which we assume exists (we will further assume a rate in Assumption 5). The (fixed) decision $\hat{z}$ can be interpreted as the *full-information* decision algorithm, i.e., the decision that the algorithm $\hat{z}_n$ would pick if it were given infinite data ($n \to \infty$).

We would like to test whether the nominal decision that our data-driven decision-making algorithm is getting at, that is, $\hat{z}$, is optimal or not. Therefore, we would like to test the following null hypothesis H_0 against the alternative H_1:

$$H_0 : R(z^\star) = R(\hat{z}), \quad H_1 : R(z^\star) > R(\hat{z}).$$

That is, we would like to test whether the nominal price that our pricing strategy would be prescribing is generating optimal profits.[3]

A test for the hypothesis H_0 can be interpreted as rejecting a pricing strategy if it *generates profits that are distinguishable from optimal to a statistically significant degree based on the data.*

[2] We may in fact simply set $r(z) = 1$ and let $Y(z)$ generically represent the random reward of decision z. The use of $r(z)$ is only helpful in some special cases such as pricing where there is a known deterministic factor to the reward.

[3] Note that our null hypothesis differs from the one considered by Besbes et al. (2010) in the definition of $R(z)$, that is, our $R(z) = \mathbb{E}[r(z)Y(z)]$ vs. Besbes et al. (2010)'s $\tilde{R}(z) = \mathbb{E}\left[r(z)Y^{\text{obs}} \middle| Z = z\right]$.

4.1. The Setting for Testing H_0

We can repeat the argument in Section 2.1 to see that in the most general settings the truth value of H_0 is not identifiable from observational data. Therefore, we can have little hope of testing it in the most general settings. Instead, we consider a setting where H_0 is identifiable. Specifically, we consider the availability additional variables that can account for the confounding factors and ensure the identifiability of H_0. What this allows us to do is evaluate the hypothesis H_0 in real, practical datasets where such features are available so to understand whether predictive approaches would work in more general settings where such features are either not available or not used.

4.1.1. The On-Line Auto Lending Case. To give an example of such a setting, we consider a case study given by the on-line auto lending data of Columbia University Center for Pricing and Revenue Management (2012) and the customized rate-setting problem considered by Besbes et al. (2010). The on-line auto lending data consists of past sale events where a customer fills out a loan application, if approved an interest rate is quoted (where rate is essentially equivalent to a price), and the customer either accepts or rejects the loan (binary demand). Besbes et al. (2010) study the problem of prescribing interest rates for automobile loans based on this data and customized to each of eight customer segments delineated by predefined ranges of FICO scores, term lengths, and season (see Example 4 for additional detail). The authors estimate predictive models for the functions $\tilde{y}(z)$ (or, $\tilde{R}(z)$) within each segment by separately regressing Y^{obs} (or, $r(z)Y^{\text{obs}}$) on Z based on the data available from each segment. The focus of their study is evaluating whether parametric demand models (specifically, logit) lead to data-driven interest rates that are optimal in the *predictive* Problem (2) by evaluating the rates in a non-parametric (Nadaraya-Watson kernel) regression estimate of the function $\tilde{R}(z)$.

The dataset description says that approval and rate was based on "credit information and other criteria." Such criteria would almost certainly also be associated with the potential likelihood of the consumer to accept a loan offer at any one particular rate (the demand curve). Even if rates are not chosen strategically in response to demand, they could be chosen based on default risk or expected loss, which may be in turn associated with the demand curve. Therefore, we argue that confounding is likely present, i.e., $\mathbb{E}[\epsilon \mid Z] \neq 0$ and hence $y(z) \neq \tilde{y}(z)$ are distinct functions. Moreover, even within each of the eight customer segments considered by Besbes et al. (2010), we would argue that confounding with FICO score is likely present even when conditioned on segment because the FICO ranges considered are wider than a single score and predefined (rather than recursively segmented to fully condition on FICO score such as using a tree model). Moreover, FICO score only account for credit information and not all of the "other criteria" that influence historical rates. Therefore, since the non-identifiability caused by confounding is not a model-specification issue, even the non-parametric estimates may estimate only the predictive objective $\tilde{R}(z)$. The question is whether the resulting interest rates nonetheless perform well in the true objective $R(z)$.

The on-line auto lending data contains much more information about each loan applicant and the associated sale event, including the precise FICO credit score of the applicant, the length of the term over which the loan is to be repaid, the dollar amount of the loan, whether the car to be purchased is new, used, or refinanced, competitors' rate, prime rate, and who referred the applicant. From here on, we let X denote these covariates. If the covariates X encompass the aforementioned "credit information and other criteria," then we may in fact be able to identify $y(z)$ and test H_0, checking the objective optimality of the predictive approaches above, whether parametric or non-parametric.

4.1.2. Identifiability of H_0. We now discuss how we might potentially test H_0. To test it, we actually need it to be identifiable. We next discuss what additional data and assumptions would be sufficient in order to identify the true function $y(z)$, the optimal decision z^*, and the truth value of H_0. Of course, when z^* is identified, one could just estimate it. However, motivated by the above auto-lending setting, we consider a situation where we have such additional necessary data but we only use it in order to test whether, had we in fact ignored it and used a purely predictive approach, would it still generate revenue indistinguishable from optimal. In particular, in the auto-lending setting, we argue such data is available but not in fact used in the setting tested by Besbes et al.

(2010) and therefore we can reuse their experimental setting and further test whether the predictive approaches used therein generate revenues indistinguishable from optimal or not, thereby testing statistically the power and limits of predictive approaches.

Let X denote some covariates observed concurrently with each historical outcome event, as in the example of the autoloan dataset, and let X_i be the observation corresponding to the i^{th} event. For example, in a pricing problem, the covariates X may include, for example, characteristics of the customer, whether a product was featured in a promotional flyer, external signals about demand used for pricing, etc. The hope is that the covariates X can help us disassociate the random variable Z and the particular idiosyncrasies of response curve $Y(z)$, the association between which is the source of confounding. One such sufficient condition for X to account for all such association is the following weak ignorability condition (Hirano and Imbens 2004):

Assumption 1 **(Weak Ignorability)**. *For every $z \in \mathcal{Z}$, we have that, conditioned on X, $Y(z)$ is independent of Z. That is, $\forall z \in \mathcal{Z},\ Y(z) \perp\!\!\!\perp Z \mid X$.*

The condition says that, historically, X accounts for all the event-specific features that may have influenced the decision Z up to idiosyncratic and independent randomness in Z. For example, we argue that X in the on-line auto lending dataset satisfies this condition.

Note that the case of experimental data is when Z is chosen without regard to any specific event (potentially conditioned on history in online learning settings), in which case Assumption 1 holds with a null X variable. This is for example the case in online demand learning and pricing as in Bertsimas and Perakis (2006), Besbes and Zeevi (2009), Harrison et al. (2012), Nambiar et al. (2019) because each sale event is assumed independent and nothing about a present sale event is considered when setting the price. In the observational setting, Assumption 1 essentially requires that the data appears experimental, or as-if random, after conditioning on X. If there is not sufficient recorded information in X to merit Assumption 1, it is said that there is residual endogeneity. There may be other conditions that enable identification such as the availability of instrumental variables (see e.g. Bijmolt et al. 2005). Here we focus on the case where Assumption 1 holds, as in the online auto-lending example.

Under Assumption 1, it is immediate and well-understood that the function $y(z)$ is identifiable:

$$y(z) = \mathbb{E}\left[\mathbb{E}\left[Y(z) \mid X\right]\right] = \mathbb{E}\left[\mathbb{E}\left[Y(z) \mid Z = z, X\right]\right] = \mathbb{E}\left[\mathbb{E}\left[Y(Z) \mid Z = z, X\right]\right] = \mathbb{E}\left[\mathbb{E}\left[Y^{\text{obs}} \mid Z = z, X\right]\right], \quad (3)$$

where the first equality is by iterated expectations, the second is by Assumption 1, the third is by $Z = z$, and the last by $Y^{\text{obs}} = Y(Z)$. Note the outer expectation in equation (3), $\mathbb{E}\left[\mathbb{E}\left[Y^{\text{obs}} \middle| Z = z, X\right]\right]$, is not conditioned on $Z = z$ and cannot be marginalized via iterated expectations. In words, it says to take the average of the conditional expectation of Y^{obs} given $X = x$, $Z = z$ over all x using the *marginal* distribution of X. This is because, per Assumption 1, in each stratum of $X = x$, we have the independence of $Y(z)$ and Z. Thus, in each stratum, we can identify the effect of the decision Z by using a regression, $\mathbb{E}\left[Y(z) \mid X\right] = \mathbb{E}\left[Y^{\text{obs}} \mid Z = z, x\right]$. To put the pieces together to get the marginal effect, we should average the within-strata estimates according to the frequency of the strata, which does not involve Z.

The right-most side of eq. (3) is expressed solely in terms of the joint distribution of (Z, X, Y^{obs}), which gives the identifiability $y(z)$ from observations thereof and hence the objective function $R(z)$. Similarly, the optimal decision is given by optimizing $R(z)$ and the truth value of H_0 is given by plugging in values to the objective function $R(z)$. We summarize this as follows.

Theorem 5. *Under Assumption 1, the mean-response curve $y(z)$, optimal decision $z^\star$, and the truth value of H_0 are all identifiable on the basis of observations of (X, Y^{obs}, Z).*

Let us consider Assumption 1 and its ramifications in an example.

Example 2 (Consulting for the MIT Coop). Consider again the hypothetical case of Example 1. Recall, Alice and Bob both came up with models for demand that completely agreed with the data but gave rise to different optimal prices. Thus, we concluded that the data observed could not possibly identify the right optimal price.

Suppose Assumption 1 holds with X being a null variable, i.e., without any extra information. This condition eliminates Bob's model – it no longer agrees with both the data and this condition.

Table 3 Data for Example 2

	$X=0$		$X=1$	
	$Z=20$	$Z=28$	$Z=20$	$Z=28$
$Y^{\text{obs}}=0$	0	4/9	0	0
$Y^{\text{obs}}=1$	4/9	2/45	0	1/90
$Y^{\text{obs}}=2$	0	0	1/18	0

On the other hand, Alice's model remains valid – in fact it turns out to be the unique model that agrees with both the data and this condition. Hence, under this condition, $z^\star = 20$ is the correct optimal price. But for Assumption 1 to hold with X being a null variable we would have needed experimental data, where prices are set at random for the sake of experiment.

Suppose the data is not experimental in that Assumption 1 does not hold with X being a null variable. Suppose instead that we recorded additional information about each sale event: whether there was a major home game that day ($X = 1$) or not ($X = 0$). On average, there is a game 2 days of each month ($\mathbb{P}(X=1) = 2/30$). Suppose tallying the historical observations led to the summary of the data shown as the frequencies in Table 3. If we assume Assumption 1 holds with this X, then it turns out that Alice's model is ruled out and Bob's model is the unique model that accommodates both this condition and the data observed, in which case $z^\star = 28$ is the correct optimal price. In this hypothetical example, we are seeking a universal price, to be set a priori without regard to whether there is a game, but the price could also potentially be customized.

With H_0 identifiable, what remains is to construct a hypothesis test. To do so, we first construct non-parametric estimates for the function $R(z)$ and the optimal decision $z^\star$, which will allow us to construct a test statistic.

4.2. Non-Parametric Estimates of $R(z)$ and $z^\star$

In this section, we study non-parametric estimates of $R(z)$ and $z^\star$ under Assumption 1 that are model-independent in that they will converge to the true objective and optimal decision regardless of the true underlying distribution, given sufficient data and some regularity assumptions. Henceforth, we assume that X is a vector of covariates taking values in $\mathbb{R}^k$.

The proof of Theorem 5 says that, under Assumption 1, the objective function can be written as $R(z) = \mathbb{E}\left[\mathbb{E}\left[r(Z)Y^{\text{obs}} \middle| Z = z, X\right]\right]$. Thus, to estimate it, one approach may be to estimate the regression function $\mathbb{E}\left[r(Z)Y^{\text{obs}}\middle|Z=z,X=x\right]$ and then average the estimated function over an estimate for the marginal distribution of X. Then, the optimizer of this estimate can be used as an estimator for the optimal decision.

First, one non-parametric estimate of the marginal distribution of X is simply the empirical distribution, which places unit mass at each of the observations X_i. Second, to estimate the regression function $\mathbb{E}\left[r(Z)Y^{\text{obs}}\middle|Z=z,X=x\right]$ non-parametrically, we can use Nadaraya-Watson kernel regression (Nadaraya 1964, Watson 1964). The estimate, based on a choice of kernel $K:\mathbb{R}^{\ell+k}\to\mathbb{R}$ and bandwidth $h_n>0$, is

$$\overline{R}_n(z,x)=\frac{\sum_{i=1}^n K(\frac{z-Z_i}{h_n},\frac{x-X_i}{h_n})r(Z_i)Y_i^{\text{obs}}}{\sum_{i=1}^n K(\frac{z-Z_i}{h_n},\frac{x-X_i}{h_n})}, \tag{4}$$

where $K(\frac{z-Z_i}{h_n},\frac{x-X_i}{h_n})=K(\frac{z-Z_i}{h_n},\frac{x_1-X_{i1}}{h_n},\ldots,\frac{x_k-X_{ik}}{h_n})$. This regression estimator arises as the conditional expectation with respect to the Parzen window density estimator (Parzen 1962) for the joint density of $(Z,X,r(Z)Y^{\text{obs}})$. There are a variety of possible kernels (Härdle 1990).

Combining the two estimators as detailed above, we arrive at the following estimate for the objective function

$$\overline{R}_n(z)=\frac{1}{n}\sum_{i=1}^n\overline{R}_n(z,X_i)=\frac{1}{n}\sum_{i=1}^n\frac{\sum_{j=1}^n K\left(\frac{z-Z_j}{h_n},\frac{X_i-X_j}{h_n}\right)r(Z_j)Y_j^{\text{obs}}}{\sum_{j=1}^n K\left(\frac{z-Z_i}{h_n},\frac{X_i-X_j}{h_n}\right)}. \tag{5}$$

Optimizing the above estimate over $\mathcal{Z}$ yields a non-parametric observational-data-driven decision-making algorithm

$$\overline{z}_n\in\arg\max_{z\in\mathcal{Z}}\overline{R}_n(z). \tag{6}$$

One question that arises is how do these decisions behave asymptotically. In particular, does this algorithm lead to a decision and objective value that converge to the optimal decision and objective value. Since the estimates are non-parametric, the hope is that this can occur under model-free assumptions. Next we show that this is indeed the case under the following regularity conditions.

Assumption 2 **(Kernel Conditions)**.

a. $0<\int_{\mathbb{R}^{1+k}}K(u)\,du<\infty$.

b. *K is zero outside a bounded set.*

c. *K is twice Lipschitz-continuously differentiable.*

d. *K has order at least $s \in \mathbb{N}$, that is, $\int K(u)u^{\alpha}du = 0 \quad \forall \alpha \in \mathbb{N}^{1+k} : |\alpha| < s$.*

e. *$h_n \to 0$ and $nh_n^{2s+3} \to 0$.*

f. *$nh_n^{k+5}/\log(n) \to \infty$ and $nh_n^{2k+1}/\log(n)^2 \to \infty$.*

Assumption 3 **(Optimality Conditions)**.

a. *$\mathcal{Z} \subseteq \mathbb{R}$ is compact.*

b. *$z^\star$ uniquely maximizes $R(z)$ on $\mathcal{Z}$.*

c. *$z^\star$ lies in the interior of $\mathcal{Z}$.*

d. *$R(z)$ is twice continuously differentiable and $R''(z^\star) < 0$.*

Assumption 4 **(Distributional Conditions)**.

a. *X and Z are continuously distributed on a compact support where the joint density, $f_{Z,X}(z,x)$, is bounded away from zero.*

b. *The marginal density of X, $f_X(x)$, is bounded and continuously differentiable.*

c. *$\mathbb{E}\left[(Y^{\text{obs}})^4\right] < \infty$ and $\mathbb{E}\left[(Y^{\text{obs}})^4 \middle| Z = z, X = x\right]$ is bounded.*

d. *$\mathbb{E}\left[(Y^{\text{obs}})^2 \middle| Z = z, X = x\right]$ is continuously differentiable.*

e. *$\mathbb{E}\left[Y^{\text{obs}} \middle| Z = z, X = x\right]$ and $f_{Z,X}(z,x)$ are $s+1$ times continuously, boundedly differentiable.*

Assumptions 2 and 4 are used directly to satisfy the various conditions of the results of Newey (1994), which establishes certain results for partial mean kernel estimators of the form of (5). Assumption 2 is a specification of how one should choose the kernel and bandwidth, rather than an assumption about the data or the problem. In particular, Assumption Assumption 4 constitutes regularity conditions on the data distribution. The assumption requires certain continuity and smoothness, which guarantee that similarity-based approaches like kernel regression work. Assumption 3 provides the necessary conditions to apply delta-method-type approaches to studying asymptotic distributions of estimators of z^* given by maximizing estimators of $R(z)$. Under these conditions, we can show the following asymptotic optimality and rates.

THEOREM 6. *Under Assumptions 1, 2, 3, and 4, we have that*

$$\sqrt{nh_n}(R(z) - \overline{R}_n(z)) \xrightarrow{d} \mathcal{N}(0, \eta_z \kappa) \quad \forall z \in \mathcal{Z},$$
$$\sqrt{nh_n^3}(z^\star - \overline{z}_n) \xrightarrow{d} \mathcal{N}\left(0, \frac{\eta_{z^\star}\kappa'}{R''(z^\star)^2}\right)$$
$$\left(nh_n^3\right)(R(z^\star) - R(\overline{z}_n)) \xrightarrow{d} \frac{-\eta_{z^\star}\kappa'}{2R''(z^\star)}\chi_1^2,$$

and, if also $nh_n^{2s+1} \to 0$*, then*

$$\sqrt{nh_n}(R(z^\star) - \overline{R}_n(\overline{z}_n)) \xrightarrow{d} \mathcal{N}(0, \eta_{z^\star}\kappa),$$

where $\mathcal{N}(0, \sigma^2)$ *denotes a centered normal distribution with variance* σ^2*,* χ_1^2 *denotes a chi-squared distribution with one degree of freedom, and* η_z*,* κ*,* κ' *are constants defined as follows*

$$\eta_z = r(z)^2\ \mathbb{E}\left[\frac{\mathrm{Var}\left(Y^{\mathrm{obs}}\middle|Z = z, X\right)}{f_{Z|X}(z|X)}\right], \quad \kappa = \int \tilde{K}(z)^2 dz, \quad \kappa' = \int \tilde{K}'(z)^2 dz,$$

where $\tilde{K}(z) = \int K(z, x)dx$ *and* $f_{Z|X}(z|x) = f_{Z,X}(z, x)/f_X(x)$ *is the conditional density of* Z*.*

The proof is given in the appendix. The constants κ, κ' capture the roughness of the effective kernel on Z, $\tilde{K}(z)$. The constant η_z captures the variance of the uncertain variable Y^{obs}, given X.

The main implication of Theorem 6 is that, under regularity conditions but without model specification, the non-parametric solution $\overline{z}_n$ has objective value that converges to the optimal z^* at a rate of $1/\sqrt{nh_n^3}$, with expected revenue shortfall converging to zero at a rate of $1/(nh_n^3)$.

4.3. A Test Statistic and Large Sample Theory

The impediment to verifying our hypothesis H_0 is that $R(z)$, $z^\star$, and $\hat{z}$ are all unknown; were they known, we would compute $\rho = R(z^\star) - R(\hat{z})$ and compare it to 0. Therefore, we must come up with an observable test statistic as a proxy to ρ. We do this by replacing the unknowns by our consistent estimates for them. We replace $R(z)$ and $z^\star$ by our non-parametric estimates $\overline{R}(z)$ as in (5) and $\overline{z}$ as in (6) and we replace $\hat{z}$ by the estimate $\hat{z}_n$. The resulting test statistic is

$$\rho_n = \overline{R}_n(\overline{z}_n) - \overline{R}_n(\hat{z}_n). \tag{7}$$

Since ρ_n estimates ρ, if ρ_n is small, we have reason to believe that $\rho = 0$, whereas if ρ_n is large, we would believe that $\rho > 0$. The question is where to draw the line.

The next result establishes the asymptotic limiting distribution of our test statistic, ρ_n, which can be used to determine what counts as small and what counts as large. To establish this result we need make one additional assumption about $\hat{z}_n$: that it converges to $\hat{z}$ just slightly faster than $\overline{z}_n$ converges to z^*.

ASSUMPTION 5 **(Convergent Decision-Making)**. $\sqrt{nh_n^3}(\hat{z}_n - \hat{z}) \xrightarrow{\mathbb{P}} 0$.

For example a parametric predictive approach that fits the regression $\tilde{y}(z)$ using maximum likelihood and plugs in the estimate into Problem (2) will have that $a_n\sqrt{n}(\hat{z}_n - \hat{z}) \xrightarrow{\mathbb{P}} 0$ for *any* $a_n \to 0$ (note that this holds regardless of correct parametric specification).[4] Since Assumption 2 requires $h_n \to 0$, setting $a_n = h_n^{3/2}$, the condition in Assumption 5 is satisfied. The same occurs for a non-parametric predictive approach. Consider the case of $s = 2$ and a method, as in Besbes et al. (2010), that fits $\tilde{R}(z)$ using a Nadaraya-Watson kernel regression with a second-order kernel and plugs in the estimate into Problem (2). Then, Ziegler (2002) shows that for any $a_n \to 0$ one can achieve $\hat{z}_n$ satisfying $a_n n^{\frac{2}{7}}(\hat{z}_n - \hat{z}) \xrightarrow{\mathbb{P}} 0$ (under some regularity; moreover $\hat{z} = \tilde{z}$). Since Assumption 2 requires $nh_n^7 \to 0$, setting $a_n = (nh_n^7)^{3/14}$, the condition in Assumption 5 is satisfied.

With this additional assumption, we have the following result.

THEOREM 7. *Suppose Assumptions 1, 2, 3, 4, and 5 hold. Let* $\Gamma = \frac{-\eta_{z^\star}\kappa'}{2R''(z^\star)}$. *Then,*

i. *under* H_0, $(nh_n^3)\,\rho_n \xrightarrow{d} \Gamma\chi_1^2$, *and*

ii. *under* H_1, $(nh_n^3)\,\rho_n \xrightarrow{d} \infty$.

The proof is given in the appendix.[5]

[4] This can bee seen as a consequence of Lemma 2 of Besbes et al. 2010 and the Lipschitz condition shown in the proof of their Lemma 3.

[5] Additionally, a more general version of the theorem and the resulting test under a relaxation of Assumption 2 is given in Section EC.2 of the appendix. The relaxation allows $\hat{z}_n$ to converge as fast as $\overline{z}_n$ rather than strictly faster, but leads to a more complicated test.

Theorem 7 says that if we only reject H_0 when $\rho_n > n^{-1}h_n^{-3}\Gamma F^{-1}_{\chi^2_1}(1-\alpha)$ (where $F^{-1}_{\chi^2_1}$ is the chi-squared quantile function), then when H_0 is true we would only falsely reject H_0 at most α fraction of the time (asymptotically). On the other hand, if H_0 is false, then we would eventually reject it using such a procedure (a property known as *consistency* of a hypothesis test). The problem is that Γ is unknown meaning that this exact procedure cannot be implemented in practice.

4.4. A Hypothesis Test

One way to implement a hypothesis is to estimate Γ and replace the estimate into the results of Theorem 7. In particular, given any estimate $\hat{\Gamma}_n$ that converges in probability to Γ, we would have as an immediate consequence of Theorem 7 that $(nh_n^3)\,\hat{\Gamma}_n^{-1}\rho_n$ converges in distribution to χ^2_1 under H_0 and to ∞ under H_1. This would give an implementable test. Non-parametric estimators for Γ can be derived from the results of Newey (1994), however, they would be convoluted and unwieldy, involving partial means of estimators of conditional variance and density as well as fragile estimates of second derivatives of partial means.

Instead, we use the bootstrap (Efron and Tibshirani 1993) and the following observation.

THEOREM 8. *Suppose Assumptions 1, 2, 3, and 4 hold. Let* $A_n = \overline{R}_n(\overline{z}_n) - \overline{R}_n(z^\star)$. *Then,* $(nh_n^3)\,A_n \xrightarrow{d} \Gamma\chi^2_1$.

The proof is given in the appendix.

So, to estimate Γ, we use a scaled estimate of the mean of A_n. In the spirit of Besbes et al. (2010), we use the bootstrap to achieve this. The fact that A_n is asymptotically pivotal suggests that a bootstrap procedure could be particularly powerful (Horowitz 2001). This bootstrap procedure is also more attractive than convoluted kernel estimates of Γ because it is less dependent on parameters and it deals more directly with the finite-sample distribution of ρ_n.

Given data $\mathcal{S}_n = \{(X_1, Y_1^{\text{obs}}, Z_1), \dots, (X_1, Y_1^{\text{obs}}, Z_1)\}$, $\hat{z}_n$, and a significance $\alpha \in (0,1)$, the full hypothesis test for H_0 proceeds as follows:

1. Compute $\overline{R}_n$ and $\overline{z}_n$ as in equations (5)–(6) based on data $\mathcal{S}_n$.
2. Compute ρ_n as in equation (7).

3. Fix B large. For $b = 1, \ldots, B$:
   a. Draw n samples with replacement from $\mathcal{S}_n$ to form the resampled dataset $\mathcal{S}_n^{(b)}$.
   b. Compute $\overline{R}_n^{(b)}$ and $\overline{z}_n^{(b)}$ as in equations (5)–(6) based on the data $\mathcal{S}_n^{(b)}$.
   c. Set $A_n^{(b)} = \overline{R}_n^{(b)}(\overline{z}_n^{(b)}) - \overline{R}_n^{(b)}(\overline{z}_n)$.
4. Let $\hat{\Gamma}_n = \frac{nh_n^3}{B} \sum_{b=B}^{n} A_n^{(b)}$.
5. Return p-value $p = 1 - F_{\chi_1^2}(\hat{\Gamma}_n^{-1} n h_n^3 \rho_n)$ and reject H_0 if $p < \alpha$.

To summarize, in this the section, we developed a hypothesis test for the hypothesis that an observational-data-driven decision $\hat{z}_n$ (or, its probability limit point), has objective value that is indistinguishable from optimal in Problem (1) to a statistically significant degree. We developed the test in the setting where Assumption 1 holds and used a non-parametric three-step optimized partial-means kernel estimator to construct a test statistic for the hypothesis. In Section 6, we apply this test to both synthetic and real data, but first we review a parametric solution to which we will also compare.

5. A Parametric Solution

In the preceding section we developed a non-parametric decision-making algorithm that converged to optimal without requiring any model to be specified. Non-parametric approaches, however, can sometimes be unwieldy because their shapelessness makes them uninterpretable and they may be slow to converge. In fact, there is a growing body of work (Besbes et al. 2010, Besbes and Zeevi 2015) arguing that parametric models are often sufficient for managerial decision-making problems, as the model may need only fit well near the optimum or even just induce an acceptable decision, whether or not the model is correct. In particular, what matters is not model fit but objective performance. In this section, we present a parametric way to estimate $y(z)$ from observational data under Assumption 1. The intention is to be able to study the power of parametric approaches that account for confounding (when possible) in the observational data setting and see if they suffice there as well for decision-making purposes.

Hirano and Imbens (2004) define the *generalized propensity score* as $Q = f_{Z|X}(Z, X)$, assuming the conditional density $f_{Z|X}(z|x)$ exists. That is, we take the conditional density $f_{Z|X}(z|x)$, which is non-random but unknown, and plug in as values the random variables Z and X. The results of Hirano and Imbens (2004) show that we can use this to identify the mean-response function $y(z)$ as follows under Assumption 1:

$$y(z) = \mathbb{E}\left[y(z, f_{Z|X}(z, X))\right], \text{ where } y(z,q) = \mathbb{E}\left[Y^{\text{obs}} \middle| Z = z, Q = q\right].$$

The implication is that it is sufficient to control just for the univariate Q rather than all of X.

This motivates the following general parametric estimation strategy:

1. Regress Z on X by fitting a generalized linear model (GLM) in order to estimate $f_{Z|X}(z,x)$. I.e., choose $\hat{\beta}_n$, $\hat{\tau}_n$ by maximum likelihood estimation, given the parametric model

$$f_{Z|X}(z|x;\beta,\tau) = h(z,\tau)\exp\left(\frac{b(\beta_0 + \beta^T x)T(z) - A(\beta_0 + \beta^T x)}{d(\tau)}\right).$$

   See McCullagh et al. (1989) for choices of b, T, A, d, h. For example, the choices $b(\mu) = \mu$, $T(z) = z,\ A(\mu) = \mu^2/2,\ d(\tau) = \tau^2$, and $h(z,\tau) = \frac{1}{\sqrt{2\pi}\tau} e^{-\frac{z^2}{2\tau^2}}$ lead to ordinary least squares (OLS). Other examples of GLMs include logistic regression, Poisson regression, Gamma regression, and loglinear regression.
2. Use the fitted GLM to impute generalized propensity scores, setting $\widehat{Q}_i = f_{Z|X}(Z_i|X_i;\hat{\beta}_n,\hat{\tau}_n)$.
3. Regress Y^{obs} on Z and $\widehat{Q}$ based on the imputed data $\{(Z_i, Y_i^{\text{obs}}, \widehat{Q}_i) : i = 1, \ldots, n\}$ using another GLM (e.g., linear or logistic regression) to produce an estimate $\hat{y}_n(z,q)$ of $y(z,q)$. For example, we can fit $Y^{\text{obs}} = b^{-1}(\alpha_0 + \alpha_1 z + \alpha_2 q + \alpha_3 q^2 + \epsilon)$ via link function b (e.g., if a log-log demand model is appropriate as in many pricing problems, we regress $\log(Y^{\text{obs}})$ on $\log(Z)$ and $\widehat{Q}$).
4. Use these to estimate the mean-response curve and prescribe the decision that optimizes the estimated objective,

$$\hat{z}_n \in \arg\max_{z\in\mathcal{Z}} \left\{ r(z) \times \tfrac{1}{n}\textstyle\sum_{i=1}^n \hat{y}_n(z, \hat{f}_{Z|X}(z|X_i;\hat{\beta}_n,\hat{\tau}_n)) \right\}.$$

The above procedure provides a flexible parametric framework for computing $\hat{z}_n$ from observational data under Assumption 1. When we apply it to examples with both real and synthetic data in Section 6, we find that it performs well and produces rewards that are often statistically indistinguishable from optimal.

Figure 3 Predictive demand and revenue in Example 3

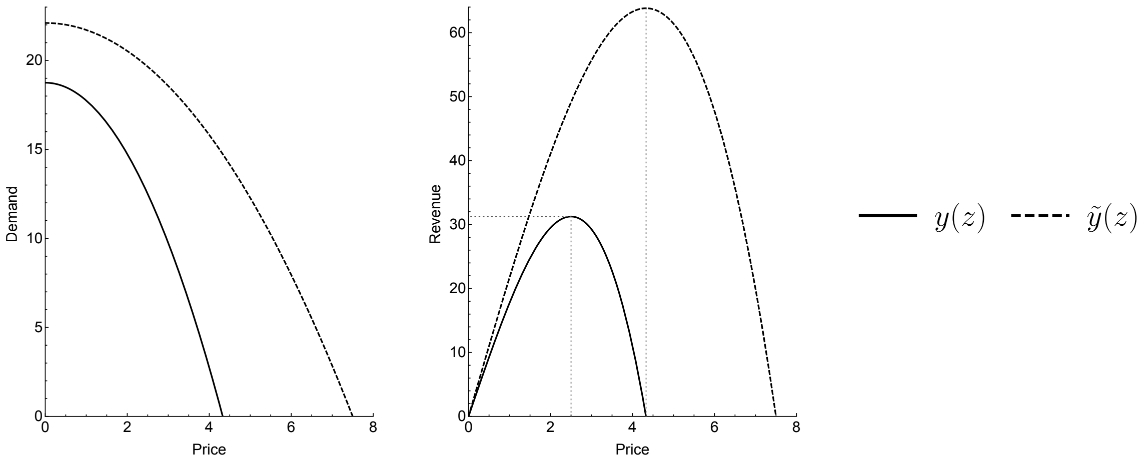

(a) Demand (b) Revenue

6. Empirical Investigation

In this section, we first use simulated data to investigate how our test performs in a controlled environment and then use our test to study the power and limits of both predictive and parametric approaches to observational-data-driven decision-making in a real example.

Example 3 (Simulated Example). Consider a pricing instance of Problem (1) with procurement cost $c=0$, potential prices $\mathcal{Z}=(0,\infty)$, and random demand curve

$$Y(z)=27.75-z^2+6Xz-9X^2+V, \tag{8}$$

where $X\sim\mathcal{N}(0,1)$, $V\sim\mathcal{N}(0,\sigma^2)$ are normal noise, and prices historically set as $Z=3X+W$ where $W\sim\mathcal{N}(0,\tau^2=15.1234)$.

The best predictor of Y^{obs}, given an observation of $Z=z$, is

$$\tilde{y}(z)=27.75-z^2+6z\mathbb{E}\left[X\middle|3X=z-W\right]-9\mathbb{E}\left[X^2\middle|3X=z-W\right]+\mathbb{E}[V]=22.108-0.393z^2,$$

which we get by plugging in Z into (8) and recognizing $(X\mid 3X=z-W)\sim\mathcal{N}\left(\frac{3z}{9+\tau^2},\,\frac{\tau^2}{9+\tau^2}\right)$. This is exactly the function we would arrive at if we used data to regress demand on price (e.g., by linear regression on Z, Z^2 or by non-parametric regression on Z). However, the expected demand when the price is set to z is a different function,

$$y(z)=27.75-z^2+6z\mathbb{E}[X]-9\mathbb{E}\left[X^2\right]+\mathbb{E}[V]=18.75-z^2,$$

Figure 4 Comparing Predictive and Prescriptive Data-Driven Pricing Strategies

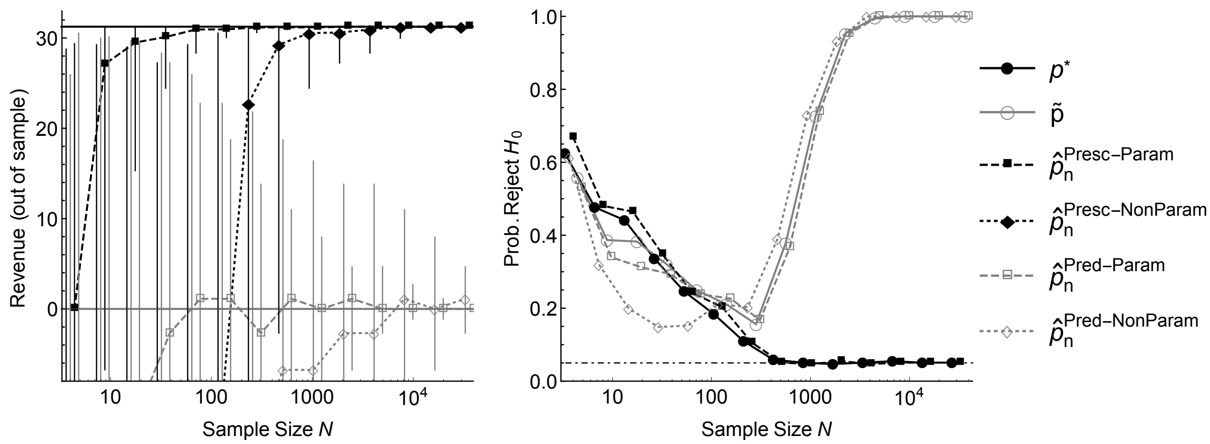

(a) Revenues (b) Test results

which we get by taking the expectation of (8). We plot these two functions in Figure 3(a).

Now consider the price optimization problem. The true profit function, $R(z)$, is optimized at $z^\star = 2.5$ with a value of $R(z^\star) = 31.25$. On the other hand, a predictive approach optimizes $\tilde{R}(z)$, leading to the price $\tilde{z} = 4.330$, which leads to exactly $R(\tilde{z}) = 0$ profit under the true profit function. We plot R, $\tilde{R}$, $z^\star$, and $\tilde{z}$ in Figure 3(b).

We next compare four different observational-data-driven pricing methods: a prescriptive non-parametric approach $\hat{z}_n^{\text{Presc-NonParam}}$ based on n observations of (X, Y^{obs}, Z), a prescriptive parametric approach $\hat{z}_n^{\text{Presc-Param}}$ based on n observations of (X, Y^{obs}, Z), a predictive non-parametric approach $\hat{z}_n^{\text{Pred-NonParam}}$ based on n observations of (Y^{obs}, Z), and a predictive parametric approach $\hat{z}_n^{\text{Pred-Param}}$ based on n observations of (Y^{obs}, Z). We also consider the (non-data-driven) true optimal price $z^\star$ of (1) and full-information predictive pricing strategy $\tilde{z}$ of (2). The prescriptive non-parametric strategy is as in (6) using a second order Gaussian kernel $K(u) = e^{-\frac{||u||_2^2}{2h_n^2}}$ and $h_n = 0.1 \times (n \log(n))^{-1/7}$, which satisfy Assumption 2 with $s = 2$, $k = 1$. For the prescriptive parametric strategy, we follow our procedure from Section 5 using OLS linear regression of Z on X for the GLM in step 1 and an OLS linear regression of Y^{obs} on Z and $\log(\widehat{Q})$ in step 3. For the predictive non-parametric strategy, we use kernel regression to regress $r(Z)Y^{\text{obs}}$ on Z (using the same kernel and bandwidth $h_n = 2.5 \times (n \log(n))^{-1/7}$). Finally, for the predictive parametric strategy, we perform OLS linear regression of Y^{obs} on Z.

First, we consider the profit performance of each of these strategies. We plot the corresponding out-of-sample profits, $R(\hat{z}_n)$, along with optimal profit $R(z^\star)$, in Figure 4(a). The plot displays the median profit (center lines) and the 10^{th} and 90^{th} percentiles (vertical lines) over 256 replicate runs of each sample size. We see that the predictive approaches, by design, have very low revenues because of significant confounding. On the other hand, the prescriptive parametric approach offers significantly better out-of-sample performance than the non-parametric approach for small samples. In this example, the parametric approach is well-specified by design.

Next, we apply our hypothesis test for profit optimality. We plot the frequencies of rejecting a pricing strategy as significantly suboptimal at a significance of 0.05 in Figure 4(b). The plot displays the fraction of times the null hypothesis is rejected out of 256 replicate runs of each sample size. We see that with sufficient data, the test can distinguish those pricing strategies that generate suboptimal profits (i.e. solely predictive strategies) from those that cannot be distinguished from optimal for all prescriptive intents and purposes. In particular, it takes a few hundred data points before the test has the desired significance of 0.05 (i.e., $z^\star$ is rejected no more than 5% of the time).

Example 4 (Auto Loan Rate Optimization). Consider again the online auto loan dataset from Section 4.1.1. In Besbes et al. (2010), the authors consider whether a parametric model suffices for the problem of fixed pricing within various customer segments of loan applicants, defined in terms of three factors:

1. FICO score: (690, 715] (range 1) or (715, 740] (range 2),
2. Loan term in months: ≤ 36 (class 1), (36, 48] (class 2), (48, 60] (class 3), or > 60 (class 4).
3. Season: first half of data (half 1) or second half (half 2).

Customers with FICO scores outside of (690, 740] are not considered (see Besbes et al. 2010, for reasoning). Term classes 2 and 4 are not considered either, but we consider these here. The authors use a per-unit profit function $r(z) = z - 2\%$. Within each segment, the authors' approach is to estimate (either parametrically or non-parametrically) the conditional expectation of demand given price and to optimize per-unit profit times this conditional expectation, i.e., solving an estimate

Table 4 Testing Revenue Optimality in the Auto Loan Rate Optimization Example

		FICO range 1 (690, 715]		FICO range 2 (715, 740]		
		Half 1	Half 2	Half 1	Half 2	
n		1359	732	1386	781	Term cl. 1
ρ_n (p-val)	Prescriptive, Param	0.37 (0.15)	0.21 (0.11)	0.24 (0.49)	0.23 (0.018*)	
	Predictive, Param	0.86 (0.030*)	0.50 (0.013*)	0.25 (0.48)	0.70 (< 0.001***)	
	Predictive, Non-Param	1.91 (0.0012**)	1.51 (< 0.001***)	1.6 (0.07)	1.35 (< 0.001***)	
n		1394	832	1327	690	Term cl. 2
ρ_n (p-val)	Prescriptive, Param	0.23 (0.21)	0.18 (0.073)	0.28 (0.053)	0.074 (0.67)	
	Predictive, Param	0.87 (0.015*)	0.24 (0.039*)	0.35 (0.033*)	0.051 (0.73)	
	Predictive, Non-Param	1.6 (0.0011**)	1.59 (< 0.001***)	1.19 (< 0.001***)	1.76 (0.040*)	
n		4495	3147	3803	2865	Term cl. 3
ρ_n (p-val)	Prescriptive, Param	0.55 (0.32)	0.26 (0.33)	0.088 (0.061)	1.4 (0.066)	
	Predictive, Param	1.19 (0.14)	0.22 (0.37)	0.28 (< 0.001***)	1.89 (0.034*)	
	Predictive, Non-Param	1.19 (0.14)	1.1 (0.046*)	0.86 (< 0.001***)	2.49 (0.015*)	
n		2347	1506	1834	1206	Term cl. 4
ρ_n (p-val)	Prescriptive, Param	0.40 (0.0071**)	0.0059 (0.63)	1.86 (0.30)	0.14 (0.46)	
	Predictive, Param	0.27 (0.026*)	0.045 (0.19)	2.19 (0.26)	0.31 (0.28)	
	Predictive, Non-Param	1.5 (< 0.001***)	1.54 (< 0.001***)	2.92 (0.19)	1.7 (0.012*)	

* denotes reject H_0 at significance 0.05, ** at 0.01, and *** at 0.001. Gray denotes p-value ≥ 0.05.

of Problem (2). Using a test that compares the parametric and non-parametric approaches, they conclude that a parametric model often suffices.

We consider the same problem again here, paying closer attention to the observational nature of the data. In Section 4.1.1 we argued that even within each segment, the data cannot be treated as experimental (i.e., satisfying Assumption 1 with respect to segment alone) and therefore that purely predictive approaches may not be estimating the true demand-price response function. We now use our hypothesis test to determine whether this distinction is moot from a profit-generated point of view. We also test whether our parametric prescriptive approach from Section 5 is successful. For the predictive approaches, we reproduce those in Besbes et al. (2010): kernel regression with the Gaussian kernel (non-parametric) and logistic regression (parametric). For our parametric prescriptive approach we fit a log-normal model for price via linear regression on X, i.e., $\left(\log(Z)\middle|X=x\right)\sim\mathcal{N}\left(\beta_0+\beta_1^Tx,\,\sigma^2\right)$, and we fit a logistic regression for demand that is linear in price and quadratic in generalized propensity score, i.e., $\hat{y}(z,q)=\left(1+e^{-\hat{\alpha}_0-\hat{\alpha}_1z-\hat{\alpha}_2q-\hat{\alpha}_3q^2}\right)^{-1}$.

We let X consist of FICO score, the loan amount, the loan term, whether the car is new or used, whether the loan is refinancing, and if so what was the previous rate (otherwise 0). In our assessment, each of these covariates has direct impact on the interest rate quoted to applicants and each can arguably impact the decision of the applicant to accept any one rate. At the same time, this summarizes all relevant data provided and thus encapsulates all customer-specific information that could have gone into a rate quote decision. Therefore, we reason that Assumption 1 holds with respect to X, while it is likely to fail with respect to any subset of X.

We run the test within each of the 16 customer segments. In Table 4, we report the estimated suboptimality ρ_n and its corresponding p-value according to our bootstrap procedure with $B = 100$ draws.

The results for the parametric predictive approach demonstrate both the power and the limits of predictive approaches. The profits are rejected as suboptimal to a statistically significant degree at $p < 0.05$ in only 9 of 16 segments, and and at $p < 0.001$ in only 2 segments. More importantly, the profit suboptimality estimate (ρ_n) is often small in magnitude. This shows that a good predictive approach can yield reasonable performance, as indicated by our results in Section 3, but not without some limits.

The non-parametric predictive approach, on the other hand, is rejected as suboptimal to a statistically significant degree at $p < 0.05$ in 13 of 16 segments, and at $p < 0.001$ in 7 segments. This highlights the difficulty of using unwieldy non-parametric models and suggests a preference for simple parametric models, even if misspecified from a statistical point of view.

Finally, the results for the parametric prescriptive approach add another dimension to the discussion on the sufficiency of parametric models for decision making. In the rare cases where it is in fact possible to identify the true average demand curve $y(z)$, parametric approaches to estimating it by controlling for covariates perform quite well from an objective-optimality standpoint. Our prescriptive parametric approach passes the test at $p \geq 0.05$ in all but 2 segments, in each of which, both predictive approaches also failed the test. Note that if the null hypothesis were actually true in *all* 16 segments (i.e., our prescriptive parametric approach is revenue optimal), we would still reject the null in two or more segments about 19% of the time. Moreover, we do not actually expect the null to be exactly true and we should expect a higher bar for optimality for large n. Importantly, even in the rejected segments, the estimated suboptimality is particularly small. Generally, we see that the estimated suboptimality of our prescriptive parametric approach is smaller than that of the non-parametric predictive approach in all segments and than the parametric predictive approach in all but 3 segments. Averaging the estimated suboptimalities of each approach over all segments (weighted appropriately by the segment's n) and comparing, we see that on average, our prescriptive parametric approach recoups 70% of the total profits lost by using a non-parametric predictive approach and 36% of those by a parametric predictive approach.

7. Conclusions

We studied the data-driven optimization problem that arises from observational data. Noting that this problem is not generally well specified by the data, we considered another problem that is in fact the problem addressed by commonly used predictive approaches. While the two problems are different in their objective functions, we argued that predictive approaches will still work to the extent that their performance in the true problem is good. To quantify the power of predictive approaches, we focused on the pricing instance of the decision-making problem and proved strong performance guarantees for predictive approaches by leveraging the special structure of the pricing problem as well as the special characteristics of pricing data.

To study the potential limits of predictive approaches, we developed a hypothesis test that can determine whether an observational-data-driven decision is indistinguishable from optimal in terms of objective to a statistically significant degree. The test was based on a partial-mean non-parametric estimate of the objective under ignorability and results on its asymptotics. Applying the test to an online auto loan dataset, we found that predictive approaches have both power and limits: the revenues generated were rejected as suboptimal only roughly half the time and when they were the suboptimality was still small. Of course, we should expect some suboptimality since the true optimal price would not be identifiable given only observational data on prices and demands. As an additional conclusion, we found that in the rare cases where everything is in fact identifiable by controlling for covariates, parametric approaches perform well, adding a new dimension to the finding that parametric approaches, even if misspecified, often suffice for decision making.

References

Bausch J (2013) On the efficient calculation of a linear combination of chi-square random variables with an application in counting string vacua. *Journal of Physics A: Mathematical and Theoretical* 46(50):505202.

Berry S, Levinsohn J, Pakes A (1995) Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society* 841–890.

Bertsimas D, Kallus N (2020) From predictive to prescriptive analytics. *Management Science* 66(3):1025–1044.

Bertsimas D, Perakis G (2006) Dynamic pricing: A learning approach. Lawphongpanich S, Hearn DW, Smith MJ, eds., *Mathematical and Computational Models for Congestion Charging*, volume 101 of *Applied Optimization*, 45–79 (Springer).

Besbes O, Phillips R, Zeevi A (2010) Testing the validity of a demand model: An operations perspective. *Manufacturing & Service Operations Management* 12(1):162–183.

Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* 57(6):1407–1420.

Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Science* 61(4):723–739.

Bijmolt TH, Heerde HJv, Pieters RG (2005) New empirical generalizations on the determinants of price elasticity. *Journal of marketing research* 42(2):141–156.

Brynjolfsson E, McElheran K (2016) The rapid adoption of data-driven decision-making. *American Economic Review* 106(5):133–39.

Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2014) The impact of linear optimization on promotion planning. *Available at SSRN 2382251* .

Columbia University Center for Pricing and Revenue Management (2012) Dataset cprm-12-001: On-line auto lending.

Efron B, Tibshirani R (1993) *An Introduction to the Bootstrap* (Chapmen and Hall).

Ferreira KJ, Lee BHA, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.

Flores CA (2005) *Estimation of Dose-Response Functions and Optimal Treatment Doses with a Continuous Treatment*. Ph.D. thesis, University of California, Berkeley.

Greblicki W, Krzyzak A, Pawlak M (1984) Distribution-free pointwise consistency of kernel regression estimate. *The Annals of Statistics* 1570–1575.

Härdle W (1990) *Applied nonparametric regression* (Cambridge Univ Press).

Harrison JM, Keskin NB, Zeevi A (2012) Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science* 58(3):570–586.

Hirano K, Imbens GW (2004) The propensity score with continuous treatments. Gelman A, Meng XL, eds., *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, chapter 7, 73–84 (New York: Wiley).

Horowitz JL (2001) The bootstrap. Heckman JJ, Leamer E, eds., *Handbook of Econometrics*, volume 5, chapter 52, 3159–3228 (Amsterdam: Elsevier).

Lee S, Homem-de Mello T, Kleywegt AJ (2012) Newsvendor-type models with decision-dependent uncertainty. *Mathematical Methods of Operations Research* 76(2):189–221.

McAfee A, Brynjolfsson E (2012) Big data: the management revolution. *Harvard Bus. Rev.* 90(10):60–66.

McCullagh P, Nelder JA, McCullagh P (1989) *Generalized linear models* (London: Chapman and Hall), 2 edition.

Nadaraya E (1964) On estimating regression. *Theory Probab. Appl.* 9(1):141–142.

Nambiar M, Simchi-Levi D, Wang H (2019) Dynamic learning and pricing with model misspecification. *Management Science* .

Newey WK (1994) Kernel estimation of partial means and a general variance estimator. *Econometric Theory* 10(02):1–21.

Parzen E (1962) On estimation of a probability density function and mode. *The annals of mathematical statistics* 1065–1076.

Phillips R (2005) *Pricing and revenue optimization* (Stanford University Press).

Phillips R, Simsek AS, VanRyzin G (2012) Endogeneity and price sensitivity in customized pricing. *Columbia University Center for Pricing and Revenue Management Working Paper* 4.

Watson G (1964) Smooth regression analysis. *Sankhyā A* 359–372.

Ziegler K (2002) On nonparametric kernel estimation of the mode of the regression function in the random design model. *Journal of Nonparametric Statistics* 14(6):749–774.

EC.1. Omitted Proofs

Proof of Theorem 2 That $y(z)$ is linear and decreasing implies that $y(z) = y_0 - \lambda(z-c)$ with $\lambda > 0$. Hence, $R(z) = y_0(z-c) - \lambda(z-c)^2$, which is unimodal and uniquely maximized at $z^\star = c + y_0/(2\lambda)$ with value $R(z^\star) = y_0^2/(4\lambda)$. Let $\delta(z) = \mathbb{E}\left[\epsilon \middle| Z=z\right]$, $\eta = y_0\gamma$. Then $|\delta(z)| \le \eta$. Note that

$$\mathbb{E}\left[Y^{\text{obs}} \middle| Z=z\right] = \mathbb{E}\left[Y(z) \middle| Z=z\right] = \mathbb{E}\left[y(z)+\epsilon(z) \middle| Z=z\right] = y(z) + \mathbb{E}\left[\epsilon \middle| Z=z\right] = y_0 - \lambda(z-c) + \delta(z).$$

Hence, the theorem is trivial if $\eta = 0$ so let us assume $\eta > 0$.

Next we ask the question, what is the largest and smallest that the maximizer $\tilde{z}$ of $\tilde{R}(z)$ can be. By assumption, $|\delta(z)| \le \eta$ for all $z \in \mathcal{Z}$. So, defining $\tilde{R}_{\delta_0}(z) := (z-c)\left(y_0 - \lambda(z-c) + \delta_0(z)\right)$, we are interested in

$$\tilde{z}_{\max} = \sup\left\{\sup\left(\operatorname*{arg\,max}_{z\in\mathcal{Z}} \tilde{R}_{\delta_0}(z)\right) : |\delta_0(z)| \le \eta\right\}, \tag{EC.1}$$

$$\tilde{z}_{\min} = \inf\left\{\inf\left(\operatorname*{arg\,max}_{z\in\mathcal{Z}} \tilde{R}_{\delta_0}(z)\right) : |\delta_0(z)| \le \eta\right\}, \tag{EC.2}$$

where we define $\sup(\varnothing) = -\infty$ and $\inf(\varnothing) = \infty$ without loss of generality because we assumed an optimizer $\tilde{z}$ exists for $\tilde{R}(z)$ so we are only interested in those functions $\delta(z)$ that induce a nonempty argmax. In what follows, define $\tilde{R}_+(z) = (z-c)\left(y_0 - \lambda(z-c) + \eta\right)$ and $\tilde{R}_-(z) = (z-c)\left(y_0 - \lambda(z-c) - \eta\right)$, which are both unimodal and uniquely maximized at $\tilde{z}_+ = c + (y_0+\eta)/(2\lambda)$ and $\tilde{z}_- = c + (y_0-\eta)/(2\lambda)$ respectively ($\tilde{z}_- < \tilde{z}_+$ because $\eta > 0$). Notice that $\tilde{R}_-(z) \le \tilde{R}_{\delta_0}(z) \le \tilde{R}_+(z)$ whenever $|\delta_0(z)| \le \eta$ with equality when $\delta_0(z) = \pm\eta$ is extremal.

First we argue that the bounds (EC.1)–(EC.2) are finite. For any $z \ge z' = c + \left(y_0 + \eta + 2\sqrt{\lambda + y_0\eta}\right)/(2\lambda)$ and $|\delta_0(z)| \le \eta$, since $\tilde{R}_+(z)$ is decreasing past $\tilde{z}_+$ and $z' \ge \tilde{z}_+$, we have that

$$\tilde{R}_{\delta_0}(z) \le \tilde{R}_+(z) \le \tilde{R}_+(z') = (y_0-\eta)^2/(4\lambda) - 1 < (y_0-\eta)^2/(4\lambda) = \tilde{R}_-(\tilde{z}_-) \le \tilde{R}_{\delta_0}(\tilde{z}_-).$$

Since $\tilde{z}_- \le z'$ we conclude that $\tilde{z}_{\max} \le z' < \infty$. Finally, since $\delta_0(z) = 0$ is feasible in (EC.1)–(EC.2), we have $c \le \tilde{z}_{\min} \le z^\star \le \tilde{z}_{\max} \le z'$.

Next we argue that in (EC.1) it is sufficient to consider functions $\delta_0(z)$ taking values in $\{-\eta, +\eta\}$ that are monotonic increasing, i.e. constant or step functions. Let $\delta_0(z)$ be feasible in (EC.1) and let $\tilde{z}_0 = \sup\left\{\arg\max_{z\in\mathcal{Z}} \tilde{R}_{\delta_0}(z)\right\}$. If at any $z_1 \geq \tilde{z}_0$ we have $\delta_0(z_1) < \eta$, then increasing the value of $\delta_0(z_1)$ to η can only increase the value of $\tilde{R}_{\delta_0}(z_1)$, which in turn may only increase the largest maximizer since $z_1 \geq \tilde{z}_0$. Moreover, if at any $z_1 < \tilde{z}_0$ we have $\delta_0(z_1) > -\eta$, then decreasing the value of $\delta_0(z_1)$ to $-\eta$ can only decrease the value of $\tilde{R}_{\delta_0}(z_1)$, which must already be at or below the maximal value and hence must leave the largest maximizer unchanged. The argument is unchanged even if $\tilde{z}_0$ is $\pm\infty$. A symmetric argument shows that in (EC.2) it is sufficient to consider functions $\delta_0(z)$ taking values in $\pm\eta$ that are monotonic decreasing.

Next we evaluate $\tilde{z}_{\max}$. Fix $\tilde{z}' = c + (\sqrt{y_0} + \sqrt{\eta})^2/(2\lambda)$ and let us consider the step function $\delta_{\max}(z) = \eta\mathbb{I}[z \geq \tilde{z}'] - \eta\mathbb{I}[z < \tilde{z}']$. Since $\tilde{z}' > z_-$, $\tilde{R}_{\delta_{\max}}(z)$ is uniquely maximized on $(c, \tilde{z}')$ at z_-, with value $\tilde{R}_{\delta_{\max}}(z_-) = \tilde{R}_-(z_-) = (y_0 - \eta)^2/(4\lambda)$. Since $\tilde{z}' > z_+$, $R_{\delta_{\max}}(z)$ is uniquely maximized on $[\tilde{z}', \infty)$ at $\tilde{z}'$, with value $\tilde{R}_{\delta_{\max}}(\tilde{z}') = \tilde{R}_+(\tilde{z}') = (y_0 - \eta)^2/(4\lambda)$. Hence, $\arg\max_{z\in\mathcal{Z}} \tilde{R}_{\delta_{\max}}(z) = \{z_-, \tilde{z}'\}$ and $\sup\{z_-, \tilde{z}'\} = \tilde{z}'$. Now we show that it is impossible to achieve a higher maximizer with $|\delta(z)| \leq \eta$, which would lead to $\tilde{z}_{\max} = \tilde{z}'$. By our previous argument we need only consider functions $\delta(z)$ taking values in $\pm\eta$ that are monotonic increasing. The constant functions taking values in $\pm\eta$ induce the maxima $\tilde{z}_-$ and $\tilde{z}_+$, both of which are smaller than $\tilde{z}'$. Next, consider any step function $\delta_0(z) = \eta\mathbb{I}[z \geq \tilde{z}_0] - \eta\mathbb{I}[z < \tilde{z}_0]$ with $\tilde{z}_0 \neq z'$. If $\tilde{z}_0 \leq \tilde{z}_+$ then, for any $z \neq z_+$, we have that $\tilde{R}_{\delta_0}(\tilde{z}_+) = \tilde{R}_+(\tilde{z}_+) > \tilde{R}_+(z) \geq \tilde{R}_{\delta_0}(z)$ since $\tilde{z}_+$ is the unique maximizer of $\tilde{R}_+(z)$; hence $\tilde{z}_+ < z'$ is the unique maximum of $\tilde{R}_{\delta_0}(z)$. Consider $\tilde{z}_0 > \tilde{z}_+$. Then, since $\tilde{z}_0 > z_+ > z_-$, $\tilde{R}_{\delta_{\max}}(z)$ is uniquely maximized on $(c, \tilde{z}_0)$ at z_-, with value $\tilde{R}_{\delta_0}(z_-) = \tilde{R}_-(z_-) = (y_0 - \eta)^2/(4\lambda)$. Since $\tilde{z}_0 > z_+$, $R_{\delta_{\max}}(z)$ is uniquely maximized on $[\tilde{z}_0, \infty)$ at $\tilde{z}_0$, with value $\tilde{R}_{\delta_0}(\tilde{z}_0) = \tilde{R}_+(\tilde{z}_0)$. If $\tilde{z}_0 < \tilde{z}'$, then either of these potential maximizers are smaller than z'. If $\tilde{z}_0 > \tilde{z}'$ then, since $\tilde{R}_+(z)$ is strictly decreasing past z_+ and $\tilde{z}' \geq z_+$, we have $\tilde{R}_{\delta_0}(\tilde{z}_0) = \tilde{R}_+(\tilde{z}_0) < \tilde{R}_+(\tilde{z}') = (y_0 - \eta)^2/(4\lambda) = \tilde{R}_-(z_-) = \tilde{R}_{\delta_{\max}}(z_-)$. Hence $\tilde{z}_- < \tilde{z}'$ is the unique maximum of $\tilde{R}_{\delta_0}(z)$. When $\eta < y_0$, a symmetric argument applied to (EC.2) shows that $\tilde{z}_{\min} = c + \left(\sqrt{y_0} - \sqrt{\eta}\right)^2/(2\lambda)$. If $\eta \geq y_0$, the lower bound $\tilde{z}_{\min} = c$ is achieved by $\delta_-(z)$. Hence, $\tilde{z}_{\min} = c + \max\left\{0, \sqrt{y_0} - \sqrt{\eta}\right\}^2/(2\lambda)$.

To summarize, we conclude that since $\left|\mathbb{E}\left[\epsilon\middle|Z\right]\right| \le \eta$, we must have

$$\tilde{z} \in [\tilde{z}_{\min}, \tilde{z}_{\max}] \quad \text{where} \quad \tilde{z}_{\min} = c + \frac{\max\left\{0, \sqrt{y_0} - \sqrt{\eta}\right\}^2}{2\lambda}, \tilde{z}_{\max} = c + \frac{(\sqrt{y_0} + \sqrt{\eta})^2}{2\lambda}.$$

Plugging these bounds into $R(z)$ we have

$$R(\tilde{z}_{\max}) = \frac{y_0^2 - 4y_0\eta - \eta^2 - 4\eta\sqrt{y_0\eta}}{4\lambda}, \qquad R(\tilde{z}_{\min}) = \begin{cases} \frac{y_0^2 - 4y_0\eta - \eta^2 + 4\eta\sqrt{y_0\eta}}{4\lambda} & \eta < y_0 \\ 0 & \eta \ge y_0 \end{cases}$$

Notice that if $\eta < y_0$ then $R(\tilde{z}_{\max}) = R(\tilde{z}_{\min}) - 2\eta\sqrt{y_0\eta}/\lambda \le R(\tilde{z}_{\min})$ and if $\eta \ge y_0$ then $R(\tilde{z}_{\max}) \le 0 = R(\tilde{z}_{\min})$. Therefore, $\min\{R(\tilde{z}_{\max}), R(\tilde{z}_{\min})\} = R(\tilde{z}_{\max})$. Since $R(z)$ is unimodal and $\tilde{z} \in [\tilde{z}_{\min}, \tilde{z}_{\max}]$, we have

$$R(\tilde{z}) \ge \min\{R(\tilde{z}_{\max}), R(\tilde{z}_{\min})\} = R(\tilde{z}_{\max}) = \frac{y_0^2 - 4y_0\eta - \eta^2 - 4\eta\sqrt{y_0\eta}}{4\lambda}.$$

Finally, using $R(z^\star) = y_0^2/(4\lambda)$,

$$\frac{R(\tilde{z})}{R(z^\star)} \ge 1 - 4\left(\frac{\eta}{y_0}\right) - 4\left(\frac{\eta}{y_0}\right)^{3/2} - \left(\frac{\eta}{y_0}\right)^2,$$

which is a univariate polynomial in $\sqrt{\eta/y_0}$. □

Proof of Theorem 3 Recall that $\mathbb{E}[\epsilon] = 0$. That ϵ and Z are jointly normal implies that

$$\mathbb{E}\left[\epsilon\middle|Z\right] = \zeta(Z - \mu), \quad \text{where } \mu = \mathbb{E}[Z], \ \zeta = \frac{\operatorname{Cov}(\epsilon, Z)}{\operatorname{Var}(Z)}.$$

By assumption of non-positive correlation, $\zeta \le 0$.

That $y(z)$ is linear and decreasing implies that $y(z) = y_0 - \lambda(z - c)$ with $\lambda > 0$. Hence, $R(z) = y_0(z - c) - \lambda(z - c)^2$, which is unimodal and uniquely maximized at $z^\star = c + y_0/(2\lambda)$ with value $R(z^\star) = y_0^2/(4\lambda)$. Also, recall from the proof of Theorem 2 that $\mathbb{E}\left[Y^{\text{obs}}\middle|Z = z\right] = \mathbb{E}\left[Y(z)\middle|Z = z\right] = \mathbb{E}\left[y(z) + \epsilon(z)\middle|Z = z\right] = y(z) + \mathbb{E}\left[\epsilon\middle|Z = z\right]$ and hence $\tilde{R}(z) = y_0(z - c) - \lambda(z - c)^2 + \zeta(z - c)(z - \mu)$, which is unimodal and uniquely maximized at its critical point $\tilde{z} = (2\lambda c + y_0 - \zeta(c + \mu))/(2\lambda - 2\zeta)$ because it is feasible since $\lambda > 0 \ge \zeta$ and $(2\lambda - 2\zeta)(\tilde{z} - c) = y_0 - \zeta(\mu - c) \ge y_0 \ge 0$.

Plugging $\tilde{z}$ into $R(z)$ we get

$$R(\tilde{z}) = \frac{(y_0 - \zeta(\mu - c))(y_0(\lambda - 2\zeta) + \lambda\zeta(\mu - c))}{4(\lambda - \zeta)^2}.$$

Rearranging and using $R(z^\star) = y_0^2/(4\lambda)$, we have

$$\frac{R(\tilde{z})}{R(z^\star)} = 1 - \left(\frac{\zeta}{\lambda - \zeta}\right)^2 \left(\frac{y_0 + \lambda(c-\mu)}{y_0}\right)^2 = 1 - \left(\frac{\zeta}{\lambda-\zeta}\right)^2 \left(\frac{\mathbb{E}\left[Y^{\text{obs}}\right]}{y_0}\right)^2,$$

where we plugged in $\mathbb{E}\left[Y^{\text{obs}}\right] = \mathbb{E}\left[Y(Z)\right] = \mathbb{E}\left[y_0 - \lambda(Z-c)\right] = y_0 - \lambda(\mu - c)$. Moreover,

$$\sup_{\zeta \le 0} \left(\frac{\zeta}{\lambda-\zeta}\right)^2 = \lim_{\zeta \to -\infty} \left(\frac{\zeta}{\lambda - \zeta}\right)^2 = 1.$$

Hence, we have the result in the statement of the theorem. □

Proof of Theorem 4 We repeat the proof of Theorem 2 but note that if $E(z)$ is non-increasing then in (EC.1) it is sufficient to consider *constant* functions $\delta_0(z)$ taking values in $\{-\eta, +\eta\}$, which implies $\tilde{z}_m ax = \tilde{z}_+$. Therefore,

$$R(\tilde{z}_{\max}) = \frac{y_0^2 - \eta^2}{4\lambda},$$

since $\eta \le y_0$ is assumed. Since $R(z)$ is unimodal and $\tilde{z} \in [\tilde{z}_{\min}, \tilde{z}_{\max}]$, we have

$$\begin{aligned} R(\tilde{z}) \ge \min\left\{R(\tilde{z}_{\max}), R(\tilde{z}_{\min})\right\} &= \min\left\{\frac{y_0^2 - \eta^2}{4\lambda}, \frac{y_0^2 - 4y_0\eta - \eta^2 + 4\eta\sqrt{y_0\eta}}{4\lambda}\right\} \\ &\ge \frac{y_0^2 - 4y_0\eta - \eta^2 + 4\eta\sqrt{y_0\eta}}{4\lambda}, \end{aligned}$$

since $\eta \le y_0$ is assumed. Finally, using $R(z^\star) = y_0^2/(4\lambda)$,

$$\frac{R(\tilde{z})}{R(z^\star)} \ge 1 - 4\left(\frac{\eta}{y_0}\right) + 4\left(\frac{\eta}{y_0}\right)^{3/2} - \left(\frac{\eta}{y_0}\right)^2,$$

completing the proof. □

Proof of Theorem 6 Assumption 1 gives $R(z) = \mathbb{E}\left[\mathbb{E}\left[r(Z)Y^{\text{obs}} \middle| Z = z, X\right]\right]$, i.e. profit is given by taking a partial mean with $Z = z$ fixed of the regression of $r(Z)Y^{\text{obs}}$ on Z and X.

By Assumption 4 part i, there exists $\delta > 0$ such that $[z^\star - \delta, z^\star + \delta]$ is contained inside the support of Z. By Assumption 4 part i, $f_{Z,X}(z,x)$ is bounded away from 0 on its support, and, by Assumption 4 part v, $\frac{\partial f_{Z,X}(z,x)}{\partial x}$ is bounded. Hence,

$$\left|\frac{\partial}{\partial x}\log(f_{Z,X}(z,x))\right| = \frac{\left|\frac{\partial f_{Z,X}(z,x)}{\partial x}\right|}{f_{Z,X}(z,x)} \le L < \infty$$

on the support of (Z,X). Therefore, we have that, for any x and $|\psi|\le\delta$,

$$\log(f_{Z,X}(z^\star+\psi,\,x))\le\log(f_{Z,X}(z^\star,\,x))+L\delta,$$

and consequently,

$$\int\sup_{|\psi|\le\delta}f_{Z,X}(z^\star+\psi,x)dx\le\int e^{L\delta}f_{Z,X}(z^\star,\,x)dx<\infty.$$

By Assumption 4 part iii, there exists M such that $\mathbb{E}\left[(Y^{\mathrm{obs}})^4\middle|Z=z,X=x\right]\le M$ for all z, x. Combined, this yields

$$\int\sup_{|\psi|\le\delta}\left(1+\mathbb{E}\left[(r(Z)Y^{\mathrm{obs}})^4\middle|Z=z^\star+\psi,X=x\right]\right)f_{Z,X}(z^\star+\psi,x)dx$$
$$\le(1+r(z^\star+\delta)^4M)\int\sup_{|\psi|\le\delta}f_{Z,X}(z^\star+\psi,x)dx<\infty. \quad \text{(EC.3)}$$

To study our profit function estimator (5) for each fixed z, we employ Theorem 4.1 of Newey (1994) (henceforth, N in this proof), which provides convergence results for two-step kernel m-estimators, a specific case of which is our profit function estimator. We let the "trimming function" of N be $\tau(x)=\mathbb{I}\left[f_X(x)>0\right]$, an indicator for the compact support of X (where its density is assumed bounded away from zero). Since our estimator (5) has K both in the numerator and denominator, it is unchanged if we rescale K by a positive constant. Similarly, the conditions of Assumption 2 remain unchanged. Hence, by Assumption 2 part i, without loss of generality we may assume $\int_{\mathbb{R}^{1+k}}K=1$. Then, Assumption K of N is satisfied with $\Delta=2$ by Assumption 2 parts i-iv. Assumption H of N is satisfied with $d=s+1$ by Assumption 4 part v. These constitute condition (ii) of N's Theorem 4.1. By Assumption 4 part i, there exists $c<z_{\max}<\infty$ such that $Z\le z_{\max}$ almost surely. Let $r_{\max}=r(z_{\max})<\infty$. Then, by Assumption 4 part iii, $\mathbb{E}\left[(r(Z)Y^{\mathrm{obs}})^4\right]\le r_{\max}\mathbb{E}\left[(Y^{\mathrm{obs}})^4\right]<\infty$ and $\mathbb{E}\left[(r(Z)Y^{\mathrm{obs}})^4\middle|Z=z,X=x\right]=r(z)\mathbb{E}\left[(Y^{\mathrm{obs}})^4\middle|Z=z,X=x\right]$ are bounded. Combined with Assumption 4 part ii, we satisfy condition (i) of N's Theorem 4.1. Condition (iii) of N's Theorem 4.1 is satisfied by our choice of $\tau(\cdot)$ and by Assumption 4 part i. The first clause of condition (iv) of N's Theorem 4.1 is satisfied by our choice of $\tau(\cdot)$ and by Assumption 4 part ii. The second clause is satisfied by Assumption 4 parts iv-v combined with the fact that for any $m>0$,

$\mathbb{E}\left[(r(Z)Y^{\text{obs}})^m \middle| Z=z, X=x\right] = r(z)^m \mathbb{E}\left[(Y^{\text{obs}})^m \middle| Z=z, X=x\right]$ and $r(z)^m$ is continuous. The third clause is satisfied by (EC.3). Since $X \in \mathbb{R}^k$ and $Z \in \mathbb{R}$, condition (v) of N's Theorem 4.1 is satisfied by Assumption 2 parts v-vi. Applying N's Theorem 4.1 for each fixed $z \in \mathcal{Z}$, we get

$$\sqrt{nh_n}(R(z) - \overline{R}_n(z)) \xrightarrow{d} \mathcal{N}(0, \eta_z \kappa) \ \forall z \in \mathcal{Z},$$

where $\eta_z \kappa$ is an algebraic simplification of the asymptotic variance in eq. (14) in N.

To study the optimizer of our profit function estimator, we employ Flores (2005) (henceforth, F in this proof). Conditions (ii-vi) of F's Theorem 3 are satisfied in a similar way to the case of N's Theorem 4.1. Condition (i) of F's Theorem 3 is satisfied by Assumption 3 parts i and iii, condition (vii) by Assumption 4 part v, condition (viii) by Assumption 3 part iv, condition (ix) by Assumption 3 parts ii and iv, and finally condition (x) by Assumption 2 parts v-vi. Applying F's Theorem 3, we get

$$\sqrt{nh_n^3}(z^\star - \overline{z}_n) \xrightarrow{d} \mathcal{N}\left(0, \frac{\eta_{z^\star}\kappa'}{R''(z^\star)^2}\right), \tag{EC.4}$$

simplifying the asymptotic variance.

By Assumption 3 part iv and using Taylor's theorem to expand $R(z)$ around $z = z^\star$, there exists $z_n \in [\min(z^\star, \overline{z}_n), \max(z^\star, \overline{z}_n)]$ such that

$$R(\overline{z}_n) = R(z^\star) + R'(z^\star)(\overline{z}_n - z^\star) + \frac{1}{2}R''(z_n)(\overline{z}_n - z^\star)^2.$$

By first order optimality conditions, $R'(z^\star) = 0$. Hence, rearranging, we have

$$R(z^\star) - R(\overline{z}_n) = -\frac{1}{2}R''(z_n)(\overline{z}_n - z^\star)^2. \tag{EC.5}$$

By continuous transformation of eq. (EC.4), we have

$$\left(nh_n^3\right)(\overline{z}_n - z^\star)^2 \xrightarrow{d} \frac{\eta_{z^\star}\kappa'}{R''(z^\star)^2}\chi_1^2. \tag{EC.6}$$

Eq. (EC.4) also implies $\overline{z}_n \xrightarrow{\mathbb{P}} z^\star$, which also implies $z_n \xrightarrow{\mathbb{P}} z^\star$ since z_n is sandwiched between $\overline{z}_n$ and $z^\star$. Since $R''(z)$ is continuous, we also get by continuous transformation that

$$R''(z_n) \xrightarrow{\mathbb{P}} R''(z^\star). \tag{EC.7}$$

Combining eqs. (EC.5)–(EC.7), we get the desired result,

$$\left(nh_n^3\right)\left(R(z^\star) - R(\overline{z}_n)\right) \xrightarrow{d} \frac{-\eta_{z^\star}\kappa'}{2R''(z^\star)}\chi_1^2.$$

If $nh_n^{2s+1} \to 0$, then we also satisfy the conditions of F's Theorem 4 with equal bandwidths. Applying F's Theorem 4, we get

$$\sqrt{nh_n}(R(z^\star) - \overline{R}_n(\overline{z}_n)) \xrightarrow{d} \mathcal{N}\left(0, \eta_{z^\star}\kappa\right),$$

simplifying the asymptotic variance. □

Proof of Theorem 7 The proof borrows the outline of the proof of Theorem 2 of Besbes et al. (2010), but applied to our new testing case and causal estimators.

Decompose the test statistic ρ_n into three terms:

$$\rho_n = \overline{R}_n(\overline{z}_n) - \overline{R}_n(\hat{z}_n) = A_n + B_n + C_n,$$

where

$$A_n = \overline{R}_n(\overline{z}_n) - \overline{R}_n(z^\star),$$

$$B_n = \overline{R}_n(z^\star) - \overline{R}_n(\hat{z}),$$

$$C_n = \overline{R}_n(\hat{z}) - \overline{R}_n(\hat{z}_n).$$

We begin by showing that $(nh_n^3)\, A_n \xrightarrow{d} \Gamma\chi_1^2$. By Assumption 2 part iii, we have that $\overline{R}_n(z)$ is twice continuously differentiable. Thus, using Taylor's theorem to expand $\overline{R}_n(z)$ around $z = \overline{z}_n$, we get that there exists $z_n \in [\min(z^\star, \overline{z}_n), \max(z^\star, \overline{z}_n)]$ such that

$$\overline{R}_n(z^\star) = \overline{R}_n(\overline{z}_n) + \overline{R}_n'(\overline{z}_n)(z^\star - \overline{z}_n) + \frac{1}{2}\overline{R}_n''(z_n)(z^\star - \overline{z}_n)^2.$$

By first order optimality conditions, $\overline{R}_n'(\overline{z}_n) = 0$. Hence, rearranging, we have

$$A_n = -\frac{1}{2}\overline{R}_n''(z_n)(z^\star - \overline{z}_n)^2. \tag{EC.8}$$

Next we show that $\overline{R}_n''(z_n) \xrightarrow{\mathbb{P}} R''(z^\star)$. Note that

$$\left|\overline{R}_n''(z_n) - R''(z^\star)\right| \leq \left|\overline{R}_n''(z_n) - R''(z_n)\right| + \left|R''(z_n) - R''(z^\star)\right|. \tag{EC.9}$$

As in the proof of Theorem 6, Assumptions 2, 3, and 4 imply the assumptions of Lemma 5.1 of Newey (1994) applied to $R''(z)$, which in turn yields the uniform convergence in probability of $\overline{R}''_n(z)$ over $\mathcal{Z}$ since $\mathcal{Z}$ is compact by Assumption 3 part i. Hence,

$$\left|\overline{R}''_n(z_n) - R''(z_n)\right| \leq \sup_{z\in\mathcal{Z}} \left|\overline{R}''_n(z) - R''(z)\right| \xrightarrow{\mathbb{P}} 0. \tag{EC.10}$$

By Theorem 6, $\overline{z}_n \xrightarrow{\mathbb{P}} z^\star$. Because z_n is sandwiched between $\overline{z}_n$ and $z^\star$, we also get $z_n \xrightarrow{\mathbb{P}} z^\star$. Since $R''(z)$ is continuous by Assumption 3 part iv, we have

$$|R''(z_n) - R''(z^\star)| \xrightarrow{\mathbb{P}} 0 \tag{EC.11}$$

by continuous transformation of the former. Combining eqs. (EC.9)–(EC.11), we get

$$\overline{R}''_n(z_n) \xrightarrow{\mathbb{P}} R''(z^\star). \tag{EC.12}$$

By continuous transformation of the result of Theorem 6 (eq. (EC.4)), we have

$$\left(nh_n^3\right)(\overline{z}_n - z^\star)^2 \xrightarrow{d} \frac{\eta_{z^\star}\kappa'}{R''(z^\star)^2}\chi_1^2. \tag{EC.13}$$

Combining eqs. (EC.8)–(EC.13), we get

$$\left(nh_n^3\right)A_n \xrightarrow{d} \frac{-\eta_{z^\star}\kappa'}{2R''(z^\star)}\chi_1^2 = \Gamma\chi_1^2. \tag{EC.14}$$

Next, we show that $(nh_n^3)\,C_n \xrightarrow{\mathbb{P}} 0$. By Assumption 2 part iii, we have that $\overline{R}_n(z)$ is twice continuously differentiable. Thus, using Taylor's theorem to expand $\overline{R}_n(z)$ around $z = \hat{z}$, we get that there exists $z'_n \in [\min(\hat{z}, \hat{z}_n), \max(\hat{z}, \hat{z}_n)]$ such that

$$\overline{R}_n(\hat{z}_n) = \overline{R}_n(\hat{z}) + \overline{R}'_n(\hat{z})(\hat{z}_n - \hat{z}) + \frac{1}{2}\overline{R}''_n(z'_n)(\hat{z}_n - \hat{z})^2.$$

Rearranging, we have

$$\left(nh_n^3\right)C_n = -\left(\sqrt{nh_n^3}\,\overline{R}'_n(\hat{z})\right)\left(\sqrt{nh_n^3}(\hat{z}_n - \hat{z})\right) - \frac{1}{2}\overline{R}''_n(z'_n)\left(\sqrt{nh_n^3}(\hat{z}_n - \hat{z})\right)^2. \tag{EC.15}$$

By Assumption 5, we have that

$$\sqrt{nh_n^3}(\hat{z}_n - \hat{z}) = o_p(1), \text{ and hence also } \left(\sqrt{nh_n^3}(\hat{z}_n - \hat{z})\right)^2 = o_p(1). \tag{EC.16}$$

Applying Theorem 4 of Newey (1994) we get the convergence in distribution of $\sqrt{nh_n}\left(\overline{R}_n'(z) - R'(z)\right)$ for any fixed z, including $\hat{z}$ and hence, since $h_n \to 0$ we have

$$\sqrt{nh_n^3}\overline{R}_n'(\hat{z}) = o_p(1). \tag{EC.17}$$

Next we show that $\overline{R}_n''(z_n') = O_p(1)$. Note that

$$\left|\overline{R}_n''(z_n') - R''(\hat{z})\right| \le \left|\overline{R}_n''(z_n') - R''(z_n')\right| + \left|R''(z_n') - R''(\hat{z})\right|. \tag{EC.18}$$

As before, $\overline{R}_n''(z)$ converges uniformly to $R''(z)$ in probability over $\mathcal{Z}$ and so

$$\left|\overline{R}_n''(z_n') - R''(z_n')\right| \le \sup_{z\in\mathcal{Z}} \left|\overline{R}_n''(z) - R''(z)\right| \xrightarrow{\mathbb{P}} 0. \tag{EC.19}$$

By Assumption 5, $\hat{z}_n \xrightarrow{\mathbb{P}} \hat{z}$. Because z_n' is sandwiched between $\hat{z}_n$ and $\hat{z}$, we also get $z_n' \xrightarrow{\mathbb{P}} \hat{z}$. Since $R''(z)$ is continuous by Assumption 3 part iv, we have

$$\left|R''(z_n') - R''(\hat{z})\right| \xrightarrow{\mathbb{P}} 0 \tag{EC.20}$$

by continuous transformation of the former. Combining eqs. (EC.18)–(EC.20), we get

$$\overline{R}_n''(z_n') \xrightarrow{\mathbb{P}} R''(\hat{z}). \tag{EC.21}$$

Combining eqs. (EC.15)–(EC.21) gives $(nh_n^3)\, C_n = -o_p(1)o_p(1) - O_p(1)o_p(1) = o_p(1)$.

Finally, we treat B_n. Under H_0, $B_n = 0$ because Assumption 3 part ii (unique optimizer) and H_0 ($R(z^\star) = R(\hat{z})$) imply that $z^\star = \hat{z}$. Next, we show that under H_1, $(nh_n^3)\, B_n \xrightarrow{\mathbb{P}} \infty$. By applying the first results of Theorem 6 to each term, we have that $B_n \xrightarrow{\mathbb{P}} R(z^\star) - R(\hat{z})$. Since $k \ge 0$, Assumption 2 part vi implies $nh_n^5/\log(n) \to \infty$, which, since we also assume $h_n \to 0$, implies $nh_n^3 \to \infty$. Hence, since $R(z^\star) - R(\hat{R}) > 0$ under H_1, we have that $(nh_n^3)\, B_n \xrightarrow{\mathbb{P}} \infty$. □

Proof of Theorem 8 Proven above. See eq. (EC.14). □

EC.2. More General Version of Theorem 7

Assumption EC.1 **(Convergent Decision-Making (relaxed))**. $\sqrt{nh_n^3}(\hat{z}_n - \hat{z}) \xrightarrow{d} \mathcal{N}(0, V)$ *for some* $V \ge 0$.

Note that Assumption 5 implies Assumption EC.1 with $V = 0$. In this sense. Assumption EC.1 is weaker and more general.

THEOREM EC.1. *Suppose Assumptions 1, 2, 3, 4, and EC.1 hold. Let* $\Gamma = \frac{-\eta_{z^\star}\kappa'}{2R''(z^\star)}$ *and* $\Gamma' = \frac{-V}{2}R''(z^\star)$. *Then,*

i. *under* H_0, $\limsup_{n\to\infty}\mathbb{P}\left(\left(nh_n^3\right)\rho_n > t\right) \le 1 - F_{\Gamma\chi_1^2+\Gamma'\chi_1^2}(t)$, *where* $F_{\Gamma\chi_1^2+\Gamma'\chi_1^2}$ *is the CDF of the weighted sum of two independent chi-squared random variables, and*

ii. *under* H_1, $\left(nh_n^3\right)\rho_n \xrightarrow{d} \infty$.

Proof. The only part of the proof of Theorem 7 that changes is the analysis of the term C_n. Following the arguments after eq. (EC.15) but using Assumption EC.1 we conclude that

$$(nh_n^3)C_n = \frac{-R''(\hat{z})}{2}(nh_n^3)(\hat{z} - \hat{z}_n)^2 + o_p(1).$$

As before, we had that

$$(nh_n^3)A_n = \frac{-R''(z^*)}{2}(nh_n^3)(z^* - \overline{z}_n)^2 + o_p(1).$$

Now, under H_0 and under Assumption 3, $R''(\hat{z}) = R''(z^*)$, which is indeed negative. Therefore, under H_1, $\left(nh_n^3\right)\rho_n \xrightarrow{d} \infty$, and, under H_0, $\left(nh_n^3\right)\rho_n \xrightarrow{d} H$, where $H = G_1^2 + G_2^2$ and (G_1, G_2) are jointly normal random variables with mean zero, variances Γ and Γ', and some covariance C. The distribution of H with some covariance C is stochastically dominated by the same with covariance 0, which yields the result. □

The implication is that if we use the p-value given by $1 - F_{\Gamma\chi_1^2+\Gamma'\chi_1^2}(nh_n^3\rho_n)$ then it would be a safe p-value in that it will still ensure at most α type-I error rate if we reject the null only when $p < \alpha$. Given Γ and Γ', we can simulate the distribution of $\Gamma\chi_1^2 + \Gamma'\chi_1^2$ to arbitrary precision or use formulae for the weighted sum of chi-squared random variables (Bausch 2013).

Finally, note that as in Theorem 8, we again have that $(nh_n^3)\mathbb{E}[C_n] \to \Gamma'$. Consequently, we can estimate Γ' using exactly the same bootstrap method but instead applied to $\hat{z}_n$.