

A Semiparametric Discrete Choice Model: An Application to Hospital Mergers*

Devesh Raval

Ted Rosenbaum

Steven A. Tenn

Federal Trade Commission

Federal Trade Commission

Charles River Associates

draval@ftc.gov

trosenbaum@ftc.gov

stenn@crai.com

October 20, 2015

Abstract

We propose a computationally simple semiparametric discrete choice estimator to model rich consumer heterogeneity. We assume groups of observably similar consumers have similar preferences, but allow preferences to vary freely across these groups. Model flexibility is easily adjusted by setting a single tuning parameter; we suggest a cross-validation method to do so. We analyze the model's properties in the context of hospital mergers, both analytically and via a Monte Carlo analysis. The model performs well for policy relevant substitution and welfare measures, even if misspecified, when the tuning parameter is set within the neighborhood of the value chosen by cross-validation.

JEL Codes: C14, D12, I11, L41

Keywords: hospitals, patient choice, demand models, semiparametric, mergers

*The views expressed in this paper are those of the authors. They do not necessarily represent those of the Federal Trade Commission, any of its Commissioners, or Charles River Associates.

Raval: Economist, Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington, DC 20580. Phone 1-202-326-2260, E-mail draval@ftc.gov

Rosenbaum: Economist, Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington, DC 20580. Phone 1-202-326-2275, E-mail trosenbaum@ftc.gov

Tenn (Corresponding Author): Vice President, Charles River Associates, 1201 F Street NW, Suite 700 Washington, DC 20004. Phone 1-202-662-3806, E-mail stenn@crai.com

1 Introduction

Individual level data on consumer choices can significantly improve predictions of consumer substitution patterns (Berry et al. (2004)). Over the past twenty years, these data have become increasingly available to researchers, a trend that seems likely to continue into the future. Individual level data facilitate the estimation of models with rich individual heterogeneity, which can yield more accurate out-of-sample predictions of individual choices than models that are less flexible along this dimension (Raval et al. (2015)).

However, the estimation of more flexible models can be fraught with statistical and computational concerns. In a context similar to ours, Ho and Pakes (2014) note that estimating models with many dummy variables via maximum likelihood can lead to an incidental parameters problem. Beyond these econometric concerns, estimating this type of model involves computational issues that are likely to consume significant researcher time (Greene (2004)).

We outline a computationally light semiparametric approach to flexibly estimate consumer substitution patterns. In our approach, we allow consumer preferences to vary freely across groups, but maintain a representative consumer assumption within each group. To form the groups, we propose an iterative procedure that matches patients with similar characteristics across multiple dimensions. In order to compute substitution patterns, we assume that agents substitute in proportion to choice shares within the group. This model can be viewed as a highly flexible multinomial logit estimator.¹

¹Our model is similar to a machine learning decision tree approach (Breiman et al. (1984)). Both approaches segment the data into a large number of groups, and estimate predicted probabilities for each group. Unlike the decision tree approach, however, our approach does not develop groups based upon in-sample model fit, and thus should be less vulnerable to over-fitting. In addition, the decision tree's use of in-sample model fit to create groups creates a pre-test bias, which invalidates conventional inference procedures.

A major advantage of our approach is that a single tuning parameter, the minimum group size, balances the trade-off between flexibly modeling preference heterogeneity and statistical power. This parameter can be adjusted easily and transparently to make the estimator more or less flexible depending on the characteristics of a given dataset. To set this parameter, we propose a validation approach to choose a minimum group size using “leave one out” cross-validation. However, we find that statistics of interest are relatively precise across a range of minimum group sizes, making it less critical to select this parameter optimally.

We then apply our approach to merger analysis in healthcare. In healthcare, prices are determined by bargaining between a supplier and intermediaries such as insurers. The relative bargaining power of the insurer intermediary depends upon the change in consumer welfare if a provider is eliminated from the choice set (Capps et al. (2003)). We use the model to compute two measures that are widely used in the provider-choice literature as proxies for insurer leverage: diversion ratios and willingness to pay (Shapiro (1996), Capps et al. (2003), Farrell et al. (2011), Gowrisankaran et al. (2015)).

To calculate these proxies, researchers and antitrust agencies have frequently used parametric discrete choice logit models (e.g., Farrell et al. (2011)). However, researchers have begun to question the predictive accuracy of these models (Doane et al. (2012), Fournier and Gai (2007), May (2013)). Raval et al. (2015) show that an implementation of the approach that we outline does well. In particular, they find that following the exogenous elimination of a choice from the choice set, our semiparametric approach does a better job of predicting consumer substitution patterns than most of the parametric specifications that are frequently used in the literature.

We first apply the semiparametric model using data from a proposed hospital merger. We

find large diversion ratios and percent changes in willingness to pay that are not sensitive to the group size tuning parameter so long as it is set within an intermediate range. However, the percent change in willingness to pay is more sensitive to the value of this parameter than the diversion ratio. We then examine the approach’s robustness to misspecification using a Monte Carlo exercise. For an intermediate range of the minimum group size, the semiparametric approach has a low rate of error for both the diversion ratio and percent change in willingness to pay across different specifications of the “true” model, including a semiparametric model with extremely large preference heterogeneity and a parametric logit model. We also find only modest increases in error for much lower sample sizes than in our main specification. Thus, the semiparametric model performs well if it is not implemented in an overly flexible or inflexible manner.

Beyond healthcare, our approach can be used in any situation where a researcher would use a discrete choice approach to model the counterfactual elimination of an option from a choice set. It can be directly applied to merger analysis in industries, such as medical devices or television, where prices are determined by bargaining between a supplier and an intermediary (Grennan (2013), Crawford and Yurukoglu (2012)). More broadly, this type of counterfactual yields a diversion ratio that can be used in an Upward Pricing Pressure (“UPP”) calculation (Farrell and Shapiro (2010), Conlon and Mortimer (2013), Jaffe and Weyl (2013)).²

The paper proceeds as follows. In Section 2 we outline our approach, in Section 3 we introduce our empirical application, and in Section 4 we show results for our Monte Carlo

²Formally, for UPP one needs to calculate the diversion ratio in response to a small price change. The diversion ratio calculated based on eliminating a choice from the choice set is equal to the diversion ratio in response to a small price change under linear demand or the representative consumer logit model (Conlon and Mortimer (2013)).

simulation. We conclude in [Section 5](#).

2 Logit Choice Models

In this section, we first detail the parametric logit choice model typically used in the hospital choice literature, and then introduce our semiparametric logit choice model. We discuss our approach in the context of a patient’s choice of provider, but it can be applied in any multinomial choice context where one has access to individual level data.³

2.1 Parametric

Each patient i chooses the specific hospital j from the set of hospitals J . The utility u_{ij} that patient i receives from choosing hospital j is specified as follows:

$$u_{ij} = \delta_{ij} + \epsilon_{ij}. \tag{1}$$

Utility u_{ij} is determined by mean utility δ_{ij} and an i.i.d. error term ϵ_{ij} that is distributed Type I extreme-value. Each patient selects the utility maximizing option from the set of hospital choices J .

While mean utility δ_{ij} may be determined by numerous factors, the literature usually assumes a functional form in which δ_{ij} depends on patient characteristics, hospital characteristics, and travel costs between patient i and hospital j . Hospital characteristics control for the set of services offered at a given hospital and the overall quality of those services.

³See [Ackerberg et al. \(2007\)](#) for further discussion of these models in general and [Gaynor et al. \(2015\)](#) for further elaboration in the healthcare context.

Travel costs capture the fact that patients are, all else equal, more likely to select hospitals close to where they live. Patient characteristics, such as an individual’s medical condition and demographics, are interacted with travel costs and hospital characteristics. For example, a patient in labor may be more likely to go to a hospital with a labor and delivery room. Under the assumed error structure, the ex-ante probability s_{ij} that patient i selects hospital j follows the familiar logistic form:

$$s_{ij} = \frac{\exp(\delta_{ij})}{\sum_{k \in J} \exp(\delta_{ik})}. \quad (2)$$

Given a set of patient and hospital interactions, a functional form for these interactions, and the observed hospital choice for each patient, the parametric logit model can be estimated via maximum likelihood. [Capps et al. \(2003\)](#), [Gowrisankaran et al. \(2015\)](#), and [Ho \(2006\)](#) all estimate parametric logit models with different specifications for δ_{ij} . After recovering the vector of δ_{ij} ’s, it is straightforward to compute substitution patterns if any option is eliminated from the choice set.

The difficulty in designing such a hospital choice model lies in how to allow for heterogeneity in preferences across patients. Any two individuals with the same δ_{ij} for all j will have the same substitution patterns following the elimination of a choice from the choice set. Therefore, the degree of flexibility in δ_{ij} determines the flexibility of the allowed substitution patterns. Even with rich individual data, as is available in the hospital context, the researcher typically specifies a parameterized function.

The degree of flexibility in δ_{ij} impacts the accuracy of predicted substitution patterns. [Raval et al. \(2015\)](#) compare parameterizations of δ_{ij} from the literature that allow for varying

levels of heterogeneity in the population. Their results show that models that richly account for heterogeneity based on patient observables (more allowed heterogeneity in δ_{ij}) provide for more accurate out-of-sample predictions of individual choices than models that are less flexible along this dimension.

2.2 Semiparametric

We propose a computationally light semiparametric estimator that allows for significant flexibility in substitution patterns. In particular, we partition individuals in the data into G “patient types” indexed by g . Our approach relies on two primary assumptions. First, within each group g , we assume preference homogeneity (i.e., that s_{ij} is equal for all i in group g). Second, we assume that patients substitute proportionally to all alternatives. Hospital preferences can freely vary, without any restrictions, across groups.

These assumptions imply that, if a given hospital j is removed from the choice set, the share of patients that choose hospital k is given by:

$$s_{k/j} = \sum_g w_g \frac{s_k^g}{1 - s_j^g}. \quad (3)$$

where w_g is the fraction of all patients that are in group g .

To group patients into G groups, we use the following iterative procedure. We start with c characteristics, and patients having the same values for all c characteristics are put into the same group.⁴ We keep groups that have at least S_{min} observations. For the remaining data sample, we then repeat this procedure using only the first $c - 1$ characteristics. We

⁴Non discrete variables, such as age, can easily be discretized into categories.

iterate on this procedure, reducing the number of characteristics by one each time, until all patients are allocated into groups.⁵

The choice of minimum group size S_{min} makes clear the tension between flexibility and tractability. A smaller value of S_{min} leads to a more flexible model with many groups, but each group may comprise a small number of people. Alternatively, one can select a higher value of S_{min} to obtain a coarser grouping. Thus, the minimum group size functions analogously to a bandwidth parameter in kernel density estimation.⁶ In our empirical application, we illustrate how to calibrate this parameter using “leave one out” cross-validation.

This approach is equivalent to estimating a multinomial logit model where there is a dummy variable for each hospital-group combination. Formally, we assume that mean utility δ_{ij} is equal to δ_j^g for all individuals i in group g . When viewed in this light, our key assumption is that observable similarity implies that patients within a group are also unobservably similar except for an i.i.d. logit error. Given the logit assumptions and a vector of δ_j^g 's, the probability that patient i in group g selects hospital j is as follows:

$$s_{i(g)j} = s_j^g = \frac{\exp(\delta_j^g)}{\sum_{k \in J} \exp(\delta_k^g)}. \quad (4)$$

As noted above, we estimate \hat{s}_j^g directly rather than estimate each $\hat{\delta}_j^g$ and using them to compute \hat{s}_j^g .

This semiparametric approach is particularly suited to large datasets. As the sample size grows, the number of people in each group may become sufficiently large such that the share

⁵A small number of ungrouped individuals may remain after iterating in this manner c times. If so, the remaining observations can either be grouped together or simply omitted from the analysis.

⁶See [Pagan and Ullah \(1999\)](#) for a comprehensive treatment of kernel density estimation.

of each group that selects a given hospital can be precisely estimated for a large number of groups. Thus, an extremely flexible model can be estimated with a large dataset.⁷

2.2.1 Cross-Validation

We propose a “leave one out” cross-validation approach (Stone (1974)) to select the minimum group size.⁸ In this approach, the researcher estimates a candidate model m on all of the data, except for one observation. Then the researcher compares the predicted value for that observation from model m to the observed value for that observation. Iterating over all observations yields a vector of predicted values from a given model and a vector of observed values corresponding to each prediction. Using these two vectors, one can compute a measure of the out-of-sample fit for model m .

For our context, we examine model fit for different minimum group sizes. For a given minimum group size, we estimate the model using all of the data except for observation i . This assigns all observations, except for i , to groups and gives a predicted share for all groups and hospitals. Then we regroup all observations, including observation i , into a new set of groups. We take the set of individuals that are in i ’s group in the “including i ” grouping and compute their predicted values from the “excluding i ” estimation. We use the average of these predicted values as individual i ’s predicted choice probability for the cross-validation.

One can use different loss functions to measure the goodness of model fit in the cross-validation. For illustrative purposes, we use root mean squared error (“RMSE”) and McFadden’s pseudo R^2 .

⁷This estimator is also easy to implement via the MapReduce algorithm for parallelization, which would reduce computational costs even further.

⁸In principle, this can also be used for variable selection.

To compute the RMSE for minimum group size m , we take the difference between the predicted and actual choices for each individual and sum over all individuals in the dataset:

$$RMSE^m = \sqrt{\frac{1}{NJ} \sum_{i=1}^N \sum_{j=1}^J (\hat{s}_{ij}^m - y_{ij})^2}. \quad (5)$$

where y_{ij} equals 1 if individual i chose hospital j and 0 if not, and \hat{s}_{ij}^m is the predicted choice probability using a minimum group size of m .

McFadden's pseudo R^2 is inversely proportional to the log likelihood of the model. To compute the log likelihood for minimum group size m , we take the log of the predicted choice probability for the actual choice:

$$E^m = \frac{1}{N} \sum_{i=1}^N \log(\hat{s}_{ij^*}^m). \quad (6)$$

where $\hat{s}_{ij^*}^m$ is the predicted choice probability of the chosen option using a minimum group size of m . This statistic is also called the relative entropy of the model.⁹ McFadden's pseudo R^2 then scales the log likelihood to range between 0 and 1 as follows:

$$R^2 = 1 - \frac{E^m}{E^{Intercept}}. \quad (7)$$

where $E^{Intercept}$ is the log likelihood of a model that only includes an intercept, so each hospital is predicted to have a $\frac{1}{J}$ share of the market. An R^2 value of zero indicates a model with the same log likelihood as an intercept only model, and an R^2 value of one perfectly

⁹Since for many minimum group sizes there are observations where $\hat{s}_{ij^*}^m$ is zero, we use a bottom code when computing the log likelihood.

predicts the data.

In our empirical application, we use two different approaches to compute goodness of fit statistics. First, we compute goodness of fit as described above, and use every observation in the dataset for cross-validation (i.e., at some point, every person in the dataset is “left out”). We also apply an alternative approach in which we leave out only a random sample of observations in the data to reduce the computational time required.¹⁰

3 Application to Hospital Mergers

3.1 Hospital Merger Setting

We examine a proposed merger between two hospital systems in a mid-sized metropolitan area, where one of the larger systems in the area (System 1) proposed acquiring one of the smaller systems (System 2). For confidentiality reasons we do not reveal the identity of the firms.

Our empirical analysis relies on inpatient discharge data for patients living in the metropolitan area where the merging parties are located.¹¹ This dataset contains 124,237 adult commercial admissions.¹² For each hospital admission, we observe patient age, gender, zip code, Diagnostic Related Group (“DRG”)¹³, Major Diagnostic Category (“MDC”)¹⁴, and whether

¹⁰In this approach, we use the full dataset, except for the excluded individual, in the estimation. The only difference is that we do not use all individuals in the dataset for validation.

¹¹This area is largely self-contained. The vast majority of patients living in the area are treated there, and few patients who are treated at area hospitals come from outside the region.

¹²Children are omitted because the merging parties rarely admit them. We also remove patients with psychiatric, substance abuse, or rehabilitation diagnoses, and patients transferred out of a hospital to another acute care facility.

¹³The DRG system is a widely employed method of classifying hospital cases which contains hundreds of different services that a hospital may offer.

¹⁴The MDC system groups DRGs into 25 mutually exclusive categories. We aggregate a small number of

the admission was an emergency.

We group admissions using the iterative procedure detailed in [Section 2.2](#). We select variable order based on two criteria. First, we put each type of variable in descending order of its likely strength in determining hospital choice. That way, to the extent necessary to maintain sufficient group sizes, individuals that differ with respect to less important types of variables are pooled together first. We assume patient location is the most important predictor, followed by admission type and patient demographics. Our second criterion is that, within each variable type, we order the characteristics from the least to most detail. This allows a finer measure to be employed when group sizes are sufficiently large (e.g., DRG), but a coarser measure for smaller groups (e.g., MDC).

In order, the variables used to group admissions are as follows:

1. Patient Location
 - (a) County
 - (b) Zip code
2. Admission Type
 - (a) MDC
 - (b) Emergency admission indicator
 - (c) DRG type (medical vs. surgical)
 - (d) DRG weight quartile¹⁵
 - (e) DRG
3. Patient Demographics
 - (a) Age group (18-45, 46-62, and 62+)
 - (b) Gender

As a robustness check, we considered alternative variable orderings. Our results do not appear to be sensitive to how variables are ordered.

MDC groups with very few admissions.

¹⁵DRG weights are a resource intensity measure used by Medicare to calculate hospital reimbursement. DRG weights are a relative measure, defined such that the resource intensity of the average admission equals one. We group highly complex tertiary admissions separately, which we define as those with a DRG weight greater than 2.

3.2 Statistics of Interest

Antitrust agencies assessing the likely competitive impact of a proposed merger use hospital choice models to calculate measures of substitution between the hospitals (Farrell et al. (2011)). We focus on two widely employed statistics. The first is a substitution measure known as the diversion ratio (Shapiro (1996)). The diversion ratio from hospital h to hospital j measures the fraction of hospital h 's patients who would switch to hospital j if hospital h were removed from the choice set. In the logit context, the diversion from hospital h to hospital j is proportional to hospital j 's share relative to the other hospitals in the market:

$$div_{ihj} = \frac{s_{ij}}{1 - s_{ih}}. \quad (8)$$

The overall diversion div_{hj} from hospital h to hospital j is obtained by computing the average patient-level diversion across the set of patients that select hospital h .

The second statistic that we consider is the post-merger percent change in a metric known as “willingness to pay” (“WTP”). Developed by Capps et al. (2003), WTP measures the reduction in consumers’ expected utility from removing a set of hospitals from the choice set. In the logit model, the ex-ante expected decline in patient i 's welfare from excluding a set of hospitals $S \subset J$ is as follows:

$$WTP_{iS} = -\ln\left(1 - \sum_{j \in S} s_{ij}\right). \quad (9)$$

A patient’s WTP is an increasing function of the probability he will select a hospital in set

S , and equals zero when that probability is zero. Overall WTP is obtained by adding up patient-level WTP across all patients.

The antitrust agencies have used WTP to assess the expected harm from a merger of two hospital systems (Farrell et al. (2011)). The combined system’s bargaining position changes post-merger, since it can now threaten to exclude both systems simultaneously from the provider’s network. Let WTP_{12} represent the WTP for the combined system, and WTP_1 and WTP_2 for System 1 and System 2 individually. If the two systems are substitutes, then the loss in welfare from simultaneously excluding both systems exceeds the sum of the losses from individually excluding each system. The percentage increase in WTP resulting from a merger between the two systems can then be calculated as follows:

$$\Delta WTP_{12} = \frac{WTP_{12}}{WTP_1 + WTP_2} - 1. \tag{10}$$

This measure has the property that it equals zero when the two systems are not substitutes, and is an increasing function of the level of substitution between the two systems.

In order to estimate a diversion ratio or obtain a finite WTP measure, there must be heterogeneity in the chosen system across group members. While generally this is not a problem for large groups, hospital choice homogeneity becomes more likely in a group consisting of only a handful of individuals. In practice, we exclude admissions where a diversion ratio cannot be calculated from estimates of the diversion ratio. We bound any measure of WTP by imposing a top code at a share of 95%; if one of the merging partners hits this bound, the bound will imply a zero change in WTP. In general, choosing S_{min} presents a bias-variance tradeoff, which we discuss in the appendix.

3.3 Merger Estimates

We use the 124,237 observations of the hospital discharge data to estimate the semiparametric choice model for minimum group size S_{min} from 3 to 50,000. Selecting such a wide range of values for S_{min} allows us to compare results for extremely flexible and inflexible models, as well as intermediate specifications. Depending on the choice of S_{min} , admissions are put in between one and 23,157 groups. The most flexible specification has a pseudo R^2 of 0.70, while the least flexible specification has a pseudo R^2 of only 0.18. We then estimate diversion ratios and the change in WTP for the two hospital systems for each value of S_{min} .

We also apply the cross-validation approaches described above in [Section 2.2.1](#) to select a value for S_{min} . [Figure 1](#) contains graphs of the cross-validation results for the RMSE and McFadden’s pseudo R^2 statistics when we use all the observations in the data for cross-validation. RMSE is minimized at a minimum group size of 25 and McFadden’s pseudo R^2 is maximized at a minimum group size of 10. For models with a minimum group size under the optimum, the benefit of the increased model flexibility is outweighed by the cost of the lower power of the \hat{s}_j^g estimates. Conversely, for a minimum group size above the optimum, the benefits of the increased model flexibility outweigh the cost associated with lower statistical power. Since the RMSE and McFadden’s pseudo R^2 approaches penalize errors differently, they weigh this tradeoff differently.¹⁶

We examine the diversion ratio from System 2 to System 1 and the percent change in WTP.¹⁷ The black line in [Figure 2a](#) is a plot of this diversion ratio for different values of

¹⁶The McFadden’s pseudo R^2 statistic reported in the text uses a “bottom code” of .05, which is consistent with our approach in setting a top code for WTP. A smaller bottom code will heavily penalize choice probabilities at or near zero and is maximized at a larger minimum group size.

¹⁷Results for the diversions from system 1 to system 2 are in [Table I](#). Those results are consistent with

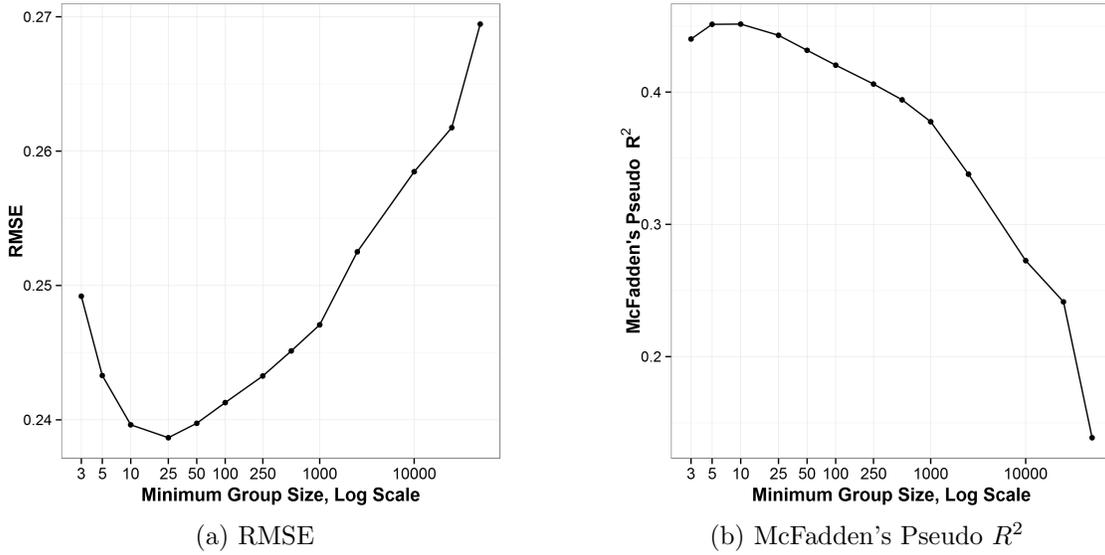


Figure 1 RMSE and McFadden's Pseudo R^2 from Leave One Out Cross-Validation by S_{min}

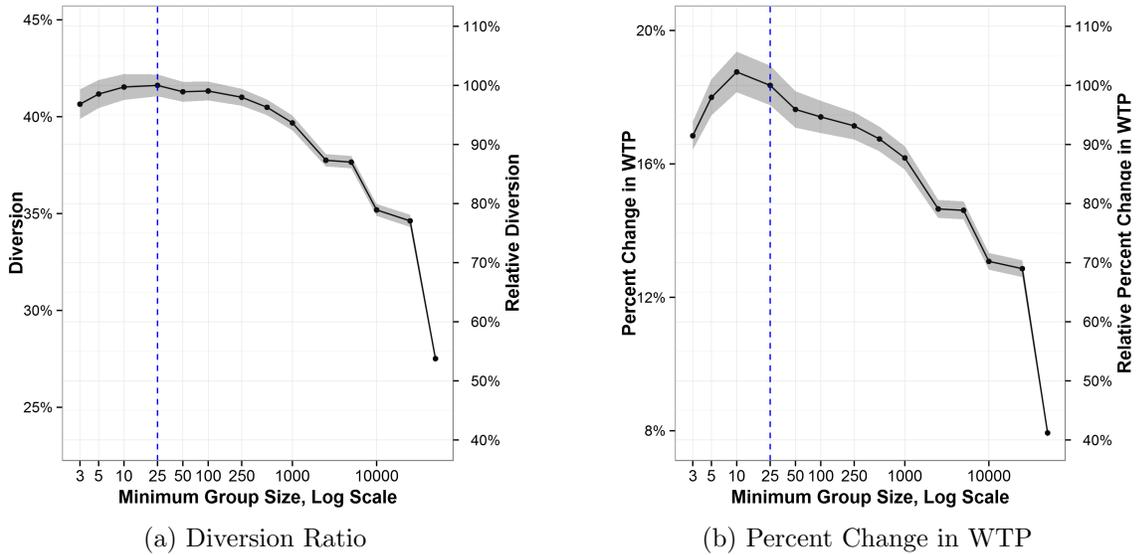


Figure 2 Estimated Diversion Ratio and Percent Change in WTP by S_{min}

Note: Diversion ratio is from System 2 to System 1. The left axis indicates the estimated value of the statistic, and the right axis the estimates scaled by the value of the statistic at a minimum group size of 25. [Table I](#) presents these estimates numerically.

the minimum group size and [Figure 2b](#) displays the estimated post-merger percent change in WTP. S_{min} is plotted on a log scale. The left axis indicates the estimated value of the statistic, while the right axis indicates the result scaled by the value of the statistic for a minimum group size of 25, the value selected by the RMSE cross-validation approach.¹⁸

Overall, these results suggest that the statistics of interest are relatively insensitive across a range of values for S_{min} in the neighborhood of 25. For a minimum group size between 5 and 500, the diversion ratio is within 4% of the value at 25, and the change in WTP is within 10% of the value at 25. The diversion ratio is less sensitive to the choice of minimum group size than the willingness to pay measure.

For both figures, the shaded region indicates 95% confidence intervals based upon bootstrapped standard errors from 50,000 draws. As expected, the use of a more flexible specification generally leads to larger standard errors for the diversion ratio. However, since fairly precise estimates are obtained even for small values of S_{min} , the loss of precision from using a more flexible model appears to be a relatively small cost. For the change in WTP, the very low S_{min} specifications also have a fairly low standard error, because the top code on WTP implies that many groups have a zero percent change in WTP.

The fraction of patients which are in groups without system choice heterogeneity is only large when S_{min} is 3, at 12%. Only 5% of patients are in groups without system choice heterogeneity when the minimum group size is 5, and 1% are when the minimum group size is 10. No patients are in such groups for a minimum group size of 25 or above. Thus, so long as S_{min} is set to at least 5 patients, groups without system choice heterogeneity do not

our findings.

¹⁸For the sake of conciseness, we selected RMSE as the goodness of fit measure to report.

seem to be a major issue.

Cross-validation on a very large dataset could take hours of computational time. Therefore, we also conduct a cross-validation using a random sample of the data, as described above in Section 2.2.1. In this alternative cross-validation, we randomly select 1000 observations from the data to serve as our validation points. This cross-validation took only minutes of computational time on a standard desktop computer and yielded similar results to cross-validation on the full dataset. This result is reassuring, since it suggests that a cross-validation exercise to suggest the minimum group size can be done quickly.¹⁹

4 Monte Carlo Analysis

The results presented in the previous section suggest that estimates of the diversion ratio and percent change in WTP are relatively robust to the minimum group size, and that the diversion ratio is more robust than the percent change in WTP. However, it is impossible to analyze the performance of the semiparametric model from the estimated results since we do not know the true value for these statistics. In this section, we use the obtained estimates to calibrate a model where we know the true value for the diversion ratios and WTP, and assess whether the semiparametric model can accurately estimate these statistics.

We undertake the following Monte Carlo analysis to assess the semiparametric model's performance in a real-world setting. Admissions are randomly sampled with replacement

¹⁹We sample one thousand people from the data fifty times. The results obtained using sampling are similar to those from using the full data set. For RMSE, in 82% of the data samples, a minimum group size of 25 minimizes the RMSE. For the remainder, the RMSE is smallest for a minimum group size of either 10 or 50 (14% and 4% of the data samples, respectively). For the pseudo R^2 measure, in 52% of the samples, a minimum group size of 10 maximizes the statistic. For the remainder, the minimum group size is either 3 (2%), 5 (40%), or 25 (6%).

from the hospital discharge data employed earlier. For each sampled admission, we randomly generate a hospital choice using the admission’s predicted choice probabilities from the semiparametric model estimated in [Section 3.3](#) for a given choice of minimum group size. This procedure results in a realistically calibrated data sample for which we know the true probability that a given patient will select any given hospital, and thus the true diversion ratios and percent change in WTP.

For various model specifications, we use this data generating process to simulate 50,000 datasets that contain the same number of admissions as the original data. Each simulated dataset is used to estimate the semiparametric model for a given choice of S_{min} , as well as the diversion ratio and the percent change in WTP. We then calculate the RMSE of the percent difference between each estimated statistic and its true value across the Monte Carlo simulations. Since we employ a large number of simulations, the estimated RMSE should accurately represent the magnitude of the semiparametric estimator’s RMSE.

We can then consider the robustness of the obtained results to model misspecification. We start with the true model calibrated based on the $S_{min} = 50$ specification. [Figure 3a](#) and [Figure 3b](#) display the RMSE for the diversion ratio and percent change in WTP for different values of S_{min} .

The percent change in WTP exhibits a U shape, with higher RMSE for very low and very high values of S_{min} . Intuitively, estimates are biased for very inflexible models, while estimates for very flexible models both have high variance as well as bias due to groups without heterogeneity. For the diversion ratio, the RMSE is still fairly low for very flexible models, with only small changes before a minimum group size of 250, but does increase considerably for sufficiently inflexible models. For both the diversion ratio and percent

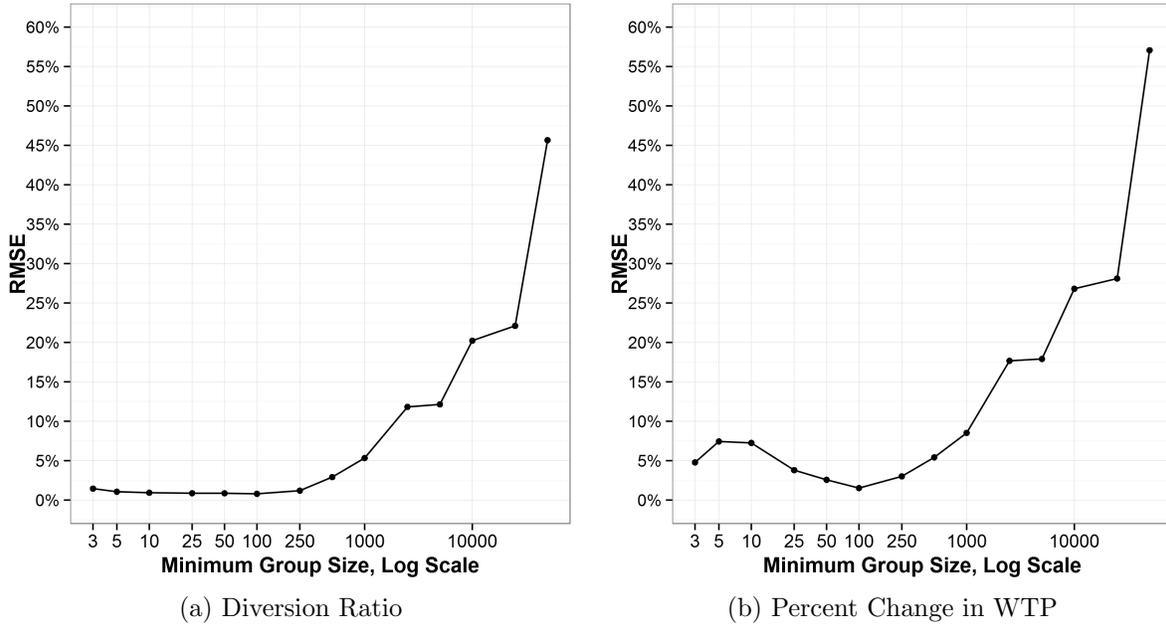


Figure 3 Estimated RMSE by S_{min} , When True Value of S_{min} is 50

Note: Diversion ratio is from System 2 to System 1. [Table II](#) presents these estimates numerically.

change in WTP, the RMSE is at its lowest when S_{min} is 100.

However, all values of S_{min} below 500 have an RMSE less than 1.5% for the diversion ratio, and all values between 3 and 500 have an RMSE below 8% for the percent change in WTP. Thus, errors may be fairly small so long as S_{min} is set within an intermediate range, although the percent change in WTP has a significantly higher error rate than the diversion ratio.

When patient heterogeneity is sufficiently prevalent, it may not be possible to choose a value for the minimum group size that is both large enough to avoid the variance associated with small groups but small enough to avoid bias from being insufficiently flexible. To examine this issue, we consider the results of a Monte Carlo analysis in which the data generating process corresponds to the $S_{min} = 3$ calibration.

The results from this analysis are displayed in [Figure 4a](#) and [Figure 4b](#). In this case, the

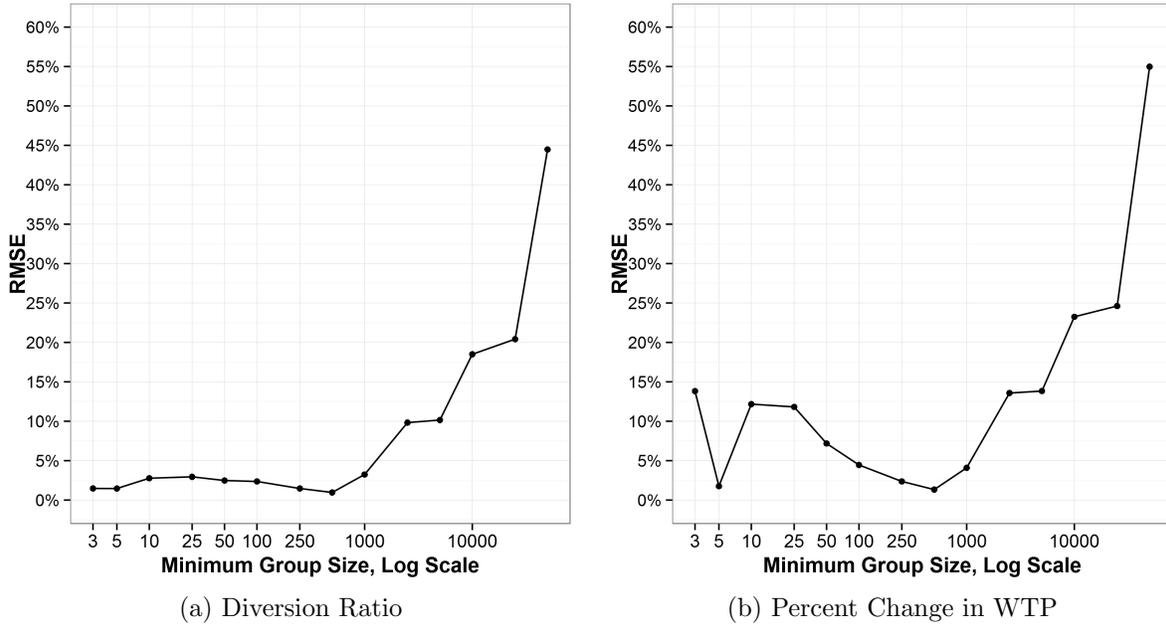


Figure 4 Estimated RMSE by S_{min} , When True Value of S_{min} is 3

Note: Diversion ratio is from System 2 to System 1. [Table III](#) presents these estimates numerically.

RMSE for both statistics has two local minima, at a minimum group size of 5 and 500, with the global minimum at 500. Again, the RMSE for the diversion ratio is fairly flat below a minimum group size of 1,000, with the RMSE below 3.5% for all values between 3 and 500. The RMSE for the percent change in WTP is less stable. It is below 15% for S_{min} between 3 and 1,000, although it is much lower, at 1.3%, at the minimum point of 500. These results suggest that, even if one is concerned that patients may have very heterogeneous hospital preferences, one may not need to use extremely flexible specifications to avoid a high degree of error. The choice of S_{min} appears to be more important for the percent change of WTP than for the diversion ratio.

Next, we analyze the efficiency loss from using a semiparametric model. The use of a properly specified parametric logit model will generally provide more precise estimates than the semiparametric model. The degree of inefficiency will depend on the functional form of

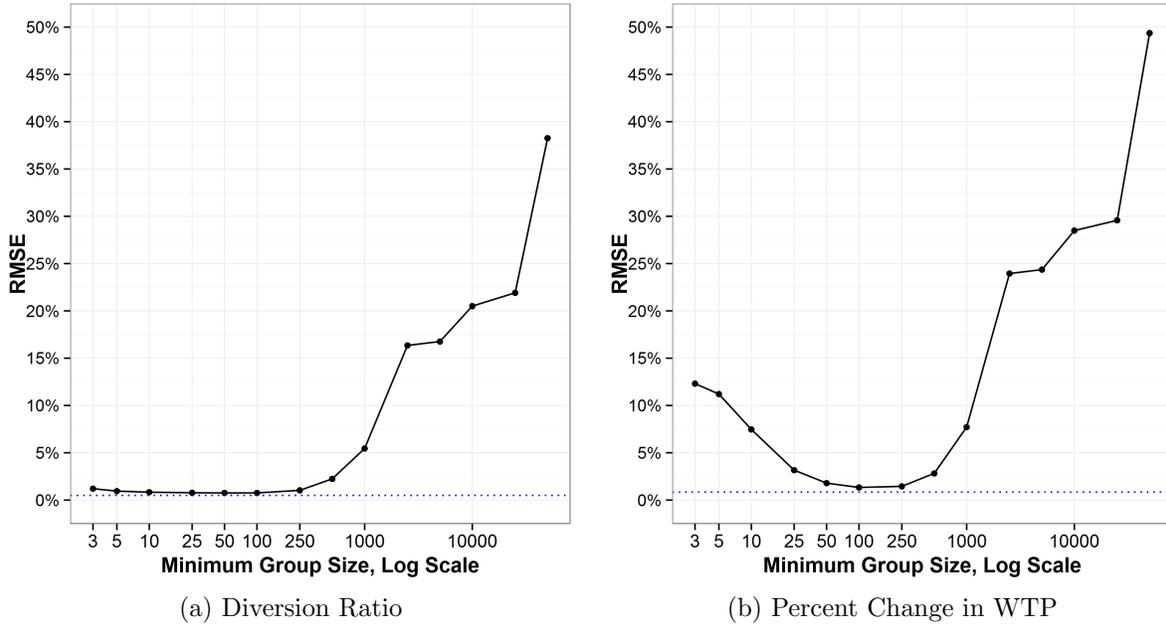


Figure 5 Estimated RMSE by S_{min} , When True Model is Parametric Logit

Note: Diversion ratio is from System 2 to System 1. Blue dotted line shows RMSE using the parametric logit, using 1,000 simulations. [Table IV](#) presents these estimates numerically.

the parametric specification. The semiparametric estimator is particularly inefficient when the true model has a simple, known parameterization. We consider a Monte Carlo analysis where the data generating process is a particularly simple specification to gain a better understanding of the “worst case” scenario for the semiparametric model. First, we use the dataset employed earlier to estimate a parametric logit model that controls only for travel time, its square, and a set of hospital fixed effects. We predict choice probabilities for each patient from the model estimates, which are then used to generate new hospital choices for admissions that are randomly sampled with replacement from the data.

The results from this analysis are presented in [Figure 5](#). We again see a U shaped curve for the RMSE for the percent change in WTP, and a flat RMSE that only rises with high values of S_{min} for the diversion ratio. The RMSE is at its lowest when the minimum group

size is 50 for the diversion ratio, and 100 for the percent change in WTP. The RMSE is approximately 1% or less for the diversion ratio for all values of S_{min} below 250. For the percent change in WTP, the RMSE is 3.5% or below for S_{min} between 25 and 500.

The blue dotted line represents the RMSE when the model is estimated using the correctly specified simple parametric logit.²⁰ The RMSE of the parametric logit is 0.5% for the diversion ratio and 0.8% for the percent change in WTP; the RMSE for the semiparametric logit is very close to the RMSE for the parametric logit for intermediate values of S_{min} . Thus, the efficiency loss from using a semiparametric model is low so long as the minimum group size is set within an intermediate range. This analysis considers the performance of the semiparametric model when the true model is an extremely simple parametric alternative. In a realistic setting where the data generating process is more complex, the inefficiency from using a semiparametric estimator is presumably smaller.

We conclude by considering the performance of the model for different sample sizes. We generate data samples between 5 percent and 100 percent of the original sample size of 124,237. In this Monte Carlo analysis, we estimate a correctly specified model where the data generating process corresponds to the $S_{min} = 50$ specification. The Monte Carlo results presented in [Figure 6](#) suggest that the semiparametric model performs relatively well even when the number of admissions is quite small. The RMSE does rise as the sample size shrinks, but the RMSE is below 4% when the sample size is at or above 10% of the original for the diversion ratio, and below 6% for the percent change in WTP.

These results suggest that the semiparametric model can be usefully applied in a wide

²⁰Due to the much higher computation burden, we employ only 1,000 simulations, rather than the 50,000 used for the semiparametric model.

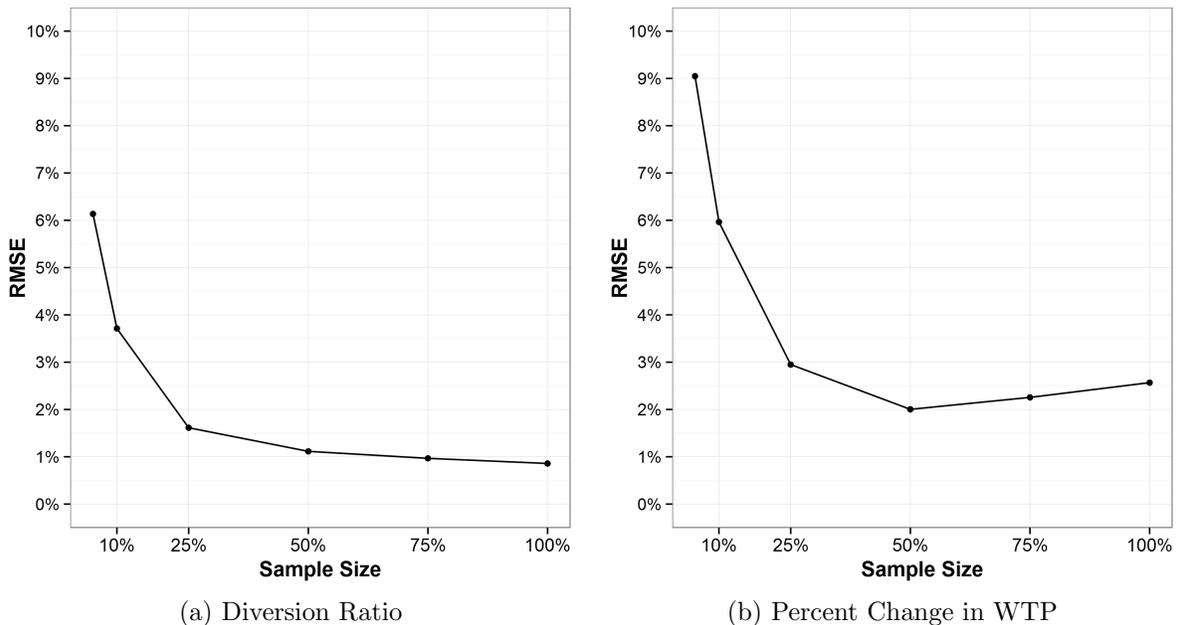


Figure 6 Estimated RMSE by Sample Size

Note: Diversion ratio is from System 2 to System 1. The true minimum group size is 50. The x axis is the fraction of the overall sample size. [Table V](#) presents these estimates numerically.

range of settings, even when the dataset contains a small number of observations. One reason why the semiparametric model performs relatively well even for small sample sizes is that the iterative grouping procedure automatically adjusts model flexibility to the size of the data sample. For a fixed value for the minimum group size S_{min} , the grouping procedure puts patients into a smaller number of groups when the sample size is smaller. This avoids potential biases associated with using very small group sizes in the semiparametric model, although it can lead to bias from model inflexibility if the sample size becomes too small.

5 Conclusion

When presented with rich microdata, researchers must balance the competing objectives of allowing for significant individual level heterogeneity while ensuring statistical power. In the

parametric logit models that are typically used, the extent of the permitted heterogeneity is limited by the parametric specification of the model. To complement these methods, we developed a semiparametric discrete choice estimator that allows for rich heterogeneity across the population. Highlighting the importance of allowing for rich heterogeneity in choice patterns, [Raval et al. \(2015\)](#) find that our proposed estimator outperforms many parametric multinomial choice models previously used in the literature in predicting choices after a change in the choice set.

In our estimator, the trade-off between heterogeneity and power is determined by a single tuning parameter, the minimum group size. We applied our semiparametric method to patient discharge data and simulated a merger of two hospital systems to test the estimators' sensitivity to this parameter. While we suggested a possible cross-validation approach to choosing this parameter, we found that the main substitution measures are relatively insensitive to the choice of minimum group size. However, of the two measures, the diversion ratio was less sensitive than willingness to pay to the choice of minimum group size.

These results should give researchers confidence to pursue this approach in healthcare, or other settings where rich microdata are available. Further, they should encourage research in other methods that relax the functional form restrictions that underlie most empirical work in discrete choice demand modeling. The increased availability of large datasets and recent research in “machine learning” approaches suggest that advances in this area may be on the horizon.

References

- Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes**, “Econometric Tools for Analyzing Market Outcomes,” *Handbook of Econometrics*, 2007, 6, 4171–4276.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 2004, 112 (1), 68–105.
- Breiman, Leo, Jerome Friedman, Charles J. Stone, and R.A. Olshen**, *Classification and Regression Trees*, Chapman and Hall, 1984.
- Capps, Cory, David Dranove, and Mark Satterthwaite**, “Competition and Market Power in Option Demand Markets,” *RAND Journal of Economics*, 2003, 34 (4), 737–763.
- Conlon, Christopher T. and Julie Holland Mortimer**, “An Experimental Approach to Merger Evaluation,” *NBER Working Paper*, 2013.
- Crawford, Gregory S. and Ali Yurukoglu**, “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, 2012, 102 (2), 643–685.
- Doane, Michael J., Luke Froeb, and R. Lawrence Van Horn**, “How Well Do Travel Cost Models Measure Competition Among Hospitals?,” *mimeo*, 2012.
- Farrell, Joseph and Carl Shapiro**, “Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition,” *The BE Journal of Theoretical Economics*, 2010, 10 (1), 1–39.
- , **David J. Balan, Keith Brand, and Brett W. Wendling**, “Economics at the FTC: Hospital Mergers, Authorized Generic Drugs, and Consumer Credit Markets,” *Review of Industrial Organization*, 2011, 39 (4), 271–296.
- Fournier, Gary M. and Yunwei Gai**, “What does Willingness-to-Pay Reveal about Hospital Market Power in Merger Cases?,” *mimeo*, 2007.
- Gaynor, Martin, Kate Ho, and Robert J. Town**, “The Industrial Organization of Health-Care Markets,” *Journal of Economic Literature*, 2015, 53 (2), 235–284.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town**, “Mergers when Prices are Negotiated: Evidence from the Hospital Industry,” *American Economic Review*, 2015, 105 (1), 172–203.
- Greene, William**, “The Behaviour of the Maximum Likelihood Estimator of Limited Dependent Variable Models in the Presence of Fixed Effects,” *Econometrics Journal*, 2004, 7 (1), 98–119.
- Grennan, Matthew**, “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, 2013, 103 (1), 145–177.
- Ho, Kate and Ariel Pakes**, “Hospital Choices, Hospital Prices, and Financial Incentives to Physicians,” *American Economic Review*, 2014, 104 (12), 3841–3884.
- Ho, Katherine**, “The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market,” *Journal of Applied Econometrics*, 2006, 21 (7), 1039–1079.

- Jaffe, Sonia and E. Glen Weyl**, “The First-Order Approach to Merger Analysis,” *American Economic Journal: Microeconomics*, 2013, 5 (4), 188–218.
- May, Sean M.**, “How Well Does Willingness-to-Pay Predict the Price Effects of Hospital Mergers?,” *mimeo*, 2013.
- Pagan, Adrian and Aman Ullah**, *Nonparametric Econometrics*, Cambridge University Press, 1999.
- Raval, Devesh, Ted Rosenbaum, and Nathan E. Wilson**, “Industrial Reorganization: Learning about Patient Substitution Patterns from Natural Experiments,” *mimeo*, 2015.
- Shapiro, Carl**, “Mergers with Differentiated Products,” *Antitrust*, 1996, 10 (2), 23–30.
- Stone, M.**, “Cross-Validatory Choice and Assessment of Statistical Predictions,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1974, 36 (2), 111–147.

Table I Semiparametric Logit Estimates for Alternative Minimum Group Sizes

	3	5	10	25	50	100	250	500	1,000	2,500	5,000	10,000	25,000	50,000
Diversion, System 1 to 2	19.7% (0.32%)	19.9% (0.30%)	20.0% (0.29%)	19.9% (0.26%)	19.7% (0.25%)	19.7% (0.23%)	19.9% (0.21%)	19.9% (0.20%)	19.6% (0.20%)	18.7% (0.18%)	18.6% (0.18%)	16.8% (0.17%)	16.4% (0.16%)	10.9% (0.10%)
Diversion, System 2 to 1	42.4% (0.56%)	43.2% (0.53%)	43.7% (0.49%)	43.8% (0.41%)	43.4% (0.38%)	43.4% (0.36%)	42.9% (0.32%)	42.2% (0.30%)	41.0% (0.28%)	38.3% (0.23%)	38.1% (0.23%)	34.6% (0.23%)	33.8% (0.23%)	23.6% (0.13%)
WTP, post-merger % chng	17.4% (0.23%)	18.7% (0.30%)	n 19.5% (0.33%)	19.0% (0.33%)	18.3% (0.30%)	18.0% (0.27%)	17.7% (0.22%)	17.3% (0.20%)	16.7% (0.19%)	15.1% (0.15%)	15.0% (0.15%)	13.4% (0.14%)	13.1% (0.14%)	7.8% (0.06%)
Number of groups	23,157	13,709	6,932	2,858	1,334	604	209	110	62	12	7	4	3	1
Average group size	5	9	18	43	93	206	594	1,129	2,004	10,353	17,748	31,059	41,412	124,237
McFadden's pseudo R^2	0.70	0.65	0.60	0.54	0.51	0.48	0.46	0.44	0.42	0.37	0.36	0.31	0.29	0.18
% observations without intra-group system choice heterogeneity	12%	5%	1%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

Notes: The semiparametric model is estimated using hospital discharge data for 124,237 admissions. Each column reports results for a separate specification that uses the indicated minimum group size. Bootstrapped standard errors based on 50,000 simulations are reported in parentheses.

Table II Monte Carlo Results, Minimum Group Size of 50 Patients is the True Model

	3	5	10	25	50	100	250	500	1,000	2,500	5,000	10,000	25,000	50,000
Diversión, System 2 to 1														
Avg estimate	42.9%	43.3%	43.4%	43.4%	43.4%	43.4%	42.9%	42.1%	41.1%	38.2%	38.1%	34.6%	33.8%	23.6%
Avg % error	-0.95%	-0.21%	0.00%	0.09%	0.16%	0.00%	-0.93%	-2.83%	-5.29%	-11.81%	-12.12%	-20.22%	-22.09%	-45.65%
SD of estimate	0.47%	0.45%	0.40%	0.37%	0.37%	0.34%	0.32%	0.30%	0.28%	0.23%	0.24%	0.23%	0.23%	0.13%
SD of % error	1.09%	1.04%	0.93%	0.86%	0.85%	0.79%	0.74%	0.68%	0.65%	0.54%	0.54%	0.52%	0.53%	0.29%
RMSE, absolute error	0.63%	0.46%	0.40%	0.37%	0.37%	0.34%	0.52%	1.26%	2.31%	5.12%	5.26%	8.77%	9.58%	19.79%
RMSE, % error	1.45%	1.06%	0.93%	0.86%	0.86%	0.79%	1.19%	2.91%	5.33%	11.82%	12.13%	20.22%	22.10%	45.65%
WTP, post-merger % chng														
Avg estimate	19.1%	19.6%	19.6%	18.9%	18.6%	18.2%	17.8%	17.3%	16.7%	15.0%	15.0%	13.4%	13.1%	7.8%
Avg % error	4.49%	7.25%	7.05%	3.44%	2.01%	-0.62%	-2.75%	-5.30%	-8.45%	-17.65%	-17.88%	-26.80%	-28.09%	-57.06%
SD of estimate	0.30%	0.30%	0.30%	0.29%	0.29%	0.25%	0.22%	0.20%	0.19%	0.15%	0.15%	0.14%	0.14%	0.06%
SD of % error	1.63%	1.63%	1.67%	1.60%	1.59%	1.38%	1.23%	1.11%	1.04%	0.81%	0.82%	0.75%	0.75%	0.33%
RMSE, absolute error	0.87%	1.36%	1.32%	0.69%	0.47%	0.28%	0.55%	0.99%	1.56%	3.23%	3.27%	4.90%	5.13%	10.42%
RMSE, % error	4.77%	7.44%	7.25%	3.79%	2.57%	1.52%	3.01%	5.42%	8.52%	17.66%	17.90%	26.81%	28.10%	57.06%

Notes: Each column reports results from a different Monte Carlo specification consisting of 50,000 simulations. A random sample of 124,237 admissions is generated for each simulation. The hospital choice for each admission is randomly generated based on estimates from the model reported in [Table I](#) for a minimum group size of 50 patients. Each simulated dataset is used to estimate the semiparametric model for the indicated minimum group size. % of estimate is the average of the statistic (i.e., % change in WTP or diversion ratio) across simulations divided by true value of the statistic in the data. SD of % error is the standard deviation of the statistic across simulations divided by the true value of the statistic in the data. RMSE, % error is the root mean squared error of the statistic, computed across simulations, also divided by the true value of the statistic in the data.

Table III Monte Carlo Results, Minimum Group Size of 3 Patients is the True Model

	3	5	10	25	50	100	250	500	1,000	2,500	5,000	10,000	25,000	50,000	
Diversions, System 2 to 1															
Avg estimate	42.1%	42.8%	43.5%	43.6%	43.4%	43.4%	43.0%	42.2%	41.1%	38.3%	38.1%	34.6%	33.8%	23.6%	
Avg % error	-0.76%	0.80%	2.52%	2.78%	2.31%	2.21%	1.25%	-0.66%	-3.18%	-9.81%	-10.14%	-18.49%	-20.40%	-44.47%	
SD of estimate	0.54%	0.52%	0.49%	0.41%	0.38%	0.35%	0.32%	0.30%	0.28%	0.24%	0.23%	0.22%	0.23%	0.13%	
SD of % error	1.26%	1.22%	1.14%	0.97%	0.89%	0.83%	0.76%	0.70%	0.67%	0.56%	0.55%	0.53%	0.55%	0.30%	
RMSE, absolute error	0.62%	0.62%	1.18%	1.25%	1.05%	1.00%	0.62%	0.41%	1.38%	4.17%	4.31%	7.85%	8.66%	18.87%	
RMSE, % error	1.47%	1.46%	2.77%	2.95%	2.48%	2.36%	1.47%	0.96%	3.24%	9.83%	10.15%	18.49%	20.41%	44.48%	
WTP, post-merger % chng															
Avg estimate	15.0%	17.3%	19.5%	19.5%	18.6%	18.1%	17.8%	17.3%	16.7%	15.1%	15.0%	13.4%	13.1%	7.8%	
Avg % error	-13.73%	-0.47%	12.02%	11.67%	6.97%	4.19%	1.98%	-0.64%	-3.94%	-13.55%	-13.81%	-23.24%	-24.60%	-54.98%	
SD of estimate	0.26%	0.30%	0.33%	0.33%	0.30%	0.26%	0.23%	0.20%	0.19%	0.15%	0.15%	0.14%	0.14%	0.06%	
SD of % error	1.52%	1.70%	1.90%	1.87%	1.71%	1.51%	1.29%	1.17%	1.10%	0.86%	0.85%	0.78%	0.79%	0.35%	
RMSE, absolute error	2.41%	0.31%	2.12%	2.06%	1.25%	0.78%	0.41%	0.23%	0.71%	2.36%	2.41%	4.05%	4.29%	9.58%	
RMSE, % error	13.82%	1.76%	12.17%	11.82%	7.18%	4.45%	2.36%	1.33%	4.09%	13.58%	13.83%	23.25%	24.61%	54.98%	

Notes: Each column reports results from a different Monte Carlo specification consisting of 50,000 simulations. A random sample of 124,237 admissions is generated for each simulation. The hospital choice for each admission is randomly generated based on estimates from the model reported in [Table I](#) for a minimum group size of 3 patients. Each simulated dataset is used to estimate the semiparametric model for the indicated minimum group size. % of estimate is the average of the statistic (i.e., % change in WTP or diversion ratio) across simulations divided by true value of the statistic in the data. SD of % error is the standard deviation of the statistic across simulations divided by the true value of the statistic in the data. RMSE, % error is the root mean squared error of the statistic, computed across simulations, also divided by the true value of the statistic in the data.

Table IV Monte Carlo Results, Simple Parametric Logit is the True Model

	Logit	3	5	10	25	50	100	250	500	1,000	2,500	5,000	10,000	25,000	50,000	
Diversion, System 2 to 1	Avg estimate	38.2%	37.9%	38.1%	38.2%	38.1%	38.1%	37.9%	37.4%	36.1%	31.9%	31.8%	30.3%	29.8%	23.6%	
	Avg % error	-0.01%	-0.58%	-0.07%	-0.03%	-0.04%	-0.09%	-0.17%	-2.12%	-5.41%	-16.34%	-16.75%	-20.50%	-21.90%	-38.26%	
	SD of estimate	0.19%	0.40%	0.36%	0.32%	0.29%	0.29%	0.28%	0.28%	0.28%	0.27%	0.19%	0.19%	0.19%	0.13%	
	SD of % error	0.49%	1.05%	0.95%	0.83%	0.77%	0.75%	0.74%	0.74%	0.72%	0.70%	0.51%	0.49%	0.50%	0.33%	
	RMSE, absolute error	0.19%	0.46%	0.36%	0.32%	0.29%	0.29%	0.29%	0.39%	0.85%	2.08%	6.24%	6.39%	7.82%	8.36%	14.60%
	RMSE, % error	0.49%	1.20%	0.95%	0.83%	0.77%	0.75%	0.76%	1.03%	2.24%	5.45%	16.35%	16.76%	20.50%	21.91%	38.26%
WTP, post-merger % chng	Avg estimate	15.5%	17.4%	17.2%	16.6%	15.9%	15.7%	15.4%	15.1%	14.3%	11.8%	11.7%	11.1%	10.9%	7.8%	
	Avg % error	-0.03%	12.18%	11.08%	7.30%	2.82%	1.19%	0.35%	-2.53%	-7.62%	-23.94%	-24.35%	-28.48%	-29.56%	-49.36%	
	SD of estimate	0.13%	0.27%	0.26%	0.24%	0.22%	0.21%	0.20%	0.20%	0.19%	0.19%	0.11%	0.10%	0.10%	0.06%	
	SD of % error	0.84%	1.76%	1.66%	1.58%	1.42%	1.33%	1.29%	1.27%	1.25%	1.20%	0.70%	0.69%	0.67%	0.39%	
	RMSE, absolute error	0.13%	1.91%	1.74%	1.16%	0.49%	0.28%	0.21%	0.22%	0.44%	1.20%	3.71%	3.77%	4.41%	4.58%	7.64%
	RMSE, % error	0.84%	12.31%	11.21%	7.47%	3.16%	1.79%	1.33%	1.45%	2.82%	7.72%	23.95%	24.36%	28.48%	29.57%	49.36%

Notes: Each column reports results from a different Monte Carlo specification consisting of 50,000 simulations (the parametric logit is based on 1,000 simulations). A random sample of 124,237 admissions is generated for each simulation. The hospital choice for each admission is randomly generated based on estimates from a parametric logit specification that controls for travel time, its square, and a set of hospital fixed effects. Each simulated dataset is used to estimate either the parametric logit or the semiparametric model for the indicated minimum group size.

% of estimate is the average of the statistic (i.e., % change in WTP or diversion ratio) across simulations divided by true value of the statistic in the data. SD of % error is the standard deviation of the statistic across simulations divided by the true value of the statistic in the data. RMSE, % error is the root mean squared error of the statistic, computed across simulations, also divided by the true value of the statistic in the data.

Table V Monte Carlo Results, Alternative Sample Sizes

		5%	10%	25%	50%	75%	100%
Diversion, System 2 to 1	Avg estimate	41.0%	42.1%	43.1%	43.4%	43.4%	43.4%
	Avg % error	-5.38%	-2.99%	-0.60%	0.05%	0.14%	0.15%
	SD of estimate	1.27%	0.95%	0.65%	0.48%	0.41%	0.37%
	SD of % error	2.94%	2.20%	1.50%	1.11%	0.96%	0.84%
	RMSE, absolute error	2.66%	1.61%	0.70%	0.48%	0.42%	0.37%
	RMSE, % error	1.65%	1.09%	0.54%	0.37%	0.41%	0.47%
WTP, post-merger % chng	Avg estimate	16.9%	17.4%	18.0%	18.3%	18.5%	18.6%
	Avg % error	-7.60%	-4.72%	-1.50%	0.19%	1.39%	2.01%
	SD of estimate	0.90%	0.67%	0.46%	0.36%	0.32%	0.29%
	SD of % error	4.91%	3.65%	2.54%	1.99%	1.78%	1.59%
	RMSE, absolute error	6.13%	3.71%	1.61%	1.11%	0.97%	0.86%
	RMSE, % error	9.05%	5.96%	2.95%	2.00%	2.26%	2.57%

Notes: Each column reports results from a different Monte Carlo specification consisting of 50,000 simulations. A random sample with the indicated number of observations is generated for each simulation. The hospital choice for each admission is randomly generated based on estimates from the model reported in [Table I](#) for a minimum group size of 50 patients. Each simulated dataset is used to estimate the semiparametric model for a minimum group size of 50 patients.

% of estimate is the average of the statistic (i.e., % change in WTP or diversion ratio) across simulations divided by true value of the statistic in the data. SD of % error is the standard deviation of the statistic across simulations divided by the true value of the statistic in the data. RMSE, % error is the root mean squared error of the statistic, computed across simulations, also divided by the true value of the statistic in the data.

A Bias-Variance Tradeoffs with S_{min}

In this section, we consider the bias-variance tradeoffs involved in setting the minimum group size, S_{min} , when the semiparametric logit specification is used to estimate diversion and WTP. Starting with a correctly specified model, we first analyze the bias from combining two groups with heterogeneous preferences. We then consider the opposite situation in which the model is unnecessarily flexible.

We start by assuming that the semiparametric model is correctly specified. The overall diversion from hospital h to hospital j across two groups A and B is a weighted average of the group-level diversions where the weight is N_h^g , the number of patients in group g that select hospital h .

$$div_{hj} = (N_h^A div_{hj}^A + N_h^B div_{hj}^B) / (N_h^A + N_h^B) \quad (11)$$

Suppose that the two groups are combined due to the (mistaken) belief that they have the same hospital preferences. The estimated diversion \hat{div}_{hj} from hospital h to j for the combined group is simply the estimated fraction of patients in A and B , after excluding those who choose hospital h , which selects hospital j . The expected value of this estimator does not equal the actual diversion div_{hj} defined in [equation \(11\)](#):

$$E(\hat{div}_{hj}) = (N_{\sim h}^A div_{hj}^A + N_{\sim h}^B div_{hj}^B) / (N_{\sim h}^A + N_{\sim h}^B) \quad (12)$$

The expected value of the estimated diversion from hospital h to j for the combined group is still a weighted average of the group-level diversions, but now the weight is $N_{\sim h}^g$, the number of patients in group g that do not select hospital h . The estimated diversion for the combined group will be unbiased only in special cases. For example, unbiased estimates are obtained when the fraction of each group that selects hospital h is the same (i.e. $\frac{N_h^A}{N_{\sim h}^A} = \frac{N_h^B}{N_{\sim h}^B}$), or when the group-level diversions are identical, so the weighting difference does not matter (i.e., $div_{hj}^A = div_{hj}^B$). In general, however, the use of an overly restrictive model leads to biased diversion estimates.

Next, consider the potential bias from using an overly flexible model. We start with the assumption that each member of group g has identical preferences, and then divide this homogeneous group into two subgroups A and B . The estimated diversion from hospital h to j across the two groups is as follows.

$$\hat{div}_{hj} = (N_h^A \hat{div}_{hj}^A + N_h^B \hat{div}_{hj}^B) / (N_h^A + N_h^B). \quad (13)$$

Since all members of the two subgroups have the same preferences, $E(\hat{div}_{hj}) = E(\hat{div}_{hj}^A) = E(\hat{div}_{hj}^B) = div_{hj}$. That is, the estimated diversion across the two groups is an unbiased estimate of the true diversion. There is a caveat, however: one must be able to estimate the diversion from hospital h to j for each subgroup. This is not possible when a group is composed solely of individuals that select hospital h .

However, diversion ratios will be less precisely estimated. Returning to the example where a homogeneous group is divided into two subgroups A and B , let α_h denote the fraction of patients, among those who choose hospital h , that are put into group A . Similarly, let $\alpha_{\sim h}$ denote the fraction of those who do not choose hospital h that are put into group A . The variance of the estimated diversion \hat{div}_{hj} defined in [equation \(13\)](#) is as follows:

$$V(\hat{div}_{hj}) = \phi_h \frac{div_{hj}(1 - div_{hj})}{N_{\sim h}}. \quad (14)$$

where $\phi_h = \frac{\alpha_h^2}{\alpha_{\sim h}} + \frac{(1-\alpha_h)^2}{1-\alpha_{\sim h}}$. When $\phi_h = 1$, the variance of the estimated diversion calculated separately for each group equals the variance of the estimated diversion when it is calculated for the combined group. Holding $\alpha_{\sim h}$ fixed, ϕ_h is a convex function of α_h that has a minimum at $\alpha_h = \alpha_{\sim h}$, at which point $\phi_h = 1$. That is, the only time that there is no efficiency loss from dividing a homogeneous group into two subgroups is when those selecting hospital h and those selecting other hospitals are allocated to the two groups in similar proportions. While this condition may approximately hold when a large group is divided, it is less likely to hold when group sizes are small due to random sampling. Thus, the efficiency cost from dividing a large group into medium sized groups is likely to be less than the loss in efficiency from dividing a medium sized group into small groups.

The use of an overly flexible model also impacts WTP estimates. As before, we start with a single group with homogeneous preferences. The group's estimated WTP (per person) for hospital h is as follows:

$$W\hat{T}P_h^g = -\ln(1 - \hat{s}_h^g). \quad (15)$$

Next, we divide the group into two. The average WTP (per person) for hospital h across the two groups is estimated as follows:

$$W\hat{T}P_h = -[N^A \ln(1 - \hat{s}_h^A) + N^B \ln(1 - \hat{s}_h^B)] / (N^A + N^B). \quad (16)$$

Both [equation \(15\)](#) and [equation \(16\)](#) provide consistent estimates of the true WTP for the group, although the estimates will not be unbiased since WTP is a nonlinear function. However, the use of a more flexible model can have a significant impact in finite samples. Since WTP is a convex function, and $\hat{s}_h = (N^A \hat{s}_h^A + N^B \hat{s}_h^B) / (N^A + N^B)$, the WTP estimate from [equation \(16\)](#) is weakly larger than the WTP estimate using [equation \(15\)](#). This can lead to an inference problem where it is unclear whether estimated WTP is high because patients strongly value a given hospital, or because an overly flexible model is being employed.

For the change in WTP, both the numerator – the WTP of the combined system – and denominator – the WTP of each individual system – increase with a more flexible model. Thus, the effect of changing the group size on the change in WTP is ambiguous.