ECONOMIC ANALYSIS GROUP DISCUSSION PAPER

Win/Loss Data, Consumer Switching Costs: Measuring Diversion Ratios and the Impact of Mergers

by

Jeff Qiu[1], Masayuki Sawada[2], and Gloria Sheu[3]

EAG 22-4 December 2022

[1] U.S. Department of Justice, Antitrust Division E-mail: yinjia.qiu@usdoj.gov

[2] Institute of Economic Research, Hitotsubashi Universit. E-mail: m-sawada@ier.hit-u.ac.jp

[3] Board of Governors of the Federal Reserve System. E-mail: gloria.sheu@frb.gov

Win/Loss Data and Consumer Switching Costs: Measuring Diversion Ratios and the Impact of Mergers*

Jeff Qiu† Masayuki Sawada‡ Gloria Sheu§

November 25, 2022

Abstract

The diversion ratio is a key input to many indicators of merger harm. Measuring the diversion ratio, however, is challenging in the presence of state dependence driven by things like consumer switching costs. We propose an identification strategy for diversion based on win/loss data. First, we show that win/loss data from the merging firms and market shares for all firms in two periods are sufficient to identify the diversion ratios between the merging partners. Second, we show that win/loss data from the merging firms alone are sufficient for partial identification, and we construct a lower bound that provides a good approximation to the diversion ratio when switching costs are high. We demonstrate the performance of our method with numerical simulations and with an application to the Anthem/Cigna merger.

*The views expressed herein are entirely those of the authors and should not be purported to reflect those of the U.S. Department of Justice. Furthermore, the analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the Board research staff or by the Federal Reserve Board of Governors.

†U.S. Department of Justice, Antitrust Division. e-mail: yinjia.qiu@usdoj.gov

‡Institute of Economic Research, Hitotsubashi University. e-mail: m-sawada@ier.hit-u.ac.jp

§Board of Governors of the Federal Reserve System. e-mail: gloria.sheu@frb.gov

1 Introduction

The diversion ratio is a key concept at the heart of work in antitrust and industrial organization. Defined as the percentage of unit sales lost by a product whose price has increased that are gained by a competing product, the ratio is important in assessing the impacts from policies that affect relative prices. The 2010 Federal Trade Commission (FTC) and U.S. Department of Justice (DOJ) Horizontal Merger Guidelines (henceforth, the "Guidelines") emphasize that diversion ratios are a component of upward pricing pressure (UPP). When substitute products come under joint ownership, the UPP created by the merger is the diversion ratio between these products times the profit margin of the good being diverted to. According to the Guidelines, diversion "can be very informative for assessing unilateral price effects, with higher diversion ratios indicating a greater likelihood of such effects."[1]

However, measuring the diversion ratio is not always straightforward, particularly when consumer preferences exhibit state dependence. Such dependence may stem from switching costs, consumer inattention, or other factors that lead to demand hysteresis. We have in mind situations where the probability of choosing a product today depends on what was chosen last period. Such behavior is common for products sold under long-term contracts, such as health insurance or telecommunications services.[2] Market share data may include sales that took place some time ago, meaning these data may reflect competitive actions that are no longer relevant. This issue can create challenges for measuring diversion from structural demand estimation, as such techniques frequently take market shares as an input.

An alternative strategy is to rely on information for customers that have recently switched products. Data of this type identifies the number of customers that previously bought from one firm and who their current supplier is, which we call loss data, and the number of customers that currently chose one firm and who their previous supplier was, which we call win data. Together we call these "win/loss" data, though they are sometimes also referred to as "switching" or "churn" data. Such

[1] See §6.1 of the Guidelines.

[2] Switching products in these instances often involves some type of disruption for the customer, such as having to change doctors, buy a new cell phone, or have new cable equipment installed.

data are frequently available to antitrust authorities, at least for the merging firms, in industries where demand exhibits state dependence.[3] In these industries, consumers often renew with the same provider repeatedly, forming a lucrative stream of income should they be enticed to switch. Therefore, in these instances it is in firms' interest to track how their customers move between suppliers over time, which incentivizes firms to expend effort on collecting win/loss data. Our methods are not meant for cases where switching data are available but consumption does not exhibit state dependence.

In this paper, our goal is to derive measures of diversion that use win/loss data in an accessible way and are consistent with a consumer demand model that includes state dependence. We consider a myopic discrete choice model, in which the utility of a consumer in the current period depends on their choice in the previous period. Keeping in mind the limited data typically available to antitrust practitioners, we propose a tiered identification strategy that relies only on market shares and win/loss information. First, we show that the diversion ratios between the merging firms are point identified with knowledge of market shares for all firms in two consecutive periods and win/loss data from the merging firms alone. Second, we show that, should market share data be unavailable, win/loss data from the merging entities are sufficient to form bounds on the diversion ratios between the merging firms. We present results for both logit and mixed logit (with observed consumer types) demands. Third, we demonstrate through numerical simulations that the lower bounds we construct are informative of the true diversion ratios. We apply our methods to data from the Anthem/Cigna merger trial (2016), thus providing a real world example.

The most common approximations to diversion from switching data are based on percentages of switching customers. One measure, which we call the "loss ratio," is the fraction of customers that switched to a given product out of all customers that switched from another specific product. Similarly, one can calculate the "win

[3] Antitrust agencies arguably have more leverage to compel information from the merging firms than from non-merging firms, as the former are beholden to the agencies for approval of (or non-objection to) their proposed transaction, whereas the latter are not.

ratio" as the fraction of customers that switched from a given product out of all customers that switched to another specific product. Such calculations were done, for instance, in the aforementioned Anthem/Cigna trial by the expert testifying on behalf of the DOJ.[4] They were also done by the Federal Communications Commission (FCC) in the AT&T/T-Mobile investigation (2011), by the DOJ expert in the H&R Block/TaxACT trial (2011), and by the FTC expert in the Staples/Office Depot trial (2016).[5]

However, Carlton and Israel (2021) point out that approximating diversion with such ratios makes potentially strong assumptions about the structure of demand. Chen and Schwartz (2016), who characterize diversion in a Hotelling oligopoly model, show that the loss ratio can frequently depart from diversion in significant ways. We have a similar finding, as we show that the loss ratio can overstate or understate diversion in our framework. Our measures of diversion are derived directly from the model and thus address this issue.

As a caveat to our analysis, we acknowledge that win/loss data may be generated by consumers switching between products for reasons other than changes in relative prices. Our baseline framework side-steps this issue by relying on the logit model, which implies that diversions due to changes in a number of factors are the same as that due to changes in price. However, this assumption may not be appropriate when customers have systematic differences in preferences, in which case the source of consumer switching can matter. One may also wonder about the appropriateness of the logit assumption generally. In order to address these concerns, we also calculate our diversion estimates when the true underlying model has random coefficients and document when our method performs well (in the presence of price and quality preference heterogeneity) and less well (in the presence of switching cost heterogeneity).[6]

[4]See Dranove (2016) at slide 46.

[5]See the Staff Analysis and Findings in WT Docket No. 11-65 at paragraph 55, available at `https://docs.fcc.gov/public/attachments/DA-11-1955A2.pdf`, the H&R Block/TaxACT Memorandum Opinion at pages 35-36, and the demonstrative slides of Carl Shapiro at slides 38-51, available at `https://www.ftc.gov/system/files/documents/cases/170216staples_redacted_shapiro_demonstrative.pdf`.

[6]Note that we do not correct for switching data that suffers from measurement error, as such

A different approach than using win/loss data is to rely on survey data asking customers about their substitution behavior. Although, in our experience, we have encountered win/loss data more often than surveys in U.S. merger cases, the same is not true of all jurisdictions. The U.K. Competition and Markets Authority (CMA) encourages merging firms to submit survey evidence and will at times commission surveys themselves.[7] Surveys can focus specifically on the likely impact of a merger, whereas win/loss data may reflect a variety of other market forces. On the other hand, if win/loss data are collected in the normal course of business in order to make competitively sensitive decisions, it seems these data are likely to be informative. Thus, we view these different sources of data as complementary. If both are available, we suggest that the resulting diversion estimates should be compared and weighed against what the practitioner knows about the reliability of the underlying data collection mechanisms. In cases where survey data are not available, perhaps because customers are difficult to contact or there are resource constraints on running a survey, win/loss data can provide an alternative.

Our paper is related to a few different strands of the industrial organization literature and to current antitrust practice. Diversion ratios, along with profit margins, are a key input into several ways of gauging the likelihood of harm from horizontal mergers. Two of the most popular methods are UPP, developed by Farrell and Shapiro (2010), and the compensating marginal cost reduction (CMCR), developed by Werden (1996).

The literature proposes several econometric estimation techniques to recover diversion ratios. A number of articles, following the seminal work of Berry et al. (1995) and Goldberg (1995), provide methods for estimating flexible demand systems and price elasticities from a variety of market-level and consumer-level datasets. Conlon and Mortimer (2021) take a complementary approach, showing how diversion estimates from second-choice data or switching surveys can be interpreted as treat-

issues have been addressed in the wider econometric literature.

[7] See the CMA guidance document "Good Practice in the Design and Presentation of Customer Survey Evidence in Merger Cases," updated as of May 2018, available at https://www.gov.uk/government/publications/mergers-consumer-survey-evidence-design-and-presentation/good-practice-in-the-design-and-presentation-of-customer-survey-evidence-in-merger-cases.

ment effects estimators. However, these techniques require sufficient data to allow for econometric estimation. In contrast, antitrust practitioners often have access to limited information, sometimes only a small set of data on market shares or on lost customers in one or two time periods. Our findings are relevant for these situations with scarce data.

The paper proceeds as follows. Section 2 explains our demand model and the associated diversion ratios. Section 3 contains our main identification results. We show numerical simulations documenting the performance of our estimators in Section 4 and results for the Anthem/Cigna example in Section 5. Section 6 concludes. The Appendix contains additional information on extensions to nested logit demand and some details on our Anthem/Cigna exercise. The Online Appendix includes supplementary numerical simulation results.

2 Theoretical Background

We begin by defining our consumer demand environment and then show what diversion ratios this setup implies. Note that we do not model the supply side of the market, in order to maintain our focus on buyer preferences, customer switching costs, and consumer diversion. Product characteristics are taken as given.

2.1 Consumer Set-Up

Our demand model follows the literature on consumer switching costs, as in, for example, Dube et al. (2009). This framework has been used to study markets where demand appears to be highly sticky, such as for healthcare (e.g. Raval and Rosenbaum (2018); Shepard (2022)). State dependence is captured through a switching cost parameter.[8] We potentially allow for heterogeneity across products, and, in the mixed logit specifications, across consumers.

[8] In the interests of allowing for flexibility in interpreting the model, we are agnostic as to what precisely generates these switching costs. Various interpretations are possible, so long as the state dependence is captured by last period's purchases. The implications for consumer welfare may depend on the specific interpretation, however.

Let $\mathcal{J} = \{0, 1, ..., J\}$ be the set of products available and the index 0 be the outside option. Consumer i's indirect utility for product $j \in \mathcal{J}$ at period t when they have previously chosen product k is

$$U_{i,j,t}^{k} = \alpha p_{j,t} + \delta_{i,j,t} + \underbrace{\eta_{i,j,k}}_{\text{switching cost}} + \varepsilon_{i,j,t}, \tag{1}$$

where $p_{j,t}$ is the price, $\delta_{i,j,t}$ is a consumer-product-time specific index, $\eta_{i,j,k}$ is the switching cost if the consumer chose product k at period $t-1$, and $\varepsilon_{i,j,t}$ is an exogenous Type 1 Extreme Value shock. Note that we do not consider price discrimination, as price does not vary by consumer.[9]

Throughout most of this paper, we illustrate our results using the pure logit case, letting $\delta_{i,j,t} = \delta_{j,t}$ and $\eta_{i,j,k} = \eta_{j,k}$. We extend to mixed logit demand in Section 3.4. Under logit, the utility in equation (1) implies that the probability of switching from product k to j is

$$s_{j|k,t} = P(a_t = j | a_{t-1} = k) = \frac{\exp(\alpha p_{j,t} + \delta_{j,t} + \eta_{j,k})}{\sum_{l \in \mathcal{J}} \exp(\alpha p_{l,t} + \delta_{l,t} + \eta_{l,k})}, \tag{2}$$

and the market share for product j at time t is

$$s_{j,t} = \sum_{k \in \mathcal{J}} s_{j|k,t} s_{k,t-1}. \tag{3}$$

The market share at time t in equation (3) depends on the market shares of all products at time $t-1$ and on the probabilities of switching from each product to product j.[10] We can interpret $s_{0,t-1}$, the share of consumers that chose the outside option last period, as the share of potential new customers.

[9]The switching data we have typically encountered in merger reviews often lacks price information, meaning it would be difficult to estimate a model with customer-specific prices. Should such data be available, allowing for price discrimination could be an interesting extension.

[10]A normalization of the utility function, typically on the outside option, is needed to estimate the parameters of the demand function. The goal of our paper, however, is to estimate the diversion ratio. Therefore, a normalization is not made explicitly.

2.2 The Diversion Ratio

For illustrative purposes, suppose that products A and B are owned by two separate single-product firms.[11] The products are thus synonymous with firms. Define the diversion ratio from firm A to firm B as

$$D_{A,B,t} = \frac{\partial s_{B,t}}{\partial p_{A,t}} \left(\left| \frac{\partial s_{A,t}}{\partial p_{A,t}} \right| \right)^{-1}. \tag{4}$$

This ratio measures the ability of firm A to shift sales to and from firm B by adjusting product A's price.

Suppose firms A and B are contemplating a merger. Without the merger, if firm A were to raise its price, that firm would consider any shift in sales to firm B a loss. After the merger, these sales are now a gain to the firm's merging partner, with a value reflected in firm B's margin. Thus, the diversion ratio is a key piece of information in assessing the incentives to raise prices after a merger.

Our logit model with switching costs implies that the diversion ratio from A to B is

$$D_{A,B,t} = \frac{\sum_{j \in \mathcal{J}} \left(s_{B|j,t} s_{A|j,t} \right) s_{j,t-1}}{\sum_{j \in \mathcal{J}} \left((1 - s_{A|j,t}) s_{A|j,t} \right) s_{j,t-1}}. \tag{5}$$

This ratio depends not only on the switching probabilities between firms A and B, but also on the market shares of all firms and on the switching probabilities from all other firms to firms A and B. Implicit in this formula is a focus on the impact of a change in the price of product A alone, holding fixed all other product attributes. Thus, this diversion does not measure feedback responses that might occur in certain fully specified supply models.

Due to state dependence, consumers are more likely to stay with their previous choice. Therefore, we expect $s_{j|j,t}$ to be much greater than $s_{j|k,t}$ for $k \neq j$, meaning that

$$s_{B|A,t} s_{A|A,t} s_{A,t-1}$$

[11] Our analysis extends to multi-product firms when switching data at the product level are available.

and

$$s_{A|B,t}s_{B|B,t}s_{B,t-1}$$

are likely to dominate the numerator of equation (5). Furthermore, these terms also tend to be larger when $s_{B|A,t}$ and $s_{A|B,t}$ are larger. Thus, as one would expect, the diversion ratio from A to B is high when consumers are frequently switching between A and B.

Rearranging our diversion ratio gives

$$D_{A,B,t} = \sum_{j\in\mathcal{J}} \frac{s_{B|j,t}}{1-s_{A|j,t}} \underbrace{\left(\frac{(1-s_{A|j,t})s_{A|j,t}s_{j,t-1}}{\sum_{j\in\mathcal{J}}(1-s_{A|j,t})s_{A|j,t}s_{j,t-1}} \right)}_{\equiv w_{j,A}} \tag{6}$$

$$= \sum_{j\in\mathcal{J}} \frac{s_{B|j,t}}{1-s_{A|j,t}} w_{j,A}, \tag{7}$$

where $w_{j,A}$ are weights that reflect what products consumers previously purchased. If there were no switching costs, the diversion ratio would simplify to that for a flat logit,

$$D^{\text{Share}}_{A,B,t} = \frac{s_{B,t}}{1-s_{A,t}}, \tag{8}$$

which is often referred to as "diversion according to share." Thus, in the spirit of Conlon and Mortimer (2021), estimating the diversion ratio for our model can be interpreted as calculating the weighted average of logit diversion ratios across each type of consumer. Here the types depend on what product a customer chose last period.

Adding switching costs to the logit allows for departures from diversion according to share by incorporating information on past purchases. However, the model still assumes that customer heterogeneity is effectively captured by the product chosen last period. Additional types of heterogeneity may be relevant in a given application. For instance, some consumers may have systematically different switching costs or preferences for product characteristics. For cases where these types of consumers can be identified in the data, we discuss a mixed logit extension in Section 3. If these consumers cannot be identified, a random coefficients model may be more appropriate. We give some examples with random coefficients preferences where our logit framework still provides useful information (and point out some instances where it does not) in Section 4.

3 Identification of the Diversion Ratio

Here we show how to combine the model with win/loss data from the merging firms in order to measure diversion ratios. Whether diversion is point- or set-identified depends on what additional data on market shares are available.

3.1 Win/Loss Data

Win/loss data from the merging firms track the number of customers switching in and out of purchasing from firms A and B. Specifically, the loss data of firm A report the number of customers who chose A last period, and what their choice this period was. That is, these data give the conditional probability of switching from A or from B to other firms,

$$P(a_t = j | a_{t-1} = A) \text{ and } P(a_t = j | a_{t-1} = B) \;\; \forall j \in \mathcal{J}. \tag{9}$$

On the other hand, the win data of firm A report the number of customers who chose A this period and what their previous choice was last period. That is,

$$P(a_{t-1} = j | a_t = A) \text{ and } P(a_{t-1} = j | a_t = B) \;\; \forall j \in \mathcal{J} \tag{10}$$

are observed from the win data of A and B. We call these the win data shares. The win data do not directly provide us with information on the conditional probability of switching from any product j to the merging firms, as they do not report information on the customers that did not switch to A or B this period. As a result,

$$P(a_t = A | a_{t-1} = j) \text{ and } P(a_t = B | a_{t-1} = j) \;\; \forall j \in \mathcal{J} \setminus \{A, B\} \tag{11}$$

are not observed.

Using these win/loss data, one can construct two common approximations of the diversion ratio from A to B,

$$D^{\text{Win}}_{A,B,t} = \frac{P(a_{t-1} = B|a_t = A)}{1 - P(a_{t-1} = A|a_t = A)} \tag{12}$$

$$D^{\text{Loss}}_{A,B,t} = \frac{P(a_t = B|a_{t-1} = A)}{1 - P(a_t = A|a_{t-1} = A)}, \tag{13}$$

where we refer to $D^{\text{Win}}_{A,B,t}$ as the conventional win ratio and $D^{\text{Loss}}_{A,B,t}$ as the conventional loss ratio. The win ratio in (12) captures the fraction of new consumers won by A that came from B; the loss ratio in (13) captures the fraction of consumers lost by A that went to B.

As can be seen by inspection, neither of the conventional measures are guaranteed to equal the diversion ratio given by equation (5). This echoes the finding of Chen and Schwartz (2016), who show that the loss ratio can overstate or understate the true diversion ratio in a linear Hotelling model. We discuss similar findings for our logit switching cost model in Section 4.

When applying the identification methods in this paper, we assume that the relevant win/loss data are generated by changes in relative prices. Whether this is true depends on the specific application. If the logit assumption holds, the diversion ratios with respect to changes in quality or switching costs are the same as that for price. In either case, the appropriate weighted average of consumer-level diversion ratios has weights given by (7). However, this equivalence does not hold when there are random coefficients. Random coefficients may be appropriate if, for instance, consumers have systematic differences in their demand for certain product characteristics or in their switching costs. The latter might be true if some customers have lower switching costs due to a change in their life circumstances, and the presence of such consumers may generate switching data that put more weight on these buyers relative to the logit case. We explore some examples of these alternative models and how they impact our simulation results in Section 4.

3.2 Point Identification of the Logit Diversion Ratio

The logit diversion ratio in equation (5) consists of the conditional winning probabilities at period t for firms A and B, $\{s_{B|j,t}, s_{A|j,t}\}$ $\forall j \in \mathcal{J}$, and the market shares of all products at period $t-1$. Assume that we have access to win/loss data for firms A and B in this period, plus market share data for all firms this period and the previous period. We show that the logit diversion ratios between A and B are then point identified.

The loss data from the merging firms report $P(a_t = A|a_{t-1} = B) = s_{A|B,t}$, $P(a_t = B|a_{t-1} = A) = s_{B|A,t}$, $P(a_t = A|a_{t-1} = A) = s_{A|A,t}$, and $P(a_t = B|a_{t-1} = B) = s_{B|B,t}$. However, these data do not report $s_{A|j,t}$ for $j \notin \{A, B\}$, $s_{B|j,t}$ for $j \notin \{A, B\}$, nor $s_{j,t-1}$ for all j. Nevertheless, the win data do provide some information about the winning probabilities, $s_{A|j,t}$ and $s_{B|j,t}$ for $j \notin \{A, B\}$. To see this, note that the win data give the win data shares as shown in (10). By Bayes rule, we know that the winning probability of A from j can be written as the product of the win data share and the ratio of present to past market shares:

$$P(a_t = A|a_{t-1} = j) = \underbrace{P(a_{t-1} = j|a_t = A)}_{\text{observed from the win data}} \frac{P(a_t = A)}{P(a_{t-1} = j)}. \tag{14}$$

Therefore, if aggregate market shares from two consecutive periods are available, we can point identify the winning probabilities.[12]

The data requirements for point identification are reasonably low, as they can be satisfied by having access to small sets of market share data from consecutive periods. There is no requirement that these share data be matched to prices, that they be voluminous enough, nor that they have variation enough to make demand estimation feasible. Antitrust authorities frequently collect market share data for the purpose of calculating the Herfindahl–Hirschman Index (HHI), as described in §5 of the Guidelines, so we believe that the data we suggest are obtainable in many cases. However, we also acknowledge that antitrust practitioners may not always

[12] In fact, this discussion shows that we require market shares for all products in period $t-1$ and market shares for only the merging parties in period t.

have share information. We address this situation in the next subsection.

3.3 Bounding the Logit Diversion Ratio

Assume now that we have access to only win/loss data for firms A and B in period t. We show that the diversion ratios between A and B are partially identified as worst case bounds (as in Manski (2003) and (2009)). Once bounds for the diversion ratio have been calculated, they can be plugged into the formulas for UPP or CMCR to generate bounds for those measures in turn.

Using Bayes rule, we know that

$$\begin{aligned} s_{A|j,t}s_{j,t-1} &= P(a_t = A|a_{t-1} = j)P(a_{t-1} = j) \\ &= P(a_{t-1} = j|a_t = A)\frac{P(a_t = A)}{P(a_{t-1} = j)}P(a_{t-1} = j) \\ &= P(a_{t-1} = j|a_t = A)P(a_t = A), \end{aligned}$$

so the logit diversion ratio can be simplified to

$$\begin{aligned} \hat{D}_{A,B,t} &= \frac{\sum_j \left(s_{B|j,t}s_{A|j,t}\right) s_{j,t-1}}{\sum_j \left((1 - s_{A|j,t})s_{A|j,t}\right) s_{j,t-1}} \\ &= \frac{\cancel{P(a_t = A)} \sum_j s_{B|j,t}P(a_{t-1} = j|a_t = A)}{\cancel{P(a_t = A)} \sum_j (1 - s_{A|j,t})P(a_{t-1} = j|a_t = A)} \\ &= \frac{\sum_j s_{B|j,t}P(a_{t-1} = j|a_t = A)}{\sum_j (1 - s_{A|j,t})P(a_{t-1} = j|a_t = A)}, \end{aligned}$$

which is a function of the winning probabilities of the merging firms and the win data shares of A. From the loss data of the merging firms, we observe $s_{A|A,t}$, $s_{A|B,t}$, $s_{B|A,t}$, and $s_{B|B,t}$, but not $s_{A|j,t}$ and $s_{B|j,t}$ for $j \notin \{A, B\}$. Since these unknown objects are probabilities, however, we know that they are bounded by zero and one.

Using the fact that $s_{A|j,t} \geq 0$ and $s_{B|j,t} \geq 0$ for all $j \notin \{A, B\}$, we can construct the lower bound of the diversion ratio:

$$\begin{aligned}\hat{D}_{A,B,t} &= \frac{\sum_j s_{B|j,t} P(a_{t-1}=j|a_t=A)}{\sum_j (1-s_{A|j,t}) P(a_{t-1}=j|a_t=A)} \\ &\geq \frac{\sum_{j\in\{A,B\}} s_{B|j,t} P(a_{t-1}=j|a_t=A)}{\sum_j (1-s_{A|j,t}) P(a_{t-1}=j|a_t=A)} \\ &\geq \frac{\sum_{j\in\{A,B\}} s_{B|j,t} P(a_{t-1}=j|a_t=A)}{1-\sum_{j\in\{A,B\}} s_{A|j,t} P(a_{t-1}=j|a_t=A)} := D^{\text{Lower}}_{A,B}. \end{aligned} \tag{15}$$

This lower bound can be calculated using only the win/loss data of the merging firms. Namely, it requires the losing probabilities of the merging firms, $s_{A|B,t}$ and $s_{B|A,t}$, the probabilities of remaining with A or B, $s_{A|A,t}$ and $s_{B|B,t}$, and the win data shares of firm A, $P(a_{t-1}=A|a_t=A)$ and $P(a_{t-1}=B|a_t=A)$.

The lower bound in (15) is likely to be informative about the diversion ratio if switching costs are high. In markets where switching costs are sizeable, $s_{A|j,t}$ and $s_{B|j,t}$ are likely to be small for $j \notin \{A,B\}$. Therefore, setting those probabilities equal to zero will not result in the loss of much information. If, however, switching costs are small, then the lower bound will be less informative. Naturally, the bound can be sharpened if one obtains additional win/loss data from other market participants besides the merging firms.

To construct an upper bound, we use the inequalities $s_{A|j,t} \leq 1$ and $s_{B|j,t} \leq 1$,

$$\begin{aligned}\hat{D}_{A,B} &= \frac{\sum_j s_{B|j,t} P(a_{t-1}=j|a_t=A)}{\sum_j (1-s_{A|j,t}) P(a_{t-1}=j|a_t=A)} \\ &\leq \frac{\sum_{j\in\{A,B\}} s_{B|j,t} P(a_{t-1}=j|a_t=A) + \sum_{j\notin\{A,B\}} P(a_{t-1}=j|a_t=A)}{\sum_j (1-s_{A|j,t}) P(a_{t-1}=j|a_t=A)} \\ &\leq \frac{\sum_{j\in\{A,B\}} s_{B|j,t} P(a_{t-1}=j|a_t=A) + \sum_{j\notin\{A,B\}} P(a_{t-1}=j|a_t=A)}{1-\sum_{j\in\{A,B\}} s_{A|j,t} P(a_{t-1}=j|a_t=A) - \sum_{j\notin\{A,B\}} P(a_{t-1}=j|a_t=A)} \\ &:= D^{\text{Worst Upper}}_{A,B}.\end{aligned}$$

We call this the “worst case” upper bound, because it approximates the winning probabilities for A and B with ones that may be much greater than their true values.

If switching costs are high, we could assume that $s_{A|j,t} \leq s_{A|A,t}$ and $s_{B|j,t} \leq s_{B|B,t}$ for $j \notin \{A, B\}$, and tighten the upper bound to

$$\hat{D}_{A,B} \leq \frac{s_{B|A,t} P(a_{t-1} = A|a_t = A) + s_{B|B,t}(1 - P(a_{t-1} = A|a_t = A))}{1 - s_{A|B,t} P(a_{t-1} = B|a_t = A) - s_{A|A,t}(1 - P(a_{t-1} = B|a_t = A))} \equiv D^{\text{Upper}}_{A,B}.$$

However, it is possible that the upper bound $D^{\text{Upper}}_{A,B}$ could exceed one. In that scenario, we would bound the diversion ratio from above by one.

Although, for the reasons we explained previously, the upper bound of the diversion ratio is likely to be loose, we also argue that the lower bound is likely to be tight when switching costs are high. As a result, the lower bound of the diversion ratio is potentially a useful tool for antitrust authorities, as it can indicate the minimum levels of UPP or CMCR one can expect from a merger. If the lower bounds of UPP or CMCR are far from zero, this would provide strong evidence that the merger is likely to be anticompetitive, unless sizeable efficiencies are also expected.[13]

3.4 Allowing for Flexible Substitution Patterns

Let $\mathcal{X}$ denote a finite set of observed consumer characteristics, such as income levels. These characteristics separate consumers into groups, each with a different "type." Assume that we have access to win/loss data from the merging firms that are also reported by consumer type. Then we can extend our results to allow for more flexible substitution patterns beyond the logit.

We build a mixed logit specification that is a discretized version of the random coefficients logit model with observed consumer characteristics. This specification has appeared elsewhere in the industrial organization literature (see, for example, Berry and Haile (2016)). The model allows switching costs to be consumer-type dependent, which is important in instances where customers vary in their propensity to switch according to meaningful patterns. For example, in the trial for the

[13] One rule of thumb is that a CMCR or UPP greater than what would be canceled out by a marginal cost efficiency of 5% is "large," under the assumption that efficiencies beyond that level are unlikely. See Carl Shapiro's remarks at the American Bar Association Section of Antitrust Law Fall Forum, November 18, 2010, available at https://www.justice.gov/atr/file/518246/download.

H&R Block/TaxAct merger, the expert testifying on behalf of the DOJ examined switching data from the Internal Revenue Service that categorized customers by the complexity of their tax returns into Simple, Intermediate, and Complex groups.[14] Such win/loss data organized by these customer types would fit well with our mixed logit framework.

We assume that the switching probability from k to j conditional on $X = x$ for each $x \in \mathcal{X}$, or $s^x_{j|k,t} = P(a_t = j|a_{t-1} = k, X = x)$ for $j, k \in \mathcal{J}$, follows the logit. Define $s^x_{k,t-1} = P(a_{t-1} = k|X = x)$ as the market share of k at $t-1$ among the type-x consumers. The market share of product j at t is

$$s_{j,t} = \sum_{x\in\mathcal{X}} \sum_{k\in\mathcal{J}} s^x_{j|k,t} s^x_{k,t-1} P(X = x), \tag{16}$$

which is a weighted sum of market shares of product j over consumer types $x \in \mathcal{X}$. The mixed logit diversion ratio is

$$\begin{aligned} \tilde{D}_{A,B} =& \frac{\sum_{x\in\mathcal{X}} \sum_{j\in\mathcal{J}} s^x_{B|j,t} s^x_{A|j,t} s^x_{j,t-1} P(X = x)}{\sum_{x\in\mathcal{X}} \sum_{j\in\mathcal{J}} (1 - s^x_{A|j,t}) s^x_{A|j,t} s^x_{j,t-1} P(X = x)} \\ =& \frac{\sum_{x\in\mathcal{X}} \sum_{j\in\mathcal{J}} s^x_{B|j,t} P(a_{t-1} = j|a_t = A, X = x) P(X = x|a_t = A)}{\sum_{x\in\mathcal{X}} \sum_{j\in\mathcal{J}} (1 - s^x_{A|j,t}) P(a_{t-1} = j|a_t = A, X = x) P(X = x|a_t = A)}, \end{aligned} \tag{17}$$

where $P(X = x|a_t = A)$ measures the fraction of customers purchasing A that are of type x and $P(a_{t-1} = j|a_t = A, X = x)$ is the win data share of A conditional on type x.[15] Due to switching costs, the diversion ratio (17) is dominated by the terms $s^x_{B|A,t} P(a_{t-1} = A|a_t = A, X = x)$ and $s^x_{B|B,t} P(a_{t-1} = B|a_t = A, X = x)$ for each $x \in \mathcal{X}$. Moreover, their contribution to the diversion ratio depends on $P(X = x|a_t = A)$, the importance of type-x consumers in firm A's customer base.

Since we assume the available win/loss data contain information on consumer characteristics, allowing for the calculation of probabilities by types, we can apply

[14] See the H&R Block/TaxACT Memorandum Opinion at page 34.

[15] We maintain the assumption that the price coefficient is constant. That is, we rule out any interaction of consumer characteristics and price. It would be possible to relax this assumption if one had rich enough switching data for demand estimation.

our identification results from the previous subsections. Thus, $P(X = x|a_t = A)$, $P(X = x|a_t = B)$, $s^x_{A|A,t}$, $s^x_{A|B,t}$, $s^x_{B|B,t}$, $s^x_{B|A,t}$, $P(a_{t-1} = j|a_t = A, X = x)$, and $P(a_{t-1} = j|a_t = B, X = x)$ for all $x \in \mathcal{X}$ and $j \in \mathcal{J}$ are identified from the win/loss data of the merging firms. Similar to the point identification results from Section 3.2, if market shares conditional on each $x \in \mathcal{X}$ over two consecutive periods are available, $s^x_{A|j,t}$ and $s^x_{B|j,t}$ for all $j \notin \{A, B\}$ are identified using Bayes rule. As a result, the diversion ratios between A and B are point identified.

If, instead, detailed market share data are not available, we can impose $s^x_{B|j,t} = s^x_{A|j,t} = 0$ for all $j \notin \{A, B\}$ and $x \in \mathcal{X}$ to construct a lower bound:

$$\tilde{D}^{\text{Lower}}_{A,B} = \frac{\sum_{x\in\mathcal{X}} \sum_{j\in\{A,B\}} s^x_{B|j,t} P(a_{t-1} = j|a_t = A, X = x)P(X = x|a_t = A)}{\sum_{x\in\mathcal{X}} \sum_{j\in\{A,B\}}(1 - s^x_{A|j,t})P(a_{t-1} = j|a_t = A, X = x)P(X = x|a_t = A)}.$$

Similarly, imposing $s^x_{B|j,t} = s^x_{B|B,t}$ and $s^x_{A|j,t} = s^x_{A|A,t}$ for $j \notin \{A, B\}$ allows us to construct an upper bound

$$\tilde{D}^{\text{Upper}}_{A,B} = \frac{\sum_{x\in\mathcal{X}} \begin{pmatrix} s^x_{B|A,t} P(a_{t-1} = A|a_t = A, X = x)+ \\ s_{B|B,t}(1 - P(a_{t-1} = A, |a_t = A, X = x)) \end{pmatrix} P(X = x|a_t = A)}{\sum_{x\in\mathcal{X}} \begin{pmatrix} 1 - s^x_{A|B,t} P(a_{t-1} = B|a_t = A, X = x)- \\ s^x_{A|A,t}(1 - P(a_{t-1} = B|a_t = A, X = x)) \end{pmatrix} P(X = x|a_t = A)}.$$

These results parallel those in Section 3.3 for the logit.

Although our method extends to models with observed heterogeneity, the key limitation is that consumers are assumed to have the same diversion ratio conditional on their past choice and observable characteristics. In Appendix A, we show that our lower bound works for a nested logit model in which the products of the merging firms are in the same nest. Furthermore, we assess the applicability of our bounds to diversion ratios derived from random coefficients models in Section 4.3 and in the Online Appendix.

4 Numerical Simulations

In this section, we illustrate numerically how our bounds relate to conventional measures of diversion (both according to share and according to the loss ratio) and to the true diversion ratio under the baseline logit switching model. In addition, we examine the performance of our bounds under a random coefficients logit model.[16]

4.1 Setup

Our numerical design resembles the substitution patterns of the Hotelling linear city model with four inside goods and one outside good, giving the set $j \in \{0, \ldots, 4\}$. This design allows for asymmetric switching costs in a systematic way, and allows us to focus on switching between adjacent pairs. Given the small number of products in the market, their linear arrangement is not cruicial to our results. Products are "located" in increasing order according to their index number. When a customer chose product 2 in the last period, we allow for the cost of switching to the left (product 1) to potentially differ from the cost of switching to the right (product 3).

The indirect utility for a product $j \in \{0, \ldots, J\}$ is given by

$$u_{i,j,t} = \begin{cases} \beta_0 + \sigma_\beta \beta_i - p_{j,t} + (1+\sigma v_i)\eta^P_{j,j} 1\{a_{i,t-1} = j\} + \eta^{N_L}_{j,j-1} 1\{a_{i,t-1} = j-1\} \\ \quad + \eta^{N_R}_{j,j+1} 1\{a_{i,t-1} = j+1\} + \epsilon_{i,j,t}, & 1 \le j \le J-1 \\ \beta_0 + \sigma_\beta \beta_i - p_{J,t} + (1+\sigma v_i)\eta^P_{J,J} 1\{a_{i,t-1} = J\} + \eta^{N_L}_{J,J-1} 1\{a_{i,t-1} = J-1\} \\ \quad + \eta^{N_R}_{J,1} 1\{a_{i,t-1} = 1\} + \epsilon_{i,J,t}, & j = J \\ \epsilon_{i,0,t}, & j = 0 \end{cases} \tag{18}$$

where $J = 4$, the $\epsilon_{i,j,t}$ are i.i.d. Type 1 extreme value shocks, v_i and β_i are i.i.d. standard normal, $\sigma \ge 0$, $\sigma_\beta \ge 0$, $p_{j,t}$ are prices, and

$$a_{i,t} = \arg \max_{j \in \{0,1,\ldots,J\}} u_{i,j,t}.$$

Prices $p_{j,t}$ are determined assuming single-product firms maximize myopic profits

[16] We leave results for the mixed logit model to the Online Appendix.

given marginal costs $c_{j,t}$.[17] The σv_i and $\sigma_\beta \beta_i$ allow for the possibility of random coefficients on switching costs and the constant term, respectively. In our baseline model, we use the flat logit ($\sigma = \sigma_\beta = 0$).[18] As a robustness check, we assess the performance of our bounds in the presence of random coefficients by varying σ and σ_β (and hence the degree of consumer heterogeneity).

The parameters η^P, η^{N_L}, and η^{N_R} allow for switching costs and switching cost heterogeneity. When all the η parameters are zero, the model collapses to the usual static logit model. When the η parameters are not zero, they allow the switching probability $P(a_{i,t} = j | a_{i,t-1} = k)$ to potentially differ from $P(a_{i,t} = j | a_{i,t-1} = l)$ for $k \neq l$. The η^P measures the baseline cost of switching away from any other previous product. The η^{N_L} adds a benefit to switching from a product to its left neighbor (e.g., it encourages switching from 2 to 3). Analogously, the η^{N_R} adds a benefit to switching from a product to its right neighbor (e.g., it encourages switching from 3 to 2).

We simulate 5000 consumers who are initially assigned to each product with equal likelihood, and we compute their switching probabilities after 19 periods of burn-in.[19] We present two sets of results. The first design sets $\eta^{N_L} = \eta^{N_R} = 2$, allowing for equal switching costs between the neighboring products; the second design sets $\eta^{N_L} = 0$ and $\eta^{N_R} = 2$, allowing for asymmetric switching patterns. The second design implies that, for example, the switching probability from product 3 to 2 is higher than that from product 2 to 3. We vary the degree of switching costs η^P and, for the random coefficients specifications, consumer heterogeneity σ in both designs.[20] Using the switching probabilities generated from the model, we compute

[17] We draw the marginal costs from a log-normal distribution. Specifically, we assume $\log(c_{j,t}) \sim N(0, 1/10)$.

[18] We set $\beta_0 = 2$ for the baseline model.

[19] Switching probabilities are generated based on new prices and the logit shocks drawn for each period.

[20] For both designs, shares of the included products are roughly equally split, but they change gradually as the parameters change. For example, $s_1 \approx 0.35$ and $s_J \approx 0.15$ when $\eta^{N_L} = 0$ and $\eta^{N_R} = 2$, whereas $s_1 \approx 0.2$ and $s_J \approx 0.2$ when $\eta^{N_L} = \eta^{N_R} = 2$ and η^P is small. In most designs, approximately 60% to 95% of consumers stay with the same product or switch to adjacent products. When $\eta^P > \eta^{N_L} + \eta^{N_R}$, staying with the same product is more likely than switching to adjacent products.

the true diversion ratio, the conventional loss ratio from equation (13), diversion by share from equation (8), and our proposed lower bound, worst upper bound, and upper bound.

4.2 Baseline Model

Figures 1 and 2 summarize the simulation results from firm 2's perspective. The figures show mean estimates out of 100 replications for each specification. In the figures, "conventional" denotes the loss ratio calculated from win/loss data, and "by share" denotes diversion calculated according to share. The lower and worst case upper bounds cover the true diversion ratio in all specifications. The "upper bound," however, may not be a proper upper bound for the diversion ratio when η^P is small, because the additional restrictions used to tighten the upper bound may fail when switching costs are low. The conventional ratio and diversion by share can over or underpredict the true diversion ratio.

Although the lower bound can be far from the truth when switching costs are low, the lower bound converges to the truth as the switching cost from η^P increases.[21] As we explained in Section 3, the lower bound is more likely to be informative when switching costs are high, as then the switching probabilities are more likely to be accurately approximated by zero. Our simulations show that the lower bound comes within 80% of the true diversion ratio from firms 2 to 3 when the probability of staying with the same product is 55%.[22] Since win/loss data are often collected in industries where one would suspect switching costs are high, our lower bound has the potential to be useful in a number of applications.[23] In Sections 1 and 2 of the Online Appendix, we show that these findings extend to measures of UPP and to the mixed logit specification, respectively.

[21]The switching costs in the simulations range from 0% to 100% of equilibrium prices (the simulated equilibrium prices range from 2.5 to 3), which are considered moderate with respect to the results in Arie and Grieco (2014).

[22]This occurs when $\eta^P = 2.8$, $\eta^{N_L} = 0$, and $\eta^{N_R} = 2$.

[23]For example, we expect switching costs to be high in products sold via long-term contracts, such as health insurance or wireless services. Merger cases in these industries have featured win/loss data, as in the Anthem/Cigna trial discussed in Section 5.

Figure 1: Diversion from Firm 2, $\eta^{N_L} = \eta^{N_R} = 2$

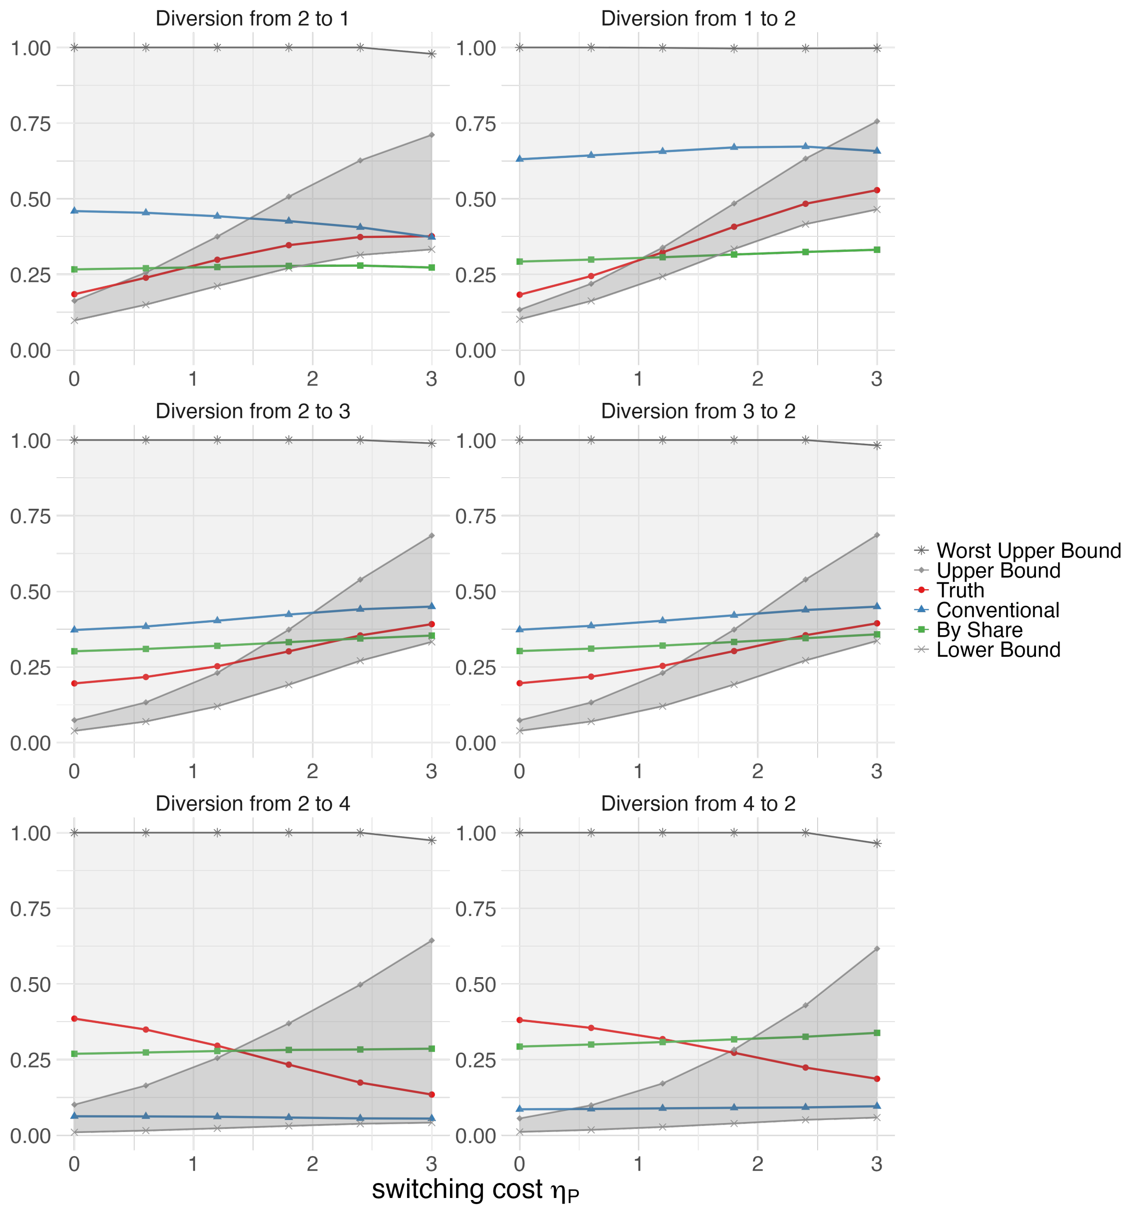

Figure 2: Diversion from Firm 2, $\eta^{N_L} = 0, \eta^{N_R} = 2$

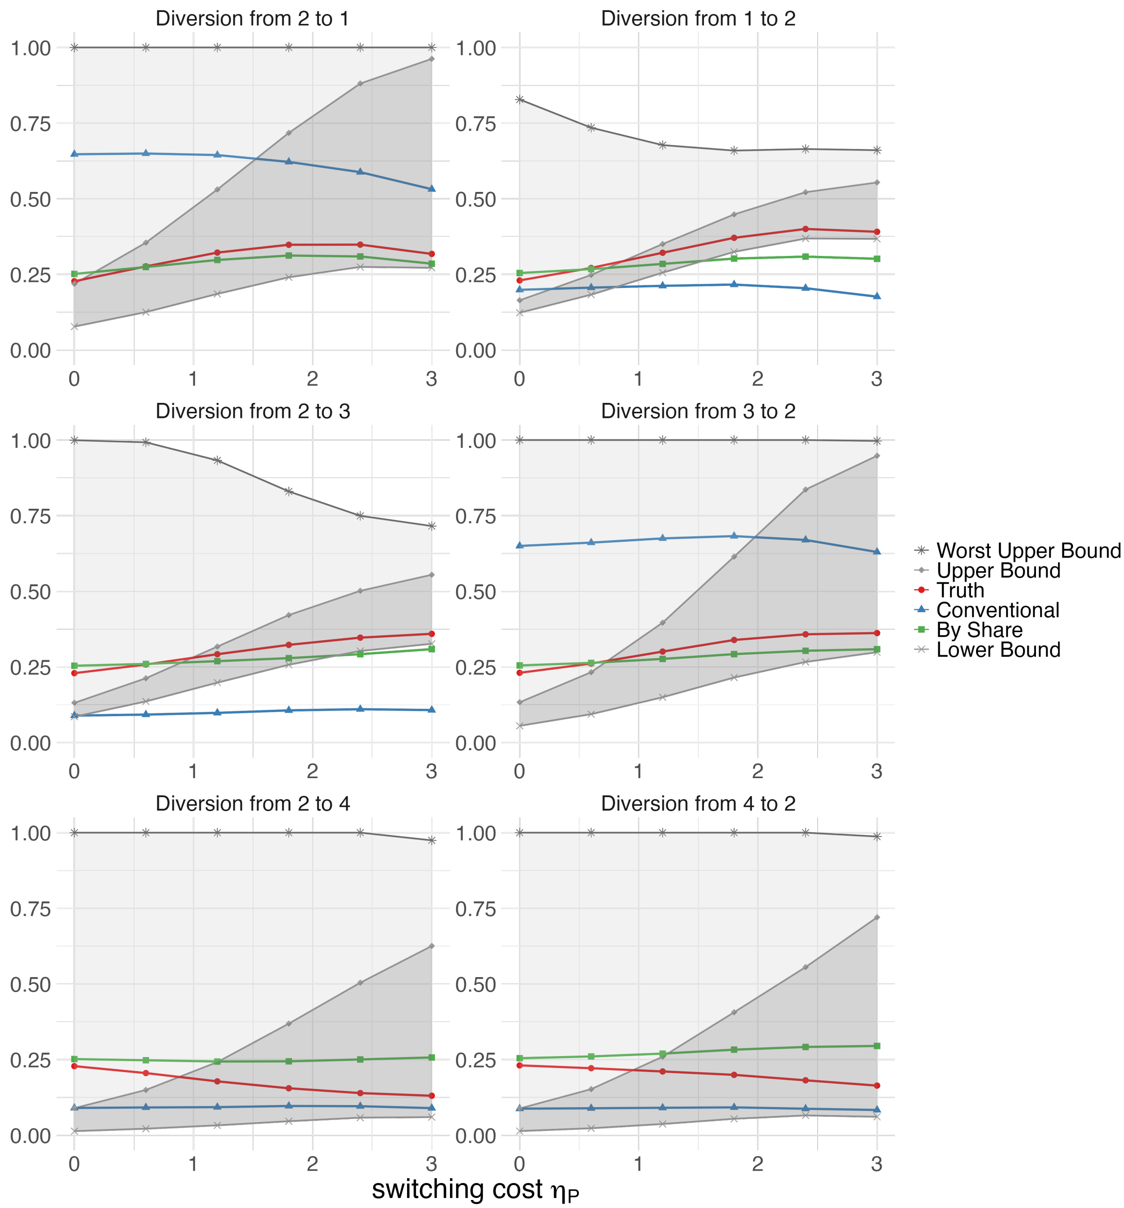

Our results for the upper bounds are more mixed. The worst case upper bound is quite loose in all instances, though it is not violated. The upper bound is violated when switching costs are low, in this instance when η^P falls below the range of 1 to 2.

Therefore, we consistently find that our bounds perform best when switching costs are high, though the lower bound may be more useful because it is never violated. Thus, when the lower bound is high, it is strong evidence that a merger may be problematic.

In Section 4 of the Online Appendix, we show that as η^P and η^{N_R} get larger, the conventional ratio becomes smaller than the true diversion ratio, while the lower bound approaches the true ratio. This means that the conventional diversion ratio may predict lower UPP than the truth when consumers exhibit strong state dependence overall and there is a tendency to switch from product 3 to product 2. Similarly, as η^P and η^{N_L} get larger, the conventional diversion ratio becomes larger than the true diversion ratio, while the lower bound still approaches the true ratio. Importantly, the true diversion ratios and the lower bounds yield similar results in both scenarios, while the conventional ratios offer opposing conclusions in terms of evaluating UPP.

4.3 Random Coefficients

Here we examine the performance of our bounds using switching probabilities generated by models with random coefficients on the constant and on switching costs, setting $\sigma_\beta = 0.5$ and varying σ in equation (18). We set $\beta_0 = 5$ for the random coefficients models. In this setup, our bounds are derived with a misspecified model, and we expect the bounds may be violated as σ, the variance of the random coefficient on switching costs, increases. We vary the magnitude of σ in equation (18) to assess the sensitivity of our bounds relative to the degree of unobserved consumer heterogeneity. We show the results from firm 2's perspective for succinct presentation.

The left panels of Figure 3 present the diversion ratios from firms 2 to 1 when $\sigma = 0, 1,$ and 3. Overall, our upper and lower bounds cover the true diversion ratios in almost all cases. The truth is slightly below the lower bound when $\sigma = 3$ and $\eta^P = 3$. Our lower bound tends to approach the truth from below as the switching cost η_P increases. The conventional ratio, on the other hand, overpredicts the actual diversion ratio.

Figure 3: Random Coefficients Model, $\eta^{N_L} = 0, \eta^{N_R} = 2$

(a) $\sigma = 0$

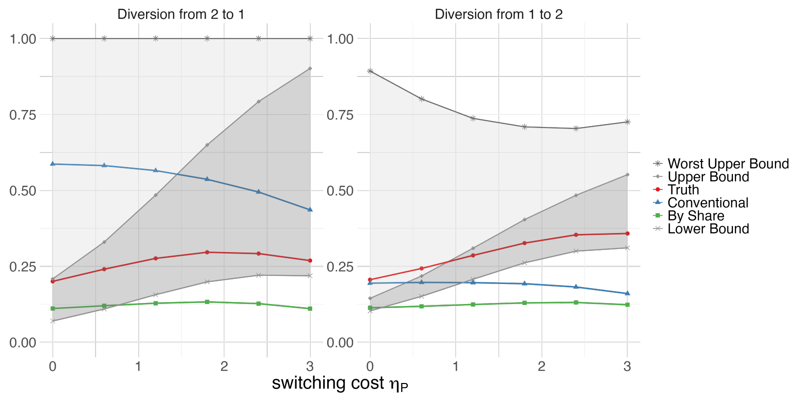

(b) $\sigma = 1$

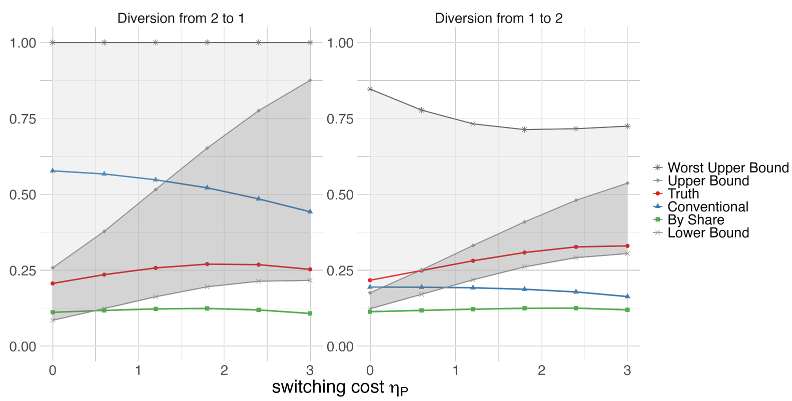

(c) $\sigma = 3$

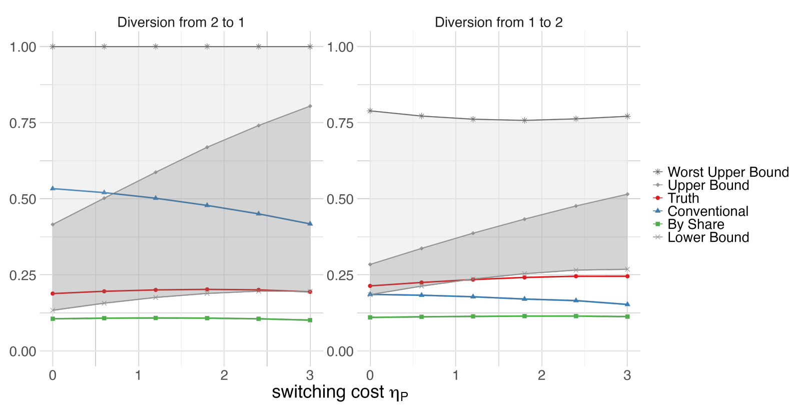

The right panels of Figure 3 present the diversion ratios from firms 1 to 2. By design, consumers are less likely to switch from firms 1 to 2 than from firms 2 to 1. Our lower bound is violated when $\sigma = 3$, particularly as η^P increases. When η_P is relatively high, consumers are less likely to switch on average. As σ increases, however, the variance of switching also increases, as more consumers are likely to experience large negative shocks to their switching costs. Since these switches are not directly generated by price changes, in this case our lower bound can be higher than the true diversion ratio. Diversion by share and the conventional ratio, on the other hand, tend to underpredict the truth.

In Section 3 of the Online Appendix, we examine the performance of our bounds with switching data generated from additional random coefficients models. When the random coefficients are on price or on exogenous product characteristics, our lower bounds perform well in the sense that they are not violated under the set of parameters we consider, even when the variance of the random coefficient is high. These exercises show that one should be cautious about applying our bounds when one believes that consumers have substantial heterogeneity in switching costs, as opposed to heterogeneity in preferences for price or quality. Heterogeneity in switching costs might be present, for example, if some consumers experience a sudden shift in their costs due to changes in their life circumstances, such as from moving residences or changing jobs. In practice, if the reasons for switching were observed, the bounds could be measured conditional on these reasons to improve the accuracy of the diversion ratio estimates.[24] Another strategy would be to focus on data collected soon after price changes, in the hope that switches driven by idiosyncratic factors would form a smaller percentage of the data.[25]

The Online Appendix also includes results where the true model is a random coefficients specification without switching costs. Our goal with these simulations

[24]The expert testifying on behalf of the DOJ in the H&R Block/TaxACT case used this strategy, by separating switching data for tax returns based on whether consumers had experienced a change in their return complexity, as categorized by the IRS. See the H&R Block/TaxACT Memorandum Opinion at page 34.

[25]The FCC used this strategy in the AT&T/T-Mobile case by examining switching data from periods where carriers introduced new wireless plan pricing. See the Economic Analysis Appendix to Staff Analysis and Findings in WT Docket No. 11-65 at pages C-10 to C-13.

is to assess the situation where consumer heterogeneity may be mistaken for state dependence. In this case our lower bounds contain the true diversions, but are typically less tight and hence less informative. Therefore, the lower bound, while still valid, tends to conservative. This result is intuitive, because our lower bound is constructed by assuming a number of switching probabilities are zero, when in reality, if the true model lacks switching costs, these probabilities are likely to be higher. Therefore, caution should be exercised if one is unsure if state dependence is present. Practitioners may have to rely on their knowledge of the industry to assess whether switching costs are relevant. However, if the lower bound happens to show a high diversion, this is evidence that substitution is high even according to a potentially conservative estimator.

5 Application to Anthem/Cigna

We now apply our bounds approach to a real world example, based on information from the Anthem/Cigna merger trial. In 2015, Anthem, the largest health insurer in the Blue Cross and Blue Shield Association, entered into an agreement to buy Cigna, another insurer, for over $54 billion. The DOJ, eleven states, and the District of Columbia filed suit to block the merger in 2016, eventually leading to a civil trial in U.S. District Court that began in November 2016 and concluded in January 2017. The products at issue were the provision of health insurance to large and national accounts employers, including both fully insured (FI) and administrative services only (ASO) accounts.[26]

As part of their case, the government presented economic expert testimony on the extent of consumer harm that was likely to result from the lessening in horizontal competition between Anthem and Cigna.[27] The economic expert testifying on behalf

[26] With FI products, the health insurance company carries the financial risk for healthcare claims. With ASO products, the employer assumes the financial risk, and instead hires the health insurance company to provide administrative services, such as billing and negotiating contracts with a healthcare network.

[27] The government expert testified four separate times, covering various aspects of the case. We focus on the first testimony, covering competition for national accounts customers, as this was

of the plaintiffs, David Dranove, argued that Anthem and Cigna were close competitors, in part by presenting diversion estimates calculated by share (as in equation (8)) alongside those calculated from loss ratios (as in equation (13)) and win ratios (as in equation (12)). Thus, we can compare our bounds estimates with these reported numbers. The plaintiffs' expert also argued that prices for FI and ASO plans sold to large employers would rise post-merger, and presented harm estimates based on (1) merger simulations assuming diversion according to share, (2) UPP assuming diversion according to share, and (3) UPP measuring diversion from win/loss data. In what follows, we compare UPP calculated with our lower bound measure to the harm estimates presented at trial.

There are some caveats to our analysis. First, we rely only on information disclosed publicly as part of the trial proceedings. We lack certain confidential data that the plaintiffs' expert had access to, and therefore this exercise should be viewed as illustrative rather than a precise measure of the merger's impact. Second, the expert used an auction model (without state dependence) for merger simulations, rather than a Bertrand logit.[28] It is not clear from the public trial materials whether the expert made any adjustments to the standard Bertrand UPP calculation in order to align with the auction specification he used for simulations. However, subsequent work published by the expert in Dranove et al. (2019) that calculates UPP for the Anthem/Cigna merger does not adjust for the auction form. Therefore, we do the same and rely on UPP derived under Bertrand competition.[29]

We use the diversion ratios by win/loss data as documented in Dranove (2016) to back out the winning probabilities and win data shares of the merging parties. See Appendix B for more details. Since market shares for two periods and the winning shares of products other than Anthem and Cigna are not reported publicly, we cannot point identify the diversion ratio. However, we can use the win/loss data of the merging parties to bound the diversion ratio.

ultimately the part of the case that was ruled on in the District Court opinion.

[28]See the discussion in the Anthem/Cigna District-level Memorandum Opinion at pages 66-67.

[29]We also do not account for any merger impacts on bargaining with upstream suppliers, which was another topic discussed extensively during the trial. Those vertical issues are examined in Sheu and Taragin (2021).

Table 1: Diversion Ratio Estimates

Expert Demonstrative	$D_{\text{Anthem,Cigna}}$	$D_{\text{Cigna,Anthem}}$
By Share	11%	43%
Loss Ratio	18%	61%
Win Ratio	35%	55%
Our Estimates		
Lower Bound	25%	55%
Upper Bound	82%	98%

Notes: The top panel of this table reports the diversion ratios from Anthem to Cigna and from Cigna to Anthem as presented by Dranove (2016) on slides 46 and 47 of Plaintiffs' Exhibit PDX005, estimated from the bar charts appearing on those slides. The bottom panel reports our bound estimates.

Table 1 reports the diversion ratios from the trial demonstrative alongside our estimates. Our lower bounds for $D_{\text{Anthem,Cigna}}$ and $D_{\text{Cigna,Anthem}}$ are roughly in the range of the diversion ratios implied by the win/loss data, but higher than the diversions implied by shares. Our results suggest that the estimated lower bounds of the diversion ratio resemble the win/loss diversions used by the plaintiffs' expert, meaning that the harm estimate presented at trial is likely a lower bound to the actual harm of the merger. Our estimated upper bounds, however, are not tight (being close to one) and therefore are in all likelihood less informative about the diversion ratios.

We now examine the implications of our diversion estimates for harm from the merger, as measured by UPP. The UPP of product A created by merging ownership with product B is defined as

$$UPP_{A,t} = (p_{B,t} - c_{B,t})D_{A,B,t}, \tag{19}$$

which is simply the margin of product B times the diversion ratio from A to B. As originally conceptualized, UPP was not meant to predict the size of the post-merger price increase, but rather whether an increase will occur. However, Miller et al. (2017) demonstrate that the value of UPP can be a reasonably accurate indicator of the magnitude of a potential price increase, particularly for log-concave demand

Table 2: Total Static Employer Harm (Millions of $)

Harm Measure	Baseline (ASO+FI)	With Claimed Variable Cost Savings (ASO + FI)	With Claimed Variable Cost Savings (ASO)
Merger Simulation	210	190	195
UPP (Diversion by Share)	390	300	310
UPP (Diversion by Win/Loss)	910	850	600
UPP (Lower Bound)	862	804	568
UPP (Upper Bound)	1940	1809	1278

Notes: The top panel of this table reports the harm calculations as presented by Dranove (2016) on slide 48 of Plaintiffs' Exhibit PDX005, estimated from the bar chart appearing on that slide. "Diversion by Share" assumes that diversion is proportional to market shares, while "Diversion by Win/Loss" assumes that diversion is measured according to the average of the loss and win ratios reported from win/loss data. The bottom panel reports our estimates assuming diversion is equal to our estimated bounds. Columns labeled "ASO+FI" cover both administrative services only and fully insured customers, while the column labeled "ASO" covers just administrative services only customers.

systems.[30]

We use our estimated bounds to calculate the harm from the merger. The top panel of Table 2 displays the harm from the merger estimated by the plaintiffs' expert using merger simulation and UPP analysis when the diversion ratio is approximated by market shares and by win/loss data. We impute the margin implied by the harm measures by assuming the expert used the average diversion ratio as an input to the UPP. We then multiply the estimated lower bound by the imputed margin to obtain a lower bound for the harm. We find that the harm of the merger is at least $862, $804, and $568 million per year under the baseline, with variable cost savings for ASO and FI, and with variable cost savings for ASO only options, respectively. The variable cost savings account for some efficiencies that Anthem claimed would be achieved by the merger. Our lower bound estimates are close to the win/loss UPP harm estimates from the expert's trial demonstrative, suggesting that the UPP estimates the expert reported are within the range of the true UPP harm. This

[30] As Miller et al. (2017) discuss, using UPP in this manner amounts to constructing a first-order approximation of the price increase (following Jaffe and Weyl (2013)), assuming that the pass-through matrix is equal to the identity matrix.

supports the argument that the merger was likely to harm consumers.

6 Conclusion

Although the importance of the diversion ratio is widely recognized in merger reviews, measuring diversion raises a number of difficulties. Antitrust authorities and practitioners frequently lack the data necessary to perform demand estimation, which may lead them to consider measures built from market shares or win/loss data. We show that these measures may perform poorly in the presence of state dependent preferences.

In response, we provide identification results for the diversion ratio that rely on the win/loss data of the merging firms. We have two main findings. First, the diversion ratio is point identified if the win/loss data of the merging parties and market share data for two consecutive periods are available. Second, the diversion ratio is partially identified if only the win/loss data of the merging parties are available. In the latter case, we derive lower and upper bounds that can be easily used in UPP and CMCR calculations. The lower bound is informative about the true diversion ratio if switching costs are high.

As a caveat to our analysis, our diversion estimates are derived from assuming consumers have logit demand with state dependence. Using numerical experiments, we show that our methods still perform well in the presence of random coefficient preferences for price and quality, but less well for cases with highly heterogeneous switching costs. An interesting area for future research would be to develop similar identification results for other models of consumer behavior.

References

Arie, Guy and Paul Grieco, "Who Pays for Switching Costs?," *Quantitative Marketing and Economics*, 2014, *12.4*, 379–419.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *65* (4), 841–890.

Berry, Steven T. and Philip A. Haile, "Identification in Differentiated Products Markets," *Annual Review of Economics*, 2016, *8*, 27–52.

Carlton, Dennis W. and Mark K. Israel, "Effects of the 2010 Horizontal Merger Guidelines on Merger Review: Based on Ten Years of Practical Experience," *Review of Industrial Organization*, 2021, *58*, 213–234.

Chen, Yongming and Marius Schwartz, "Churn vs. Diversion: An Illustrative Model," *Economica*, 2016, *83* (332), 564–583.

Conlon, Christopher and Julie Mortimer, "Empirical Properties of Diversion Ratios," *RAND Journal of Economics*, 2021, *52* (4), 693–726.

Dranove, David, "U.S., et al. v. Anthem, Inc, and Cigna Corp. Plaintiff's Exhibit PDX005," 2016. https://www.justice.gov/atr/page/file/914606/download.

_ , Dov Rothman, and David Toniatti, "Up or Down? The Price Effects of Mergers of Intermediaries," *Antitrust Law Journal*, 2019, *82* (2), 643–677.

Dube, Jean-Pierre, Gunter J. Hitcsch, and Peter E. Rossi, "Do Switching Costs Make Markets Less Competitive?," *Journal of Marketing Research*, 2009, *46*, 435–445.

Farrell, Joseph and Carl Shapiro, "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition," *B.E. Journal of Theoretical Economics*, 2010, *10*, 1–39.

Goldberg, Pinelopi Koujianou, "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 1995, *63* (4), 891–951.

Jaffe, Sonia and E. Glen Weyl, "The First Order Approach to Merger Analysis," *American Economic Journal: Microeconomics*, 2013, *5* (4), 188–218.

Manski, Charles F., *Partial Identification of Probability Distributions*, Springer

Science & Business Media, May 2003.

— , *Identification for Prediction and Decision*, Harvard University Press, June 2009.

Miller, Nathan H., Marc Remer, Conor Ryan, and Gloria Sheu, "Upward Pricing Pressure as a Predictor of Merger Price Effects," *International Journal of Industrial Organization*, 2017, *52* (4), 683–709.

Raval, Devesh and Ted Rosenbaum, "Why Do Previous Choices Matter for Hospital Demand? Decomposing Swithing Costs from Unobserved Preferences," *The Review of Economics and Statistics*, 2018, *100*, 906–915.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," *American Economic Review*, 2022, *112*, 578–615.

Sheu, Gloria and Charles Taragin, "Simulating Mergers in a Vertical Supply Chain with Bargaining," *The RAND Journal of Economics*, 2021, *52* (3), 596–632.

Werden, Gregory J., "A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products," *Journal of Industrial Economics*, 1996, *44* (4), 409–413.

Appendix

A Bounds for Nested Logit

Suppose product j is in nest g. Let $s_{j,t}(A)$ denote the probability of switching from product A to j, $\bar{s}_{j|g,t}(A)$ denote product j's within nest g probability, and $\bar{s}_{g,t}(A)$ be the nest g probability, where both $\bar{s}_{j|g,t}(A)$ and $\bar{s}_{g,t}(A)$ are conditional on being a consumer who had chosen A in period $t-1$. The probability $s_{j,t}(k)$ is

$$s_{j,t}(A) = \bar{s}_{j|g,t}(A)\bar{s}_{g,t}(A), \tag{A.1}$$

where

$$\bar{s}_{j|g,t}(A) = \frac{\exp((\delta_{j,t} - \eta_{j,A})/(1-\sigma))}{\sum_{j \in g} \exp((\delta_{j,t} - \eta_{j,A})/(1-\sigma))} \text{ and} \tag{A.2}$$

$$\bar{s}_{g,t}(A) = \frac{\left(\sum_{j \in g} \exp((\delta_{j,t} - \eta_{j,A})/(1-\sigma))\right)^{(1-\sigma)}}{\sum_{g \in \mathcal{G}} \left(\sum_{j \in g} \exp((\delta_{j,t} - \eta_{j,A})/(1-\sigma))\right)^{(1-\sigma)}}. \tag{A.3}$$

The derivatives of $s_{j,t}(A)$ with respect to the mean utility are

$$\frac{\partial s_{j,t}(A)}{\partial \delta_j} = \frac{1}{1-\sigma} s_{j,t}(A)(1 - \sigma \bar{s}_{j|g,t}(A) - (1-\sigma)s_{j,t}(A)), \tag{A.4}$$

$$\frac{\partial s_{j,t}(A)}{\partial \delta_k} = \frac{-1}{1-\sigma} s_{j,t}(A)(\sigma \bar{s}_{k|g,t}(A) + (1-\sigma)s_{k,t}(A)) \text{ when } j \text{ and } k \text{ are in nest } g, \text{ and} \tag{A.5}$$

$$\frac{\partial s_{j,t}(A)}{\partial \delta_k} = -s_{k,t}(A)s_{j,t}(A) \text{ when } j \text{ is in nest } g, \text{ but } k \text{ is not in } g. \tag{A.6}$$

Suppose products A and B are in nest g. The diversion ratio is

$$\begin{aligned} D_{A,B,t}(\sigma) &= \frac{\sum_{j \in \mathcal{J}} \frac{1}{1-\sigma} s_{A,t}(j)(\sigma \bar{s}_{B|g,t}(j) + (1-\sigma)s_{B,t}(j))s_{j,t-1}}{\sum_{j \in \mathcal{J}} \frac{1}{1-\sigma} s_{A,t}(j)(1 - \sigma \bar{s}_{A|g,t}(j) - (1-\sigma)s_{A,t}(j))s_{j,t-1}} \\ &= \frac{\sum_{j \in \mathcal{J}} s_{A,t}(j)(\sigma \bar{s}_{B|g,t}(j) + (1-\sigma)s_{B,t}(j))s_{j,t-1}}{\sum_{j \in \mathcal{J}} s_{A,t}(j)(1 - \sigma \bar{s}_{A|g,t}(j) - (1-\sigma)s_{A,t}(j))s_{j,t-1}}. \end{aligned}$$

It is easy to show that

$$D_{A,B,t}(0) \leq D_{A,B,t}(\sigma) \leq D_{A,B,t}(1), \tag{A.7}$$

where $D_{A,B,t}(0)$ is equivalent to equation (15) in the main text. Therefore, the same lower bound also applies to the nested logit diversion ratio, provided that products A and B are in the same nest.

B Anthen Cigna Win/Loss Data

We assume that the expert used the conventional method to derive diversion ratios based on win/loss data. Specifically, the diversion ratios implied by the loss data are

$$D^{\text{Loss}}_{\text{C,A}} = \frac{s_{A|C,t}}{1 - s_{C|C,t}} = 61\%$$
$$D^{\text{Loss}}_{\text{A,C}} = \frac{s_{C|A,t}}{1 - s_{A|A,t}} = 18\%$$

and the diversion ratios implied by the win data are

$$D^{\text{Win}}_{\text{C,A}} = \frac{P(a_{t-1} = A|a_t = C)}{1 - P(a_{t-1} = C|a_t = C)} = 55\%$$
$$D^{\text{Win}}_{\text{A,C}} = \frac{P(a_{t-1} = C|a_t = A)}{1 - P(a_{t-1} = A|a_t = A)} = 35\%.$$

The report does not disclose the number of Anthem and Cigna customers who chose not to switch. Therefore, we cannot back out the switching probabilities and the win data shares of the merging parties. We assume $s_{A|A,t} = s_{C|C,t} = p(a_{t-1} = A|a_t = A) = p(a_{t-1} = C|a_t = C) = \alpha = 0.1$, and vary α as robustness checks. Our results are robust for $\alpha \in [0, 0.5]$.