

Upward Pricing Pressure as a Predictor of Merger Price Effects

Nathan H. Miller*
Georgetown University

Marc Remer[†]
Swarthmore College

Conor Ryan[‡]
University of Minnesota

Gloria Sheu[§]
Department of Justice

January 9, 2017

Abstract

We use Monte Carlo experiments to evaluate whether “upward pricing pressure” (UPP) accurately predicts the price effects of mergers, motivated by the observation that UPP is a restricted form of the first order approximation derived in ? (?). Results indicate that UPP is quite accurate with standard log-concave demand systems, but understates price effects if demand exhibits greater convexity. Prediction error does not systematically exceed that of misspecified simulation models, nor is it much greater than that of correctly-specified models simulated with imprecise demand elasticities. The results also support that UPP provides accurate screens for anticompetitive mergers.

Keywords: upward pricing pressure; UPP; merger simulation; merger enforcement; Herfindahl-Hirschman Index; HHI; unilateral effects; antitrust
JEL classification: K21; L13; L41

*Georgetown University, McDonough School of Business, 37th and O Streets NW, Washington DC 20057. Email: nhm27@georgetown.edu.

[†]Swarthmore College, Department of Economics, 500 College Avenue, Swarthmore, PA, 19081. Email: mremer1@swarthmore.edu.

[‡]University of Minnesota, Department of Economics. Email: ryan0463@umn.edu.

[§]U.S. Department of Justice, Antitrust Division, Economic Analysis Group, 450 5th St. NW, Washington DC 20530. Email: gloria.sheu@usdoj.gov. 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.

1 Introduction

In a number of recent antitrust enforcement actions, the U.S. Department of Justice (DOJ) and the Federal Trade Commission (FTC) have alleged that mergers between producers of competing differentiated products would adversely affect unilateral pricing incentives.¹ This follows a decades-long trend that has both spurred on and been informed by academic research on how mergers affect prices (e.g., ? (?); ? (?); ? (?); ? (?); ? (?); ? (?); ? (?)). Continuing this evolution, the DOJ and the FTC updated its Horizontal Merger Guidelines in 2010, in part motivated by a desire to better align the document with economic theory and antitrust practice as they relate to markets with differentiated products (? (?)).

One point of emphasis in the 2010 Horizontal Merger Guidelines is that mergers between competitors create opportunity costs, which in turn place upward pricing pressure (or “UPP”) on the combining firms. This principle is easily derived from basic economic models, and the magnitude of the opportunity costs often can be quantified with information from only the merging parties. This combination of theoretical and practical simplicity make UPP a useful diagnostic tool. Referring to UPP as the value of diverted sales, the Guidelines state that “[t]he Agencies rely more on the value of diverted sales than on the level of the HHI for diagnosing unilateral price effects in markets with differentiated products.”² The FTC has employed UPP calculations to support arguments in court (FTC v. Sysco Corporation, et al.) and to justify enforcement decisions (Family Dollar/Dollar Tree).³

Although UPP has a direct relationship to firms’ pricing incentives, antitrust economists have been wary about using it as a prediction of price effects. UPP does not incorporate how the pass-through of costs to prices depends on the higher-order properties of the underlying demand system. Nor does it account for the possibility that non-merging competitors change prices as the market shifts to a new equilibrium. Two of the principal authors of the 2010 Horizontal Merger Guidelines, Joseph Farrell and Carl Shapiro, emphasize in their academic work that “UPP does not predict post-merger prices, but only predicts the sign of changes in price” (? (?)).⁴ Furthermore, ? (?) show that UPP must be scaled by an

¹E.g., U.S. v. H& R Block Inc., et al.; U.S. v. AT&T Inc., et al.; U.S. v. Bazaarvoice, Inc.; FTC v. Sysco Corporation, et al.; U.S. v. AB Electrolux, et al.

²See Section 6.1.

³The FTC Press Release is available online: https://www.ftc.gov/system/files/documents/public_statements/681781/150713dollartree-jdwstmt.pdf.

⁴See also ? (?), who writes that:

The value of diverted sales, taken alone, does not purport to quantify the magnitude of any post-merger price increase.... The value of diverted sales is a measure of the extra (opportunity) cost the merged firm bears in selling units of Product 1. Higher costs give the merged firm an incentive to raise the price of Product 1. But further analysis is needed to determine how that

appropriate measure of pass-through to provide a first order approximation (or “FOA”) to price effects. Obtaining estimates of the requisite pass-through information can be difficult even in advantageous empirical settings (e.g., ? (?)).

In this paper, we revisit whether UPP accurately predicts the magnitude of price effects. We begin with a theoretical discussion in which we develop that using an identity matrix to proxy for pass-through introduces only limited misspecification error in some standard settings. It follows that UPP may provide a reasonable approximation to the true price effects. We then explore this possibility using a large-scale Monte Carlo experiment that simulates the unilateral effects of mergers in markets with differentiated products. Results indicate that UPP is quite accurate with standard log-concave demand systems but understates price effects if demand exhibits greater convexity. To put the magnitude of prediction error in context, we compare the UPP results to merger simulations conducted with either a functional form misspecification or inaccurate structural parameters (with either problem merger simulation is inexact). We find that the prediction error of UPP does not systematically exceed the prediction error that occurs due to functional form misspecification in simulation models, nor is it much greater than the prediction error that arises in correctly specified simulation models that rely on inaccurate structural parameters.

The Monte Carlo experiments follow the data generating process developed in ? (?). We repeatedly draw randomized market shares and a single price-cost margin, and use these data to calibrate the parameters of a logit demand system. We then calibrate the less restrictive linear, almost ideal, and log-linear demand systems to match the elasticities of the logit model, sometimes incorporating a degree of measurement error in the elasticities. The analysis thus features two log-concave demand systems (linear and logit demand) alongside two demand systems that exhibit greater convexity (almost ideal and log-linear demand). These four demand systems are commonly employed in antitrust analyses (? (?); ? (?)), and also have been used in academic studies that examine the effect of demand curvature on the precision of counterfactual simulations (e.g., ? (?); ? (?)). This research design complements that of ? (?), which compares the accuracy of UPP and a structurally estimated merger simulation model within the specific context of the U.S. airline industry.

One important feature of the Monte Carlo experiments is that consumer substitution in the pre-merger equilibrium is proportional to market shares. This property of logit demand transmits to the other demand systems in calibration. We accept the feature in our experiments because it simplifies the process of obtaining sensible structural parameters.

cost increase translated into a price increase.

Further, the accuracy of UPP should hold in settings with different substitution patterns because UPP adjusts with the underlying demand derivatives. To confirm this conjecture, we conduct a robustness check in which we introduce an additional source of randomness into the demand cross-derivatives. The accuracy of UPP is maintained in these alternative Monte Carlo experiments, which suggests that the logit restriction in the baseline results does not drive the main results.

Our experiments also allow us to evaluate merger screens based on the Herfindahl-Hirschman Index (HHI). In the baseline Monte Carlo experiments, we find that the change in HHI (or “ ΔHHI ”) correlates strongly with merger price effects but that the level of HHI is less predictive. This result has been anticipated in the antitrust literature because there is a theoretical connection between ΔHHI and unilateral price effects in differentiated products settings under the logit restriction of substitution-by-share (e.g., ? (?)). Unlike UPP, however, the ΔHHI formula does not account for the underlying demand derivatives. Consistent with this, the correlation between ΔHHI and price effects is weaker in the alternative Monte Carlo experiments that introduce an additional source of randomness into the demand cross-derivatives. Overall, we characterize the results as indicating that screens based on ΔHHI can be useful if consumer substitution is roughly proportional to shares.

The analysis is subject to a number of caveats and limitations. As with much Monte Carlo research, the most serious of these pertain to external validity. First, we focus exclusively on pricing in differentiated products markets, although the UPP calculation itself can be generalized to other settings (e.g., ? (?)). Second, we impose the Nash-Bertrand equilibrium concept throughout the data generating process in order to focus the analysis on the “unilateral effects” of mergers. UPP is unlikely to perform as well for mergers that create coordinated effects. Lastly, we make a number specific assumptions about the demand systems and marginal cost functions that are necessitated by the Monte Carlo approach. We align these assumptions with those typically made in merger simulations commonly performed by or presented to antitrust agencies; one notable exclusion is that we do not examine the random coefficients logit model, which requires an entirely different data generating process. In light of these limitations, we do not seek to provide the most general results available and instead seek to establish certain relationships that advance the dialog on UPP and motivate future research.

The remainder of the paper proceeds as follows. Section 2 details the theoretical connection between UPP, first order approximation, and the price effects of mergers. Section 3 describes the Monte Carlo experimental design and provides summary statistics. Section 4 presents the results. There we plot the raw data, compare the prediction error that arises

from UPP with that from merger simulation, and also discuss the use of UPP and HHI-based measures as early-stage screens. Section 5 describes the alternative data generating process in which substitution is not proportional to share and provides results. Section 6 concludes with a summary and a discussion of the appropriate scope of application for UPP.

2 Theoretical Framework

2.1 Merger price effects and UPP

We examine the connection between different methods of merger price prediction within the context of Nash-Bertrand price competition between multi-product firms. Assume that each firm, i , produces a subset of products available to consumers, faces a twice-differentiable demand function, and maximizes the following profit function:

$$\pi_i = P_i^T Q_i(P) - C_i(Q_i(P))$$

where P_i is a vector of firm i 's prices, Q_i is a vector of firm i 's unit sales, P is a vector containing the prices of every product, and C_i is the cost function. The superscript T denotes the vector/matrix transpose. Profit-maximizing prices are characterized by first-order conditions:

$$f_i(P) \equiv - \left[\frac{\partial Q_i(P)^T}{\partial P_i} \right]^{-1} Q_i(P) - (P_i - MC_i) = 0, \quad (1)$$

where $MC_i = \partial C_i / \partial Q_i$ is a vector of firm i 's marginal costs. Now consider a merger between two firms j and k that, for simplicity, does not affect the cost functions. The post-merger first-order conditions are given by

$$h_i(P) \equiv f_i(P) + g_i(P) = 0 \quad \forall i \in I \quad (2)$$

where

$$g_j(P) = - \underbrace{\left(\frac{\partial Q_j(P)^T}{\partial P_j} \right)^{-1} \left(\frac{\partial Q_k(P)^T}{\partial P_j} \right)}_{\text{Matrix of Diversion from } j \text{ to } k} \underbrace{(P_k - MC_k)}_{\text{Markup of } k} \quad (3)$$

and $g_k(P)$ is defined analogously, while $g_i(P) = 0$ for all $i \neq j, k$. Prices that satisfy the post-merger first order conditions can be computed given sufficient information on the demand

system and marginal costs. Directly computing post-merger prices using this information is referred to as merger simulation and has been a main focus of research spanning more than two decades; numerous literature reviews summarize the topic (e.g., ? (?); ? (?); ? (?)).

The merger can be interpreted as creating an opportunity cost within the joined firm. Aggressive pricing from one merging partner creates forgone profits that otherwise would be earned by the other. The magnitude of these opportunity costs – given by the $g(P)$ function – depends multiplicatively on the customer diversion rates between the merging firms and their markups. Reinforcing this interpretation is the fact that both marginal costs and $g_i(P)$ are additively separable in the post-merger first order conditions. ? (?) refer to these opportunity costs as the UPP due to the merger. They propose UPP as an initial screen in merger investigations, on the basis that higher marginal costs tend to be associated with higher prices.

2.2 UPP as a price predictor

The equilibrium post-merger price effects in this model depend upon how pricing pressure is passed through to consumers. In merger simulation models, pass-through behavior is determined by the demand system, and there is existing research that explores how functional form restrictions on demand affect the accuracy of simulation (e.g., ? (?); ? (?)). A somewhat more general solution to calculating merger price effects is provided by ? (2013). There it is shown that first order approximation (FOA) to the price change is

$$\Delta P = - \left(\frac{\partial h(P)}{\partial P} \right)^{-1} \Bigg|_{P=P^0} g(P^0) \quad (4)$$

where P^0 is the vector of pre-merger prices. The FOA equals UPP pre-multiplied by the opposite inverse Jacobian of $h(P)$, which ? refer to as the merger pass-through matrix. By inspection, merger pass-through depends on the first and second derivatives of demand, but not higher order derivatives. ? (?) provide Monte Carlo evidence that FOA is an accurate predictor of true price effects provided that the pass-through is known.

With this foundation in place, it follows that UPP itself may provide a useful prediction of the price effect, insofar as the identity matrix can reasonably proxy for the merger pass-through matrix. We provide a simple numerical example to fix ideas. Consider three firms, each of which has a margin of 0.50 and a 30% market share (the outside good has a 10% share). Consumer behavior is given by the logit demand system. With a merger between

the first two firms, equation (4) becomes

$$\begin{bmatrix} 0.204 \\ 0.204 \\ 0.052 \end{bmatrix} = \begin{bmatrix} 0.771 & 0.180 & 0.297 \\ 0.180 & 0.771 & 0.297 \\ 0.122 & 0.122 & 0.776 \end{bmatrix} \begin{bmatrix} .214 \\ .214 \\ 0 \end{bmatrix} \quad (5)$$

Here the value of UPP (0.214) nearly equals the first order approximation (0.204) for the merging firms. This happens because the diagonal elements of the merger pass-through matrix are somewhat below one, while the off-diagonal elements are positive. Thus, using an identity matrix to proxy merger pass-through overstates some effects and understates others; the balance is a prediction close to the first order approximation. Further, again in this example, it is worth noting that both UPP and the first order approximation are close to the true price effects (0.190 for the merging firms).

This idea extends beyond the example provided. Countervailing biases arise provided that (i) the diagonal elements of the merger pass-through matrix are below unity, and (ii) prices are strategic complements so the off-diagonal elements are positive. In such settings, UPP may provide a reasonable approximation to the true price effects. The two demand systems we consider that exhibit log-concavity (linear and logit) always satisfy both conditions. Note that log-concavity is sufficient to ensure incomplete *cost* pass-through (e.g., ? (?)), but that cost pass-through differs somewhat from merger pass-through. Following ? (?), pre-merger cost pass-through equals $-(\partial f(P)/\partial P)^{-1}|_{P=P^0}$ and post-merger cost pass-through equals $-(\partial h(P)/\partial P)^{-1}|_{P=P^1}$ where P^1 is the vector of post-merger prices. Thus, merger pass-through is based on post-merger cost pass-through equations evaluated at pre-merger prices.

There are at least three specific scenarios in which the countervailing biases do not arise. First, if demand exhibits enough convexity, then the diagonal elements of the merger pass-through matrix can exceed one, and UPP should understate price effects. As we develop in the Monte Carlo experiments, this is typically the case with almost ideal and log-linear demand. Second, if prices are strategic substitutes then the off-diagonal elements of the merger pass-through matrix are likely to be negative, which would lead UPP to overstate price effects.⁵ This does not occur in the models we consider but can occur in other models, such as the random coefficients logit, if the consumers that switch in response to price changes are particularly price sensitive. We are not aware of empirical research that systematically

⁵The off-diagonal elements of the cost pass-through matrix are negative if and only if prices are strategic substitutes, but the relationship between merger pass-through and strategic substitutability is not as precise.

investigates the prevalence with which prices are strategic substitutes.

Third, suppose that the second merging firm has a marginal cost efficiency that is sufficiently large to make second element in the UPP vector negative. Then UPP should overstate the price effects on the first merging firm, and return a price prediction for the second firm that is too negative. To see this, consider an adjusted version of the numerical example above, in which the sign of the second element of the UPP matrix is flipped:

$$\begin{bmatrix} 0.126 \\ -0.126 \\ 0.000 \end{bmatrix} = \begin{bmatrix} 0.771 & 0.180 & 0.297 \\ 0.180 & 0.771 & 0.297 \\ 0.122 & 0.122 & 0.776 \end{bmatrix} \begin{bmatrix} .214 \\ -.214 \\ 0 \end{bmatrix} \quad (6)$$

The UPP and FOA equal 0.214 and 0.126 for the first firm, respectively, and equal -0.214 and -0.126 for the second firm. Thus, there are identifiable scenarios for which the logic of countervailing biases does not apply. In many applications it may be possible to evaluate whether cost pass-through exceeds unity, prices are strategic substitutes, or cost efficiencies generate downward pricing pressure. Such evaluations could inform priors about the existence of countervailing bias and the accuracy of UPP as a price predictor.

The appeal of UPP to antitrust authorities derives in part from its limited informational requirements. It can be calculated with diversion and markups for only the merging firms, and such information often becomes available during the course of merger investigations. By contrast, merger simulation models typically require a full set of demand elasticities, encompassing consumer responses to the prices of all firms in the model. FOA requires these demand elasticities along with pass-through. To the extent that UPP provides accurate price predictions, the importance of obtaining elasticities and pass-through would be diminished, and this motivates the Monte Carlo experiments developed below. We note that it is possible to conduct simulation using only information from the merging firms if one holds the prices of non-merging firms fixed, and we explore that method in our experiments as well.

2.3 Market shares and HHI

The economic theory outlined above demonstrates that diversion ratios and markups are directly related to unilateral price effects in markets with differentiated products. Except in a special case that we describe below, no such direct theoretical connection exists between unilateral price effects and market concentration. Nevertheless, the 2010 Guidelines maintain that the HHI is useful in informing competitive effects, at least in the broad sense of “identify[ing] some mergers unlikely to raise competitive concerns and some others for which

it is particularly important to examine whether other competitive factors confirm, reinforce, or counteract the potentially harmful effects of increased concentration.”⁶ We calculate HHI as the sum of squared market shares:

$$HHI = \sum_i s_i^2 \tag{7}$$

where s_i is between 0 and 100 and represents the market share of firm i . The change in HHI due to a merger between firms j and k requires only the merging parties’ market shares:

$$\Delta HHI = 2s_j s_k \tag{8}$$

The Monte Carlo experiments allow us to evaluate the accuracy of the HHI statistics as a screening device. Because HHI and ΔHHI are merger-specific statistics (unlike UPP which is firm-specific), we compare them to the average price change of the merging firms.

The direct theoretical connection between unilateral price effects and market concentration arises if consumer diversion is proportional to market share. Then diversion from product j to product k equals $s_k/(1 - s_j)$, and can be approximated by $s_k(1 + s_j)$ for small s_j . Diversion from k to j is analogous, meaning that the sum of the approximate diversion ratios is $s_j + s_k + 2s_j s_k$. This can be expressed, using equation (8), as $s_j + s_k + \Delta HHI$. These mathematics, due to ? (?), provide a theoretical connection between ΔHHI and price effects because diversion is one factor that enters the first order approximation of equation (4). Further, they provide a theoretical foundation for the many empirical studies that relate merger price effects to the predicted change in HHI (e.g., ? (?); ? (?)). Lastly, we note that the *level* of HHI often is influenced greatly by the shares of non-merging firms. While the strategic reactions of rivals affect post-merger equilibrium, typically they are of secondary importance for unilateral effects. It follows that there is not as tight a relationship between the level of the HHI, as opposed to the ΔHHI , and the extent of unilateral price effects in differentiated product markets.

⁶See the 2010 Horizontal Merger Guidelines, Section 5.3.

3 Design of the Monte Carlo Experiments

3.1 Data generation

We generate data that are consistent with the theoretical model outlined in the previous section. Each draw of simulated data is independent and characterizes the pre-merger equilibrium conditions of a single market. Together, the data cover a wide range of competitive conditions. The markets contain six firms that produce differentiated products at constant marginal cost. Firms compete in prices and equilibrium is Nash-Bertrand. All pre-merger prices are normalized to one, which results in price effects that are identical in levels and percentages. The specific data generating process is as follows:

1. Obtain market shares for six firms and an outside good, and the first firm's margin. The market shares are drawn from independent uniform distributions and then normalized to sum to one. That is, we draw $\hat{s}_i \sim U[0, 1]$ for $i = 0, \dots, 6$, and calculate the market share of firm i as $s_i = \hat{s}_i / \sum_{j=0}^6 \hat{s}_j$. The margin is the percent markup of price over marginal cost, $m = \frac{p-c}{p}$, and is drawn from a uniform distribution bounded between 0.20 and 0.80.
2. Calibrate the parameters of a logit demand system based on the margin and market shares, and calculate the demand elasticities that arise in the pre-merger equilibrium. This entails selecting demand parameters that rationalize the random data. The parameters are exactly identified given market shares, prices, and a single margin.
3. Calibrate linear, almost ideal, and log-linear demand systems based on the logit demand elasticities. The parameters of these systems are exactly identified given market shares, prices, and the logit elasticities.
4. Calculate UPP for a merger of the first and second firms in each market, based on the elasticities and margins for each draw of data. UPP is invariant to the demand system. Also simulate the merger under each of the demand systems.
5. Repeat steps (1) - (4) until 4,500 draws of data are obtained.

The algorithm generates 18,000 mergers to be examined, each defined by a draw of data and a demand system. We provide mathematical details in Appendix A. Our objective in drawing the data the way we do is to cover a wide portion of the parameter space, rather to mimic the distribution of market shares and margins across industries.

The data generating process imposes the restriction that diversion is proportional to share in the pre-merger equilibrium because the structural parameters in all the models are calibrated to be consistent with the elasticities that arise with logit demand. This reduces the dimensionality of the random data that must be drawn, enforces that the elasticities are identical across the different demand systems, and helps ensure that the structural parameters make economic sense. We have found that the alternative approach of drawing structural parameters directly (e.g., as in ? (?)) often produces substitution patterns that are difficult to rationalize, which in turn requires one to “filter out” the bad draws. Among the four demand systems, only with logit demand does diversion remain proportional to share away from the pre-merger equilibrium.

We see no reason that imposing diversion by share in the pre-merger equilibrium pre-disposes the Monte Carlo exercise to favor UPP because the logic of the countervailing biases developed in equations (4) and (5) is more general. To see this, note that the magnitude of UPP adjusts explicitly to settings with non-proportional diversion because the UPP formula in equation (3) incorporates the underlying demand derivatives explicitly. This is not the case for ΔHHI , however, which retains its theoretical connection to unilateral price effects only in the specific case of proportional diversion. Thus, our baseline Monte Carlo exercise can be interpreted as providing something of a best case scenario for ΔHHI . We return to these topics in Section 5, which summarizes the results obtained from an alternative experiment that incorporates non-proportional diversion.

The data generating process allows us to assess easily the accuracy of UPP in absolute terms. Suppose that the true demand system is logit. Then prediction error can be calculated as the difference between UPP and the logit merger simulation.

We are also interested in how this prediction error compares with results from alternative predictors. For our first comparison, we evaluate UPP against merger simulation conducted with a functional form misspecification. We note that the elasticities of each demand system are identical in the pre-merger equilibrium, for a given draw of data, so that differences in price effects arise solely due to functional form. Thus, for example, if the true demand system is logit then the prediction error of UPP can be compared against the prediction error that arises from simulations using almost ideal, linear, and log-linear demand. This mimics the position of an antitrust authority that knows margins and diversion but does not have pass-through or other information that informs demand curvature.

For the second comparison, we incorporate imprecision into the observed demand elasticities, and evaluate how the predictive accuracy of UPP and simulation degrade. Specifically, we add a uniformly distributed error to each product’s own-price elasticity of demand.

To do so in a manner that preserves the property that own-price elasticities are less than negative one, we define the observed own-price elasticities to be

$$\tilde{\epsilon}_{kk} = \epsilon_{kk} + \nu \quad \text{where} \quad \nu \sim U(-t(\epsilon_{kk} + 1), t(\epsilon_{kk} + 1)) \quad (9)$$

The support of the error is element-specific and depends on $t \in [0, 1]$. We examine three levels of error: $t = (0.2, 0.5, 0.8)$. We then scale each product’s cross-price elasticity according to the percent error of that products’ own-price elasticity, i.e. $\tilde{\epsilon}_{jk} = \epsilon_{jk} \frac{\tilde{\epsilon}_{kk}}{\epsilon_{kk}}$. This restriction eliminates economically unlikely scenarios in which a substitution away from a given product is exceeded by substitution to other products; when we generate data without such a restriction, even modest amounts of error often result in negative price predictions. An interesting implication of the restriction is that diversion is unaffected (margins and UPP are affected). This second comparison uses only the linear, almost-ideal, and log-linear demand systems because logit demand cannot accommodate changes in the elasticity matrix with a fixed set of market shares.

3.2 Summary statistics

Table 1 summarizes the empirical distributions of the data. The distribution of firm 1’s share is centered around 15 percent, which reflects that shares are allocated among six products and the outside good. The margin distribution is determined by the uniform draws with support over $(0.20, 0.80)$. The own-price elasticity of demand, which equals the inverse margin, has a distribution centered around two, and 80 percent of the own-price elasticities fall between 1.37 and 3.87. The diversion ratios have a distribution centered at 0.17, and 0.80 percent of diversion ratios fall between 0.04 and 0.29. The distribution of market shares, margins, elasticity, and diversion are nearly identical for the other firms due to the design of the data generation process. The median pre-merger and post-merger HHI are 1,562 and 1,931, respectively.⁷ The median Δ HHI is 317, and the median UPP is 0.07.

[Table 1 about here.]

The table reveals that the demand systems have very different pass-through properties. The diagonal elements of the merger pass-through matrix (“own merger pass-through”) are less than unity for logit and linear demand, consistent with log-concavity, but often exceed

⁷In calculating HHI, we use the shares of the six firms that strategically react to the merger and ignore the share of the outside good. This is equivalent to an assumption that the outside good is sold by an infinite number of atomistic firms, and leads to a conservative estimate of the HHI level.

unity with almost ideal and log-linear demand. The off-diagonal elements (“cross merger pass-through”) are positive with logit, almost ideal and linear demand, but negative with log-linear demand. The own and cross pass-through terms also exhibit a fair amount of variation within each demand system. A quick examination of the data shows that the countervailing biases of UPP exist for every draw with logit and linear demand.

The median merger price effects are 0.06, 0.11, 0.05, and 0.18 for the logit, almost ideal, linear, and log-linear demand systems, respectively. Because pre-merger prices are normalized to one, these statistics reflect both the median level change and median percentage change. The relative sizes of price increases across demand systems match Monte Carlo results in ? (?). Also notice that price increases tend to be larger for demand systems that exhibit greater own pass-through. This connection between pass-through and merger effects also is documented in ? (?) and makes theoretical sense given the first order approximation of equation (4). Dispersion of price effects within demand systems reflects the range of market conditions that arise from the data generating process. Figure 1 summarizes the merger price effects graphically; the distributions are approximately exponential and have long tails.

[Figure 1 about here.]

4 Results

4.1 Graphical analysis

We begin by plotting the data. Figure 2 depicts the accuracy of UPP in predicting post-merger price increases under each of the demand systems considered. Each dot represents the predicted and true changes in firm 1’s price for a given draw of data; its vertical position is the prediction of simulation and its horizontal position is the true price effect. Dots that fall along the 45-degree line represent exact predictions while dots that fall above (below) the line represent over (under) predictions. If the underlying demand is logit or linear, UPP appears to be quite accurate, albeit somewhat larger in magnitude than the true price effect. These systems are log-concave and the two biases developed in Section 2 are countervailing. UPP exceeds the actual price increases because using ones to proxy the diagonal pass-through elements (which amplifies predictions with incomplete pass-through) has a somewhat larger affect on results than suppressing the cross terms (which damps predictions with strategic complements). By contrast, UPP understates price increases with almost ideal and log-linear demand, again consistent with the theoretical discussion.

[Figure 2 about here.]

Figure 3 depicts the accuracy of merger simulation conducted with incorrect functional form assumptions. The scatter plots show data sorted by the underlying demand system: logit (column 1), almost ideal (column 2), linear (column 3), and log-linear (column 4), and by the merger simulation model: logit (row 1), almost ideal (row 2), linear (row 3), and log-linear (row 4). In many instances, the prediction error that arises due to functional form misspecification visibly exceeds the prediction error of UPP. This sensitivity of merger simulation to functional form assumptions is well known (e.g., ? (?); ? (?)) and, in antitrust settings, it is standard practice to generate predictions under multiple assumptions as a way to evaluate the scope of potential price changes.

[Figure 3 about here.]

Figure 4 depicts the accuracy of merger simulation conducted with imprecisely measured demand elasticities. As one might expect, predictions are centered around the true effects but prediction error increases as elasticities lose precision. Interestingly, predictions are relatively robust to imprecision in the elasticities with linear demand. The explanation for this begins with equation (4), which shows that price changes are driven by pass-through, diversion, and margins (the latter two of which enter through UPP). The elasticity error that we introduce does not affect diversion, and it affects margins equally in all demand systems. Thus, the differential effect of the elasticity error on prediction accuracy across demand systems arises due to the way that elasticity error changes pass-through. With linear demand, the support of pass-through is relatively limited (e.g., see Table 1) so there is less scope for elasticity errors to transmit into pass-through. Theoretical derivations of cost pass-through in each of the four demand systems are provided in ? (?).

[Figure 4 about here.]

4.2 Numerical analysis

Table 2 presents the median absolute prediction error (“MAPE”) of UPP when the true underlying demand system is logit, almost ideal, linear, or log-linear. UPP is quite accurate if demand is logit; the MAPE of 0.006 is roughly 10 percent of the median price effect. With the other demand systems, the MAPEs are somewhat larger and range from 38 to 60 percent of the median price effect. The table also shows the results of “partial simulations” in which we hold fixed the prices on non-merging firms. These simulations have the same informational

requirements as UPP, and they tend to be more accurate than UPP for each demand system. It would be inappropriate to conclude from the analysis that partial simulation dominates UPP, however, because it is assumed that partial simulation is conducted with the correct functional form of demand, and that substantially reduces the scope for prediction error.⁸

[Table 2 about here.]

Table 2 also provides the MAPEs that arise with misspecified merger simulations, to help put the prediction error of UPP in context. The table shows, for example, that a simulation based on logit demand has a MAPE of 0.049 if the true underlying demand system is linear. The results suggest that functional form misspecification tends to introduce more prediction error than UPP. Looking within each column, the MAPE of UPP is smaller than at least two of the three misspecified simulations. Focusing on the two log-concave demand functions, UPP is more accurate than linear simulation if the true underlying demand is logit, and UPP is nearly as accurate as logit simulation if the truth is linear.

Table 3 pushes this comparison farther, showing the frequency with which UPP provides a more accurate prediction than misspecified simulation. Among the twelve comparison groups the data generating process allows, UPP dominates misspecified simulation in all but two instances (the exceptions being logit simulation with linear demand and almost ideal simulation with log-linear demand). The results summarized in Tables 2 and 3 suggest that the accuracy of UPP can exceed the accuracy of misspecified merger simulation models.

[Table 3 about here.]

Table 4 reports MAPEs separately for mergers with price effects less than 10 percent (Panel A) and greater than 10 percent (Panel B), in order to examine whether the accuracy of UPP arises over only a limited support of the data. The accuracy of all predictors is better for mergers with smaller price effects, at least in absolute terms. More notable is that the accuracy of UPP relative to the other predictors is basically unchanged for mergers with small and large price effects, supporting the generality of the main results.

[Table 4 about here.]

We next consider merger simulation conducted with the correct demand system but imprecise demand elasticities. In this exercise, we recalculate UPP based on the imprecise

⁸Partial simulation is exact in the case of log-linear because prices are neither strategic complements nor substitutes with that demand system. Partial simulation conducted with the incorrect demand system yields results very similar to those in rows 3-6 in the table.

elasticities in order to help facilitate an “apples-to-apples” comparison. (UPP is affected through the margins but not diversion, which is unchanged by the elasticity errors.) Panel A of Table 5 summarizes the prediction errors that arise with merger simulation. The MAPEs increase significantly with the amount of imprecision, but remain small relative to the median merger price effects for each underlying demand system. Panel B shows that the MAPEs of UPP, by contrast, do not increase much with imprecision. Despite this robustness, the prediction error of UPP exceeds that of merger simulation even with a substantial amount of measurement error. Together with the prior results, this suggests that the relative value of UPP as a price predictor is diminished if a reasonable functional form of demand can be selected (e.g., based on pass-through information or other information), even if some of the structural demand parameters are not precisely estimated. By contrast, UPP has relatively greater value if there is uncertainty about the functional forms of demand.

[Table 5 about here.]

Table 5 also suggests that the accuracy of merger simulation may depend more on the accuracy of diversion than on the magnitudes of the underlying elasticities. This is relevant for empirical researchers because estimation error often may “cancel out” as parameters are converted to diversion. This is especially stark in the case of logit demand: The own-price derivative for product i is $\alpha s_i(1 - s_i)$, for a price coefficient α , and a cross-price derivative is $\alpha s_i s_k$. Thus the diversion ratio, $s_k/(1 - s_i)$, is free of parameters, meaning that estimation error in α does not affect diversion. Our results suggest that simulation results are somewhat robust to estimation error in that context. Of course, logit is a restrictive model for empirical research. If instead a random coefficients logit model is estimated, our results suggest that obtaining precise estimates of the nonlinear parameters (which drive diversion) may be more important than obtaining a precise estimate of the price coefficient. Further, if a bootstrap on the counterfactual prediction is too computationally intensive then checking the precision of implied diversion ratios may be a reasonable substitute.

4.3 Preliminary screens in merger analysis

4.3.1 Upward Pricing Pressure

The Monte Carlo experiments also allow us to assess the properties of UPP as a preliminary screen in merger analysis. The evidence shown thus far – most visibly in Figure 2 – demonstrates that UPP is strongly correlated with price effects, a property that is highly desirable in a screen. Indeed, UPP is almost perfectly correlated with price changes under logit and

linear demand, and highly correlated with price changes under almost ideal and log-linear demand. The correlation coefficients are 0.996, 0.955, 0.857 and 0.895, respectively. To extend the analysis, suppose that the objective of the antitrust authority is to block mergers that increase price more than 10 percent, and employs a screen in which it investigates if and only if UPP exceeds 10 percent. How well would the antitrust authority sort mergers?

Table 6 provides the frequency of the two possible errors: “false positives” and “false negatives.” We define false positives as benign mergers that are investigated, and false negatives as anticompetitive mergers that are not investigated. False positives during an initial screen may be acceptable because such mergers can be identified and cleared in the subsequent investigation. False negatives are more consequential because no such *ex post* correction is possible. It follows that it may be appropriate to place more weight on false negatives than false positives in this evaluation. The results are broken out by the true underlying demand system. If demand is log-concave (i.e., linear or logit) then false positives are much more likely than false negatives, consistent with UPP being an effective screen. With almost ideal and log-linear demand, false negatives exceed false positives. Thus, if substantial demand convexity is deemed a realistic feature of many industries then the threshold level used in the UPP screen should be revised downward. For instance, a UPP of two percent possibly could be used to screen out mergers with price effects under five percent, reducing the number of false negatives.

[Table 6 about here.]

4.3.2 Herfindahl-Hirschman Index

The 2010 Horizontal Merger Guidelines define a set of HHI levels and changes to help stratify mergers into those that are unlikely to pose a problem and those that warrant a close investigation. Five categories are defined as follows:

- (i) Post-merger $HHI > 2500$ and $\Delta HHI > 200$.
- (ii) Post-merger $HHI > 2500$ and $\Delta HHI \in (100, 200]$.
- (iii) Post-merger $HHI \in (1500, 2500]$ and $\Delta HHI > 100$.
- (iv) Post-merger $HHI \leq 1500$.
- (v) $\Delta HHI < 100$.

The Guidelines state that categories (i)-(iii) are likely to raise competitive concern and lead to further investigation, while mergers in categories (iv) and (v) are unlikely to create competitive problems. Many in the antitrust community view these latter categories as providing safe harbors, although this is not specifically endorsed in the Guidelines. Because (iv) and (v) are not mutually exclusive, a single merger may slot into both, and we conduct our analysis accordingly.

The Monte Carlo experiments allow us to evaluate these HHI thresholds. The necessary caveat is that the baseline data generating process imposes the restriction that diversion is proportional to share in the pre-merger equilibrium. This is precisely the scenario in which there is a theoretical relationship between ΔHHI and price effects. It follows that experiments provide the best case scenario for concentration-based screens. We revisit the performance of HHI with non-proportionate diversion in the next section.

Table 7 shows the fraction of mergers that result in price increases of at least 5% (Panel A) and 10% (Panel B), sorted by HHI category. Mergers in categories (i) and (iii) frequently produce substantial price increases. This is especially true of mergers in category (i), which generate price elevations above 5 percent in around 90 percent of the mergers with logit and linear demand, and in more than 95 percent of the mergers with almost ideal and log-linear demand. Perhaps more surprising, mergers in category (ii) appear relatively benign and never produce a 5 percent price increase with log-concave demand. The results for category (iv) indicate that a nontrivial minority of mergers in markets that are not deemed “moderately concentrated” nonetheless result in price increases of 5 percent or higher. By contrast, virtually no mergers in category (v) result in such price increases if demand is log-concave.

[Table 7 about here.]

These results suggest that ΔHHI is more directly connected to unilateral effects theory than the post-merger HHI given proportional diversion. Accordingly, we investigate whether the ΔHHI could be used effectively as a screen. The data indicate a strong correlation between ΔHHI and the price change. The correlation coefficients range between 0.519 and 0.846 depending on the demand system. Figure 5 plots post-merger prices against the ΔHHI to illustrate this correlation. Note that relationships shown are noisier than the UPP-price relationships shown previously. This is because UPP accounts for both diversion and markups (through their interaction), whereas ΔHHI is imperfectly related to diversion and does not account for the markup. Also, the distribution of ΔHHI 's is the same across all

four demand systems, so the results confirm that AIDS and log-linear demand produce larger price effects than linear and logit demand for a given ΔHHI .

[Figure 5 about here.]

Table 8 shows the fraction of mergers that result in prices increases of at least 5 percent (Panel A) and 10 percent (Panel B), sorted by ΔHHI . Results are provided for (i) change in HHI greater than 200, (ii) change in HHI between 100 and 200 and (iii) change in HHI less than 100. Most mergers with $\Delta\text{HHI}>200$ produce substantial price increases, as do a many mergers with $\Delta\text{HHI}\in(100,200)$. Some mergers with $\Delta\text{HHI}<100$ also cause substantial price increases if demand is almost ideal or log-linear, but virtually no mergers with $\Delta\text{HHI}<100$ result in substantial price increases if demand is log-concave. This corresponds the scenario (v) of the previous table.

[Table 8 about here.]

5 Alternative Data Generating Process

We also produce results using an alternative data generating process that does not impose the logit restriction that diversion is proportional to market shares. This introduces additional complications to the calibration process, and data generation is incompatible with logit demand. Nonetheless, it helps assess the extent to which the main results are driven by the logit restrictions. The alternative data generating process is as follows:

1. Randomly draw (i) market shares for six firms, s_i , and an outside good, (ii) diversion shares for each firm, \hat{s}_i , and the outside good, \hat{s}_o , and (iii) a calibration parameter \hat{m} drawn from a uniform distribution bounded between 0.2 and 0.8. Calculate HHI and ΔHHI based on the market shares. Normalize prices to one.
2. Calculate diversion to be proportional to diversion shares, unrelated to market share.
3. Define the own price elasticity of demand for each product as $\epsilon_{ii} = -\frac{1}{\hat{m}} \frac{\hat{s}_i (1 - \hat{s}_i)}{s_i (1 - \hat{s}_o)}$. The entire demand elasticity matrix is implied from the own price elasticities, the diversion matrix, and an assumption that the matrix of demand first derivatives is symmetric.
4. Calibrate linear, almost ideal, and log-linear demand systems based on the demand elasticities. The parameters of these systems are exactly identified given market shares, prices, and the elasticities.

5. Calculate UPP for a merger of the first and second firms in each market, based on the elasticities and margins for each draw of data. UPP is invariant to the demand system. Also, simulate the effect of the merger under each of the demand systems.
6. Repeat steps (1) - (5) until 4,500 draws of data are obtained.

The algorithm generates 13,500 mergers to be examined, each defined by a draw of data and a demand system. Appendix B provides mathematical details on steps 2 and 3. In order to maintain symmetry in the distribution of margins and elasticities across firms, we cannot draw margins directly. Instead, we throw out and replace any markets that generate price-cost margins greater than one or less than zero. The process is then able to generate sensible demand parameters for the almost ideal, linear, and log-linear demand system.⁹

The empirical distributions in the alternative data are similar to those of the baseline data, with a few differences. The median diversion rates are 0.17 in the alternative data, just as in the baseline data. The median difference between diversion and diversion proportional to share is -0.03 , with the 10th and 90th percentiles being -0.18 and 0.13 . The median elasticity is 2.03 in the baseline data and 2.92 in the alternative data. The median price effects are somewhat smaller with the alternative data, at 0.06, 0.03, and 0.10 for almost ideal, linear, and log-linear demand, respectively.

Table 9 provides the MAPEs for UPP and the misspecified merger simulations, and none of the predictors are noticeably less accurate than with baseline data. Indeed, the MAPEs tend to be somewhat smaller with non-proportional diversion, though the improved accuracy is due to the fact that the merger price effects are slightly smaller on average with the alternative data. As an additional check, in Figure 6 we plot the prediction error of UPP against the difference between the actual non-proportional diversion and the hypothetical diversion that would arise under proportional diversion. The graphics show visually that prediction error does not increase for markets with highly non-proportional diversion. Together, these analyses confirm that the main results on the accuracy of UPP are not driven by the logit assumption employed in the baseline data generating process.

[Table 9 about here.]

[Figure 6 about here.]

⁹A small number of draws still cannot be rationalized with one of the three demand systems. This arises, for instance, if a firm has both an unusually small market share and an unusually high price-cost margin. We replace these to obtain the 4,500 draws.

Finally, we find that HHI is somewhat less accurate as a screen if diversion is not proportional to share, as shown in Table 10. In the alternative data, the screen for markets with a change in HHI greater than 200 is less accurate by more than 10 percentage points in nearly every case than the baseline data in Table 8. This should be expected as the only explicit theoretical connection between HHI and unilateral price effects is when diversion is proportional to share. The degradation is not dramatic because the departure from diversion by share is not too great, which is partially a result of the data generation process. If a firm has a large market share but an especially small diversion share, the implied margin may be greater than one and left out of our analysis. Thus, firms with large market share tend to have large diversion as well. We do find that the correlation between change in HHI and price effect is related to how well actual diversion approximates diversion by share. For example, when demand is linear, the correlation coefficient is 0.75 when diversion is within two percentage points of diversion by share and 0.68 otherwise. Appendix B elaborates on the implications of the calibration process on the empirical distribution of diversion.

[Table 10 about here.]

6 Conclusion

This research evaluates the accuracy of UPP in predicting post-merger price changes, using a large-scale Monte Carlo experiment. The results are supportive overall: we find that UPP is quite accurate with standard log-concave demand systems but understates price effects if demand exhibits greater convexity. Prediction error does not systematically exceed that of misspecified simulation models, nor is it much greater than that of correctly-specified models simulated with imprecise demand elasticities. We provide a theoretical basis for these results by observing that UPP is a restricted version of the first order approximation derived in ? (?). We conclude that UPP has greater utility than is currently recognized.

That UPP often outperforms simulation models in our Monte Carlo experiments raises a question about the appropriate scope of application. In our view, the value of UPP as a price predictor is greatest in merger investigations and similar policy endeavors, due to its expediency and simplicity. The academic literature provides an array of methodologies that are capable of both limiting functional-form misspecification in simulation models and reducing standard errors in structural estimation. These methodologies (which we do not examine in the Monte Carlo experiments) may well allow simulation to produce more robust and accurate predictions than are available from UPP. Thus, we are skeptical that our results

have significant bearing on empirical industrial economics. By contrast, because state-of-art academic methodologies often may be too time-consuming to be used within the constraints of merger investigations, our results are immediately relevant for antitrust practice.

References

Appendix

A Mathematical Details of the Calibration Process

We provide mathematical details on the calibration process in this appendix. To distinguish the notation from that of Section 2, we move to lower cases and let, for example, s_i and p_i be the market share and price of firm i 's product, respectively.¹⁰ Recall that in the data generating process we randomly assign market shares among the six single-product firms and the outside good, draw the price-cost margin of the first firm's product from a uniform distribution with support over $(0.2, 0.8)$, and normalize all prices to unity. The calibration process then obtains parameters for the logit, almost ideal, linear and log-linear demand systems that reproduce these draws of data.

Calibration starts with multinomial logit demand, the basic workhorse model of the discrete choice literature. The system is defined by the share equation

$$s_i = \frac{e^{(\delta_i - \alpha p_i)}}{1 + \sum_{j=1}^N e^{(\delta_j - \alpha p_j)}}. \quad (\text{A.1})$$

The parameters to be calibrated include the price coefficient α and the product-specific quality terms δ_i . We recover the price coefficient by combining the data with the first order conditions of the first firm. Under the assumption of Nash-Bertrand competition this yields:

$$\alpha = \frac{1}{m_1 p_1 (1 - s_1)} \quad (\text{A.2})$$

where m_1 is the price-cost margin of firm 1. We then identify the quality terms that reproduce the market shares:

$$\delta_i = \log(s_i) - \log(s_0) + \alpha p_i, \quad (\text{A.3})$$

for $i = 1 \dots N$. We follow convention with the normalization $\delta_0 = 0$. Occasionally, a set of randomly-drawn data cannot be rationalized with logit demand and we replace it with a set that can be rationalized. This tends to occur when the first firm has both an unusually small market share and an unusually high price-cost margin.

The logit demand system sometimes is criticized for its inflexible demand elasticities. Here, the restrictions on substitution are advantageous and allow us to obtain a full matrix of elasticities with a tractable amount of randomly drawn data. The derivatives of demand

¹⁰We define market share $s_i = q_i / \sum_{j=1}^N q_j$, where q_i represents unit sales.

with respect to prices, as is well known, take the form

$$\frac{\partial q_i}{\partial p_j} = \begin{cases} \alpha s_i(1 - s_i) & \text{if } i = j \\ -\alpha s_i s_j & \text{if } i \neq j \end{cases} \quad (\text{A.4})$$

We use the logit derivatives to calibrate the more flexible almost ideal, linear and log-linear demand systems. This ensures that each demand system has the same first order properties in the pre-merger equilibrium, for a given draw of data.

The AIDS is written in terms of expenditure shares instead of quantity shares ? (?). The expenditure share of product i takes the form

$$w_i = \alpha_i + \sum_{j=0}^N \gamma_{ij} \log p_j + \beta_i \log(x/P) \quad (\text{A.5})$$

where x is total expenditure and P is a price index. We incorporate the outside good as product $i = 0$ and normalize its price to one; this reduces to N^2 the number of price coefficients in the system that must be identified (i.e., γ_{ij} for $i, j \neq 0$). We further set $\beta_i = 0$ for all i , a restriction that imposes in income elasticity of unity. Under this restriction, total expenditures are given by

$$\log(x) = (\tilde{\alpha} + u\tilde{\beta}) + \sum_{k=1}^N \alpha_k \log(p_k) + \frac{1}{2} \sum_{k=1}^N \sum_{j=1}^N \gamma_{kj} \log(p_k) \log(p_j) \quad (\text{A.6})$$

for some utility u . We identify the sum $\tilde{\alpha} + u\tilde{\beta}$ rather than $\tilde{\alpha}$, u and $\tilde{\beta}$ individually.¹¹ Given this structure, product i 's unit sales are given by $q_i = xw_i/p_i$ and the first derivatives of demand take the form

$$\frac{\partial q_i}{\partial p_j} = \begin{cases} \frac{x}{p_i^2}(\gamma_{ii} - w_i + w_i^2) & \text{if } i = j \\ \frac{x}{p_i p_j}(\gamma_{ij} + w_i w_j) & \text{if } i \neq j \end{cases} \quad (\text{A.7})$$

The calibration process for the AIDS then takes the following four steps:

1. Calculate x and w_i from the randomly drawn data on market shares, using a market size of one to translate market shares into quantities.
2. Recover the price coefficients γ_{ij} for $i, j \neq 0$ that equate the AIDS derivatives given in

¹¹The price index P is defined implicitly by equation (A.6) as the combination of prices that obtains utility u given expenditure x . A formulation is provided in ? (?).

equation (A.7) and the logit derivatives given in equation (A.4). Symmetry is satisfied because consumer substitution is proportional to share in the logit model. The outside good price coefficients, γ_{i0} and γ_{0i} for all i , are not identified and do not affect outcomes under the normalization the $p_0 = 1$. Nonetheless, they can be conceptualized as taking values such that the adding up restrictions $\sum_{i=0}^N \gamma_{ij} = 0$ hold for all j .

3. Recover the expenditure share intercepts α_i from equation (A.5), leveraging the normalization that $\beta_i = 0$. The outside good intercept α_0 is not identified and does not affect outcomes, but can be conceptualized as taking a value such that the adding up restriction $\sum_{i=0}^N \alpha_i = 1$ holds.
4. Recover the composite term $(\tilde{\alpha} + u\tilde{\beta})$ from equation (A.6).

This process creates an AIDS that, for any given set of data, has quantities and elasticities that are identical in the pre-merger equilibrium to those that arise under logit demand. The system possesses all the desirable properties defined in ? (?). Our approach to calibration differs from ? (?), which does not model the price index as a function of the parameters, and from ? (?), which assumes total expenditures are fixed.

We turn now to the linear and log-linear demand systems. Linear demand takes the form

$$q_i = \alpha_i + \sum_j \beta_{ij} p_j \quad (\text{A.8})$$

The parameters to be calibrated include the firm specific intercepts α_i and the price coefficients β_{ij} . We recover the price coefficients directly from the logit derivatives in equation (A.4). We then recover the intercepts to equate the implied quantities in equation (A.8) with the randomly drawn market shares, again using a market size of one. Of similar form is the log-linear demand system:

$$\log(q_i) = \gamma_i + \sum_j \epsilon_{ij} \log p_j \quad (\text{A.9})$$

where the parameters to be calibrated are the intercepts γ_i and the price coefficients ϵ_{ij} . Again we recover the price coefficients from the logit derivatives (converting first the derivatives into elasticities). We then recover the intercepts to equate the implied quantities with the market share data. This process creates linear and log-linear demand systems that, for any given set of data, has quantities and elasticities that are identical to those of the calibrated logit and almost ideal demand systems, in the pre-merger equilibrium.

B Mathematical Details with Alternative Diversion

This appendix section contains mathematical details on the data generation process when diversion is not proportional to market share. Let s_i and p_i be the market share and price of firm i 's product, respectively. Let s_o be the share of the outside good. Just as in the calibration process when diversion is by share, we randomly assign market shares among the six single-product firms and the outside good, normalize all prices to unity, but do not directly draw a price-cost margin. We randomly draw a diversion matrix and an additional calibration parameter \hat{m} , which can be used to generate the matrix of first derivatives of demand under the assumption of symmetry. The parameters for the almost ideal, linear and log-linear demand systems are then obtained to reproduce these draws of data using the methodology described Appendix A.

We generate a diversion matrix by randomly assigning “diversion shares” to each firms’ product, \hat{s}_i , and the outside good, \hat{s}_o . The diversion matrix is calculated according to these diversion shares:

$$\hat{d}_{ij} = \frac{\hat{s}_j}{1 - \hat{s}_i} \quad (\text{B.1})$$

where \hat{d}_{ij} represents the diversion from firm i 's product to firm j 's product. This calculation is intentionally similar to the calculation of diversion by share, in which case $d_{ij} = \frac{s_j}{1-s_i}$. By calculating diversion in this way, we randomly generate a diversion matrix which is guaranteed to sum to no greater than one across each row and imply a symmetric matrix of first derivatives of demand, properties which make the calibration of each demand system tractable.

One challenge of generating random diversion is obtaining markets with reasonable price-cost margins. In the calibration process for diversion by share, we randomly generate one margin, and the other firms’ products’ margins are implied by the diversion matrix. Under the assumption of Nash-Bertrand competition, the price-cost margin is related to the own-derivative of demand by

$$\frac{\partial q_1}{\partial p_1} = -\frac{1}{m_1} \frac{s_1}{p_1} \quad (\text{B.2})$$

Under the assumption of symmetric derivatives of demand,

$$\frac{\partial q_i}{\partial p_i} = \frac{\partial q_1}{\partial p_1} * \frac{\partial q_i / \partial p_1}{\partial q_1 / \partial p_1} * \frac{\partial q_i / \partial p_i}{\partial q_1 / \partial p_i} = \frac{\partial q_1}{\partial p_1} * \frac{\hat{d}_{1i}}{\hat{d}_{i1}} \quad (\text{B.3})$$

When diversion is by share and prices are normalized to one, this results in,

$$\frac{\partial q_i}{\partial p_i} = -\frac{s_1}{m_1} \frac{s_i}{1-s_1} \frac{1-s_i}{s_1} \quad (\text{B.4})$$

$$\begin{aligned} \frac{\partial q_i}{\partial p_i} &= -\frac{s_i}{m_1} \frac{1-s_i}{1-s_1} \\ m_i &= m_1 \frac{1-s_1}{1-s_i} \end{aligned} \quad (\text{B.5})$$

This relationship breaks down when diversion is not by share and results in uneven empirical distributions for the margins of each firm. To remedy this issue, we instead draw a random calibration parameter $\hat{m} \sim U[0.2, 0.8]$ and calculate the own derivative of demand for product i as

$$\frac{\partial q_i}{\partial p_i} = -\frac{1}{\hat{m}} \frac{\hat{s}_i}{p_i} \frac{1-\hat{s}_i}{1-\hat{s}_o} \quad (\text{B.6})$$

This formulation satisfies the relationship in (B.3) and is symmetric across all firms. We then find the cross-price derivatives of demand using the diversion ratios.

$$\frac{\partial q_i}{\partial p_j} = \frac{\partial q_j}{\partial p_j} * \frac{\partial q_i / \partial p_j}{\partial q_i / \partial p_j} = \frac{\partial q_j}{\partial p_j} * -\hat{d}_{ji} \quad (\text{B.7})$$

This provides a matrix of demand derivatives that is similar in functional form to those found through the logit calibration process described in Appendix A:

$$\frac{\partial q_i}{\partial p_j} = \begin{cases} -\hat{\alpha} \hat{s}_i (1 - \hat{s}_i) & \text{if } i = j \\ \hat{\alpha} \hat{s}_i \hat{s}_j & \text{if } i \neq j \end{cases} \quad (\text{B.8})$$

where

$$\hat{\alpha} = \frac{1}{\hat{m}(1 - \hat{s}_o)} \quad (\text{B.9})$$

The price-cost margins of each firm can be characterized similarly to (B.5), but with an additional term that represents the ratio of a firms' product's market share to its diversion share.

$$m_i = \hat{m} \frac{1 - \hat{s}_o}{1 - \hat{s}_i} \frac{s_i}{\hat{s}_i} \quad (\text{B.10})$$

This demonstrates how the randomly drawn data may generate margins that exceed one. If one firms' product has a large market share and an exceptionally low diversion share, the margin can be large. Indeed, the empirical distribution of margins is not bounded above.

To deal with this complication, we remove and replace any markets in which margins exceed one. The implication of this filtering is that diversion cannot be too far away from actual market share, especially for products with large market shares.

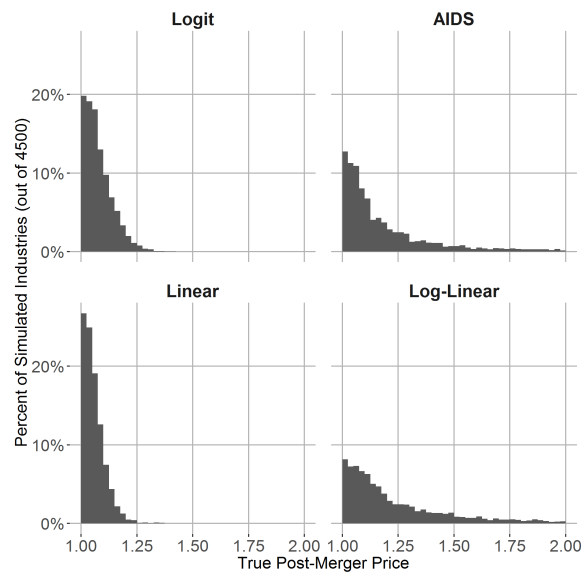


Figure 1: Distribution of Post-Merger Prices

Notes: The histograms characterize the distribution of post-merger prices when the underlying demand system is (from top left to bottom right) logit, almost ideal, linear, and log-linear.




Figure 2: Graphical Illustration of UPP as a Price Predictor

Notes: The scatter plots characterize the accuracy of UPP as a price prediction when the underlying demand system is logit, almost ideal, linear, and log-linear. Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data.




Figure 3: Prediction Error from Standard Merger Simulations

Notes: The scatter plots characterize the accuracy of merger simulations when the underlying demand system is logit (column 1), almost ideal (column 2), linear (column 3), and log-linear (column 4). Merger simulations are conducted assuming demand is logit (row 1), almost ideal (row 2), linear (row 3), and log-linear (row 4). Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data.

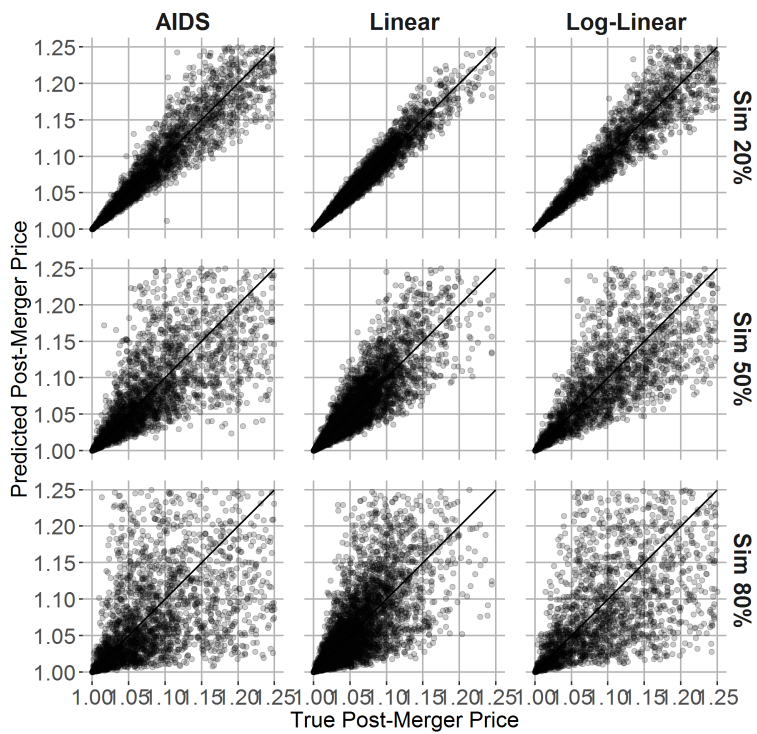


Figure 4: Prediction Error from Standard Merger Simulations

Notes: The scatter plots characterize the accuracy of merger simulations when there is error in the observed elasticities of demand. Merger simulations are conducted assuming the true demand system is known: almost ideal (column 1), linear (column 2), and log-linear (column 3). Each dot represents the first firm's predicted and actual post-merger prices for a given draw of data.

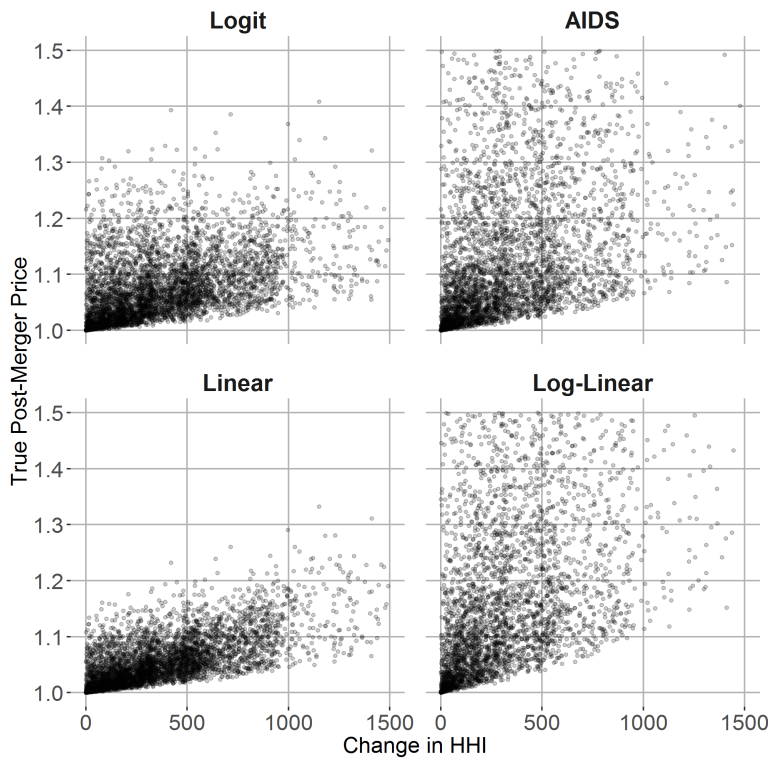


Figure 5: Post-Merger Prices and ΔHHI

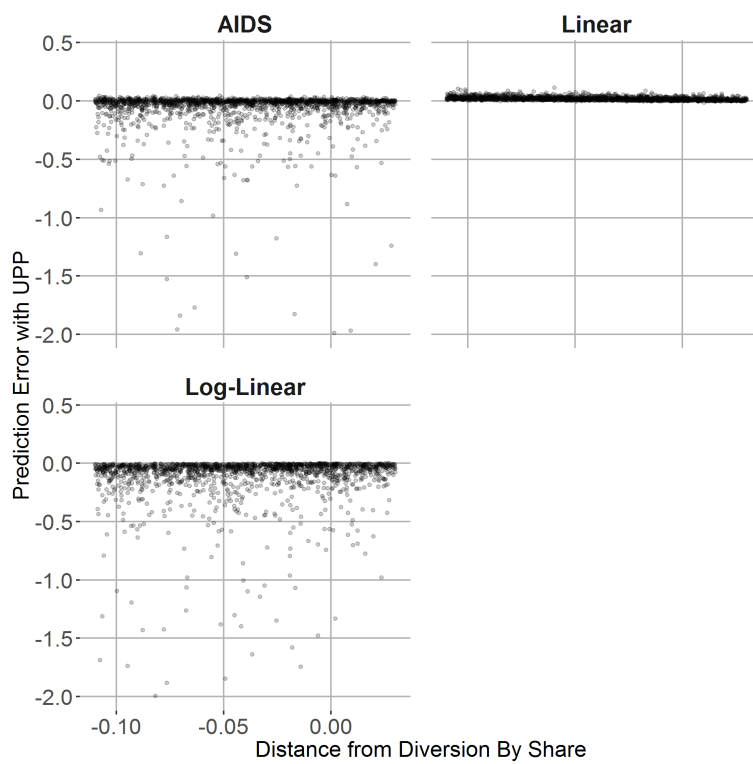


Figure 6: Prediction Error of UPP without Diversion by Share

Notes: The scatter plots characterize the prediction error of UPP with the alternative data. The “distance” from diversion-by-share is calculated by subtracting a hypothetical proportional diversion ratio from the actual non-proportional diversion ratio. Diversion from firm 1 to firm 2 is used, without loss of generality.

Table 1: Order Statistics

	Median	10%	25%	75%	90%
<i>Market Conditions</i>					
Market share	0.15	0.04	0.08	0.20	0.25
Margin	0.49	0.26	0.35	0.64	0.73
Elasticity	2.03	1.37	1.56	2.89	3.87
Diversion	0.17	0.04	0.09	0.23	0.29
<i>Herfindahl–Hirschman Index (HHI)</i>					
Pre-Merger	1562	1231	1361	1816	2073
Post-Merger	1931	1452	1642	2277	2667
Δ HHI	317	48	139	555	833
<i>Upward Pricing Pressure</i>					
UPP	0.07	0.01	0.04	0.12	0.18
<i>Own Merger Pass-Through</i>					
Logit	0.86	0.77	0.81	0.92	0.97
AIDS	1.43	0.83	0.99	2.41	4.16
Linear	0.54	0.51	0.52	0.56	0.59
Log-linear	2.72	1.40	1.89	4.67	9.47
<i>Cross Merger Pass-Through</i>					
Logit	0.03	0.00	0.01	0.04	0.06
AIDS	0.32	0.06	0.16	0.67	1.71
Linear	0.12	0.03	0.06	0.16	0.21
Log-linear	-0.17	-3.88	-0.77	-0.02	-0.00
<i>Merger Price Effects</i>					
Logit	0.06	0.01	0.03	0.11	0.16
AIDS	0.11	0.02	0.05	0.28	0.69
Linear	0.05	0.01	0.02	0.08	0.12
Log-Linear	0.18	0.03	0.08	0.46	1.18

Notes: Summary statistics are based on 4,500 randomly-drawn sets of data on the pre-merger equilibria. The market share, margin and elasticity are for the first firm. The cross-diversion is diversion from firm 1 to firm 2 (the two merging firms). Market share and margin are drawn randomly in the data generating process. The elasticity is the own-price elasticity of demand and equals the inverse margin. Own merger pass-through is the first element of the diagonal of $h(P)$, and cross merger pass-through is the first off-diagonal element of $h(P)$. The merger price effects are the change in firm 1's equilibrium price.

Table 2: Median Absolute Prediction Error

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
UPP	0.006	0.042	0.022	0.110
Partial Simulation	0.001	0.013	0.004	0.000
Logit Simulation	0.000	0.049	0.014	0.117
AIDS Simulation	0.050	0.000	0.068	0.065
Linear Simulation	0.014	0.066	0.000	0.132
Log-Linear Simulation	0.123	0.065	0.139	0.000

Notes: The table provides the median absolute prediction error of UPP, a partial merger simulation holding non-merging prices constant, and standard simulations when the true underlying demand system is logit, almost ideal, linear, and log-linear.

Table 3: Frequency with Which UPP Improves Accuracy

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
Logit Simulation	.	92.2%	3.2%	100%
AIDS Simulation	95.1%	.	90.8%	10.6%
Linear Simulation	69.0%	98.5%	.	99.0%
Log-Lin Simulation	100%	74.6%	100%	.

Notes: This table shows the fraction of mergers for which UPP has a smaller absolute prediction error than standard merger simulations in predicting the price change.

Table 4: Median Absolute Prediction Error for Small and Big Mergers

Panel A: Less than 10% True Price Increase				
<u>Underlying Demand System:</u>				
	Logit	AIDS	Linear	Log-Linear
UPP	0.005	0.007	0.018	0.022
Partial Simulation	0.000	0.004	0.003	0.000
Logit Simulation	0.000	0.010	0.011	0.025
AIDS Simulation	0.025	0.000	0.049	0.017
Linear Simulation	0.009	0.019	0.000	0.031
Log-Linear Simulation	0.071	0.025	0.109	0.000

Panel B: Greater than 10% True Price Increase				
<u>Underlying Demand System:</u>				
	Logit	AIDS	Linear	Log-Linear
UPP	0.014	0.154	0.054	0.221
Partial Simulation	0.002	0.041	0.013	0.000
Logit Simulation	0.000	0.164	0.036	0.230
AIDS Simulation	0.320	0.000	0.628	0.117
Linear Simulation	0.042	0.190	0.000	0.253
Log-Linear Simulation	0.728	0.156	1.450	0.000

Notes: The table provides the median absolute prediction error of UPP and standard simulations when the true underlying demand system is logit, almost ideal, linear, and log-linear.

Table 5: MAPE with Imprecise Demand Elasticities

Panel A: Simulation			
<u>Underlying Demand System:</u>			
	AIDS	Linear	Log-Linear
20% Error	0.007	0.001	0.012
50% Error	0.019	0.004	0.030
80% Error	0.031	0.006	0.049

Panel B: Upward Pricing Pressure			
<u>Underlying Demand System:</u>			
	AIDS	Linear	Log-Linear
20% Error	0.042	0.022	0.110
50% Error	0.041	0.022	0.108
80% Error	0.045	0.023	0.104

Notes: Panel A shows the median absolute prediction error of simulation when elasticities are observed with 20%, 50%, and 80% error and the true demand system is known. Panel B shows the median absolute prediction error of UPP when elasticities are observed with 20%, 50%, and 80% error.

Table 6: UPP as a Screen

	Underlying Demand System:			
	Logit	AIDS	Linear	Log-Linear
	<i>Baseline Calculation</i>			
False Positives (Type I Error)	5.0%	0.2%	18.4%	0.0%
False Negatives (Type II Error)	0.0%	22.4%	0.0%	36.6%
	<i>Alternative Calculation</i>			
False Positives (Type I Error)	7.1%	0.4%	22.0%	0.0%
False Negatives (Type II Error)	0.0%	41.3%	0.3%	53.7%

Notes: The first row shows the fraction of all mergers for which UPP exceeds 10% but the true price change is less than 10%. The second row shows the fraction of all mergers for which UPP is less than 10% but the true price change exceeds 10%. The third row shows the fraction of mergers for which the price change is less than 10% but UPP exceeds 10%. The fourth row shows the fraction of mergers for which the price change exceeds 10% but UPP is less than 10%.

Table 7: HHI Category Screens

Panel A: Frequency of 5% Price Increase				
Underlying Demand System:				
	Logit	AIDS	Linear	Log-Linear
Category (i)	90.9%	95.7%	86.7%	98.8%
Category (ii)	0.0%	33.3%	0.0%	61.9%
Category (iii)	63.8%	79.8%	45.1%	92.4%
Category (iv)	19.3%	46.6%	6.0%	61.5%
Category (v)	0.2%	20.7%	0.0%	30.2%

Panel B: Frequency of 10% Price Increase				
Underlying Demand System:				
	Logit	AIDS	Linear	Log-Linear
Category (i)	53.2%	80.2%	49.5%	93.6%
Category (ii)	0.0%	9.5%	0.0%	33.3%
Category (iii)	17.7%	57.1%	8.0%	73.9%
Category (iv)	0.7%	27.0%	0.0%	36.7%
Category (v)	0.0%	5.8%	0.0%	9.9%

Notes: Panel A shows the fraction of mergers in each HHI category for which the weighted-average change in the merging firms' prices is greater than 5%. Panel B shows the the same statistic for a price increase greater than 10%. Category (i): Post-merger $HHI > 2500$ and $\Delta HHI > 200$. Category (ii): Post-merger $HHI > 2500$ and $\Delta HHI \in (100, 200]$. Category (iii): Post-merger $HHI \in (1500, 2500]$ and $\Delta HHI > 100$. Category (iv): Post-merger $HHI \leq 1500$. Category (v): $\Delta HHI < 100$.

Table 8: Screens Based on ΔHHI

Panel A: Frequency of 5% Price Increase

Underlying Demand System:

	Logit	AIDS	Linear	Log-Linear
$\Delta\text{HHI} > 200$	76.1%	87.0%	60.3%	96.9%
$\Delta\text{HHI} \in (100, 200)$	20.2%	53.4%	0.3%	71.7%
$\Delta\text{HHI} < 100$	0.2%	20.7%	0.0%	30.2%

Panel B: Frequency of 10% Price Increase

Underlying Demand System:

	Logit	AIDS	Linear	Log-Linear
$\Delta\text{HHI} > 200$	27.4%	65.8%	17.6%	81.9%
$\Delta\text{HHI} \in (100, 200)$	0.0%	30.9%	0.0%	45.5%
$\Delta\text{HHI} < 100$	0.0%	5.8%	0.0%	9.9%

Notes: Panel A shows the fraction of mergers in each ΔHHI category for which the weighted-average change in the merging firms' prices is greater than 5%. Panel B shows the same statistic for a price increase greater than 10%.

Table 9: MAPEs with Alternative Data

	Underlying Demand System:		
	AIDS	Linear	Log-Linear
UPP	0.017	0.018	0.046
AIDS Simulation	0.000	0.027	0.039
Linear Simulation	0.025	0.000	0.064
Log-Linear Simulation	0.039	0.067	0.000

Notes: The table provides the median absolute prediction error of UPP and standard simulations when the true underlying demand system is almost ideal, linear, and log-linear.

Table 10: Screens Based on Δ HHI with Alternative Data

Panel A: Frequency of 5% Price Increase

Underlying Demand System:

	AIDS	Linear	Log-Linear
Δ HHI>200	78.5%	32.6%	96.1%
Δ HHI \in (100,200)	35.3%	1.8%	58.4%
Δ HHI<100	11.2%	0.1%	22.0%

Panel B: Frequency of 10% Price Increase

Underlying Demand System:

	AIDS	Linear	Log-Linear
Δ HHI>200	45.7%	3.6%	69.1%
Δ HHI \in (100,200)	14.1%	0.1%	21.0%
Δ HHI<100	2.8%	0.0%	6.7%

Notes: Panel A shows the fraction of mergers in each Δ HHI category for which the weighted-average change in the merging firms' prices is greater than 5%. Panel B shows the same statistic for a price increase greater than 10%.