International Journal of Industrial Organization 28 (2010) 695–707

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Mergers when firms compete by choosing both price and promotion[☆]

Steven Tenn[a], Luke Froeb[b,*], Steven Tschantz[b]

[a] Federal Trade Commission, Washington, DC 20580, United States
[b] Vanderbilt University, Nashville, TN 37203, United States

ARTICLE INFO

Article history:
Received 6 July 2009
Accepted 7 April 2010
Available online 14 April 2010

JEL classification:
L1: Market Structure, Firm Strategy, and Market Performance
L4: Antitrust Policy

Keywords:
Merger analysis
Structural oligopoly model
Promotion

ABSTRACT

We analyze the bias from predicting merger effects using structural models of price competition when firms actually compete using both price and promotion. We extend the standard merger simulation framework to allow for competition over both price and promotion and ask what happens if we ignore promotional competition. This model is applied to the super-premium ice cream industry, where a merger between Nestlé and Dreyer's was challenged by the Federal Trade Commission. We find that ignoring promotional competition significantly biases the predicted price effects of a merger to monopoly (5% instead of 12%). About three-fourths of the difference can be attributed to estimation bias (estimated demand is too elastic), with the remainder due to extrapolation bias from assuming post-merger promotional activity stays constant (instead it declines by 31%).

1. Introduction

Antitrust laws prohibit mergers that substantially lessen competition, and most often this means mergers that raise price. Enforcement agencies have a relatively good understanding of price competition, and how to measure the loss of price competition caused by merger. Less well understood are the effects of mergers when firms compete in multiple dimensions. In this paper we study mergers in industries where firms compete by setting both price and promotion, and ask what happens if we mistakenly assume that price is the only dimension of competition.

To answer the question, we build a structural merger model where firms compete using both price and promotion. We find two sources of potential bias from ignoring promotional competition. The first is what we call "estimation bias," a type of omitted variables bias. If promotion is correlated with price, then observed price changes will proxy for unobserved (or ignored) changes in promotional activity. As a consequence, price elasticity estimates will be biased. Bias in estimated own-price elasticities affects the post-merger price prediction because a merged firm facing a more elastic demand would not raise price as high as a merged firm facing a less elastic demand, all else equal.

The second source of potential bias is what we call "extrapolation bias." Following a merger, we would expect the merged firm to internalize both price and promotional competition among its commonly owned products. In price-only merger models, promotional activity is implicitly held constant at pre-merger levels when the post-merger equilibrium is calculated. This leads to extrapolation bias when optimal price depends on the level of promotional activity.

Interestingly, both types of bias depend, in part, on whether promotional activity makes demand more or less elastic. Our analysis suggests a simple heuristic: if promotional activity makes demand more (less) elastic, price-only merger models will under- (over-) estimate the magnitude of the post-merger price increase. Only in the special case where optimal price and promotion are independent of each other will a price-only model provide an accurate prediction of post-merger prices.

To empirically assess the magnitude of this bias, we extend the standard merger simulation framework by allowing firms to compete over both price and promotional activity. This framework is applied to the super-premium ice cream industry, where the Federal Trade Commission (FTC) challenged a 2003 merger between Nestlé and Dreyer's.[1] The results from our price-plus-promotion merger simulation framework are compared to counterpart estimates from a price-only model. The structural model which ignores promotion substantially under-predicts the price effects of a merger when firms actually compete via both price and promotion. This comparison demonstrates

☆ DISCLAIMER: the views expressed in this paper do not purport to represent the views of the U.S. Federal Trade Commission, nor any of its Commissioners. We wish to acknowledge useful discussions with Hajime Hadeishi, Dawn Iacobucci, David Schmidt, Timothy Muris, David Scheffman, Robert McMillan, Craig Peters, Henry Schneider, Shawn Ulrick, Matthew Weinberg, David Weiskopf, Gregory Werden, and participants at Cornell, Duke, and FTC workshops.

* Corresponding author.
E-mail addresses: stenn@ftc.gov (S. Tenn), luke.froeb@owen.vanderbilt.edu (L. Froeb), tschantz@math.vanderbilt.edu (S. Tschantz).

[1] See Section 4 for background on the Nestlé–Dreyer's transaction.

0167-7187/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijindorg.2010.04.003

that the standard merger simulation framework can be extended to accommodate situations where firms compete in multiple dimensions, and that doing so may significantly alter merger predictions.

2. How reliable are structural merger models?

To determine whether a proposed merger is anticompetitive, enforcement agencies must be able to predict its effects. Merger predictions from structural oligopoly models are made in two distinct steps. First, a model of firm and consumer behavior is estimated from observed pre-merger data. Then, the estimated model parameters are used to forecast post-merger prices as the merged firm internalizes competition among commonly owned products. The difference between observed pre-merger prices and predicted post-merger prices is an estimate of the merger effect. The popularity of the methodology and its use to inform enforcement decisions has raised questions about its reliability (Werden et al., 2004; Hosken et al., 2005).

Tests of model reliability have taken three forms. The first is to test over-identifying model restrictions. For example, Werden (2000) and Pinske and Slade (2004) compare pre-merger price-cost margins to demand elasticities in the Chicago bread industry and the UK brewing industry, respectively. They find that the familiar margin–elasticity relationship of Bertrand models holds in the pre-merger data. However, Kim and Knittel (2006) find that the margin–elasticity relationship does not hold in the California electricity market.

A second test of model reliability is to determine how well structural models can predict actual events. Nevo (2000) finds that predicted price changes are close to actual price changes for two mergers in the ready-to-eat cereal industry. However, Peters (2006) finds that simulation methods "do not generally provide an accurate forecast" of post-merger prices for five airline mergers. Using information from the post-merger period, he finds that the merger prediction error is caused by unobserved cost changes or firm behavioral changes. Hadeishi and Schmidt (2004) find that a structural model could not predict the effects of a merger between two complements (peanut butter and jelly) because demand shifted post-merger. Lastly, Weinberg and Hosken (2008) find that a structural model slightly under-predicts the price effects of a merger among producers of motor oil, but over-predicts the price effects of a merger among breakfast syrup producers.

Testing whether models can accurately predict real events is difficult because it requires estimating both a structural model as well as the merger effect. The latter requires controls for confounding factors that could cause post-merger prices to change. Typically, a difference-in-difference estimator is used to compare prices of the merging firms to a set of control prices. The difficulties of this approach were the subject of recent FTC hearings (Federal Trade Commission, 2005) on whether the oil mergers in the late 1990s increased the price of gasoline. Panelists raised questions about the choice of control group, the choice of variables, the maintained assumptions behind the difference-in-difference estimator, and the selection bias that follows from antitrust enforcement as agencies challenge only mergers that are thought to be anticompetitive. The panel concluded that more study is needed to reliably answer the question that motivated the hearings.

The third approach to testing model reliability is to determine the sensitivity of post-merger predictions to the behavioral assumptions on which the models are built. Bertrand models of competition, for example, assume that consumers choose products based on the price differences between them, and that firms compete solely on the basis of price (e.g., Hausman et al., 1994; Werden and Froeb, 1994). This is a simplification of the way consumers and firms actually behave, but models necessarily abstract away from reality. Finding that a model is unrealistic is neither surprising nor interesting. Rather, we want to know whether a model is unrealistic in ways that make its predictions misleading. There are a number of studies that try to address this issue.

Gandhi et al. (2008) find that if firms can change their "locations" in product space, in addition to price, the merged firm will move its products apart to avoid cannibalizing sales. Since this reduces the incentive to raise price, ignoring such repositioning overstates the effects of a merger. Similarly, Froeb et al. (2003) find that capacity constraints on the merging firms attenuate merger effects just as those on the non-merging firms amplify them. Since the former effect is typically bigger than the latter, ignoring the effects of capacity constraints likely overstates the effects of a merger. Crooke et al. (1999) and Froeb et al. (2005) find that demand functional form determines, to an extent, the predicted merger effect. They recommend either sensitivity analysis, or computing what they call "compensating marginal cost reductions," the reductions in marginal cost necessary to offset the merger effect. These depend only on elasticities, and not on demand functional form (Werden, 1996). Finally, several authors have found that the downstream price effects of upstream manufacturing mergers are determined, in part, by the vertical relationship between manufacturers and retailers. Depending on the types of agreements between manufacturers and retailers, the retail sector can attenuate, amplify, or completely hide the effects of upstream manufacturing mergers (Froeb et al., 2007; O'Brien and Shaffer, 2005).

This paper follows this third strand of research by asking whether ignoring competition based on promotional activity biases the predictions of structural merger models.

3. When does promotion matter for merger analysis?

In this section, we specify a simple canonical model of price and promotional competition as a static non-cooperative game that is typical of the kinds of structural models used to predict merger effects. We discover a correspondence between firm behavior in a price-plus-promotion model and firm behavior in a price-only model that allows us to characterize the estimation bias that comes from ignoring promotion. We also find that, under general conditions, firms have an incentive to change promotion and price in the post-merger equilibrium. Extrapolation bias arises in price-only models that mistakenly hold promotion fixed at pre-merger levels.

3.1. Mergers that eliminate price competition between the merging firms

We begin by reviewing the determination of Nash equilibria in static Bertrand price-only models. The industry is composed of n products, with product j having price p_j and quantity demanded $q_j(\mathbf{p})$ that is a function of the vector $\mathbf{p}$ containing each product's price. The cost of producing product j is denoted by $c_j(q_j)$. Define $\text{profit}_j = p_j q_j - c_j$ as the profit associated with product j. If each product is owned by a different firm which sets price so as to maximize its profit, then the first-order condition for optimal pricing is given by

$$0 = \frac{\partial \text{profit}_j}{\partial p_j} = q_j + \left(p_j - c_j'\right)\frac{\partial q_j}{\partial p_j}, \tag{3.1}$$

where c_j' is the marginal cost for product j. This gives rise to the familiar margin–elasticity relationship, or Lerner equation, characteristic of single-product Bertrand models,

$$\frac{p_j - c_j'}{p_j} = -\frac{1}{\varepsilon_{jj}}, \tag{3.2}$$

where $\varepsilon_{kj} = (\partial q_k/\partial p_j)(p_j/q_k)$ is the cross-price elasticity of demand for product k with respect to price j. If margins are observed and demand elasticities can be estimated, then Eq. (3.2) is an over-identifying