Economics at the FTC: Horizontal Mergers and Data Security

Ginger Zhe Jin
Dan Hanner
Marc Luppino
Ted Rosenbaum

Abstract: This article describes aspects of the economic analyses that were performed on three matters on which economists in the Bureau of Economics at the Federal Trade Commission have worked during this past year. The first two matters were merger investigations: They (separately) involved tobacco manufacturers and food distributors. While these investigations shared some common attributes, such as the importance of the proposed divestitures, this essay demonstrates how our analysis must vary based on the types of information and the competitive concerns presented by specific transactions. The third section discusses a general economic approach to estimating consumer harm from data breaches, which is illustrated with an example of an application to a recent case.

Address for all authors but Jin:
Federal Trade Commission, Bureau of Economics
600 Pennsylvania Ave., N.W.
Washington, DC 20580, U.S.A

Address for Jin:
Department of Economics
University of Maryland
College Park, MD 20742

determined in this instance that a proposed divestiture would be insufficient to maintain the competition that would likely be lost via the merger. In June of 2015, after the U.S. District Court for the District of Columbia granted the Federal Trade Commission's request for a preliminary injunction, Sysco and US Foods abandoned their proposed merger. A number of challenging topics had to be addressed in this investigation. This article will discuss several of those, including price discrimination, bidding markets, and the analysis of the net effect of competitive harms and cost savings likely to be passed-through.

The final section is a discussion of data security issues and focuses on an economic analysis that can be used to estimate injuries from data breaches. The courts have held that the FTC has authority to bring data security cases under the provision of Section 5 of the FTC Act that outlaws unfair acts or practices in or affecting commerce.[13] Rather than addressing the contours of the legal responsibilities of companies to safeguard information, this section instead focuses exclusively on how to use an economic framework to assess the harm that is done when such precautions are not taken and information is compromised. As an illustrative example, this section discusses the potential application of this framework in the recent FTC case against Wyndham Worldwide Corporation for allegedly employing insufficient security protocols.

II.   Tobacco Merger: Reynolds and Lorillard

During 2014-2015, the FTC investigated Reynolds American's proposed acquisition of Lorillard. At the time, Reynolds and Lorillard were the second and third largest U.S. cigarette manufacturers, respectively. On the same day that the Reynolds/Lorillard transaction was announced, Reynolds reached an agreement to sell several cigarette brands and other assets to Imperial, the fifth largest cigarette manufacturer in the U.S. The result of these two transactions would be that Reynolds would acquire Lorillard's Newport brand, while Imperial would gain three of Reynolds' brands (Winston, Kool, and Salem) and two Lorillard brands (Maverick and blu eCigs – an electronic cigarette brand). Reynolds's other major brands are Camel and Pall Mall. The Commission accepted a consent decree that required that the firms divest the five brands to Imperial, in lieu of a challenge to the transaction.[14]

---

[13] FTC v. Wyndham Worldwide Corp., 799 F.3d 236 (3d Cir. 2015).
[14] See https://www.ftc.gov/system/files/documents/public_statements/644971/150526reynoldscommstatement.pdf.

As part of the investigation, BE economists used detailed data, document analysis, and testimony from the merging parties to assess the impact of the series of transactions. In this article, we discuss how the merger would change firms' unilateral pricing incentives ("unilateral effects"). However, as part of our investigation, we also considered how the merger might change firms' ability to collude tacitly ("coordinated effects"),[15] and whether the merger might create efficiencies.

A. The Theory of Harm

In many ways, the theory of harm from this merger follows that of any merger of differentiated substitutes. However, measuring the welfare impacts from this transaction presented a few additional complexities relative to those of a "typical" case.

First, the notion of "welfare effects" of a merger of firms that sell cigarettes is a potentially fraught topic. To the extent that there are negative externalities (e.g., via second hand

assess harm from this transaction. Both of these approaches yielded estimates of state-by-state changes in aggregate consumer surplus: First, we utilized first-order approximation methods (upward pricing pressure or "UPP") to compute a range of consumer harm estimates, each of which relied on different sets of assumptions about consumer substitution patterns. Second, by calibrating and estimating a simple demand system, we conducted a merger simulation to measure consumer and total welfare changes. The remainder of this section will discuss the UPP approach in detail, and conclude with how a merger simulation complemented our other analyses.[17]

B. Using First-Order Approximations to Compute Harm

1. Overview of method

To compute estimates of price changes and consumer harm without specifying a demand system and computing equilibrium, we utilized a first-order approximation approach that is outlined in Jaffe and Weyl (2010) and Miller et al. (2016).[18] For this analysis, we assumed the firms were engaging in Nash-Bertrand price competition.

Within that context, consider a firm f that faces the following twice-differentiable profit function:

$$\pi_f = q_f(P)'(p_f - mc_f) \,,$$

where p_f is a $n_f \times 1$ vector of prices; $q_f(P)$ is a $n_f \times 1$ vector of quantities that are evaluated at the market price vector P; mc_f is a $n_f \times 1$ vector of firm f's marginal cost for each product; and n_f is the number of products that are controlled by firm f prior to the transactions.

We can write a first-order approximation to the equilibrium price change that follows a transaction that modifies firm f's pre-merger product portfolio. In particular, consider an acquisition where firm f acquired a set of products from firm k and divested a set of products to firm l. As a result of the products from firm k that enter firm f's portfolio and the products that are divested to firm l (and thus leave firm f's portfolio), there is an impact on the equilibrium prices of all of the other products.

---

[17]

The first-order approximation of the post-transaction change in the price of the products that were originally in firm f

In the context of the investigation, we computed the $\approx z_i$ vector for the pre-merger products of Reynolds, Lorillard, and Imperial. These were the products that would have a change in upward or downward pricing pressure because of the series of transactions.

Further, one can compute the change in consumer surplus using the expression $\approx\% 5N \; \approx z^i \; \grave{q}$ where $\approx z^i$ and $\grave{q}$ are vectors that approximate the post-merger price change and each product's quantity respectively. For $\grave{q}$, there can be theoretical justifications for using pre-merger or post-merger quantities or a geometric average of the two (Jaffe and Weyl, 2010).

Therefore, this analysis of a first-order approximation using the 4r(ae-6r(a)-dn (ng)10 (prue)4 (t)-2 (.tp0 Tc

In the widely used logit demand system, the diversion ratio is proportional to the market share of the two products (Conlon and Mortimer, 2013), which can be expressed as $D_{jm} = \frac{s_m}{1-s_j}$. However, this finding is also true at the group level – in other words, we can combine individuals into small groups, and then assume proportional diversion ratios within each group – which allows for substantial flexibility across groups (Raval et al., 2015; Farrell et al., 2011). The group-level diversion ratio is expressed as $D_{jm}^g = \frac{s_m^g}{1-s_j^g}$ where s_m^g and s_j^g are m and j's shares within group g. Therefore, this approach allows for flexible substitution patterns across groups, while imposing a specific structure on substitution within each group. The aggregate diversion ratio is a weighted average of the group diversion ratios, where the weights are given by the share of each group in the share of product j.[22]

With the proprietary datasets that were obtained in the course of our investigation, we used three distinct approaches to define groups: geography, consumer type, and demography. Ideally, we would have a dataset that contains information on all three of these together in order to stratify consumers by all of these types jointly. However, as is often the case, this type of data was not available.

Therefore, given the data constraints, we looked at each of the three approaches to define groups in isolation from the others: For geography, we used store-level information on product shipments and assumed proportional substitution within the store. For consumer types, we used survey data on individuals' brand preferences and assumed proportional substitution within these brand preferences.[23] For demography, we assumed proportional substitution within demographic groups and used survey data on age, sex, and ethnicity, for example.

While each of these approaches has drawbacks -- especially when viewed in isolation -- when viewed in concert they helped to paint a clear picture of the substitution patterns in the market. Further, due to the simplicity of computing each of these approaches, we were able to present a range of harm estimates each based on a different set of plausible, but imperfect, assumptions.

---

form assumptions – including logit -- imply that the diversion ratio is constant or close to constant over a range of prices, in which case a diversion ratio from a natural experiment may be informative (see Conlon and Mortimer, 2013).

[22] Formally, $D_{jm} = \sum_g S_j^g D_{jm}^g$ where S_j^g is the share of demand for product j that is purchased by group g.

[23] In this case, we put a weight of one on the diversions from the regular users of that product. For example, to compute the diversion ratio from j to m, we only used regular users of j.

b. Markups

The second major ingredients needed are the gross markups, which require information on prices and costs. Prices (which included both cross-sectional and time variation within products) and costs were estimated using information that was obtained in the investigation. We included taxes in our estimates of marginal costs, since that portion of price is not captured by the firm as profit.

c. Pass-through

The final necessary ingredient is the pass-through matrix. As our baseline harm estimate, we used the identity matrix for the pass-through matrix, which implies that the net pricing pressure is the prediction of the post-merger price change. In the absence of marginal cost shocks or a flexible demand system, it is not possible to know the "true" pass-through matrix. Therefore, we relied on results that show that for many common demand systems net pricing pressure is frequently a better predictor of post-merger price changes than a mis-specified demand system (Miller et al., 2016).

3. Potential Limitations of the UPP Analysis

This UPP analysis suffered from two limitations: one based upon our data, and the other endemic to the approach.

The nature of a first-order approximation method is that it is a local analysis. In other words, this approach computes a linear approximation to a small change from the pre-merger equilibrium. If there is significant curvature in the demand system – such that the shape of demand is very different at pre-merger and post-merger equilibrium prices -- this may not be a good approximation.

Second, when we estimated diversion ratios, we lacked information on the extent to which a price increase would cause consumers to stop smoking or reduce the number of cigarettes that they smoke per day or per week. We assumed that in response to a relatively small post-merger price increase, this would be a small effect.

C. Full Merger Simulation

To address concerns as to the limitations of the UPP analysis, we also estimated and calibrated a parametrized demand system and conducted a merger simulation using it. Within

that framework, we were able to conduct a non-local analysis and estimate substitution to the outside option. Moreover, in our welfare calculations we were able to account for changes in quantities purchased as a result of the changes in prices. The cost of the full simulation approach is that it required stronger assumptions about the nature of the substitution patterns between the products than were required for the UPP analysis.

D. Broader Lessons

In our investigation, the UPP and full simulation approaches and the differing approaches to computing diversion ratios for UPP proved to be complementary. By evaluating the predicted

Overall, the combination of individual negotiations and the use of sales reps to evaluate each customer's specific circumstances gives i   [(i)-6inovanm2 6v tc 0.004 Tw -11.68 -[eru10 (et-1 (-6 (c tta4 (

This analysis was complicated by the fact that the merging parties signed a divestiture agreement with PFG ("Performance Food Group", a regional distributor) during the investigation to sell eleven of USF's distribution centers to PFG. While the expanded PFG network improved PFG's geographic footprint, even with the divestiture PFG remained substantially weaker than the existing Sysco and USF networks along many dimensions. For example, for Sysco's national account customers, the average distance from their locations to Sysco's distribution network and the average distance from their locations to USFs distribution network were nearly identical. In contrast, the average distance to the network that PFG would have had even after the divestiture would have been approximately 50% greater. And since distance mattered for both delivery costs and customer service (errors were easier to remedy for nearby customers), for these customers the value of the next best alternative appeared to be substantially reduced by the merger.[29]

Similar analyses compared PFG to Sysco and USF using other measures, such as: the variety of SKUs carried; customer service ratings; distributor costs (based on scale and other factors); the flexibility of distribution schedules; and the implications of current market shares for the overall quality of the next best alternative.

BE also evaluated the potential efficiencies from the transaction. As noted above, in addition to economists, the bureau staff also includes financial analysts who worked with the economists to evaluate different aspects of the efficiency claims. The economists and financial analysts focused on a number of questions, such as whether or not efficiencies are likely to be passed through to consumers and whether the efficiencies are merger-specific.

While the simple bidding model suggests that pass-through is more likely when the merging firms are the customer's second- and third-best options, and that firms would not have an incentive to pa 0 Tw9C   /P (is222 (m)-6 (p)-4 (l)-6 (i)1 (Tc 74)]TJ   -20.8d   [(a)6 0.004 Tw 0.A[(a)6)2 (ti

natural to expect that pass-

$$2k\$_\gamma o\, R\frac{2k\,\$_s = 1\,\hat{e}\,\dagger_a = 1\,o}{2k\dagger_a\,\$_\gamma = 1\,o},$$

which one could calculate and use as a benchmark in assessing the data security practices of the company and the likelihood of consumer injury. It is possible that a data security expert may be able to provide a point or bounded estimate of the likelihood that a company with similar data security practices would eventually experience a data breach. If such an estimate is not available, we can rely on the fact that this probability could be at most one and proceed with presenting an analysis with a range of plausible data breach probabilities.

One can combine the aforementioned estimated inputs as suggested by the price premium formula to produce a point or bounded estimate of the price premium that results from a given firm's deficient data security practices. In turn, this estimate can be multiplied by the total number of customers to produce an estimate of the extent of consumer injury that is attributable to a firm's data security violation.

B.   Illustrative Example in a Data Security Case

1.   Background

In this section, we consider how the economic framework that was laid out above could be applied to real world cases; we use

Over the course of these legal proceedings, the Third Circuit Court of Appeals affirmed the federal district court's ruling that the FTC could use the prohibition on unfair practices in Section 5 of the FTC Act to challenge the alleged data security lapses that were outlined in its complaint against Wyndham.[33] Wyndham ultimately reached a settlement with the FTC under which the company agreed to establish a comprehensive information security program that is designed to protect cardholder data, conduct annual information security audits, and maintain safeguards in the connections to its franchisees' servers.

2.   Analysis

We will briefly discuss various sources of information that can be used to estimate specific inputs in the potential consumer injury calculation described above, although the specific details underlying how one would produce these estimates in a particular case would typically be confidential. The Bureau of Justice Statistics has intermittently included Identity Theft Supplements in its National Crime Victimization Survey, and the corresponding data from these surveys is a very useful starting point for assessing potential consumer injury that results from deficient data security.

For example, the 2012 Identity Theft Supplement of the National Crime Victimization Survey includes information on out-of-pocket costs and time costs that are associated with debit and credit card fraud, which can be used to estimate the expected cost that a consumer would incur as a result of identity theft. This type of data can also be used to estimate the incidence of successful credit card and debit card fraud, which -- along with information about the percentage of the population that holds debit or credit cards -- can be used to calculate the baseline identity theft rate.[34]

It is difficult to find publicly available information to estimate the likelihood of experiencing identity theft, conditional on being exposed in a data breach. One possible source of information is data from major banks and credit card companies on the transaction histories of cards that have been exposed in identified data breaches, including whether fraud was observed on the exposed cards. Estimating the probability that a given firm would experience a data

---

[33] See https://www.ftc.gov/system/files/documents/cases/140407wyndhamopinion.pdf.
[34] For example, the Federal Reserve Bank of Boston's Survey of Consumer Payment Choice is one potential source of information on the percentage of households with access to different payment instruments, such as debit and credit cards. See https://www.bostonfed.org/publications/survey-of-consumer-payment-choice.aspx.

breach is also potentially difficult, but one can calculate consumer injury estimates for ranges of

Farrell, J. & Shapiro, C. (2010). Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition. The B.E. Journal of Theoretical Economics, 10(1), DOI 10.2202/1935-1704.1563

Federal Trade Commission & Department of Justice. (2016). Hart-Scott