

**A more recent version of this article has been accepted for publication in:
Journal of Competition Law and Economics, 2008**

THE ANALYSIS OF MERGERS THAT INVOLVE MULTI-SIDED PLATFORM BUSINESSES

David S. Evans*

Michael D. Noel[†]

April 2008

* University College London and University of Chicago, and LECG, LLC. devans@lecg.com.

[†] University of California at San Diego, mdnoel@ucsd.edu.

This paper is an extensive revision of a paper that previously circulated under the title “Analyzing Market Power and Market Definition in Multi-Sided Platform Markets.” We would like to thank Howard Chang, Kirsten Edwards, Andrei Hagiu, Alison Oldale, Michael Salinger, Richard Schmalensee, and Vanessa Yanhua Zhang for valuable comments and Murtaza Akhter, Eri Budo, Marina Danilevsky and Nese Nasif for exceptional research assistance.

ABSTRACT

A multi-sided platform (MSP) serves as an intermediary for two or more groups of customers who are linked by indirect network effects. Recent research has found that MSPs are significant in many industries and that some standard economic results—such as the Lerner Index—do not apply to them, in material ways, without some significant modification to take linkages between the multiple sides into account. This article extends several key tools used for the analysis of mergers to situations in which one or more of the suppliers are MSPs. It shows that the application of traditional tools to mergers involving MSPs results in biases the direction of which depends on the particular tool being used and other conditions. It also extends these tools to the analysis of the merger of MSPs. The techniques are illustrated with an application to an acquisition involving the multi-sided online advertising industry.

I. INTRODUCTION

This article presents an empirical framework for examining market definition and unilateral effects in mergers in which one or more of the businesses that may be considered for the hypothetical market are multi-sided platforms (MSPs). MSPs provide goods or services to several distinct groups of customers who need each other in some way and who rely on the platform to intermediate transactions between them (Evans (2003a, b), Rochet and Tirole (2003, 2006)).¹ They typically reduce transaction costs and thereby permit value-creating exchanges to take place that otherwise would not occur (Evans and Schmalensee (2007a, b)). In particular, they facilitate the realization of indirect network externalities, and externalities in use, between the members of distinct customer groups (Rochet and Tirole (2003)).

Many old industries are based on MSPs, ranging from village matchmakers that date from ancient times to advertising-supported newspapers introduced in the 17th century to payment cards introduced in the mid 20th century. However, an increasing number of significant modern businesses are MSPs as a result of technological changes that have drastically lowered the costs and increased the benefits of connecting diverse customer groups on a single platform. These include most internet-based businesses such as eBay, Facebook, and Google. These businesses are creating new products and services such as social networking platforms and are disrupting existing industries such as advertising-supported media (Evans and Schmalensee (2007b)).

The standard tools of antitrust and merger analysis, which were developed based on the economics of single-sided businesses, do not necessarily apply in ways that are material to the analysis of competition that involves multi-sided businesses. Each side of the MSP's business influences and constrains its strategies on the other side. Antitrust analysis that focuses on one side of the business in isolation from the

¹ Many multi-sided platforms have two primary sides such as advertising and readers, and much of the economic literature has focused on the case of two-sided platforms (2SPs).

other side is incorrect as a matter of economics, and can lead to the wrong answer when indirect network effects are significant and are relevant for assessing the practice at issue (Evans (2003a), Wright (2004)). This article shows how the standard tools used for analyzing market definition and unilateral effects for mergers need to be modified when the parties are MSPs. While we caution against in relying too heavily on mechanical market definition exercises as a general matter, as long as practitioners and courts use these techniques, they should be done correctly when MSPs are involved. The analysis of market definition and power for mergers has obvious extensions to other areas of antitrust.

We present an empirical framework that can be used to handle situations in which one or more MSPs may be the subject of the merger analysis. Section II provides an informal discussion of the analysis of market definition, market power, and unilateral effects for situations that involve MSPs and how that analysis differs from that for situations that only involve single-sided firms. Section III then presents our formal analysis for the special case of two-sided markets. Section IV considers a series of examples and simulations to highlight the benefits of pursuing the correct analysis. Section V applies the analysis to an example based on Google's acquisition of DoubleClick. Section **Error! Reference source not found.** presents brief conclusions.

II. ANALYSIS OF MARKETS WITH MULTI-SIDED PLATFORMS

Consider profit maximization for a platform that serves customer groups A and B . Suppose the platform has already established prices for both groups and is considering changing them.² If it raises the price for members of group A , fewer A 's will join. If nothing else changed, the relationship between price and the number of A 's would depend on the price elasticity of demand for A 's. Since, however, members of group B value the platform more if there are more A 's, fewer B 's will

² To keep matters simple we consider the case where each side is charged a membership fee as in Armstrong and Wright (2007). MSPs generally involve platforms on which interactions take place and customers face an access fee and a usage fee although they may choose to make some of those fees zero.

join the platform at the current price for B 's. That drop-off depends on the indirect network externality which is measured by the value that B 's place on A 's. But with fewer B 's on the platform, A 's also value the platform less leading to a further drop in their demand. There is a feedback loop between the two sides. Once this is taken into account the effect of an increase in price on one side is a decrease in demand on the first side because of the direct effect of the price elasticity of demand and on both sides as a result of the indirect effects from the externalities. The change in revenue from a change in the price for A 's therefore depends on the price elasticity of demand for A 's and the indirect network effects between the two sides. Costs necessarily go down so long as marginal costs are positive since the number of customers has dropped on both sides. As is always the case with profit maximization, the price increase is profitable if revenues do not decline more than costs decline.

The platform would like to find the prices for each side that maximize its profits by taking these considerations into account. As Rochet and Tirole (2003) observe, one can think of these as determining the absolute and relative levels of prices. Three key results hold for two-sided platforms based only on the assumptions that there are two distinct customer groups, there are positive externalities between members of those groups, and a two-sided platform provides a good or service that facilitates exchange of value between the two customer groups in the face of these externalities:

- Each optimal price depends on the price elasticities of demand for both sides; the nature and intensity of the indirect network effects between each side; and the marginal costs that result from changing output of each side.
- An increase in marginal cost on one side does not necessarily result in an increase in price on that side relative to the price of the other side. The price ratio between two sides depends only on the *ratio* of elasticities (not inverse elasticities), and not on marginal cost.
- The profit-maximizing price for one side may be below the marginal cost of supply for that side or even negative. A common situation analyzed by

Armstrong (2006) is when the platform in effect buys A 's who are valued by B 's. The relationship between price and cost is complex, and the simple formulas that have been derived for single-sided markets do not apply.

Several results that are relevant for the analysis of market definition and unilateral effects follow immediately from these results. Significantly for our purposes, the widely used Lerner Index

$$\frac{p - c}{p} = \frac{1}{\eta} \quad (1)$$

where η is the usual own-price elasticity of demand, does not accurately summarize the profit-maximizing equilibrium for MSPs when applied to a particular side. This condition does not consider the linkage between the two sides and as a result does not reflect the profit-maximizing equilibrium condition for a two-sided platform. For the special case considered by Rochet and Tirole (2003) the two-sided version of the Lerner Index is

$$\frac{p^A + p^B - c}{p^A + p^B} = \frac{1}{\eta^A + \eta^B} \quad (2)$$

Merger analyses tools such as the Critical Loss and Diversion Ratio that are based on the one-sided Lerner Index are therefore not correct when applied to multi-sided businesses.³

Moreover, merger simulation models are misspecified when they fail to consider the multi-sided nature of the business.⁴ Standard models, when applied to one side of a two-sided industry, fail to consider the feedback effects between the two sides—in other words that the demand by side A depends on the number of customers on side B —and one must consider the demand by sides A and B simultaneously to properly account for all the feedback effects.

³ We derive the two-sided Lerner Index for our model below.

⁴ See Werden, Gregory J & Froeb, Luke M (1994); Baker, J., and T. Bresnahan (1985); Feldman, R. (1994); and Nevo, A. (2000).

A further implication is that the SSNIP test for defining a relevant market does not apply without significant modifications when any of products involved in the analysis are produced by an MSP. Consider the case of a merger between two symmetric MSPs that serve the same customer groups *A* and *B*. To define the market an analyst proceeds by starting with the merger of the products that serve demand for, say, side *A* because that is the focus of the competition concern.⁵ The set of products is expanded until a hypothetical monopolist over that set of products could raise price by, say, 5 percent or more on each of those products. That set of products then defines the market for analysis.

However, by ignoring side *B* the analyst fails to consider that the hypothetical price increase reduces the number of side *A* customers available to side *B*, which thereby reduces the prices that side *B* customers will pay, and furthermore reduces the number of side *B* customers available to side *A*, which in turn reduces the prices that side *A* customers will pay. The link between sides *A* and *B* reduces the profitability of any price increase. Therefore, the market would be drawn too narrowly and estimates of market concentration too high, because the standard approach fails to consider the tempering effects on price coming from the other side. However, we will also show that the use of common one-sided calibration techniques to obtain elasticity estimates which could not otherwise be directly estimated can cause a reverse bias instead, with the market drawn too broadly. So the overall bias will depend on the nature of the MSPs and the estimation technique used.

The mistake though is more profound. The purpose of market definition is, in part, to help focus the economic analysis on a relevant but finite set of products and competitive relationships for analysis. For industries in which the multi-sided effects are sufficiently strong, market definition that excludes one side of an MSP results in the failure to consider multi-sided strategies and market linkages. Failure to consider those multi-sided relationships can result in Type I and Type II errors: failing to

⁵ For a discussion of the SSNIP test see US Department of Justice Merger Guidelines available at http://www.usdoj.gov/atr/public/guidelines/horiz_book/hmg1.html. In particular see Section 1 on Market Definition, Measurement and Concentration.

recognize practices that may be harmful because of the two-sided relationships and condemning practices that are innocuous in a two-sided context.

The remainder of this paper provides the toolkit for conducting merger and unilateral effects analyses for MSPs.

III. CRITICAL LOSS ANALYSIS FOR MULTI-SIDED PLATFORMS

Economists have developed a number of techniques to assess market definition and the competitive consequences of a merger following seminal contributions by Farrell and Shapiro (1990), Willig (1991), and Werden and Froeb (1994). This work and subsequent contributions build on the fundamental insight that one can infer own-price elasticities of demand from price-cost margins (the Lerner Index) and use these estimated elasticities to conduct a variety of simple analyses. Given enough data, one can also evaluate the effects of a merger by estimating a demand system under some assumptions (such as differentiated-market Bertrand) about strategic interactions (see Ivaldi and Verboven (2005)).

This section uses one of the more popular parsimonious techniques—Critical Loss Analysis—to show how these techniques can be extended and modified for mergers involving MSPs. We begin with a brief overview of one-sided Critical Loss Analysis and then introduce the two-sided variant.

A. One-Sided Critical Loss

Critical Loss Analysis was introduced by Harris and Simons (1989) as a user-friendly implementation of the SSNIP test.⁶ It is a simple calculation intended to provide a first look at potential market size. It compares “Critical Loss” (*CL*)—the percentage loss in quantity of a hypothetical monopolist’s products that would be exactly enough to make an *X*% price increase in the price of all its products unprofitable—to “Actual Loss” (*AL*)—the predicted percentage loss in quantity that

⁶ See Harris and Simons (1989), O’Brien and Wickelgren (2003) and Katz & Shapiro (2003). See Harris and Veljanovski (2003) for discussion of the use of critical loss analysis in antitrust proceedings.

the monopolist would suffer if it did increase prices on all its products by $X\%$. With symmetric one-sided firms, the well-known formulas are given by $CL = X/(X + M)$ and $AL = X(\epsilon^{OWN} - \epsilon^{CROSS})$, where M is the percentage markup, and ϵ^{OWN} and ϵ^{CROSS} are the own and cross price elasticities respectively. A relevant market is found when Actual Loss equals Critical Loss for a hypothetical monopolist of the given set of products in the proposed antitrust market. If Actual Loss exceeds Critical Loss, the relevant market is expanded to include more substitutes. Otherwise, it is contracted.

Critical Loss Analysis has also been used to estimate potential unilateral price effects from proposed mergers. In this context, one inquires whether the merger of the parties would lead the merged firm to raise prices by $X\%$ or more. If the test shows that Actual Loss exceeds Critical Loss, the merged entity would not find an $X\%$ price increase in all its products profitable. Otherwise, it would find it profitable. Used in this way, it provides a first look at the potential for unilateral price effects, although more sophisticated simulation techniques that take into account the intricacies of the industry and businesses involved would ideally be used to confirm this initial finding.

The one-sided calculations are trivial, and the technique has won significant appeal both because of its simplicity and because of its easy measurement of inputs.⁷

B. Two-Sided Critical Loss

We now show the one-sided formulas given above are analytically wrong whenever they are applied to markets involving MSPs. The reason is that when an MSP increases the price to customers on one side of its business, it results not only in a loss of customers on that side but also in a loss of customers on the other side of the market. This in turn causes a shift in the relative and absolute sizes of all platforms in

⁷ One could also calculate and compare profit totals directly across hypothetical and real worlds. The sum of firm profits before is $\pi = 2Mpq$, and the joint profit after is $\pi' = 2pq(M+X)(1+X(\epsilon^{OWN} - \epsilon^{CROSS}))$. We, however, maintain the Critical Loss formulation throughout the paper because of its appeal, its transparency (over black box profit comparisons), and its tangible metric – namely, how many sales a firm can afford to lose when it raises price.

the market, giving rise to further implications. The direction of the bias from using one-sided formulas depends upon the analytical method being used and the nature of the differences between the MSPs, but can be non-trivial as a general rule.

For our exposition of the biases from the use of one-sided formulas when MSPs are involved, we consider the case of two symmetric platforms serving two distinct customer groups. Symmetry implies that each elasticity and indirect network externality are the same across the two platforms. While less commonly encountered in practice, the symmetric model best conveys the key insights in the analysis. We present the more complex equations for the case of asymmetric platforms (which also includes the case of platforms that do not have all sides in common) and the case of MSPs with $N > 2$ in detail in Appendix 1.A. There are two opposite biases that can occur—“Estimation Bias” and “Lerner Bias”—depending on the estimation technique. In the symmetric case, the direction of each bias is unambiguous.

Estimation Bias. Suppose the analyst estimates a demand system for one of the products offered by two-sided platforms using data and following the techniques that give an unbiased estimate of the short-run own-price elasticity of demand. Here we define short-run to represent the length of time it takes for customers who experience a price increase directly to respond, but before any feedback effects commence. (The feedback effects, we know, will cause additional indirect responses over time as relative and absolute platform sizes change.⁸) Since the analyst’s estimate does not account for feedback effects, the full impact of the price increase on demand is underestimated. As a result, antitrust markets necessarily will be defined too narrowly, and merger analysis will overstate the increase in market power of merging parties and overstate the predicted unilateral price effects of the transaction.

Lerner Bias. An opposite bias can occur when the analyst uses observed markups to calibrate the own elasticity of demand based on the one-sided Lerner

⁸ In practice, most likely the analyst will estimate some unknown mix of a short-run and long-run elasticity (i.e. once all feedback effects have been worked out), since indirect network effects from platform size changes are not properly accounted for. The direction of bias is the same.

Index. This bias overstates the true short-run own-price elasticity of demand. In fact, it yields an elasticity estimate that is even larger than the true long-run own-price elasticity of demand. We define long-run to be as long as necessary for all meaningful feedback effects to be worked out, from the consequent changes in platform size.⁹ In this case, the resulting market definition would be too broad and predicted unilateral effects of a transaction would be too small.

The situation is more complex with asymmetric platforms. When the test considers equal price changes (e.g. $X^i = 5\%$) across all of the hypothetical monopolist's products, the biases are signed as above. However, if the asymmetry is significant *and* if the analyst considers differential price changes across products, the direction of the overall bias depends on the analytical method and also the relative and absolute sizes of the price increase(s) considered. Differential price changes are more likely to be a consideration in the analysis of unilateral effects, where antitrust concern may center around particular products, than in market definition analyses where the price increases deemed significant are generally taken to be uniform (e.g. 5% or 10%), as suggested by the DOJ/FTC merger guidelines. These additional complications become pertinent in our analysis of the Google-DoubleClick merger in Section V, below.

We now return to the symmetric case. Call the platforms firm 1 and firm 2, and the two groups of customers A and B . Note that platform symmetry implies $P^s = P_1^s = P_2^s$ for $s = A, B$, but even with "symmetric platforms" we are not assuming symmetry across the two sides, $P^A \neq P^B$. The sides are typically quite different in practice. Also, define total quantity across platforms on side s as $Q^s = Q_1^s + Q_2^s$.

Imagine that a hypothetical monopolist of both two-sided platforms increased prices to the consumers on side s by ΔP^s and to consumers on side r by ΔP^r , where $r \neq s$. Similar to the one-sided case, the gain on inframarginal sales due to the price increase on side s would be:

⁹ Our definitions of "short" and "long" are defined to be the periods in which the direct price responses and meaningful indirect price effects take place, respectively. They do not relate to calendar time and are different from the concepts of short- and long-run (+/- one or two years) generally used by antitrust practitioners.

$$\Delta P^s (Q^s + \Delta Q^s) \quad (3)$$

However, unlike the one-sided case, the marginal loss due to the price increase on side s now consists of two components:

$$-(P^s - C^s)\Delta Q^s - (P^r - C^r)\Delta Q^r \quad (4)$$

the loss on side s and the loss on side r due to the price increase on s .¹⁰ Here, marginal costs are denoted by C^s and C^r , and ΔQ^s represents the loss in total quantity on side s over the time period that is relevant for the merger review (typically one or two years). Summing up and equating the gains to the losses, we get the Two-Sided Critical Loss Formula in response to changes in the prices on side A and B of $X^A\%$ and $X^B\%$ respectively:

$$\sum_{s=A,B} \left[R^s (X^s + M^s) \left(\frac{\Delta Q^s}{Q^s} \right) + R^s X^s \right] = 0 \quad (5)$$

where $R^s = Q^s P^s$, the revenue earned from side s . Critical Loss is the set of percentage quantity reductions on side A , $\Delta Q^A / Q^A$, and side B , $\Delta Q^B / Q^B$, that would leave the hypothetical monopolist profits unchanged.¹¹

The Actual Loss to the hypothetical MSP monopolist depends on several factors. As with one-sided firms, Actual Loss depends on the (short-run) own-price and cross-price elasticities of demand. Higher own price elasticities tend to increase Actual Loss since a price increase at a given platform results in relatively more customers switching away from that platform. On the other hand, higher (short-run) cross price elasticities between the monopolist's platforms tend to decrease Actual

¹⁰ The merger guidelines state that when determining the relevant antitrust market "the terms of sale of all other products [outside the proposed antitrust market] are held constant". In general, we would expect a hypothetical monopolist to adjust prices on side B as well as on side A . This and subsequent formulas account for both cases.

¹¹ A special case is when the two sides are tied together in a fixed proportion, e.g. a transaction market like credit card services, where a transaction takes place between a customer on side A (a card holder) and a customer on side B (a merchant). This is the case considered by Rochet and Tirole (2003) and Emch and Thompson (2006) among others. Suppose the proposed antitrust market is the transaction, which by definition includes both sides. Given $Q = Q^A = Q^B$ and $X^A = X^B = 5\%$, equation (8) collapses to: $|\Delta Q / Q| = X / (X + M^{AB})$ where $M^{AB} = (P^A + P^B - C^A - C^B) / (P^A + P^B)$. This is the familiar one-sided Critical Loss formula where the "product" is the transaction with a composite price of $p^A + p^B$.

Loss because relatively more of the customers that switched away from one platform can be recaptured with the monopolist's other platform.

Unlike one-sided firms, Actual Loss also depends on the strength of the indirect network externalities that customer groups provide to one another. Recall that an MSP is only successful because it is able to bring two distinct customer groups together in significant numbers. When an MSP increases its price on side s for platform i , there is the usual contraction in demand on side s , as with one-sided firms. But now because there are fewer side s customers, the platform is less valuable to side r customers. This causes a contraction on side r as well. The feedback effects take over, causing another contraction of side s , then side r again, and so on. The stronger the externality across groups, the greater the demand contractions will be after a price increase and the greater is the Actual Loss, all else equal.

The exact Actual Loss formula depends on the specific demand form chosen. We consider, as an example, an isoelastic demand function adapted to include the special features of MSPs.¹² Let

$$q_i^s = \alpha_s q_i^r - \delta_s q_j^r + \theta_i^s \text{ where } \theta_i^s = \mu_s - \beta_s p_i^s + \gamma_i^s p_j^s \quad (6)$$

$i \neq j, s \neq r$, where $q_i^s = \ln(Q_i^s)$ is the log quantity demanded by side s customers at platform i , $p_i^s = \ln(P_i^s)$ is log price to side s on platform i , and parameters α_s , β_s , and δ_s are non-negative.¹³ Since each platform is offering a competing service to the other platform on each side of its business, γ_s is non-negative as well. Given the log-log form, the β s and γ s are the usual short-run own and cross price elasticities for side s customers respectively. Recall that we define short-run to include the time needed for the direct response of customers impacted by the price increase but too short to include any of the indirect network feedback effects that follow. The indirect network externalities (the marginal value one side puts on the presence of the other at a platform) are captured by α_s , the within-platform cross-side externality, and δ_s , the

¹² The biases described here are not exclusive to the particular demand function chosen.

¹³ For expositional purposes, we take the indirect network effects to be positive across the two sides in the main text, though this needs not hold in general.

cross-platform cross-side externality.¹⁴ The α_s , for example, is the percentage change in quantity demanded by side s at platform i in response to a one percent change in the quantity of platform i 's side r customers. We assume $\alpha_s > \delta_s$, $\beta_s > \gamma_s$, and $\alpha_s + \delta_s < 1$ to ensure stability of the system.¹⁵

Solving for the reduced form equations of the q_i^s 's (as a function only of prices), and taking the total derivative of each q_i^s , we derive the formula for Actual Loss to side s customers after price increases to side A and side B of $X^A\%$ and $X^B\%$ respectively:

$$L^s = \frac{(\gamma_s - \beta_s)X^s + (\alpha_s - \delta_s)(\gamma_r - \beta_r)X^r}{1 - (\alpha_s - \delta_s)(\alpha_r - \delta_r)} \quad (7)$$

L^s is the percentage Actual Loss on side s at each platform and at both platforms collectively, $L^s = \Delta q_i^A = \Delta \ln Q^A = \Delta Q^A / Q^A$ (since platforms are symmetric), once all feedback effects have been worked out.

We see the two-sided Actual Loss formula for side s is a generalization of the one-sided formula. When all cross-side externalities are zero, so that each group's demands do not depend on the numbers of customers on the other side at either platform, the formulas for L^A and L^B coincide with their one-sided counterparts. When two-sided externalities are present, however, it is necessarily the case with symmetric platforms that the correct two-sided formulas give larger Actual Loss estimates than the one-sided versions, given otherwise unbiased estimates of the short-run price elasticities.

There are two reasons for the "Estimation Bias" here. First, an increase in price on side A alone causes not only a reduction in demand on side A but also a reduction in demand on side B , since the B group values a platform less when there

¹⁴ The same formulas apply whether firms multi-home or single-home. While it does not change the structure of the demand system, it would be expected to impact estimates of the parameters.

¹⁵ The first condition ensures demand on the other side of both platforms does not increase after a price increase and demand contraction at both platforms on the first side. The second ensures demand does not increase on both sides after a price increase on one side by both platforms. The third prevents exploding demand following a price increase on one side by one platform.

are fewer A 's. So, unlike the one-sided case, there are losses on both sides ($L^A < 0$ and $L^B < 0$).

Second, and importantly, there is a multiplier effect in two-sided industries that magnifies the immediate loss on side A and on side B over time. As discussed above, when side A contracts after its price increases, the platform is less valuable to its B customers and the B side contracts. But fewer B 's decreases the value of the platform to the A customers and now the A side contracts further. Both sides contract in turn, and Actual Losses are larger than they would be in a one-sided world where absolute and relative platform size does not matter.

The multiplier effect can be seen in the denominator of the two-sided Actual Loss formulas. The denominator is necessarily less than one, so the two-sided Actual Loss formulas necessarily yield larger estimates of loss than the one-sided formulas, for a given set of parameters. The biased one-sided versions overestimate the hypothetical monopolist's ability to raise prices.

To calculate whether a given price increase or increases are likely to be profitable, one simply compares Critical Loss and Actual Loss, by substituting the values of L^s into $\Delta Q^s / Q^s$ in the Critical Loss equation (5). If the left hand side of (5) is negative, the price increase(s) will not be profitable. If positive, the price increase(s) will be profitable.

In merger analyses, the test would be performed using just the set of products or services that the newly merged firm would control. The merged entity is *de facto* the hypothetical monopolist and assumptions about possible cost efficiencies post-merger are easily worked into the analysis. In market definition exercises this test is repeated many times, each time expanding or contracting the proposed antitrust market, until a market is found where Critical Loss would equal Actual Loss for a hypothetical monopolist of all the products in that market. Given a market definition, the market shares of the "in" firms are calculated and then used as proxies for market power.

While presenting these corrected formulas, we continue to caution that there are many reasons to be wary of mechanical market definition exercises such as the SSNIP test and these concerns do not go away when MSPs are involved. Some may even be exacerbated by their presence.

For example, should one include both sides of the business of an MSP in the market definition or just one side? The answer can be difficult and depends on the situation at hand. The merger guidelines, in discussing the SSNIP test, refer to adding or removing “substitute products”. The difficulty is that the sides of a given MSP represent highly complementary products, and pricing decisions on both sides critically affect the MSP’s profitability. Yet the SSNIP test is silent on handling complementary products¹⁶. We believe if the two sides are very highly complementary and closely linked – for example, if the MSPs facilitates transactions between the groups that occur in fixed proportions – and MSPs in an industry all tend to serve the same two sides, then it can be reasonable to include both sides in the market definition and the “transaction” as the product. However, in other industries MSPs may all cater to the same side *A* customers but cater to very different kinds of side *B* customers. If the antitrust concern centered around the side *A* business, including both sides of all MSPs that share the *A* side in the market definition would open a Pandora’s box of unrelated “*B*” types that make no sense under a single coherent market definition. Then the market may need to be defined on the basis of side *A* only, but with the critical understanding that the *B* sides are an important constraint on behavior and the formulas presented here must be used to account for its influence.

A second issue with the SSNIP test is that we would expect a hypothetical monopolist of several platforms to reoptimize prices across sides and across platforms just as we would expect a hypothetical monopolist to reoptimize across products in a one-sided world. The DOJ/FTC merger guidelines suggest that the price increase deemed significant for market definition purposes will generally be uniform

¹⁶ Weisman (200r) also notes that the presence of complementary goods, if ignored by the analyst, can lead to an overestimation of equilibrium markups and of market power.

across all products within the definition (e.g. 5%), and “the terms of sale of all other products [outside the proposed antitrust market] are held constant” (i.e. 0% price increase). In the case of MSPs, one could instead imagine price increases that differ across sides or platforms, or consider a kind of 5% quantity-weighted “average” price increase across sides, allowing the hypothetical platform monopolist to reoptimize relative prices.

Settling these issues, which are also concerns with one-sided SSNIP tests, lies beyond the scope of our paper. Our goal here is simply to show that the one-sided Critical Loss formulas for conducting the SSNIP test are wrong when applied to two-sided markets, and also to provide a useable framework for how to conduct a SSNIP test when MSPs are involved. Whether to include one or both sides, and how to choose X^A and X^B , will depend on the situation. The equations we present—which allow for differential X^i 's—handle all these situations.

As an alternative to the SSNIP test, one could instead perform a full-scale simulation allowing full relative and absolute reoptimization of prices by the hypothetical monopolist, and such an exercise is helpful in an advanced analysis. However, the Critical Loss implementation remains a simple and useful first pass, in spite of its drawbacks, and it is important, if one wishes to apply it to MSPs, to use the correct formulas.

With the SSNIP test, as in any estimation, it is generally preferable to estimate all the parameters of the demand system whenever reliable data exist.¹⁷ In practice, data availability is often such that at least some parameters must be calibrated with limited information. Theoretical restrictions, like the Lerner Index

¹⁷ There are several empirical papers that estimate two-sided demand systems in particular industries in order to establish the existence of indirect network effects. Although generally not in the context of merger or market definition analysis, these highlight the importance and challenges of proper demand identification in a two-sided context. For example, see Rysman (2004) establishing a link between readers and advertisers in the yellow pages market, Rysman (2007) for cardholders and merchants in the payment card industry, Argentesi and Filistrucchi (2005) for readers and advertisers in the Italian magazine industry with an estimation of market power, Wright (2004) and Kaiser and Wright (2006) for linking readers and advertisers in the German magazine industry and Dubois, Hernandez-Perez, and Ivaldi (2007) linking readers and authors in the academic publishing industry. Argentesi and Ivaldi (2005) recap antitrust cases involving market definition in (two-sided) media industries. Using limited data for the French magazine market, they include an estimation of subscriber demand elasticities using the price of advertising as an instrument, suggesting a two-sided linkage. Elasticity estimates differ with and without the instrument.

that relate markups and elasticities, can be used to calibrate parameters that otherwise may be difficult or impossible to estimate. Other theoretical restrictions, such as the Slutsky Symmetry rule that relates cross-price elasticities, can be used to ensure logical consistency across certain parameters.

The relationship between markups and elasticities is especially important here because markups appear in the Critical Loss formula and the elasticities appear in the Actual Loss formulas. This relationship should be consistent with economic theory.¹⁸ In calibration exercises, it is now common practice to use the (one-sided) Lerner Index to estimate the own price elasticity from the observed markup when the latter cannot be estimated from data, perhaps because data are lacking or the analyst has insufficient time. This turns out to be a problem in two-sided settings. The one-sided Lerner Index is incorrect for MSPs and relying on it significantly overestimates the true short-run own price elasticity of demand and overstates Actual Loss.

It is well known in the one-sided case that the percentage markup equals the inverse of the own price elasticity at the profit maximizing output. In the two-sided case, we can derive the first order conditions for profit maximization. Each platform i chooses prices on each side to maximize:

$$\max_{p_i^A, p_i^B} \Pi_i = \sum_{s=A,B} q_i^s(\vec{p})^* (p_i^s - c_i^s) \quad (8)$$

The two first order conditions simplify to:

$$R_i^s M_i^s \epsilon_{ii}^{ss} + R_i^r M_i^r \epsilon_{ii}^{sr} = R_i^s \quad (9)$$

for $s = \{A,B\}$, where:

$$\epsilon_{ii}^{ss} = \frac{(\alpha_A \delta_B + \alpha_B \delta_A) \gamma_s + (1 - \alpha_A \alpha_B - \delta_A \delta_B) \beta_s}{(1 - (\alpha_A - \delta_A)(\alpha_B - \delta_B))(1 - (\alpha_A + \delta_A)(\alpha_B + \delta_B))}$$

¹⁸ Some authors disregard the Lerner Index as too simple and unrealistic for application in specific real-world antitrust cases (see Scheffman and Simons (2003) and Harris (2003)). We agree with Katz and Shapiro (2003) that there should be a presumption of the theoretical markup-elasticity relationship pending evidence to the contrary. It is possible, however, to estimate both markups and elasticities independently of one another directly from the data and sidestep these issues.

$$\varepsilon_{ii}^{sr} = \frac{(\alpha_A \delta_B + \alpha_B \delta_A)(\alpha_r \gamma_s + \beta_s \delta_r) + (1 - \alpha_A \alpha_B - \delta_A \delta_B)(\alpha_r \beta_s + \delta_r \gamma_s)}{(1 - (\alpha_A - \delta_A)(\alpha_B - \delta_B))(1 - (\alpha_A + \delta_A)(\alpha_B + \delta_B))}$$

We call ε_{ii}^{ss} and ε_{ii}^{sr} the long-run own-price elasticity and the long-run cross-side price elasticity, in absolute value, once all feedback effects have been worked out. The i subscripts on the ε 's indicate these are elasticities with respect to changes in the price of the side s good *by firm i only*.

The first order conditions do not reproduce the one-sided Lerner Index formula. They do not even uniquely identify the values of the own price elasticities from the markups, but rather only constrain the relationship between the own and cross price elasticities and the indirect network externalities.

Using the one-sided Lerner Index to calibrate own price elasticities when MSPs are involved creates a “Lerner Bias”. To see the Lerner Bias most easily, imagine for a moment symmetry across sides $\alpha_A = \alpha_B$ and $\delta_A = \delta_B$. The markup equation for side s then simplifies to $M_i^S = (\varepsilon_{ii}^{ss} + \varepsilon_{ii}^{sr})^{-1}$. Now note that the long-run own price elasticity, ε_{ii}^{ss} , is necessarily larger than β (the usual short-run own price elasticity), due to the indirect network effects. The long-run cross-side price elasticity, ε_{ii}^{sr} , which is zero in a truly one-sided industry, is greater than zero in absolute value here.¹⁹ Consequently, for a given set of elasticity estimates, equilibrium markups are lower than the one-sided Lerner Index would imply. The converse is that for a given set of markups observed in the data, the true short-run own-price elasticities of demand are lower than the one-sided Lerner Index would suggest. The reason is that the indirect network externalities penalize price increases more, so the short-run own price elasticity must be especially low in order to support markups at a given level.

¹⁹ Recall this is the long run change in quantity on side r at platform i due to a change in the price of side s at platform i , once all feedback effects have been worked out. There is no direct effect of side s prices on side r demand, but there is an indirect effect of side s prices on side r demand for two-sided platforms, via changes in side s demand.

The bias is strong – since $\varepsilon_{ii}^{sr} > 0$, the one-sided estimate is even larger than the true long-run own-price elasticity, and this results in a bias in the opposite direction. Market definitions will be set too large, or the expected price increases from a merger will be assumed too small.

The comparative statics of the two-sided markup equations give insight into the problem. As with the one-sided calculation, a higher short-run own elasticity β (which increased both ε_{ii}^{ss} and ε_{ii}^{sr}) lowers markups. But unlike the one-sided case, a higher short-run cross price elasticity of demand γ also results in lower markups (again increasing ε_{ii}^{ss} and ε_{ii}^{sr}). The reason is that a higher cross price elasticity causes a higher proportion of the side A customers who switch after a side A price increase to buy from platform j rather than buy nothing instead. Since platforms are two-sided, the relative value of firm i 's platform to the side B customers falls when more A 's move to j , causing a demand contraction on the B side as well. The feedback effects begin, and platform i contracts further.

Greater indirect network externalities also work to decrease markups. First, by raising price a platform triggers the market shrinking feedback effect causing its demand to repeatedly contract on each side. It is easiest to see this by noting $M_i^s = (1 - \alpha)/\beta$ when $\delta = 0$ in the symmetric case. A greater cross-side own-platform externality α penalizes price increases more and acts to lower equilibrium markups.

Second, and surprisingly, a higher cross-side cross-platform externality δ also acts to lower markups. To see this, note that $M_i^s = (1 - (\alpha + \delta))/\beta$ as $\gamma \rightarrow \beta$ in the limit in the symmetric case. The parameters α and δ now appear as a sum, rather than a difference (as they did in the Actual Loss formulas). Only platform i is changing price, not the hypothetical monopolist of both products. A higher δ means that in response to the shrinking of side A at platform i after its i raises p^A , more side B customers switch to platform j rather than to nothing at all. This makes j more attractive to side A , then more attractive to side B , then side A and so on. Thus, higher

values of both α and δ penalize price increases more and lead to lower markups in equilibrium.

Overall, for a given set of short-run own-price and cross-price elasticities, two-sided pressures lower equilibrium markups below that which the one-sided Lerner Index would predict. One-sided estimates of short-run own-price elasticity derived from markups overstate the true elasticity and understate the profitability of potential price effects.

We now turn to a demonstration of the potential “Estimation” and “Lerner” Biases from using one-sided formulas in situations involving symmetric MSPs.

IV. COMPARISON OF ONE AND TWO-SIDED CRITICAL LOSS: SOME SIMULATIONS

Assume that two platforms each serve sides A and B , and as an example, assume a relevant antitrust market is proposed that would include the products on side A of these two platforms. The test would be the same if we were considering a merger between the two platforms and wanted to know if a given price increase on side A would be profitable for the merged entity.²⁰ We maintain the assumption of symmetric platforms (though not necessarily symmetric sides).

It is useful to rearrange the Critical Loss equation (5), substituting in L^A and L^B for $\Delta Q^A / Q^A$ and $\Delta Q^B / Q^B$, to get:

$$\omega \frac{X^A + M^A}{X^A} L^A + (1 - \omega) \frac{X^B + M^B}{X^B} L^B = -1 \quad (10)$$

which holds when Critical Loss equals Actual Loss exactly. It is clear that the equation is really a weighted sum of the ratios of Actual Loss to Critical Loss on each side, where the weight is given by

$$\omega = \frac{R^A X^A}{R^A X^A + R^B X^B}$$

²⁰ If the market were proposed to include both sides A and B , we would set $X^A = X^B = 5\%$ and the same qualitative results still hold.

We refer to the left hand side of this equation, in absolute value, as the Actual-Critical Loss Ratio (ACR). If the ACR is greater than 1, Actual Loss exceeds Critical Loss, and the price increase would not be profitable. The relevant market would be expanded. If ACR is less than 1, the price increase is more than profitable and the relevant market contracted. The ACR represents a measure of the closeness of the Actual and Critical Loss.²¹ Note that if the industry were truly one-sided, so that $L^B = 0$ after a price increase on side A , this equation reduces to the ratio of Actual Loss to Critical Loss that would be applicable with one-sided firms.

First, assume that the short-run own-price elasticities are well estimated and the analyst does not depend on the (one-sided) Lerner Index. We see the one-sided calculations are biased towards a lower ACR, and the bias grows as the indirect network externalities become more important. Figure 1 plots the ACR against the own-price elasticity β for several different values of the cross-side own-platform externality α . Recall the parameter α represents the short run percentage change in a platform's side A business due to a 1% change in demand on its side B business, and vice versa. The figure includes both the ACR calculated using the one-sided formula (which assumes α and δ are both zero), and also using the correct two-sided formula for MSPs. Other parameters are held constant, and the price increase to side s is taken to be 5%.²²

First, and as we would expect, the ACR increases as the short-run own-price elasticity increases in both calculations for a given α . There is an indirect effect of higher β through markups that increases Critical Loss, but the direct effect of β that increases Actual Loss dominates, and the ACR rises with β .

The calculated effect of the own-platform cross-side externality α on the ACR, however, depends on the calculation being used. In the two-sided calculation, a higher α is associated with stronger feedback effects, which result in greater Actual Loss and greater ACR following a price increase, all else equal. In contrast, the ACR

²¹ The derivation assumes $X^A \neq 0$ and $X^B \neq 0$. When $X^B = 0$, the second term reduces to $(R^B M^B / R^A X^A) L^B$.

²² The assumption of 5% is often used but arbitrary.

falls under the one-sided calculation. This is because the Actual Loss calculated with the one-sided formulas $((\beta_A - \gamma_A)X_A)$ is independent of α , but the Critical Loss is greater since observed equilibrium markups are lower with higher α . The gap between the ACRs under the two calculations grow with higher α , and with higher β .

Figure 2 plots the “critical beta” under each calculation against the cross-side own-platform externality α . The critical beta is the value of the own price elasticity, conditional on the other parameters, that would cause Actual Loss to exactly equal Critical Loss. The area between the one and two-sided critical β s is the region of error: the combinations of α and β for which a contraction of the proposed antitrust market would be called for under the one-sided calculation when an expansion of the proposed market should be called instead. In merger analyses, these would be the combinations of α and β for which the analyst would conclude from the one-sided calculation that the hypothetical monopolist in a market definition analysis or the merged entity in the unilateral effects analysis would significantly increase prices when in fact it could not.

The figure shows that as the indirect network effects become stronger, the region of error grows wider. The corresponding figure for the “critical gamma”, with β held fixed (not shown), yields the same conclusion.

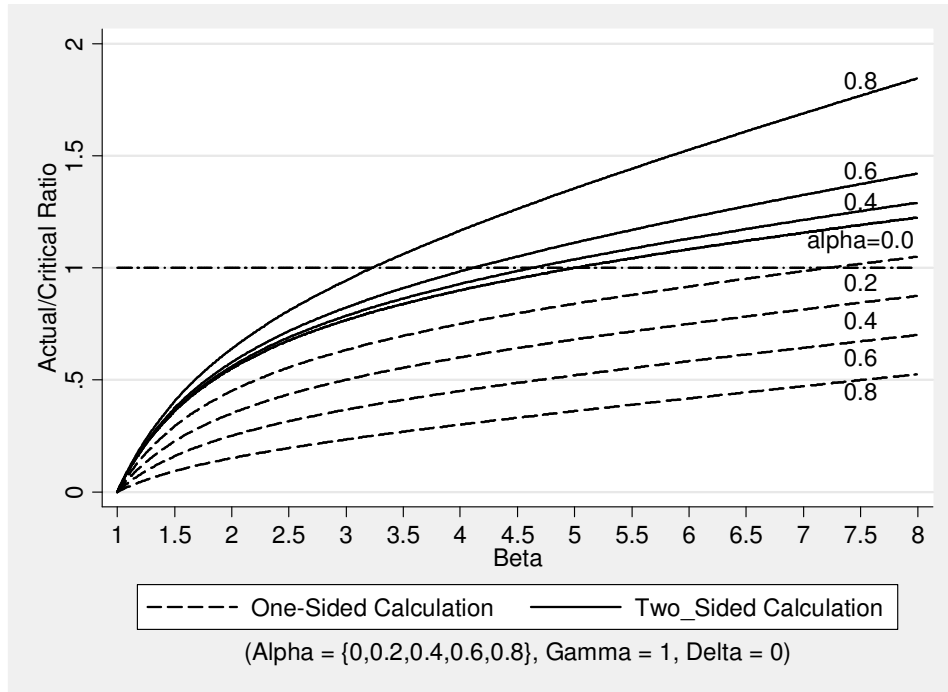


Figure 1. Actual/Critical Loss Ratios, One- and Two-Sided Calculations

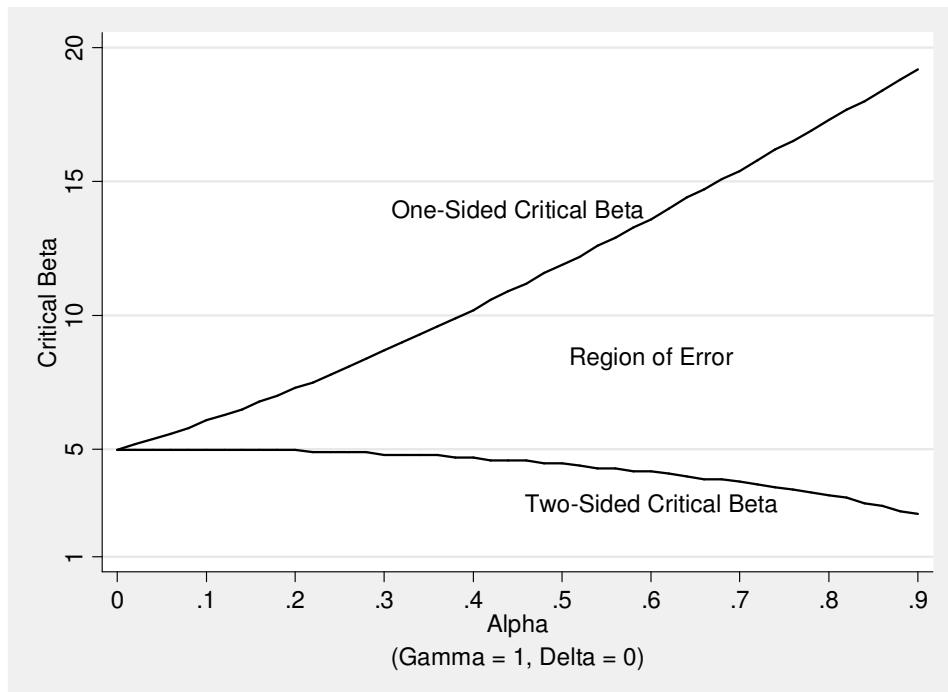


Figure 2. Region of Error, One- and Two-Sided Calculations

In these examples, it is assumed that the short-run elasticities are properly estimated; only the formulas used differ. However, if the analyst relies on the one-sided Lerner Index to estimate the own price elasticity of demand, and then uses the usual one-sided formulas, the bias goes in the other direction. The analyst will overestimate the short-run and long-run own price elasticities of the MSPs and will overestimate Actual Loss. The one-sided Critical Loss calculation is unchanged (because it is based on observed markups) and so the ACR will be erroneously high. In other words, profitable price increases will not be expected for even relatively low values of the true short-run own-price elasticity β . The critical betas and the region of error are plotted in Figure 3, and this time the critical betas are lower under the one-sided calculation than under the two-sided one. In the region of error, the analyst who relies on the one-sided calculations would recommend an expansion of the antitrust market when a contraction should be called for, or would conclude that merging platforms could not profitably raise prices after the merger when in fact they could.

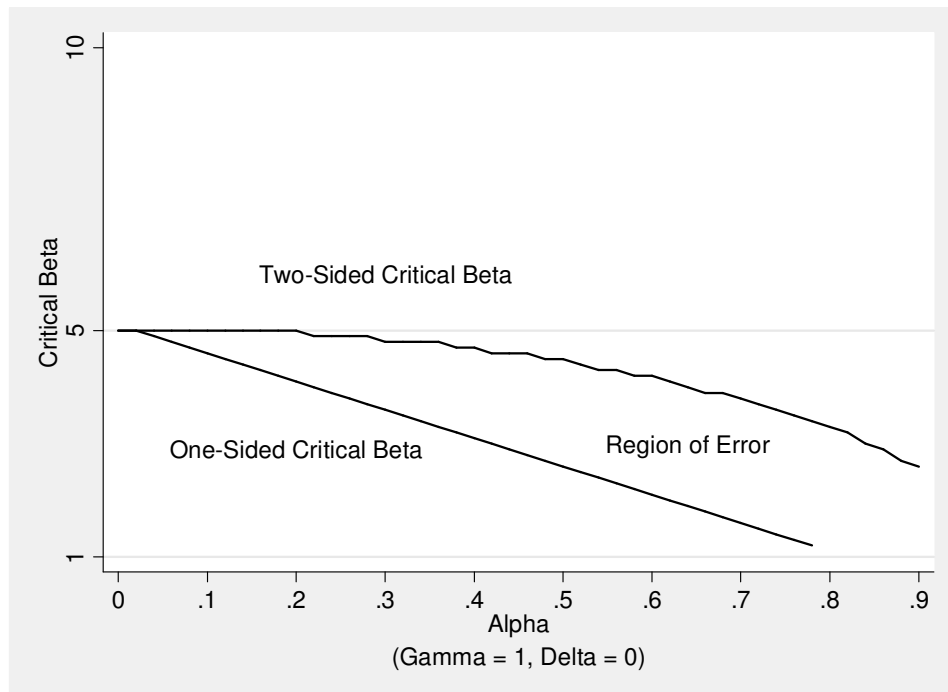


Figure 3. Region of Error, Mismeasured β in One-Sided Calculations

The asymmetric platform case, given by formulas set out in Appendix 1.A, is more involved. When platforms are asymmetric but the price increases considered on any side are equal across platforms, the biases move in the same directions as in the symmetric case. However, when platforms are asymmetric *and* price increases considered are also allowed to differ across products, the direction of the bias also depends on the degree of asymmetry and the particular mix of price increases. The reason differential price increases are important is because they result in a shift in the relative sizes of the platforms and generate feedback effects. If the platform is asymmetric enough, the impact of certain price increases will be overstated, and others understated.²³

As an example, consider a proposed merger where the platforms differ significantly in initial profitability. A price increase at the low-margin platform will be beneficial for the hypothetical monopolist if network effects are strong and doing so convinces enough customers to switch to its high-margin platform. One-sided calculations do not account for this possibility, and as a result, understate the profitability of this price increase. If we consider a price change at only the high-margin platform, we can get an opposite bias. The profitability of the price increase would be overstated by a one-sided calculation as long as the gain from expansion of the low-margin platform cannot compensate for the contraction of demand at the high-margin platform. Thus, when asymmetric platforms and asymmetric price changes are considered, the direction of the biases depends on the specific question being asked as well as the method used.

²³ Across all four potential price changes of a given equal amount, the “average” bias from those price changes collectively goes in the same direction as that in the symmetric case, which in turn depends on the analytical method used. However, in the asymmetric case there is variation in the size of these biases and with sufficiently asymmetric firms, some may have the reverse sign. Because differential price changes act as weights in Actual Loss (equation A.4), the overall bias can go in the opposite direction of the symmetric case if the reverse-signed biases are sufficiently highly weighted.

V. ANALYSIS OF UNILATERAL EFFECTS: AN APPLICATION TO GOOGLE'S ACQUISITION OF DOUBLECLICK

Google is a provider of online advertising services for web publishers and advertisers in addition to operating a search-engine that sells advertising space. DoubleClick is principally a provider of server-based software tools and services for managing online advertising for web publishers and advertisers. Google agreed to purchase DoubleClick for \$3.1 billion and sought clearance from several competition authorities around the world. The merger was cleared by the Federal Trade Commission, the European Commission, and other competition authorities after lengthy investigations. This section uses data on margins for Google and DoubleClick along with other assumptions on their competitive relationships to show how the techniques used above can be employed to estimate the unilateral effects from a merger of MSPs. The data and calculations are illustrative. We take no position on whether the merger should or should not have been cleared by the competition authorities; these authorities had access to more reliable data than what are used here and addressed fuller range of issues in their assessments.²⁴ Our purpose is only to show how the tools developed in the previous sections can be used in practice.

Web publishers make space available on their web pages for advertisements. Their ad inventory corresponds roughly to the number of people that will view each space over some given time period.²⁵ Advertisers are on the other side of the market. They buy advertising inventory to reach consumers. Like all sellers and buyers, publishers and advertisers require ways to identify optimal trading opportunities and to establish transaction prices. This “intermediation” occurs *directly* through bilateral exchanges between publishers and advertisers and *indirectly* through multilateral exchanges between publishers and advertisers using advertising

²⁴ The authors were consultants to Microsoft and authored submissions to the competition authorities that assessed the possible competitive effects of the acquisition based on information that was publicly available.

²⁵ Advertisers value the demographic and other characteristics of viewers in addition to their sheer number. Thus views by high-income residents of Boston are more valuable to the local BMW dealer than views by low-income residents or high-income residents of Topeka.

networks. Larger publishers use more direct methods while smaller publishers use more indirect methods.

This section concentrates on large web publishers, which account for the preponderance of advertising revenue and large advertisers, which account for the preponderance of online spending.²⁶ A significant part of the advertising inventory bought and sold by these large advertisers and publishers involves bilateral exchange. Large publishers either have direct sales forces or hire third-party sales reps to sell their ad inventory. Likewise, large advertisers have purchasing agents or, more often, use media buyers at their advertising agency to purchase ad inventory. Advertising inventory sold this way is said to be “reserved”. Large publishers often sell their “premium ad space” this way.

Large publishers often also rely on other intermediaries to sell ad inventory indirectly that they have not “reserved” for advertisers directly. These intermediaries are called “advertising networks”. Publishers may use ad networks because they are more efficient than a direct sales force for some, or all, of their ad inventory; or because they have excess inventory that they have not sold directly, perhaps because of spikes in viewers. Advertisers use ad networks because it is another way to reach viewers. Advertising inventory sold indirectly costs less on average than advertising inventory sold directly because the space is less desirable than the directly sold space.

Advertisers and publishers require management, reporting, and technology solutions such as those offered by DoubleClick and aQuantive. These tend to be server-based software that can help manage advertising inventory and campaigns that may involve millions of ad impressions (that is, views of an ad by an individual) a day. These server-based software tools are highly sophisticated mission-critical applications

²⁶ A large publisher refers to one which is sufficiently large to use stand-alone tools to serve their directly sold ads and any remnant ads sold through non-integrated and/or integrated ad networks. This includes at least the top 500 publishers, and accounts for the majority of advertising revenues. Smaller publishers do not find it economically efficient to hire direct sales forces. They usually rely on an ad network to sell their ad inventory.

Large publishers usually use a publisher tool such as DoubleClick's Dart for Publishers (DFP).^{27,28} This tool is typically hosted on a web server maintained by the provider. The publisher hardcodes links to the publisher tool to fill the ad space for which it wants to use the management, reporting, and serving capabilities of the publisher tool. It also typically integrates the publisher tool into many other aspects of the website technology and business practices. Publishers will typically only use one publisher tool (i.e., they "single-home" on tools). Surveys of large websites show that DFP is used by somewhat more than half of these websites. Larger sites are more likely than smaller sites to have DFP. DFP has about 63 percent share based on page views.²⁹

Large advertisers and advertising agencies often have an advertiser tool such as DoubleClick's Dart for Advertisers (DFA).³⁰ Large advertisers typically manage advertising on hundreds of websites and across numerous products using many methods of online advertising. This tool helps them manage these various campaigns. Advertisers usually use one advertiser tool although advertising agencies may use several. However, single-homing does appear to be the norm. DFA holds roughly a 26 percent share of non-search advertiser tools.³¹ Google is the leading integrated ad platform.

A few providers offer more or less complete solutions for advertisers and publishers. Google's AdSense/AdWords platform is one of these. Publishers can hardcode ad space to Google's AdSense³², which takes care of everything – selling the inventory, managing the ad space, serving ads to the viewer, and sending the publisher a portion of the proceeds after taking a commission.³³ Advertisers,

²⁷ See <http://www.doubleclick.com/products/dfp/index.aspx>

²⁸ A handful of mega-large publishers such as MSN have their own proprietary tools but most others use a third-party tool.

²⁹ The survey conducted by LECG shows that AdSense is used on 50% of all websites sampled. DoubleClick is used on 63% of all websites sampled. A second survey conducted by Keystone Strategy for American viewers shows that AdSense is used on 45% of surveyed sites while DoubleClick is used on 61%.

³⁰ See <http://www.doubleclick.com/products/dfa/index.aspx>

³¹ This is based on aQuantive's estimates of shares for itself and DoubleClick in advertiser tools (among providers that sell those as standalone non-integrated products), adjusted for those advertiser tools that are sold on an integrated basis (with the ads), by providers such as AdSense.

³² See https://www.google.com/adsense/login/en_US/

³³ In the online advertising Traffic Acquisitions Costs (TAC) refers to what an advertising platform pays for traffic. Google pays TAC to publishers in return for contributing their advertising inventory to the Google

likewise, can buy space from the Google Content Network³⁴ through AdWords³⁵ (which bundles Google's search-based and contextual-based advertising products). Yahoo! and Microsoft offer similar all-inclusive solutions. These solutions have all resulted from leveraging the technologies developed for search-based advertising – especially the keyword bidding auctions – to the buying and selling of publisher ad inventory. Some other ad networks are also integrated to lesser degrees; they may offer publishers serving technologies so that publishers can hardcode ad network into particular space. Many large publishers use an integrated platform for contextual ads and an unintegrated platform (based on a particular publisher tool) for non-contextual ads; the unintegrated platform is used to access multiple standalone ad networks, as well as ads sold directly. Estimates suggest that Google accounts for 51 percent of the ad revenue through the indirect channel and 27 percent overall. The unintegrated ad networks account for an estimated 45 percent of ad revenue through the indirect channel and 25 percent of web publisher ad spending overall.³⁶

Following the acquisition, Google controls an input into the unintegrated channel for selling ads directly and indirectly. It would therefore be expected to coordinate the price of DFP, and through DFP, the overall price of distribution through the direct sales and through the ad networks.³⁷ The question we consider is whether the combined entity will have the incentive to alter prices significantly and, more importantly for our purposes, how the results differ between the one-sided and two-sided calculations. Incorporating two-sided effects is important for analyzing this question because changes in the prices of the inputs alter the advertiser's demand for space as well as the publisher's demand for selling through a particular method.

Table 1 reports the data and sources we have relied on. Google and DoubleClick appear to earn roughly 75 percent gross margins based on publicly

Content Network. The publisher receives TAC. In this case 1-TAC is the percent commission paid by the publisher to Google for selling its ad inventory. If Google pays the publisher 80% of the revenue that Google receives from the advertiser the publisher has paid a commission of 20% to Google.

³⁴ See <https://adwords.google.com/select/afc.html>

³⁵ See <http://adwords.google.com/select/Login>

³⁶ Source: Keystone Strategy.

³⁷ We focus on price effects as is traditional in this sort of analysis. Google could engage in a variety of other strategies post-acquisition that could involve the exercise of acquired market power but that do not necessarily involve increasing prices.

available data. Ad inventory sells on average, with great dispersion, for about \$2 per thousand impressions (that is, views by users) based on estimates reported to us by industry observers. Google sells on a cost per click basis but obtains the equivalent of approximately 40¢ per thousand impressions based on the evidence we have seen.³⁸ Publishers typically pay 40¢ per thousand impressions on average to the standalone ad networks and 5¢ per thousand impressions to the publisher tool provider.

	DFP	DFA	AdSense	Source
Unit price (cpm)	5¢	7.5¢	40¢	Discussions with publishers. The revenue paid to Google by advertisers is estimated at around \$2 CPM of which publishers keep approximately 80%.
Percentage gross margin	75%	75%	75%	DoubleClick's gross margin according to its 2004 Form 10K was 74%. According to Google's 2006 Form 10K its costs of earning revenue (net of traffic acquisition costs paid to publishers) are around 10% of total revenues. This implies its gross margin on AdSense is 90%. Its overall gross margin (for search and non-search ad programs) has been between 50-70% in the past five years. We use 75% as reasonable.
Revenues	\$325m		\$196m	Google's 2004 revenues on third party ads after paying publishers. DoubleClick's 2004 revenues for "Tech Solutions" which appears to mainly be comprised of DART products.

Table 1. Data and Sources.

We want to investigate how the one-sided and two-sided techniques differ in their prediction of whether the merged Google-DoubleClick entity will have the ability to increase the price of DFP by $X\%$, holding all else constant. The potential

³⁸ The revenue paid by advertisers by publishers is estimated at around \$2 CPM of which publishers keep approximately 80%. See <http://www.mydigitallife.info/2006/10/21/google-adsense-giving-publishers-average-of-78-revenue-share/> and <http://www.webmasterworld.com/forum89/6913.htm>

for an increase in the price of DFP was a central point of concern with the merger. Google’s AdSense is an integrated product, and so charges higher markups than DoubleClick’s DFP. The concern was that it might be advantageous for Google to raise the price of DFP and shift publishers to its AdSense network.³⁹ In the analysis, we take X to be either 5% or 10%, and since the platforms are inherently asymmetric, we use the formulas set out in Appendix A.1. Because of asymmetric platforms and asymmetric price changes, the direction of bias from using the one-sided calculation in the two-sided setting is unknown *ex ante*.

Given the limited data available, calibration is the most feasible approach.⁴⁰

For the one-sided calculation, there are four parameters to estimate (all on the publisher side), using a combination of data and restrictions based in economic theory. In the two-sided model there are sixteen parameters total in the demand system that we need to estimate using a combination of data and economic theory. Appendix A.2 discusses the specifics of the theoretical restrictions used in calibrating the parameters under both one-sided and two-sided calculations. The formulas are in Appendix A.1, which extends the model to the case of asymmetric platforms, as is relevant here. Because we calibrate, the one-sided calculation will be subject to the Lerner bias.

Following common practice, we report the more easily interpretable “critical diversion ratios” and “critical switching levels” rather than the critical elasticities (β ’s or γ ’s themselves).⁴¹ The critical diversion ratio is defined as the fraction of the quantities lost at DoubleClick’s DFP that needs to be immediately recaptured by Google’s AdSense platform, for a price increase in DFP to be profitable once the indirect network feedback effects work themselves out. It is effectively the quantity-weighted ratio of the cross-price elasticity (of AdSense’s quantity with respect to the

³⁹ In the two-sided calculation, we attribute 2/5th of the Google unit price and revenues to the publishing side and 3/5th to the advertising side, to match the DoubleClick ratio. These are the expected side-specific prices that Google would charge had advertisers paid publishers directly for the advertising space and each paid Google a commission for the value of the tools and intermediation services provided. This avoids the need for negative prices and negative externality parameters in the calibration, which adds transparency.

⁴⁰ Given enough time a competition authority could possibly obtain sufficient data from members of the industry and through surveys to obtain estimates from demand-system estimation.

⁴¹ See Shapiro (1996).

price of DFP) to the own-price elasticity (of DFP). The Critical Switching Level is defined as the fraction of the total quantities of DFP that needs to be immediately recaptured by AdSense after a price increase for the increase to be profitable once all feedback effects are worked out. It is calculated as the AdSense-quantity-weighted cross-price elasticity to total DFP quantity.

	5% Critical Price Increase		10% Critical Price Increase	
	One-sided Calculation	Two-sided Calculation	One-sided Calculation	Two-sided Calculation
Critical diversion ratios	0.83%	0.00%	1.67%	0.01%
Critical switching levels	0.06%	0.00%	0.22%	0.00%

Table 2. Critical Diversion Ratios and Switching Levels

Table 2 reports the critical diversion ratios and critical switching levels under our best data assumptions, for both the one-sided and two-sided calculations, and for critical price increases (X) of 5% and 10%.⁴²

While all the numbers are low given the concentrated nature of the proposed merger, the critical diversion ratios and switching levels implied by the one-sided results turn out to significantly overestimate the two-sided calculations. In other words, the one-sided calculation understates the potential for unilateral price effects from the merger in this circumstance. Whereas the one-sided numbers allow for a small possibility that the price increase would not be profitable, the two-sided calculations suggest that Google is likely to raise prices on DFP even if virtually no customers switched to AdSense.

⁴² We present this analysis to demonstrate the effect of two-sided considerations on standard diversion ratio analysis. Whether and under what circumstances it is appropriate to conclude from this sort of analysis that a merger should be blocked is a much more complicated issue which we do not address in this article.

This creates a “region of error” between the critical diversion ratios. If the actual diversion ratio is below the one-sided figure and above the two-sided figure, it would result in different conclusions depending on the calculation used. The one-sided calculation would suggest profitable price increases, the two-sided calculation would not. The same is true for actual switching levels.

Generally speaking, the finding of profitable price increases with an effectively zero cross-price elasticity in the two-sided calculation might be surprising. Merging firms in a one-sided world only raise price above its formerly profit-maximizing level if doing so directly causes customers to switch to its other products, i.e. a large enough cross-price elasticity. The reason why zero switching levels are plausible in the two-sided case is because of the indirect network externalities. By raising the price of DFP, DoubleClick loses publishers and, via the feedback effects, also loses advertisers, then more publishers, more advertisers and so on. As the DoubleClick platform shrinks and its value deteriorates, advertisers and publishers switch to other platforms, including Google’s AdSense.⁴³ It turns out that the increase in profits to Google from this switching, which is driven by changes in platform size rather than changes in relative prices directly, makes the price increase on DFP profitable. Hence, Google benefits from contracting and devaluing the low (absolute) margin DoubleClick platform to increase the relative value of its high (absolute) margin AdSense platform.⁴⁴

In summary, the one-sided and two-sided calculations yield different critical diversion ratios (and critical switching levels) and open up a region of error between them. The size of this region of error depends on the data and externality assumptions, but is non-trivial as a general rule. If the actual diversion ratio falls in between the one- and two-sided calculations, the analyst would give an opposite recommendation about unilateral price effects. This highlights the need for incorporating two-sided effects into any economic analysis involving MSPs.

⁴³ Customers who are sensitive to platform size, even if perfectly price inelastic, make this switch.

⁴⁴ The fact that Google can profit from DFP price increases that are purely contractionary suggests that other non-price methods of shrinking the relative size of the DoubleClick platform can also be profitable.

This section has highlighted the errors that arise from using approaches based on the one-sided Lerner Index – i.e. subject to the Lerner bias. Suppose, however, that an analyst had used econometric techniques for assessing market definition and competitive effects in single-sided markets to assess the Google-DoubleClick transaction. In that case, our conclusions would be reversed (the Estimation bias). The analyst would tend to define markets too narrowly and overstate the competitive consequences of the transaction.

VI. CONCLUSIONS

Traditional methods for evaluating mergers between single-sided firms can lead to erroneous conclusions when mergers involve MSPs. The direction of the bias depends on the analytical technique used, with traditional demand estimation tending to draw markets too narrowly and overstating competitive effects, and Lerner Index based calibration methods drawing markets too broadly and understating competitive effects. The direction of bias can also vary with the degree of platform asymmetry and the particular set of products that are subject to price increases. When mergers involve MSPs, the correct analysis must account for the indirect network effects between the multiple sides and the consequent effect on prices and output for the multiple sides. Failing to do so can lead to material mistakes as we have shown through the simulations reported above.

BIBLIOGRAPHY

- Argentesi, E. and L. Filistrucchi (2005). "Estimating Market Power in a Two-Sided Market: The Case of Newspapers," *Journal of Applied Econometrics*, *forthcoming*.
- Argentesi, E. and M. Ivaldi (2005). "Market Definition in Printed Media Industry: Theory and Practice," CEPR Discussion Paper No. 5096, Available at SSRN: <http://ssrn.com/abstract=779107>.
- Armstrong, Mark (2006). "Competition in Two-Sided Markets," *Rand Journal of Economics*, Vol. 37, No. 3, pp. 668-691.
- Armstrong, M. and J. Wright (2007). "Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts," *Economic Theory*, Vol. 32, No. 2, pp. 353-380.
- Baker, J., and T. Bresnahan (1985) "The Gains from Merger or Collusion in Product-Differentiated Industries," *Journal of Industrial Economics*, 33(4), pp. 427-444.
- Carlton, Dennis (2007). "Market Definition: Use and Abuse," *Competition Policy International*, Vol. 3, No 1.
- DOJ and FTC (1997). "U.S. Horizontal Merger Guidelines." Available at: http://www.usdoj.gov/atr/public/guidelines/horiz_book/hmg1.html.
- Dubois, P., A. Hernandez -Perez and M. Ivaldi (2007). "The Market of Academic Journals: Empirical Evidence from Data on French Libraries," *Journal of the European Economic Association*, Vol. 5, No. 2, April-May 2007, pp. 390-399
- Emch, E. and T.S. Thompson (2006). "Market Definition and Market Power in Payment Card Industries," *Review of Network Economics*, Vol. 5, No. 1, pp. 45-60.
- Evans, D. (2003a). "The Antitrust Economics of Two-Sided Markets," *Yale Journal of Regulation*, Vol. 20, No. 2, pp. 325-382.
- Evans, D. (2003b). "Some Empirical Aspects of Multi-Sided Platform Industries," *Review of Network Economics*, Vol. 2, No. 3, pp. 191-209.
- Evans, D., A. Hagiu and R. Schmalensee (2006). *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*. MIT Press, 2006.
- Evans, David and Richard Schmalensee (2005a). *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, 2d Ed. MIT Press, Cambridge, MA.

- Evans, D. and R. Schmalensee (2005b). "The Economics of Interchange Fees and Their Regulation: An Overview." Proceedings – Payment Systems Research Conferences, Federal Reserve Bank of Kansas City, May 2005, pp. 73-120.
- Evans D. and R. Schmalensee (2007a). "The Industrial Organization of Two-Sided Platforms," Competition Policy International, Vol. 3, No. 1.
- Evans, D. and R. Schmalensee (2007b). *Catalyst Code: The Strategies behind the World's Most Dynamic Companies*. Harvard Business School Press, 2007.
- Farrell, J and C. Shapiro (1990). "Horizontal Mergers: An Equilibrium Analysis," The American Economic Review, Vol. 80, No. 1, pp. 107-126.
- Feldman, R. (1994) "The Welfare Economics of a Health Plan Merger," Journal of Regulatory Economics, 6(1), pp. 67-86.
- Harris, B. (2003). "Recent Observations about Critical Loss Analysis." Available at: <http://www.usdoj.gov/atr/public/workshops/docs/202599.htm>
- Harris, B. and J.J. Simons (1989). "Focusing Market Definition: How Much Substitution is Necessary?" Research in Law and Economics, Vol. 207, pp. 207-226.
- Harris, B. and C. Veljanovski (2003). "Critical Loss: Its Growing Use in Competition Law," European Competition Law Review, Vol. 5, pp. 213-218.
- Ivaldi, M. and F. Verboven (2005). "Quantifying the Effects from Horizontal Mergers in European Competition Policy," International Journal of Industrial Organization, Vol. 23, pp. 669-691.
- Kaiser, U. and J. Wright (2006). "Price Structure in Two-Sided Markets: Evidence from the Magazine Industry," International Journal of Industrial Organization, Vol. 24, Issue 1, pp. 1-28.
- Katz, M. and C. Shapiro (2003). "Critical Loss: Let's Tell the Whole Story," Antitrust Magazine, Vol. 17, No. 2, pp. 49-56.
- Klein, B., Andres V. Lerner, Kevin M. Murphy and Lacey L. Plache (2006). "Competition in Two-Sided Markets: The Antitrust Economics of Payment Card Interchange Fees," Antitrust Law Journal, 73(3), pp. 571-626.
- Nevo, A. (2000). "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," RAND Journal of Economics, 31(3), pp. 395-421.
- O'Brien, D. and A. Wickelgren. (2003). "A Critical Analysis of Critical Loss," FTC Working Paper 254.

Rochet, J-C. and J. Tirole (2003). "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1:4, pp. 990-1029.

Rochet, J-C. and J. Tirole (2006). "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, Vol. 37, Issue. 3, pp. 645-667.

Rysman, M. (2004). "Competition Between Networks: A Study of the Market for Yellow Pages," *Review of Economic Studies*, 71, pp. 483-512.

Rysman, M. (2007). "An Empirical Analysis of Payment Card Usage," *Journal of Industrial Economics*, 55, pp. 1-36.

Scheffman, D. and J. Simons (2003). "The State of Critical Loss: Let's Make Sure We Understand the Whole Story," *The Antitrust Source*, Vol. 3, No. 2, pp. 1-9.

Shapiro, Carl (1996). "Mergers with Differentiated Products," *Antitrust*, 10, pp. 23-30.

Weisman, Dennis L. (2005). "Assessing Market Power: The Trade-off Between Market Concentration and Multi-Market Participation," *Journal of Competition Law and Economics*, 1 (2), pp. 339-354.

Werden, G.J. and L. M. Froeb (1994). "The Effects of Mergers in Differentiated Products Industries: Logit Demand and Merger Policy," *Journal of Law, Economics, & Organization*, 194, pp. 407-426.

Willig, R. (1991). "Merger Analysis, Industrial Organization Theory, and Merger Guidelines," *Brookings Papers on Economic Activity: Microeconomics*, 281, pp. 292-293.

Wright, J. (2004). "The Determinants of Optimal Interchange Fees in Payment Systems," *Journal of Industrial Economics*, Vol. 52, No. 1, March, pp. 1-26.

A1. APPENDIX. N-PLATFORM AND ASYMMETRIC PLATFORM CASES

We extend the results in the text to the case of asymmetric platforms and to the case of more than two platforms. The Critical Loss equation for two asymmetric platforms is given by:

$$\sum_{s=A,B} \sum_{i=1,2} \left[R_i^s (X_i^s + M_i^s) \frac{\Delta Q_i^s}{Q_i^s} + R_i^s X_i^s \right] = 0 \quad (\text{A.1})$$

The N platform formula differs only in the number of summed terms. This formulation allows different price changes at each platform on each side for use in a variety of questions.

For Actual Loss, consider again the isoelastic demand system given by

$$q_i^s = \mu_i^s + \alpha_i^s q_i^r - \delta_i^s q_j^r - \beta_i^s p_i^s + \gamma_i^s p_j^s \quad (\text{A.2})$$

for $i, j = 1, 2, i \neq j$ and $s = A, B$, where lower case q 's and p 's represent log values. Parameters are positive (though this need not be the case) and restricted such to prevent the two-sided effects from exploding following price changes.

Since this a demand system, ownership patterns do not matter for estimation. Also, if there are two platforms that both sell to side A , it is not necessary that they both sell to side B . The two "side B 's" can refer to completely different goods. It may also be that one platform sells to sides A and B , whereas another firm sells only to side A , and a third firm only sells to side B .

The reduced form value of q_j^s as a function only of parameters and prices, accounting for all the feedbacks in the system, is:

$$q_i^s = [\Gamma_i^s (\alpha_j^s \theta_j^r - \delta_j^s \theta_i^r + \theta_j^s) + \Psi_i^s (\alpha_i^s \theta_i^r - \delta_i^s \theta_j^r + \theta_i^s)] / \Omega_i^s \quad (\text{A.3})$$

where

$$\begin{aligned} \theta_i^s &= \mu_i^s - \beta_i^s p_i^s + \gamma_i^s p_j^s \\ \Gamma_i^s &= -\alpha_i^s \delta_i^r - \delta_i^s \alpha_j^r \\ \Psi_i^s &= (1 - \alpha_j^s \alpha_j^r - \delta_j^s \delta_i^r) \\ \Omega_i^s &= \Psi_i^s (1 - \alpha_i^s \alpha_i^r - \delta_i^s \delta_j^r) + \Gamma_i^s (\alpha_j^s \delta_j^r + \delta_j^s \alpha_i^r) \end{aligned}$$

Totally differentiating, and replacing all dp_i^s with X_i^s and dq_i^s with L_i^s , we derive the Actual Loss formula:

$$\begin{aligned}
L_i^s &= (\Gamma_i^s \gamma_j^s - \Psi_i^s \beta_i^s) X_i^s / \Omega_i^s \\
&\quad + (-\Gamma_i^s \beta_j^s + \Psi_i^s \gamma_i^s) X_j^s / \Omega_i^s \\
&\quad + [\Gamma_i^s (\alpha_j^s \gamma_j^r + \delta_j^s \beta_i^r) + \Psi_i^s (-\alpha_i^s \beta_i^r - \delta_i^s \gamma_j^r)] X_i^r / \Omega_i^s \\
&\quad + [\Gamma_i^s (-\alpha_j^s \beta_j^r - \delta_j^s \gamma_i^r) + \Psi_i^s (\alpha_i^s \gamma_i^r + \delta_i^s \beta_j^r)] X_j^r / \Omega_i^s
\end{aligned} \tag{A.4}$$

The analyst then compares Actual Loss and Critical Loss as described in the text.

The analysis can further be extended to N platforms. The Actual Loss formula is derived from the following system of $2N$ equations:

$$q_i^s = \sum_{k=1}^N \delta_{ik}^s q_k^r + \theta_i^s \tag{A.5}$$

for $i = 1..N$; $s, r = A, B$; $s \neq r$, where

$$\theta_i^s = \mu_i^s + \sum_{k=1}^N \gamma_{ik}^s p_k^r$$

Quantity at platform i on side s depends on the quantity at each platform k (including i) on the other side r (the δ_{ik} 's), and the price at each platform (including i) on side s (the γ_{ik} 's).⁴⁵ Substituting each q_i^r equation into each q_i^s equation, the system can be rewritten in matrix form: $\Delta^s \vartheta^s = \Phi^s$ for each side s , where Δ^s is an $N \times N$ matrix with ij^{th} element λ_{ij}^s equal to

$$\lambda_{ij}^s = 1(i = j) - \sum_{k=1}^N \delta_{ik}^s \delta_{kj}^s \tag{A.6}$$

⁴⁵ We have replaced the notation α_i from the two-platform version with δ_i^r 's and the coefficient on own price from β_i with γ_i^s 's.

where $1(i = j)$ is an indicator function equal to one when $i = j$ and zero otherwise. Also, ϑ^s is a N -vector with i^{th} element equal to q_i^s , and Φ^s is an N -vector with i^{th} element equal to ϕ_i^s , where:

$$\phi_i^s = \theta_i^s + \sum_{k=1}^N \delta_{ik}^s \theta_k^r \quad (\text{A.7})$$

Actual Loss is calculated for each q_i^s by total differentiation, i.e. $L^s = (\Delta^s)^{-1} d\Phi^s$, for $s = A, B$, replacing each dp_i^s with X_i^s in the $d\phi_i^s$ terms.

To close the model, we note that the first order conditions defining markups in the two-platform symmetric case are the same as those in the 2-platform and N -platform asymmetric cases, since each set of equations is platform specific.

As before, the calculated Actual Loss amounts X_j^s are substituted into the Critical Loss equation and the ACR is checked. (Alternatively, one can fix all parameters except one, and use the $\text{ACR} = 1$ identity to back out the “critical” value of the remaining parameter.) The comparative statics from the symmetric and two-platform cases carry through to the N firm asymmetric case.

A2. APPENDIX. CALIBRATION OF ELASTICITIES FOR GOOGLE-DOUBLECLICK MERGER ANALYSIS

We consider the question of whether it may be profitable for Google to raise the price of DFP post-merger by 5% or 10%. Due to the sparse nature of the available data, we use a calibration technique to get a first look at the potential unilateral effects of the merger.

For the one-sided calculation, there are four parameters to be calibrated, all on the publisher side ($\beta_{\text{GOOGLE}}^{\text{PUBLISHER}}$, $\beta_{\text{DOUBLECLICK}}^{\text{PUBLISHER}}$, $\gamma_{\text{GOOGLE}}^{\text{PUBLISHER}}$, $\gamma_{\text{DOUBLECLICK}}^{\text{PUBLISHER}}$) since all indirect network effects are erroneously assumed to be zero. First, we use the theoretical one-sided Lerner Index to fix the own-price elasticities, given the observed markups (two restrictions). This yields a biased estimate of the true short-

run price elasticities. We then assume Slutsky symmetry, which equates revenue-weighted cross-price elasticities across platforms (one restriction), which is commonly done. This essentially equates the cross-partials $\partial Q_i^s / \partial P_j^s = \partial Q_j^s / \partial P_i^s$, written in levels instead of logs.⁴⁶ We then just need to estimate or calibrate one of the cross-price elasticities and then we could calculate Critical Loss and Actual Loss and compare them. However, we instead employ the common and useful technique in calibration exercises to set Critical Loss equal to Actual Loss (the fourth restriction) and then back out the values of the cross-price elasticities (connected together by Slutsky symmetry) that would exactly equate *CL* and *AL*. Rather than report these “critical” cross price elasticities directly, as explained in the text, we follow the common practice of reporting critical diversion ratios and critical switching levels, which are derived from the critical cross price elasticities.

For the two-sided calculation, whose formulas are in Appendix A.1, we must calibrate sixteen parameters, either directly or using restrictions guided by economic theory. First, we use the true two-sided formulas for the Lerner Index given by equation (9). Conditional on the externality parameters (α 's and δ 's), this yields a linear relation between each own-price elasticity and its corresponding cross-price elasticity (on each side for each platform), a total of four restrictions. We again assume Slutsky symmetry to equate the revenue-weighted cross-price elasticities across platforms on each side (two more restrictions), though results are similar for other reasonable assumptions.

We also use a similar type of Slutsky symmetry for the relative size of the indirect network effects. Given the absence of direct data on the size of the network effects, this reasonably relates pairs of externality elasticities and yields four more restrictions. Effectively, we set for each side $\partial Q_i^s / \partial Q_i^r = \partial Q_j^s / \partial Q_j^r$ and $\partial Q_i^s / \partial Q_j^r = \partial Q_j^s / \partial Q_i^r$, written in levels instead of logs. Using other reasonable assumptions does not affect results qualitatively. We directly set Google's own-platform cross-side externality to 0.1 (10%) on each side, and its cross-side cross-

⁴⁶ This must be exactly true with Hicksian demand curves or if there were no income effects, and is an approximation for the Marshallian (or regular) demand curves estimated here.

platform externality to 0.05 (5%) on each side (four more restrictions). Slutsky symmetry pins down the corresponding externality parameters for DoubleClick.

These restrictions add to 14 and equating Actual Loss equal to Critical Loss counts for a 15th restriction. We need only specify one own-price elasticity or one cross-price elasticity to complete the calibration. As we do not have direct data on any of these elasticities, it is more prudent to seek out the “most symmetric” solution. Technically speaking, many combinations of β (own) and γ (cross) are possible on a given side and still satisfy its markup condition. However, note that the impact on markups of a change in β is about $1/\alpha$ times the impact of a change in γ . So if $\alpha = 0.1$, the impact of β is ten times that of γ . Therefore, a relatively low β could exist on one side as long as its corresponding γ is extremely high to compensate. But since AL must equal CL , this causes an opposite stress on the other side. It tends to force the other β to be high, and the other γ to be extremely small – and the latter can easily come out negative (with a similarly large magnitude as the γ on the first side) which is not possible. The most reasonable solution, absent good data, is that which yields the most symmetric β 's and γ 's, conditional on satisfying all the other constraints. In a fully symmetric model this ensures a purely symmetric solution; in an asymmetric model the β 's and γ 's still differ across sides but only to the extent the externality parameters require this. We implement the most symmetric solution by minimizing the sum of squares of the γ 's across sides (alternatively, the sum of squares of the β 's across sides.) This is the 16th and final restriction.

We are now able to back out the critical β 's and γ 's (tied together by the above constraints) that would result in CL equaling AL . Of course, in cases where more and better data is available, one need not rely as heavily on theoretical restrictions and instead estimate elasticities directly from the data itself. This is always preferred.