MERGER ANTITRUST LAW

LAW 1469
Georgetown University Law Center
Fall 2025

Tuesdays and Thursdays, 3:30 pm – 5:30 pm
Dale Collins
wdc30@georgetown.edu
www.appliedantitrust.com

Class 11 (September 26, 2025): H&R Block/TaxACT (Unit 5)[1]

In this class, we will continue our examination of market definition. So far, we have developed the hypothetical monopolist test when discussing Sanford Health/MDC and followed it with the critical loss implementation in the prior class. We will finish market definition in *H&R Block* with the one-product SSNIP recapture implementation of the hypothetical monopolist test.

The critical loss implementation primarily applies to homogeneous product markets, where quantity is the control variable and there is no price discrimination. Empirically, this occurs relatively infrequently in the real world. Instead, competing products tend to be differentiated from a consumer perspective (either by product characteristics or by location), giving each firm some ability to set its own price to maximize profits against the firm's (downward-sloping) residual demand curve.

In addition, although the original hypothetical monopolist test in the 1982 Merger Guidelines required the hypothetical monopolist to increase the prices of all of the products in the candidate market by a uniform percentage SSNIP, the 2010 and 2023 Merger Guidelines allow the hypothetical monopolist to raise the prices of one or more products in a differentiated candidate market, selectively raising their prices while leaving the others constant. Under this change, the hypothetical monopolist test only requires that the hypothetical monopolist be able to profitably *raise the price of a single product* in the product group for the product grouping to be a relevant market. The 2010 Guidelines required at least one of the products subject to the SSNIP to be a product of one of the merging firms. The 2023 Guidelines allowed the product subject to the SSNIP to be any product in the candidate market, but courts have sensibly retained the requirement that the SSNIP apply to at least one of the products of the merging firms.

When a hypothetical monopolist raises the price on only one product (say, Product 1), three things happen:

1. The monopolist gains the $SSNIP on each of the inframarginal sales of Product 1
2. The monopolist loses the profit margin on the lost marginal sales of Product 1
3. The monopolist gains the profit margin on any lost marginal sales "recaptured" by other products in the candidate market (which the hypothetical monopolist controls).

The first two factors characterize the critical loss tests covered in the previous class since there is no recapture of lost marginal sales by other products in the candidate market. Instead, all lost sales divert to products outside the candidate market. In a differentiated products market,

[1] A reasonably complete set of the most important filings in the litigation (including the trial transcript) may be found here on AppliedAntitrust.com.

however, some of Product 1's lost marginal sales may divert to other products in the candidate market, and the hypothetical monopolist will earn incremental profits on those recaptured sales.

An example may be helpful to compare a uniform SSNIP test with a one-product SSNIP test. Consider a candidate market consisting of four different types of gourmet yogurt with the following characteristics:

	Price	$Margin	Unit Sales
Siggi's	$2.50	$1.00	1000
Noosa	$2.50	$1.25	800
La Fermiete	$3.50	$1.75	600
Oikos Triple Zero	$4.00	$2.00	500
			2900

When the prices of all products are increased by five percent, each product loses ten percent of its sales. None of these sales divert to another product in the group. Does the product group satisfy the HMT under a uniform SSNIP of five percent?

We can set up the following worksheet for a "brute force" accounting:

	Price	$Margin	Unit Sales	Marginal Sales	%Δq	Inframarginal Gain	Marginal Loss	Net
Siggi's	$2.50	$1.00	1000	100	10.00%	112.50	100.00	12.50
Noosa	$2.50	$1.25	800	80	10.00%	90.00	100.00	-10.00
La Fermiete	$3.50	$1.75	600	60	10.00%	94.50	105.00	-10.50
Oikos Triple	$4.00	$2.00	500	50	10.00%	90.00	100.00	-10.00
						387.00	405.00	-18.00

Inframarginal gain for each product is 5% of the product's price times the inframarginal unit sales. Marginal loss is the dollar margin of the product times the marginal unit sales. The sum of the net incremental profits is -$18.00, so a uniform 5% SSNIP is unprofitable.

Now, let's perform a one-product SSNIP test. Say Siggi's is a product of one of the merging firms and subject it alone to a five percent SSNIP. Suppose Siggi's loses 120 sales (12%) due to the price increase. Siggi's loses one half of these unit sales to products outside the candidate market while the other half diverts to the other products in the candidate market at the levels shown in the Recaptured Sales column:

	Price	$Margin	Unit Sales	Marginal Sales	Recaptured Sales	Inframarginal Gain	Marginal Loss	Recapture Gain	Net
Siggi's	$2.50	$1.00	1000	120		$110.00	-$120.00		-$10.00
Noosa	$2.50	$1.25	800		30			$37.50	$37.50
La Fermiete	$3.50	$1.75	600		20			$35.00	$35.00
Oikos Triple	$4.00	$2.00	500		10			$20.00	$20.00
			2900		60	$110.00	-$120.00	$92.50	$82.50

Two points are worth noting here. First, Siggi's loses more sales when only its price increases than when the (percentage) price increase is applied to all four products because the other products in the group are more attractive to Siggi's marginal customers at the lower prices. Second, notwithstanding the greater unit and dollar losses on Siggi's when only its price is subject to the SSNIP, the hypothetical monopolist makes a positive profit of $82.50 because of the profits earned on Siggi's lost marginal sales that are diverted (recaptured) by the other three products in the candidate market. Hence, although the candidate market fails the uniform SSNIP test, it passes the one-product SSNIP test.

As this example suggests, one-product SSNIP tests typically yield relevant product markets that are smaller (have fewer products) than uniform SSNIP tests, which in turn often result in market shares and market concentration levels that are higher than in those relevant markets that must satisfy a uniform SSNIP test.[2] This tendency often makes one-product SSNIP tests appealing to plaintiffs.

All of this should feel familiar from our study of unilateral effects. If you already understand the basic principle of unilateral effects—where a merged firm profits from raising price on only one product because it recaptures enough diverted sales on its other product—you will recognize the logic of the one-product SSNIP recapture test. Here the same principle applies, except that the hypothetical monopolist recaptures diverted sales across all of the products in the candidate market that do not experience the SSNIP.

With this introduction, read the class notes summarizing the basic principles of the one-product SSNIP recapture test (slides 38–45). While the following class notes develop formulas that allow the test to be applied more efficiently, remember that brute-force accounting—as in the examples here and on slides 43 and 44—will always work.

The next section of the class notes (slides 46–53) develops diversion ratios, which are central to applying the one-product SSNIP recapture test. We introduced diversion ratios earlier in the course, but this section adds considerable detail to the concept. A diversion ratio is always specific to the particular SSNIP applied to the product whose price increases: it measures the fraction of that product's lost sales that divert to another product whose price remains unchanged. In practice, diversion ratios may be estimated from a firm's own in-loss records, other internal company documents, or relevant empirical evidence such as switching data from insurers, customers, or marketing studies. Where no more probative evidence is available, the agencies and courts sometimes resort to the relative share method, which allocates diversion in proportion to the shares of the potential destination products in the candidate market. As you review these slides, focus on how diversion ratios are defined, how they can be estimated from the available evidence, and how they directly inform the profitability condition for a one-product SSNIP.

The class notes then develop the formal concepts and tools for applying the one-product SSNIP recapture test. They begin by defining the aggregate diversion (or recapture) ratio, which measures the share of sales lost by the product with the SSNIP that are recaptured by other products in the candidate market (slide 55). You will see a number of synonymous labels in the literature and case law—aggregate diversion ratio, aggregate diversion test, one-product SSNIP recapture test—but they all point to the same basic idea. The class notes then explain how the

[2] A one-product SSNIP also is likely to yield smaller geographic markets around the locations of the merging firms.

Merger Guidelines employ this framework: the HMT is framed in terms of whether a hypothetical monopolist could profitably impose a SSNIP on one product by recapturing enough diverted sales on the other products in the candidate market, an approach that courts have increasingly adopted. (slides 56).

The minimum recapture ratio required to make the SSNIP profitable is called the *critical recapture ratio*. This ratio is specific to the product on which the SSNIP is imposed and the magnitude of the SSNIP. The one-product SSNIP recapture test then holds:

> *If the actual recapture ratio exceeds the critical recapture ratio, the candidate market satisfies the one-product SSNIP implementation of the HMT.*
>
> *Otherwise, the candidate market fails this implementation of the HMT.*

Two points to note here.

First, whereas the critical loss test looks to whether actual loss is less than the critical loss, the one-product SSNIP recapture test looks to whether the actual recapture ratio exceeds the critical recapture ratio. This makes sense: critical loss asks whether the marginal sales of the hypothetical monopolist are too large to make the SSNIP profitable, while the one-product SSNIP test asks whether the hypothetical monopolist recaptures enough of the lost marginal sales to make the SSNIP profitable.

Second, it is essential to remember that a candidate market that fails one implementation of the HMT but satisfies a different implementation still passes the HMT. So, a candidate market that fails a uniform SSNIP critical loss test but satisfies a one-product SSNIP test still passes the HMT. Likewise, a candidate market that fails a one-product SSNIP test when the SSNIP is applied to one product, but satisfies the test when the SSNIP is applied to another product, passes the HMT. Do not make the mistake of concluding that a candidate market fails the HMT just because it fails one implementation.

Next, the class notes set out the fundamental formula for determining the critical recapture ratio:

$$R^{1}_{Critical} = \frac{\delta p_1}{\$m_{RAve}} \quad \left(= \frac{\$\text{SSNIP}_1}{\$m_{RAve}} \right),$$

where δ is the percentage SSNIP applied to Product 1[3]

p_1 is the pre-SSNIP price of Product 1

\$SSNIP is the dollar SSNIP applied to Product 1

$\$m_{Rave}$ is the recapture share-weighted average of the other products in the candidate market that are not subject to the SSNIP (see slide 59)

Don't worry too much about how to calculate the $\$m_{Rave}$. The key thing to remember is that it is the weighted average of the dollar margins of the products in the candidate market, where the weights are each product's relative recapture share. The effect is that products that recapture more of Product 1's lost marginal sales count more heavily in the average than those that recapture fewer.

The proof of this proposition is quite involved and therefore optional. You will not be tested on the details of the derivation, but I encourage you to work through it if you want to see how the

[3] By convention, unless otherwise specified, Product 1 is the product on which the hypothetical monopolist applies the one-product SSNIP.

algebra anchors the economic logic. For most purposes, it is enough to understand the result: the profitability of a one-product SSNIP turns on whether the actual recapture ratio exceeds the critical recapture ratio.

The class notes then turn to several corollaries and simplifications of the one-product SSNIP recapture test (slides 64-72). These results show how the general profitability condition can be expressed more simply under common assumptions. For example, when all products in the candidate market are symmetric or share the same margins, the test reduces to a straightforward comparison that is easier to apply. Note, however, that most of these corollaries (Corollaries 1-3) rely on the use of $\$m_{RAve}$. You will *not* be responsible for calculating $\$m_{RAve}$ or for knowing or using Corollaries 1-3 (slides 64, 66–70, and 72). The exception is Corollary 4 for symmetric products (slides 65 and 71), which you are expected to know and be able to use. It will also play a role when we later look at estimating the magnitude of profit-maximizing unilateral effects price increases.

The final slides in this section add two important cautions. First, do not assume that an economist's calculations are always correct. In practice, even experts can make mistakes or overlook details, so you should always double-check the logic and arithmetic before relying on their conclusions. For example, two well-known antitrust economists once derived and published an equation for calculating a one-product SSNIP critical recapture ratio that turned out to be incorrect; yet, it was relied upon by other economists and even the courts until the error was discovered (slide 74). Second, while it may be tempting to assume that a 100% recapture ratio guarantees a profitable SSNIP, that is not always the case. Whether the SSNIP is profitable also depends on the size of the margins and the magnitude of the SSNIP itself. A merged product might recapture all of the lost sales, but if margins are low or the price increase too large, the price increase may still be unprofitable (slides 75-79, including an optional proof).

The class notes then cover the uniform SSNIP recapture test for differentiated product markets (slides 80-89). Unlike the one-product SSNIP test, here the hypothetical monopolist increases the prices of all products in the candidate market simultaneously. The profitability condition is framed in terms of whether enough of the lost sales are recaptured on these products at their new, higher prices. This test parallels the one-product SSNIP logic but is applied across the entire candidate market.

Which test is appropriate depends on the nature of the diversion ratios available. If diversion ratios are estimated from one-product price changes, the proper tool is the one-product SSNIP test. If only switching data is available, then a uniform SSNIP test is more usually appropriate, as observed switching may be driven by factors other than price.

The class notes then introduce the uniform SSNIP recapture sufficiency test (slides 85-88). This test can be used when the available data allows an estimate of a uniform SSNIP recapture ratio but not a one-product SSNIP recapture ratio. It provides a sufficient condition for profitability: if the estimated uniform recapture ratio is high enough, then a uniform SSNIP would certainly be profitable. But it is important to remember that this is only a sufficient test, not a necessary one. A candidate market that fails the sufficiency test may still satisfy the HMT under the general

uniform SSNIP framework; the sufficiency condition simply offers a convenient shortcut in those limited circumstances where only uniform recapture data can be estimated.[4]

Finally, some commentators have suggested that in a uniform SSNIP test, one can use the single-product SSNIP diversion and recapture rates to substitute them into the sufficiency test, thereby creating a presumption that the profitability condition holds and the candidate market therefore satisfies the HMT (slide 89). In other words, they treat the one-product diversion evidence as if it were sufficient for the uniform framework. I am very skeptical of this approach. It is not supported by a meaningful economic argument, and it risks significantly overstating the extent of within-market diversion in response to a uniform SSNIP. For that reason, I discourage its use. If one-product diversion ratios are available, the test to use is the one-product SSNIP recapture test, not a uniform SSNIP recapture test. However, you should understand the suggestion, as it has been mentioned in the commentary, but also be clear on why it is unreliable.

The class notes conclude with a summary of the different HMT implementations (slides 90-94). This section synthesizes the key principles covered in the class notes, including the various ways the HMT can be implemented: uniform SSNIP critical loss, one-product SSNIP recapture, uniform SSNIP recapture, and their associated corollaries. The summary illustrates how each implementation is connected to the same fundamental economic logic and explains how your choice of implementation should depend on the nature of the available diversion evidence. Importantly, the section emphasizes a critical point to remember: failing one implementation of the test does not mean the candidate market fails the HMT entirely. Read this summary section carefully. It serves as a capstone that distills the central insights from the more technical material into a clear restatement of the essential concepts you need to retain and apply going forward.

With this background, read the section of the *H&R Block* opinion discussing the economists' testimony and evidence on the HMT (pp. 51–65). In this section, Dr. Warren-Boulton, testifying for the government, applied a critical loss analysis and a merger simulation. He concluded that a 10% SSNIP in the digital do-it-yourself ("DDIY") tax preparation market would be profitable, since his estimates of diversion ratios from IRS switching data all exceeded the critical loss threshold, and his simulation likewise predicted significant post-merger price increases.

As you read the court's description of Warren-Boulton's HMT analysis, ask yourself the following questions: (1) What HMT implementation did Warren-Boulton use? (2) How did the court describe his analysis, and did it get it right? (3) Was this the right implementation to apply in this case? (4) What assumptions underlie Warren-Boulton's reliance on switching data, and are they valid? (5) What role did his analysis play in the court's ultimate decision on market definition—was it central, corroborative, or peripheral?

Dr. Meyer, testifying for the defendants, sought to rebut this analysis on several fronts. She argued that switching data were unreliable proxies for diversion because consumers switch for reasons other than price, and she presented alternative diversion estimates derived from a pricing simulator and an email survey. But the court found both of these data sources deeply flawed—most notably because the simulator assigned no prices to key alternatives, leading to anomalous and counterintuitive results.

[4] I suspect that if we reverse the inequality signs in Proposition 2 (slide 85), we get a necessary test. But I have not tried to prove that yet.

As you read the court's description of Dr. Meyer's testimony, ask yourself the following questions: (1) What methods did Meyer use to rebut Warren-Boulton's HMT analysis? (2) The court found her data flawed, but do you agree with that assessment? (3) Did her criticisms of Warren-Boulton's reliance on switching data highlight valid weaknesses, or did they fall short? (4) How did the court evaluate Meyer's methods and conclusions? (5) If you were the court, would you have given her testimony more, less, or the same weight as the court ultimately did?

If you have any questions, send me an email. See you in class.