



**A CRITICAL ASSESSMENT OF THE LATEST
CHARGE OF GOOGLE'S ANTICOMPETITIVE
BIAS FROM YELP AND TIM WU**

Geoffrey A. Manne, R. Ben Sperry & Kristian Stout

*ICLE Antitrust & Consumer Protection Research Program
White Paper 2016-3*

A CRITICAL ASSESSMENT OF THE LATEST CHARGE OF GOOGLE'S ANTICOMPETITIVE BIAS FROM YELP AND TIM WU

Geoffrey A. Manne, R. Ben Sperry & Kristian Stout*

Introduction

Late last year, Tim Wu of Columbia Law School (and now the White House Office of Management and Budget), Michael Luca of Harvard Business School (and a consultant for Yelp), and a group of Yelp data scientists released a study claiming that Google has been purposefully degrading search results from its more-specialized competitors in the area of local search.¹ The authors' claim is that Google is leveraging its dominant position in general search to thwart competition from specialized search engines by favoring its own, less-popular, less-relevant results over those of its competitors:

To improve the popularity of its specialized search features, Google has used the power of its dominant general search engine. The primary means for doing so is what is called the "universal search" or the "OneBox."²

* Geoffrey Manne is the founder and Executive Director of the International Center for Law & Economics (ICLE); Ben Sperry is the Associate Director at ICLE; Kristian Stout is the Associate Director for Innovation Policy at ICLE. ICLE has received financial support from numerous companies and individuals, including Google, as well as several of its competitors. Unless otherwise noted, all ICLE support is in the form of unrestricted, general support. The ideas expressed here are the authors' own and do not necessarily reflect the views of ICLE's advisors, affiliates or supporters. Please contact the authors with questions or comments at icle@laweconcenter.org.

¹ Michael Luca, Timothy Wu, Sebastian Couvidat, Daniel Frank, & William Seltzer, *Is Google degrading search? Consumer harm from Universal Search* (Harvard Business School Working Paper No. 16-035), available at <http://people.hbs.edu/mluca/SearchDegradation.pdf> (hereinafter "Wu, *et al.*").

² *Id.* at 10.

This is not a new claim, and researchers have been attempting³ (and failing⁴) to prove Google's "bias" for some time. Likewise, these critics have drawn consistent policy conclusions from their claims, asserting that antitrust violations lie at the heart of the perceived bias.

But the studies are systematically marred by questionable methodology and bad economics.⁵ The primary difference now is the saliency of the "Father of Net Neutrality,"⁶ Tim Wu, along with a cadre of researchers employed by Yelp (one of Google's competitors and one of its chief antitrust provocateurs⁷), saying the same thing in a new research paper, with *slightly different* but equally questionable methodology, bad economics, and a smattering of new, but weak, social science.⁸

The basic thesis of the study is that Google purposefully degrades its local searches (e.g., for restaurants, hotels, services, etc.) to the detriment of its

³ See, e.g., Benjamin Edelman & Benjamin Lockwood, *Measuring Bias in "Organic" Web Search*, BENEDELMAN (last updated Jan. 19, 2011), <http://www.benedelman.org/searchbias/>.

⁴ See Joshua D. Wright, *Defining and Measuring Search Bias: Some Preliminary Evidence* (George Mason University Law and Economics Research Paper Series 12-14, Nov. 3, 2011), available at http://www.law.gmu.edu/assets/files/publications/working_papers/1214DefiningandMeasuringSearchBias.pdf.

⁵ See, e.g., Josh Wright, *Search Bias and Antitrust*, TRUTH ON THE MARKET (Mar. 24, 2011), <https://truthonthemarket.com/2011/03/24/search-bias-and-antitrust/>; Josh Wright, *Sacrificing Consumer Welfare in the Search Bias Debate*, TRUTH ON THE MARKET (Apr. 22, 2011), Josh Wright, *Sacrificing Consumer Welfare in the Search Bias Debate, Part II*, TRUTH ON THE MARKET (Jun. 28, 2011), <https://truthonthemarket.com/2011/06/28/sacrificing-consumer-welfare-in-the-search-bias-debate-part-ii/>.

⁶ See Emil Guillermo, *Father of Net Neutrality, Tim Wu, Hails FCC Decision*, NBCNews (Feb. 26, 2015), <http://www.nbcnews.com/news/asian-america/father-net-neutrality-tim-wu-hails-fcc-decision-n313656>.

⁷ See, e.g., Kelly Fiveash, *Yelp files competition complaint against Google search biz in EU*, The Register (Jul. 9, 2014), http://www.theregister.co.uk/2014/07/09/yelp_files_official_competition_complaint_against_google_search_biz_in_eu/.

⁸ For a thorough criticism of the inherent weaknesses of Wu, *et al.*'s social science methodology, see Miguel de la Mano, Stephen Lewis, and Andrew Leyden, *Focus on the Evidence A Brief Rebuttal of Wu, Luca, et al* (2016), available at SSRN: <http://ssrn.com/abstract=2825006> (hereinafter "de la Mano, *et al.*").

specialized search competitors, local businesses, consumers, and even Google's bottom line — and that this is an actionable antitrust violation.⁹

The study approaches its analysis by employing A/B testing to determine whether consumers prefer Google's "local specialized universal results" or Yelp's preferred alternative as delivered with the "Focus on the User"¹⁰ ("FOTU") plugin:

Our main treatment uses data from a Chrome Browser Extension called Focus on the User – Local (FOTU), which was designed to detect Local OneBoxes and perform alternative searches for results from online review websites. Upon detecting a Local OneBox in Google search results, FOTU conducted a search for links to third party local review websites (such as Yelp, TripAdvisor, and ZocDoc). FOTU then extracted and ranked results from these websites (as well as from Google's review content), according to a combination of Google's organic ranking, the business's average star rating, and the number of reviews. FOTU essentially constructed an alternative method of presenting local results, based on the algorithm that powers Google organic search....

To create a control condition, we captured a screenshot of results for searches for the phrase "coffee Louisville KY". Throughout the experiment, this was used as a control display.... [The FOTU] treatment is identical to the control except for the content of the OneBox. Instead of containing Google content, the OneBox presents content from FOTU. By comparing the control to this treatment, we can explore consumer preferences across the two sets of content.¹¹

The study's authors conclude that the allegedly "neutral" search results derived from the FOTU plugin are more relevant to users than are Google's unadulterated Universal Search results — and thus that Google must be harming consumers by showing its own results instead of Yelp's preferred results (the FOTU-generated results).

⁹ Wu, *et al.*, *supra* note 1 at 6.

¹⁰ *Id.* at 4-5. See also *Focus on the User* (last accessed Sept. 13, 2016), <http://www.focusontheuser.eu/>.

¹¹ Wu, *et al.*, *supra* note 1, at 23-25.

But in fact the study shows nothing of the kind. Instead, the study is marred by methodological problems that, in the first instance, make it impossible to draw *any* reliable conclusions. Nor does the study show that Google's conduct creates any antitrust-relevant problems. Rather, the construction of the study and the analysis of its results reflect a superficial and inherently biased conception of consumer welfare that completely undermines the study's purported legal and economic conclusions.

Methodological Problems

The conclusions of any such study are only as reliable as the methodology used to generate them. Unfortunately for Wu and his coauthors, the study suffers from several critical methodological flaws.

I. The study is based on an insufficiently small and narrow data set

As de la Mano, *et al.* note:

[The p]aper's user testing methodology is woefully invalid. It relies on a single example and does not provide any data on other queries featured prominently on the FOTUL website. Participants from across the United States were asked to evaluate static screenshots with search results for a hyper-local query (coffee shops in Louisville, Kentucky). In addition to this query likely being irrelevant to the overwhelming majority of testers, participants could not see the destination webpages of the results displayed (which is essential to evaluate the quality of a set of search results). Based on a recreation of the FOTUL Paper's methodology, the very same experiment yields exactly the opposite outcome for another query from FOTUL's website, i.e., users preferred Google's supposedly "degraded" results over FOTUL's. In these circumstances, the FOTUL experiment is plainly unreliable, and its results ought to be disregarded as empirical evidence.¹²

It is impossible to derive statistically relevant results from the limited sample in the Wu, *et al.* study. In fact, the methodological flaws of the study are so severe

¹² de la Mano, *et al.*, *supra*, note 8 at 4.

that when it was re-run using just a single, slightly different search term, the researchers obtained contradictory results.¹³

In addition, the fact that “[t]esters were not allowed to visit the referenced pages... further limited participants’ ability to determine whether FOTUL results were of any use.”¹⁴ As we discuss below, it is impossible to draw conclusions about the consumer welfare effects of different search results without considering users’ behavior when they click on the results. Among other things, whatever their first impressions, if users would have ended up consistently returning to the initial results after having found the destination webpages wanting, this would demonstrate considerable consumer dissatisfaction. Yet nothing in the Wu, *et al.* paper tests this.

2. The study is based on a flawed measure of consumer welfare

To begin with, the entire study is premised on the assertion that click-through rates are an accurate measure of consumer welfare:

We find that users are roughly 40% more likely to engage with universal search results (which receive favored placement) when the results are organically determined relative to when they contain only Google content. To shed further light on the underlying mechanisms, we show that users are more likely to engage with the OneBox when there are more reviews, holding content constant. This suggests that Google is reducing consumer welfare by excluding reviews from other platforms in the OneBox.¹⁵

In the summary quoted above, “engage with” means “click-through.” The more likely users are to click on the Universal Search results to visit the linked pages from which they are derived, according to the study, the more consumer welfare gains are to be realized from those results.

¹³ *Id.* at 9-12 (“Based on a recreation of the FOTUL Paper’s methodology, the very same experiment yields exactly the opposite outcome for another query from FOTUL’s website, *i.e.*, users preferred Google’s supposedly “degraded” results over FOTUL’s. In these circumstances, the FOTUL experiment is plainly unreliable, and its results ought to be disregarded as empirical evidence.”).

¹⁴ *Id.* at 10.

¹⁵ Wu, *et al.*, *supra* note 1, at 1.

But there is nothing (other than the authors' assertion) to suggest that the methodology of measuring click-through rates accurately reflects the extent — or even the direction — of consumer welfare effects.

Information provided without requiring click-throughs may be more valuable than links

First, if a consumer is using a search engine in order to find a direct answer to a query rather than a link to another site to answer it, click-through would actually represent a *decrease* in consumer welfare, not an increase.

In fact, the study fails to incorporate this dynamic even though it is precisely what the authors claim the study is measuring:

According to Google, a principal difference between the earlier cases and its current conduct is that universal search represents a pro-competitive, user-serving innovation. By deploying universal search, Google argues, it has made search better. As Eric Schmidt argues, “if we know the answer it is better for us to answer that question so [the user] doesn’t have to click anywhere, and in that sense we... use data sources that are our own because we can’t engineer it any other way.”¹⁶

Wu, *et al.* acknowledge that universal search (i.e., direct answers) in concept can be welfare enhancing, but question, based on their study, whether it is benefiting consumers in this case:

No one truly disagrees that universal search, in concept, can be an important innovation that can serve consumers. The more challenging question arises not from the deployment of universal search to tell people the weather, but its intentional degradation for exclusionary in areas where Google faces the most serious long-term competition.¹⁷

¹⁶ *Id.* at 12 (quoting *Testimony of Eric Schmidt, Executive Chairman, Google Inc.*, before the Senate Committee on the Judiciary Subcommittee on Antitrust, Competition Policy, and Consumer Rights (September 21, 2011)).

¹⁷ *Id.* at 13.

According to the authors, “[t]he goal of this paper is to test, whether in fact the user experience has been improved.”¹⁸

The problem, however, is that Wu, *et al.* employ a methodology that equates *more clicks* with consumer benefit in order to test whether something that exists “so [the user] doesn’t have to click anywhere” confers consumer benefit. This makes no sense, of course.

One would think, perhaps, that testing whether Universal Search improves the user experience would entail a methodology that equates a *smaller* number of clicks (and hence greater satisfaction with the results as presented) with greater consumer welfare — not the opposite.

But the disconnect between the study and any sensible measure of consumer welfare from Google’s Universal Search is actually even more substantial. By measuring only click-throughs and not the direct, informational aspects of Google’s results, the study systematically misses the most important and valuable source of consumer benefit from universal search — the very thing that Eric Schmidt indicated in the quote cited by Wu, *et al.* was the *intention* of Google’s Universal Search.

Users looking for a phone number, or an address, or any number of other pieces of information don’t need to click *anywhere* to realize value from Google’s Universal Search results. It is hardly surprising that click-throughs would be less common than for less-information-rich organic results (as the study tests), when simply looking at the Universal Search box will often provide the answer a user was searching for.

In one of the two sample Google Universal Search results pages shown in the study, for example, *all* of the Google-only results for the query “pediatrician toronto on” provide telephone numbers.¹⁹ The current version of the study

¹⁸ Wu, *et al.*, *supra* note 1, at 17. Although nothing other than the framing has changed, an older version of the paper even more clearly tied the purported purpose of the paper to the specific claims made by Eric Schmidt above by placing the following directly after the quote from Eric Schmidt: “The object of this paper is to test the truth of that proposition.” See Michael Luca, Tim Wu, and the Yelp Data Science Team, *Is Google degrading search? Consumer Harm from Universal Search*, Working Paper, available at <https://www.law.berkeley.edu/wp-content/uploads/2015/04/Luca-Wu-Yelp-Is-Google-Degrading-Search-2015.pdf> (hereinafter, “Wu, *et al.*, Older Version.”).

¹⁹ Wu, *et al.*, *supra* note 1, at 40.

doesn't include the corresponding results derived from the FOTU plugin, but a previous version of the paper did include such results for an analogous query, "pediatrician nyc." In that case, while, every OneBox result contained a phone number, only a *single* FOTU result contained a phone number (see Figure 1, below).²⁰

For this search in particular, Google's inclusion of results with phone numbers in the OneBox is surely not accidental. The phone number is the critical piece of information that users need in order to make a pediatrician appointment for their child. If that's the information a user was searching for (as it certainly will frequently be, in the real world), it would not at all accurately measure the true consumer value of the results to look only at click-through rates. Yet that is precisely what the study does: It constructs test scenarios and adopts such a crabbed measure of consumer welfare that arguably the most important aspect of the results it purports to be measuring is simply ignored.

Figure 1



Even on its own terms, the study fails to properly measure click-throughs

Even on its own terms, the methodology of the study undermines its claims. The study did not test *actual* click-throughs, but only *potential* click-throughs.²¹ On the study's theory, the measure of value of the links is in the (presumed) content of the destination pages — not the information presented directly in the Universal Search results. But to accurately assess such value, the study would have to consider what happens when a user actually clicks through and judges the relevance of the content on the other side of the links.

²⁰ Wu, *et al.*, Older Version, *supra* note 18.

²¹ de la Mano, *et al.*, *supra* note 8 at 10.

Yet users in the Wu, *et al.* study were shown only a set of static screenshots and asked to compare which set of links they would be likely to click on. This is a critical point: *They were never shown the results of their clicks.* This entirely removes the feedback mechanism that would occur had users been allowed to actually interact with the system, and makes it impossible to assess relative consumer satisfaction with different results.

Moreover, as de la Mano, *et al.* found, many of the searches performed via FOTU contained results that were either completely irrelevant or not within the searched geography.²² For example, in one instance they searched for local plumbers and were given FOTU results that included a hotel:

Evidently, this search result about a hotel does not belong to the category of local business that the user is searching for (plumbers). It just so happens that the webpage about the hotel contains a review where a customer complained about plumbing issues. This is precisely the type of problem that specialized local search results such as Google's are designed to avoid, by understanding local businesses as individual entities with their own database identities, rather than as flat webpages.²³

Thus, even within the cherry picked sample that Wu, *et al.* selected, the results seem likely to be highly irrelevant in one manner or another. But without allowing users to actually access that content, the true relevance of the results could not be determined.

At best, the results of Wu *et al.*'s study could be said to support the conclusion that sometimes FOTU search results might superficially *look* better, to some users, in some stylized situations. But, again, not only does this completely miss out on the power of large data processing that entity search methodology leverages, it has also been refuted as a useful study design.

Click-throughs don't measure consumer satisfaction with chosen results

Critically, this failure to assess user interaction with the links or to otherwise assess how the user would react to universal search results means that the study also makes no effort to actually measure consumer satisfaction (*i.e.*, consumer welfare).

²² *Id.* at 7-9.

²³ de la Mano, *et al.*, *supra*, note 8 at 7-8.

In particular, the study does not assess how many searchers might have regretted their click when they found themselves faced not with a site for their chosen doctor's office or local restaurant, for example, but rather another collection of links on Yelp. Nor does it indicate how many users would have returned to the original results page and clicked on another result, found that useless, returned, and kept clicking until they finally found something useful. In such a scenario more click-throughs would correlate with *decreased* consumer welfare. Similarly, the study also fails to attempt to assess whether users were happy to avoid additional clicking by getting their desired answer from Google in the first place.

For instance, if a user searches in Google for "restaurants near me," she might want (and Google seems to think she wants) to see a map with addresses, phone numbers, some indication of others' ratings, and a set of links to the restaurants' own websites. The user doesn't necessarily want a link to Yelp or another website with similar information generated by a different algorithm or based on different data. And, crucially, if she did, she could easily use the Yelp app, navigate directly to Yelp.com, or likely find Yelp links at or near the top of the organic search results.

In this respect the study is effectively comparing apples and oranges, assuming that answers to queries (with direct links to the results) are the same thing as (or worse than) links to third party sites that offer nothing more than another site's presentation of the results, likely requiring more click-throughs — and more time spent online — in order to answer users' queries.

Rather than analyze the entire user experience to determine whether consumer welfare was enhanced by their alternative results, Wu, *et al.* use a partial snapshot in a way that skews the analysis in favor of only one kind of result. The FOTU plugin (and the study) presumes that searchers are looking for results incorporating more reviews and links to vertical search sites (presumably precisely because that's the business these competing search engines are in). But, again, for searches performed on a general search engine like Google, it is at least as likely that a user is seeking a direct answer as it is that she is seeking a comprehensive set of reviews.

Fundamentally, the sort of computer intelligence necessary to determine the answer to a question is completely distinct from that necessary to determine which set of linked pages may be most relevant to a query. And where a query suggests a direct question, looking for a simple answer, fewer click-throughs

indicate a *better, more relevant* set of search results than do results that generate more clicks.

Unfortunately, the study is constructed in such a way that it simply doesn't reflect how users use Google in the real world, and how Google has responded to these users.

3. The study is designed with biased search terms and biased search intent that doesn't accurately represent the way Google is used

Unlike some previous critics,²⁴ the authors of the study attempt to measure purported harm rather than just simply presuming that prioritization of Google-owned properties is anticompetitive. Nevertheless, the study is based upon cherry-picked terms (a *single* cherry-picked term, in fact) for local services, rather than on the broad range of informational queries that make up the majority of Google searches.²⁵ Moreover, it tests only scenarios designed to direct test subjects to vertical search engine links — scenarios that don't adequately represent the range of searches for a general search engine like Google.

Biased search terms and biased test scenarios skew the results

For the entirety of its data set, the study relies upon a single search string, "coffee louisville ky,"²⁶ meant to represent the sort of search used to find local products and services by someone not familiar with an area — a fact built into the study but not, of course, otherwise known to a search engine interpreting search queries in the abstract. In addition, to prime the users in the study, the researchers asked what results they would click in the following scenario:

You're visiting Louisville, Kentucky for the first time with a friend. Your friend suggests you both get a cup of coffee. You type "coffee louisville ky" into Google and click "Search", which brings you to the following screen. What do you click first?²⁷

²⁴ Cf. Edelman & Lockwood, *supra* note 3.

²⁵ See Bernard J. Jansen, Danielle L. Booth, & Amanda Spink, *Determining the informational, navigational, and transactional intent of Web queries*, 44 INFO. PROCESSING & MGMT. 1251 (2008).

²⁶ Wu, *et al.*, *supra* note 1, at 24.

²⁷ *Id.* at 41, Figure 3.

Unfortunately, by using such a limited and non-representative sample of terms and by priming the participants in this way, the results derived from such a sample will necessarily fail to present a representative picture of Google's conduct or that of its users.

First, by designing the study around only one local search term, the authors incorporate a misleading bias into the study by imbuing testers with a pre-determined (but not necessarily transparent) search intent that doesn't necessarily represent the intent of typical Google users who may want the answers provided in Google's Universal Search results. As a consequence, the authors import their bias regarding what the "correct" answer to such terms should look like, effectively assuming, in this case, that results are accurate only when they present the sort of specialized information found on competing local search sites like Yelp:

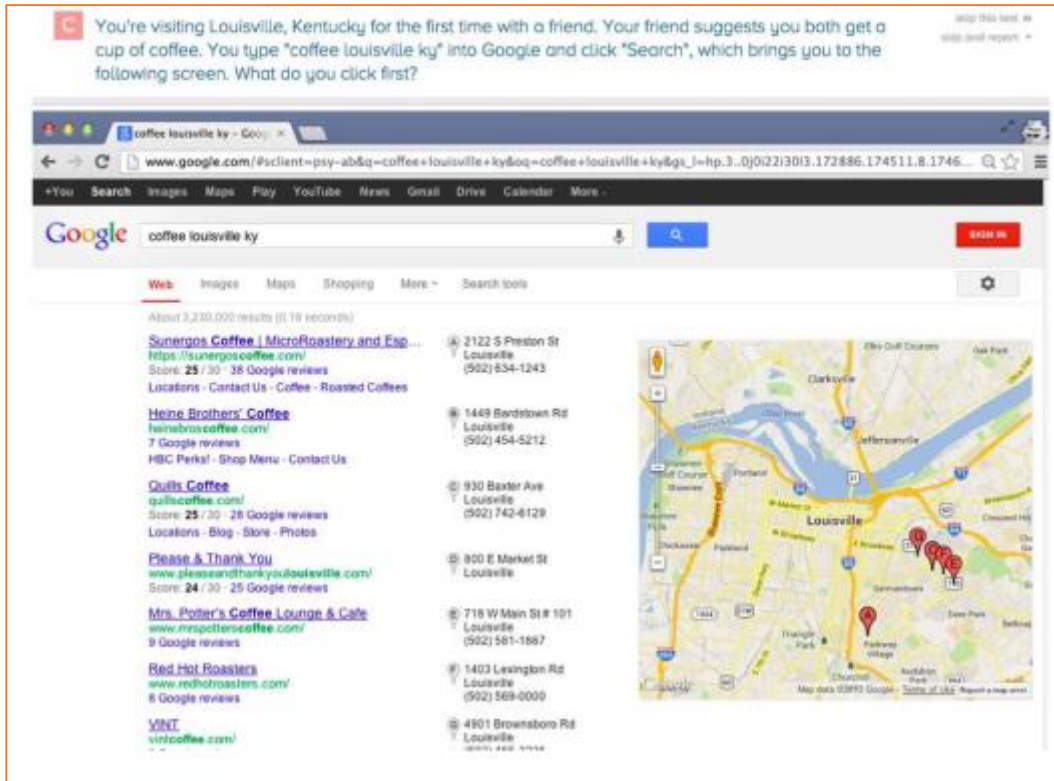
However, Google's organic search employs a merit-based algorithm that can easily be used to identify **better candidates to populate its local search boxes**... (Emphasis added).²⁸

"Better" here is presumed to be results incorporating more user reviews, for example. But, as mentioned above, Google, as a general search site, may well be optimizing different things in its Universal Search results.

Having set up the study to feed testers only with screenshots of search results for which Yelp-like results are more likely to be appropriate, it is a foregone conclusion that some testers may not prefer Google's Universal Search results. But this doesn't mean that Google is offering less-relevant results in its Universal Search box; it means only that the study is flawed.

²⁸ Wu, *et al.*, *supra* note 1, at 4.

Figure 2



The problem is exacerbated by priming the user into thinking they were visiting Louisville for the first time, and looking for a coffee shop in an unfamiliar city. Such a scenario virtually ensures that a searcher is looking to find one thing: a set of options from which she can *pick* a place to get coffee, and, most likely, a list of *ranked* options (given local unfamiliarity).

Unlike the authors of the study, Google didn't already know that the tester was primed in this way. Operating in relative ignorance and based only on inferences derived from search terms or other cues, Google tailors its results to provide answers to the primed question, *but also to enable the user to obtain other information*. For instance, a user already familiar with Louisville may search for "coffee louisville ky" because she wants: a list or map of coffee shops to assist in remembering nearby options and making a quick choice; the phone number of a particular coffee shop to check operating hours; a quick address check before driving off; or any of a wide range of other things. Further, a user not primed to select a coffee shop for a social occasion may also care less about the store ratings in which Yelp specializes — for example, if she just wants to quickly grab a coffee nearby before leaving for the airport.

And, crucially, the predictive analytics that function well for answering a question — for instance using learning algorithms to understand how previous users have interacted with the results for a particular search query — are wholly distinct from the data techniques (like PageRank) used to analyze a static web page for keyword relevance.

In other words, Google, as a general search engine, optimizes results for users searching for a much wider range of answers than the limited experience of the primed test user in the Wu, *et al.* study. In contrast, Yelp, TripAdvisor, and the rest of the specialized search services provide exactly what the test user has been told to prefer. So it's hardly surprising that Google's results don't map perfectly onto the testers' preferences when they are essentially told to look for the kind of results they might find on Yelp, rather than on Google.

Instead, all this demonstrates is that users who know that they are looking for Yelp-like results may prefer Yelp-like results. But presumably users who know this either start such searches on Yelp in the first place, or else click on Yelp (or one of its competitors) in their Google search results.

As it happens, this occurs quite frequently.²⁹ And far from demonstrating that Google limits access for Google users to Yelp and similar sites, Yelp itself acknowledges that Google searches are among the most significant sources of traffic for Yelp.³⁰

²⁹ See Danny Sullivan, *Leaked Documents Show How Yelp Thinks It's Not Getting Screwed by Google*, Search Engine Land (Jul. 9, 2014), <http://searchengineland.com/leaked-documents-yelp-google-196208>.

³⁰ See, e.g., Jason Kincaid, *Stoppelman: 75% Of Yelp's Traffic Comes From Google*, TECHCRUNCH (Sept. 21, 2011), <https://techcrunch.com/2011/09/21/stoppelman-75-of-yelps-traffic-comes-from-google/>. Indeed, a core theory that motivates Google's antitrust provocateurs is that it functions as something like an essential facility upon which their entire business depends. But between the continuing relevance of traffic from Google and the increasing relevance of traffic from Yelp's own app, this claim is unsustainable:

Approximately 21 million unique devices accessed Yelp via the mobile app on a monthly average basis in the first quarter of 2016, an increase of 32% compared to the same period in 2015. App users viewed approximately 70% of page views in the first quarter and were more than 10 times as engaged as website users, as measured by number of pages viewed.

Yelp Announces First Quarter 2016 Financial Results, YELP (May 5, 2016), <http://www.yelp-ir.com/phoenix.zhtml?c=250809&p=irol-newsArticle&ID=2165708>.

The FOTU plug-in systematically directs users away from the direct answers Google seeks to provide

Moreover, the FOTU plugin at the heart of the study systematically redirects users *away from* direct links to local businesses and toward third party search engines like TripAdvisor and Yelp instead. In many cases, this is unlikely to yield results corresponding to what searchers who have affirmatively chosen to use Google instead of going directly to one of those sites are looking for in their Universal Search results.

To see this, recall the “coffee louisville ky” example. The prompt was engineered, first, to *require* a (hypothetical) click-through (thus excluding the possibility that merely viewing a phone number in a Universal Search box was optimal), and, second, to encourage test subjects to prefer results that would likely offer a full list of options rather than a link to any specific coffee shop. This is what FOTU does — but it hardly represents the sum total of what real-world searchers are looking for on Google.

Moreover, because it biases “correct” results toward sites that offer such lists, FOTU necessarily contemplates multiple clicks (on multiple sites). To the extent that Google seeks to *minimize* clicks (as Schmidt indicated³¹), there is a clear disconnect between what Google is optimizing and what Wu and his co-authors assert promotes consumer welfare. The study simply *assumes* that conclusion; it does nothing to actually demonstrate it.

4. The authors’ interpretation of the study’s results is unsupported by the data

Perhaps reflecting their bias toward Yelp-like results, the authors write that:

³¹ *Testimony of Eric Schmidt, Executive Chairman, Google Inc.*, before the Senate Committee on the Judiciary Subcommittee on Antitrust, Competition Policy, and Consumer Rights (September 21, 2011), *available at* <http://www.judiciary.senate.gov/meetings/the-power-of-google-serving-consumers-or-threatening-competition>.

Suppose you are planning a trip to Louisville, Kentucky and are searching for a coffee shop through Google. Clearly, there is a wide variety of content that might facilitate this search. Competing ratings and reviews ranging from Yelp to TripAdvisor to Food & Wine invest heavily in developing such content. In this situation, Google's content may be more or less useful to users than other content. If Google provides favorable placement to Google content in a world in which Food & Wine is — on average — more useful, then this creates harm to consumers.³²

Overall, these findings show that users prefer FOTU content to Google content, and that this is mediated by the fact that FOTU contains more reviews. By intentionally excluding reviews from other platforms, Google is hence degrading the quality of its product.³³

But, as we discuss below, the study doesn't consider that users may be misled by the FOTU results in the context in which they are presented in the study. That is, even if the study seems to show that users prefer FOTU results to OneBox results, that conclusion may be limited to the extremely stylized and narrow circumstances presented to the study's test subjects.

More doesn't necessarily equal better

To begin with, the conclusion that more reviews equates with higher quality search results cannot be correct. It is surely the case that specialized search engines like Yelp believe that their own collections of reviews are more accurate and more useful than their competitors'. But it isn't clear why it should necessarily be the case that incorporating hundreds of reviews from other sites (some of which users may never have heard of) yields "better" results than incorporating dozens of reviews from a single, more-familiar and possibly more-reliable source. Whatever the test subjects' initial propensity to prefer OneBox results with more reviews demonstrates, it doesn't necessarily demonstrate that users are well-served by doing so. As we discussed above, because the study didn't evaluate testers' actual preference for the FOTU results, and because of the known inaccuracies of those results (which, presumably, users would have discovered if they were permitted to actually

³² Wu, *et al.*, *supra* note 1 at 3-4.

³³ *Id.* at 26-27.

click-through), the assertion that consumer welfare is being harmed by Google's unadulterated results because they offer fewer reviews is unsupported.

The study's design may mislead users

Of course, all of this still begs the question of why users, as represented by the study's small cohort of test subjects and single test search, nonetheless systematically chose the screen shots that showed results with more reviews than the screen shots that had results presenting only Google's own review data.

One possibility is that the test subjects in the study, expecting to find Google results in the Universal Search box, may think it is in fact *Google* directing them to the FOTU-generated links (instead of the links being generated in the same fashion as the organic results below). This could indicate any number of things to users. Most obviously, it could suggest that Google has determined that it doesn't have information on the specific query sufficient to present its typical Universal Search results, and that the best option for that query is to obtain more information by clicking through to another search engine. Users who know and trust Google are likely to follow this suggestion.

On the other hand, it could simply suggest that the appropriate Universal Search results are no different than the organic results below, and thus that Google has determined that putting the organic results into the Universal Search box is the most efficient thing to do. In fact, this is exactly what the FOTU results are designed do: to mimic the organic results that follow. But Google users who realize this would quite logically be less likely to click on the subsequent, identical presentation of organic results instead of the ones in the Universal Search box.

The study fails to measure the appropriate baseline

It is also important to note what else the study fails to measure. The authors didn't look into Bing or other general search engines to determine whether Google's practices were uniquely biased, or even perhaps *less biased* than the practices of competitors without any market power. If Bing or another search engine with less market power were actually *more* biased, for instance,³⁴ then a

³⁴ See Wright, *Defining and Measuring Search Bias*, *supra* note 4, at 50 ("When search engines appear to rank their own content more prominently in these samples, it is clear both Bing and Google do so. Indeed, Bing appears to be more biased than Google.").

quite plausible inference would be that this conduct is actually procompetitive product differentiation rather than anticompetitive exclusion.

The appropriate baseline comparison for Google's actual conduct can't simply be an artificially constructed world that might not exist even in the most textbook-perfect competitive market (and which others have found not to even reliably exist in repetitions of the study³⁵). To make any antitrust sense at all, the relevant baseline (particularly when measuring anticompetitive foreclosure) should be what would *in fact* prevail in a competitive market. Without even assessing what other search engines without large market shares are doing, the authors can't establish with any certainty that Google's conduct deviates at all from what one would expect to see in the absence of market power or anticompetitive incentives.

Moreover, the study glosses over the fact that users *do* in fact engage with the organic search results that appear on the same page as the Universal Search results they criticize. Nor does it mention the fact that many consumers directly navigate to these other sites, or that an increasing number of mobile users access their content through apps and the mobile web.³⁶

The reality is that Google doesn't tinker with the list of organic links generated by its algorithm to show the Universal Search results; it actually includes them on the same page as its Universal Search results. What Google *does* do is present an additional, adjunct collection of information in its Universal Search results box. From an antitrust perspective, this means that Google is not — at least not in any way based on the evidence the study presents — foreclosing these other information sources from access to Google's users (or vice versa). In fact, it is even possible that the presence of the OneBox information actually *increases* users' incentives and ability to engage with the organic search results below by providing first-cut information that may help users evaluate the organic results.

Economic and Legal Problems

As the above begins to suggest, aside from its serious methodological problems, the paper also reflects poor analysis of antitrust law and economics.

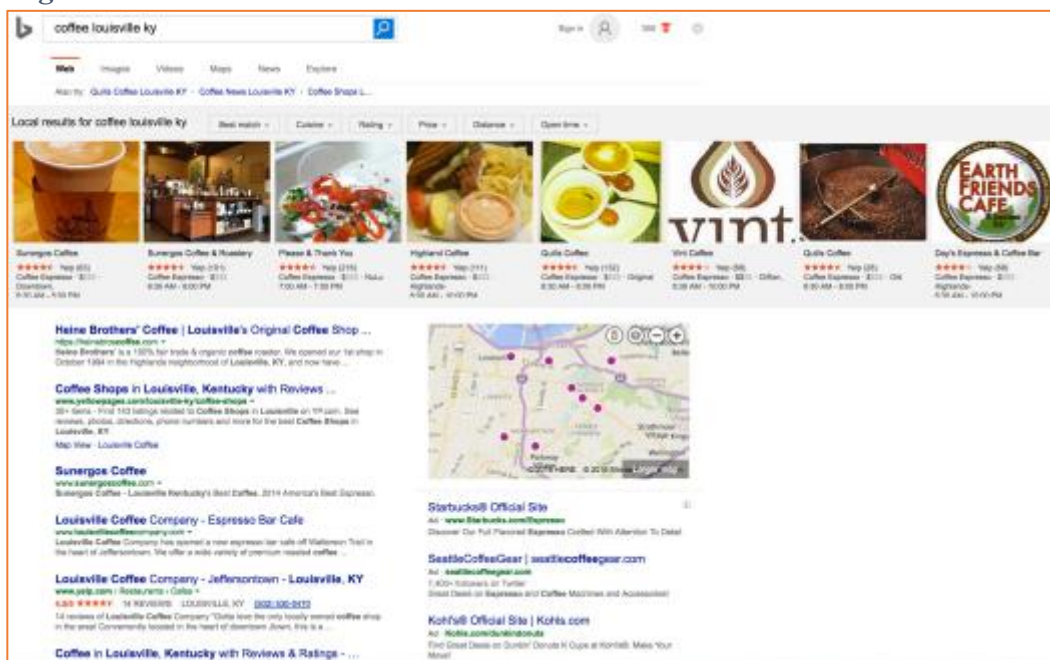
³⁵ See generally de la Mano, *et al.*, *supra*, note 8.

³⁶ For example, Yelp recently reported to investors that the vast majority of its page views now come via its mobile app —hardly the fate of a stifled competitor. See *Yelp Announces First Quarter 2016 Financial Results*, *supra* note 30.

I. “Bias” is not inherently anticompetitive

If companies that lack any market power in search engage in similar practices, it suggests that those practices may be procompetitive. And, indeed, Bing presents its preferred “Carousel” results at the top of its search results pages, and Yahoo! offers similar universal search results at the top of its results pages, as well (Figures 3 and 4, below) — neither of which perfectly tracks the organic results that follow, and neither of which is exactly the same as the other or Google. What’s more, as it happens, because of special deals with Yelp, each prioritizes Yelp content in its universal search results.³⁷ At minimum it is problematic to infer abuse of dominance from conduct if non-dominant firms also engage in it.

Figure 3



³⁷ See, e.g., Jason Abbruzzese, *Yahoo Taps Yelp to Boost Search Results*, Mashable (Mar. 12, 2014), <http://mashable.com/2014/03/12/yahoo-yelp-partnership/#Dtpk5e.tQ8qn>.

Figure 4

YAHOO! coffee louisville ky Search

Web Images Video More Anytime

Also try: john conti coffee louisville ky, heine bros coffee louisville ky

related to: coffee louisville

Coffee near Louisville, KY

1 **Quills Coffee**
Coffee House · \$ · Closed · ★★★★★ 132 on Yelp
930 Baxter Ave, Louisville, KY · (502) 742-6129

2 **Highland Coffee**
Coffee House · \$ · Closed · ★★★★★ 111 on Yelp
1140 Bardstown Rd, Louisville, KY · (502) 451-4545

3 **Graehaus**
Coffee House · \$\$ · Closed · ★★★★★ 147 on Yelp
1001 Baxter Ave, Louisville, KY · (502) 454-7075

4 **Sunergos Coffee**
Coffee House · \$ · Closed · ★★★★★ 102 on Yelp
2122 S Preston St, Louisville, KY · (502) 634-1243

5 **Sunergos Coffee**
Coffee House · \$ · Closed · ★★★★★ 63 on Yelp
231 S 6th St, Louisville, KY · (502) 589-3222

See all 325 results Add or update your business

Heine Brothers' Coffee | Louisville's Original ...
heinebroscoffee.com
Heine Brothers' is a 100% fair trade & organic coffee roaster. We opened our 1st shop in October 1994 in the Highlands neighborhood of Louisville, KY, and now have ...

Coffee in Louisville, Kentucky with Reviews &...
www.yellowpages.com/louisville-ky/coffee
Find 171 listings related to Coffee in Louisville on YP.com. See reviews, photos, directions, phone numbers and more for the best Coffee & Espresso Restaurants in ...

Sunergos Coffee
www.sunergoscoffee.com
Sunergos Coffee - Louisville Kentucky's Best Coffee. 2014 America's Best Espresso

Louisville Kentucky
www.Amazon.com/books
Low Prices on Millions of Books. Free Shipping on Qualified Orders!

Coffee Louisville Ky
Shopping.net/Coffee Louisville Ky
Find Our Lowest Possible Price! Coffee Louisville Ky for Sale

Louisville Kentucky
www.TripAdvisor.com/Louisville
What's the #1 Hotel in Louisville? See Deals for Louisville Hotels!

Louisville Kentucky
Info.com/Louisville Kentucky
Get Louisville Kentucky Info. Your Search & Social Results Now.

John Conti Coffee -70% - Buy cheap John Conti Coffee & save.
www.John-Conti-Coffee.best-deal.com
Buy cheap John Conti Coffee & save. Compare offers at BEST-DEAL.com!

The inclusion of Google-specific (or Yelp-specific, for that matter) search results in a Universal Search box might occur because of (procompetitive) product differentiation. If real-world consumers use Google or Bing or another general search engine because they prefer the type of tailoring that search engine applies — because, in other words, that search engine offers more consistently relevant results — then it is difficult to make the case that the tailoring is harmful to consumer welfare, regardless of what A/B tests suggest.

Former FTC Commissioner Josh Wright’s analysis of a previous study purporting to find search “bias” in Google’s results applies with equal force here:

[I]t is critical to recognize that bias alone is not evidence of competitive harm and it must be evaluated in the appropriate antitrust economic context of competition and consumers, rather individual competitors and websites... [but these results] are not useful from an antitrust policy perspective because they erroneously—and contrary to economic theory and evidence—presume natural and procompetitive product differentiation in search rankings to be inherently harmful.³⁸

Google is a vertically integrated company that offers general search, but also a host of other products, including email, calendars, cloud sharing services, social networking, maps, translation services, browsing, mobile operating systems, rating systems, online payment systems, video streaming, targeted advertising, and Internet access services, among other things.³⁹ With its well-developed algorithm and wide range of products, it is hardly surprising that Google can provide not only direct answers to factual questions, but also a wide range of its *own* products and services that meet users' needs. If consumers choose Google not randomly, but precisely because they seek to take advantage of the direct answers and other options that Google can provide, then removing the sort of "bias" alleged by Wu, *et al.* would affirmatively hurt, not help, these users.

2. There is no legal basis for imposing upon Google a duty to deal with competitors like Yelp

The study's legal conclusions make a demand of Google that is rarely made of any company (even vertically integrated ones): they essentially amount to a *de facto* demand that it must provide access to its competitors on equal terms. The way the study's authors hope to accomplish this is by enjoining Google to include only those results in its OneBox that are ranked as relevant by Google's PageRank algorithm.⁴⁰ Under his theory, if Google were forced to use

³⁸ Wright, *Defining and Measuring Search Bias*, *supra* note 4, at 19.

³⁹ See Google>About Google>Products, <http://www.google.com/about/products/> (last accessed Sept. 13, 2016).

⁴⁰ Wu, *et al.*, *supra* note 1 at 34-35. It bears noting that this would essentially amount to an injunction to remove the Universal Search results altogether. After all, if Google were forced to show results determined by PageRank, and those same results already occur in the organic results on the page, what would be the point of maintaining a duplicative list of results in the OneBox?

PageRank as the metric for the One Box, Google's competitors would be better able to reach consumers with their (purportedly) superior results.

But one can hardly imagine a valid antitrust complaint arising because McDonald's refuses to list Burger King's menu alongside its own.⁴¹ Nor can one conceive of Burger King being liable for harming consumer welfare because A/B tests suggest that consumers would prefer McDonald's french fries with their Whopper.⁴² The law on duties to deal is heavily circumscribed for good reason. As the Supreme Court observed in *Trinko*, the imposition of a duty to deal would threaten to "lessen the incentive for the monopolist, the rival, or both to invest in... economically beneficial facilities."⁴³

Requiring Google to link to other powerful and sophisticated online search companies on the precise terms it links to its own information, or instead of providing direct answers in its search results, would reduce the incentives of everyone to invest in their underlying businesses to begin with.⁴⁴

⁴¹ See Ben Sperry, *Google and Antitrust: Economic Liberty in the Balance*, OPENMARKET (Jun. 21, 2012), <https://cei.org/blog/google-and-antitrust-economic-liberty-balance>.

⁴² Joshua Wright (@ProfWrightGMU), Twitter (Jun. 30, 2015, 8:43 PM), <https://twitter.com/ProfWrightGMU/status/616089693864726529> ("Would anyone find it relevant to an antitrust claim against Burger King if a survey showed consumers love when u replace BK fries with McDs?").

⁴³ *Verizon Communications, Inc. v. Law Offices of Curtis v. Trinko, LLP*, 540 U.S. 398, 407-08 (2004). See also, Geoffrey A. Manne & Joshua D. Wright, *If Search Neutrality Is the Answer, What's the Question?*, 2012 COLUM. BUS. L. REV. 151, available at <http://ssrn.com/abstract=1807951> ("Because imposition of a duty to deal with rivals threatens to decrease the incentive to innovate by creating new ways of producing goods at lower costs, satisfying consumer demand, or creating new markets altogether, courts and antitrust agencies have been reluctant to expand the duty.").

⁴⁴ Moreover, as a doctrinal legal matter the "outer boundary" of liability (that is, for treating a company as an "essential facility," meriting a rare exception to the principle described above) referred to in *Aspen Skiing* (*Aspen Skiing Co. v. Aspen Highlands Skiing Corp.*, 472 U.S. 585 (1985)) and *Trinko* likely doesn't apply here, anyway. While the study's authors suggest that Google has lessened product quality in order to exclude rivals, it isn't clear that Google has actually given up a previously profitable course of conduct with Yelp or any other vertical search engine, as the case law would require. See *Trinko*, 540 U.S. at 409 ("The Court [in *Aspen Skiing*] found significance in the defendant's decision to cease participation in a cooperative venture... The unilateral termination of a voluntary (and thus presumably profitable) course of dealing suggested a willingness to forsake short-term profits to achieve an anticompetitive end."). In this case, in fact, it is Yelp that ended its (presumably profitable) arrangement with Google — and that now, ironically, wants to use the antitrust laws to force Google to comply with its further change of heart. See, e.g., MG Siegler, *Not Only Is Google*

Moreover, and fatal to Yelp's case, Google *still* shows organic search results on pages that contain the OneBox. And those search results are ranked using the PageRank algorithm (among other things). All Google has done is add an additional, valuable tool for users, albeit one that may not prominently feature competitors in the specialized search space. Even assuming some users were "duped" into using "degraded" search results in the OneBox one or two times, however, presumably those users would easily learn to skip over inferior results in the OneBox and resume relying on the organic results lower down on the page, just as they generally skip over paid results when they appear and offer less relevant links. At root, this is not so much a case of exclusion sufficient to merit the imposition of a duty to deal as it is one of churlish complaining over Google's specific (and procompetitive) placement of particular results within Google's own property.⁴⁵

3. Antitrust law protects competition, not competitors' idiosyncratic wish lists

Antitrust law is concerned with preserving competition (*i.e.*, access to competitors), not with satisfying every competitor's demands regarding *how* exactly that competition should be preserved. Yelp may prefer to dictate to Google the precise terms of how and when Yelp's content should appear on Google's site, but this isn't an antitrust-relevant concern. As long as *consumers* are not being unreasonably denied access to Yelp (and, indeed, consumers are accessing Yelp via *myriad* sources in addition to Google, as noted above), Yelp's idiosyncratic preferences for precisely how *Google* furthers that access is of no antitrust moment. Google continues to prominently display Yelp in organic and paid results. Nothing in the study (or anywhere else) demonstrates that Yelp is being anticompetitively excluded just because Google directly

Places Going After Yelp, They're Doing So With Yelp's Content, TECHCRUNCH (Jul. 26, 2010), <https://techcrunch.com/2010/07/26/google-yelp/>.

⁴⁵ In other words, this is analogous to a grocery store slotting arrangement. But the availability of access via other channels, and Yelp's overall (and increasing) success, render claims of foreclosure hollow. See Benjamin Klein & Joshua D. Wright, *The Economics of Slotting Contracts*, 50 J. L. & ECON. 421 (2007) ("The primary competitive concern with slotting arrangements is the claim that they may be used by manufacturers to foreclose or otherwise disadvantage rivals, raising the costs of entry and consequently increasing prices. It is now well established in both economics and antitrust law that the possibility of this type of anticompetitive effect depends on whether a dominant manufacturer can control a sufficient amount of distribution so that rivals are effectively prevented from reaching minimum efficient scale.").

answers some queries using its own properties in its Universal Search results box.

Wu and his co-authors struggle to make a point out of the FTC's closed investigation into Google's allegedly anticompetitive conduct:

The FTC, in its closing of the case, stated, based on the evidence it had, that “the documents, testimony and quantitative evidence the Commission examined are largely consistent with the conclusion that Google likely benefited consumers by prominently displaying its vertical content on its search results page.” ... [But w]hile the full scope of the evidence reviewed is not public, it is evident that the Commission did not have the benefit of randomized controlled trials when it offered that conclusion in 2013. The goal of this paper is to test, whether in fact the user experience has been improved. As we stressed, in some areas, the results suggest the opposite.⁴⁶

Despite their efforts to downplay the FTC decision not to pursue Google on antitrust charges, the FTC's view of the law and facts clearly cut against Wu, *et al.*, regardless of how their study came out. On the one hand, the FTC staff (note, not the Commission itself, but an internal staff memorandum⁴⁷ that was not adopted by the Commission) recommended the FTC *not* sue Google on these claims and concluded that, at worst, Google was *both* providing users with the right answer, *as well as* driving traffic to its own sites.⁴⁸ In the language of antitrust that's called a “procompetitive justification,” meaning that, even if there might be identifiable negative effects for some rivals, they wouldn't be enough to make out an antitrust case if they were incidental to legitimate and consumer-welfare-enhancing business justifications — like getting the “right” result.

Moreover, driving traffic to one's own sites isn't inherently anticompetitive. Just as grocery stores may stock their store brands in key shelf spaces, and just as Coke isn't required to carry Pepsi in its vending machines, companies are free to try to bolster their own products; in fact, that's precisely what vigorous competition anticipates.

⁴⁶ Wu, *et al.*, *supra* note 1 at 17.

⁴⁷ See FTC Staff Report, Google Inc., File No. 111-0163, *available at* <http://graphics.wsj.com/google-ftc-report/>.

⁴⁸ Wu, *et al.*, *supra* note 1 at 17.

And one final point on legality: There is no viable remedy here. Finding a duty to deal would require central planning of Google's algorithms by antitrust agencies or courts — remedies both impractical and outside of normal antitrust practice.

Wu, *et al.*'s suggested remedy (assuming a finding of liability) is to enjoin Google as follows:

For those searches which trigger its local "OneBox" result, the firm must henceforth use the results of its own PageRank algorithm to populate that OneBox, as opposed to arbitrarily populating the OneBox with its own properties.⁴⁹

But, contrary to the authors' claim, even this remedy would end up requiring intrusive and impractical judicial or regulatory oversight, would be technologically problematic, or, at best, would amount to a total prohibition on Google's trying to optimize the user experience in any way that deviated from PageRank-based results.

Moreover, such a remedy would be unworkable because, as soon as Google decides to change any part of the PageRank algorithm (as it regularly does), the exact same sort of complaints would arise, so long as some searches end up valuing Google's properties over rivals'. Courts would have to get into the business of reviewing Google's programming code in order to ensure compliance with extended legal remedies. In the end, such an injunction would have to lock in *today's* PageRank algorithm, effectively demanding an end to Google's local search OneBox results as well as any other, future innovations that might in any way preference Google's properties — an outcome that would certainly leave consumers worse off.

4. Yelp hasn't been harmed by Google — quite the opposite, in fact

Regardless, it isn't remotely clear that Yelp (or its users) has been significantly — or in any way — harmed by Google's conduct. If anything, Google is one

⁴⁹ Wu, *et al.*, *supra* note 1 at 35.

of Yelp's greatest benefactors: Yelp's traffic from Google has increased faster than Google's own traffic, in fact.⁵⁰

Yelp, meanwhile, has done quite well for itself by engaging in its *own* forms of "bias."⁵¹ It has successfully struck deals for its content to be displayed (to the exclusion of, or on more preferable terms than, other content) on a number of services including Apple Maps, Siri, Bing, and Yahoo!.⁵² And Yelp CEO Jerry Stoppelman has repeatedly asserted that Yelp's market position is robust because of the strength of its mobile app and the company's position as the "de facto local search engine."⁵³ It is not clear that Google has the ability to affect Yelp's competitive position any more than Yelp has the ability to affect Google's.

Further, and ironically, the entire exercise undertaken by the study's authors actually shows there are tools available (until recently the FOTU plug-in, for example) to consumers to circumvent Google's tailoring, if they so choose, while still taking advantage of the company's algorithm. This suggests that there are, in fact, many choices available in the marketplace, and that users choose Google because they generally like the service they are receiving.

⁵⁰ See, e.g., Eric Schmidt, *The New Grundergeist*, GOOGLE EUROPE BLOG (Oct. 13, 2014), <https://europe.googleblog.com/2014/10/the-new-grundergeist.html> ("This issue of providing direct answers to questions is at the heart of complaints being made about Google to the European Commission. Companies like Expedia, Yelp, and TripAdvisor argue that it deprives their websites of valuable traffic and disadvantages their businesses. They'd rather go back to 10 blue links. What's interesting is that the traffic these websites get from Google has increased significantly — faster in fact than our own traffic — since we started showing direct answers to questions.").

⁵¹ See, e.g., *Yelp Announces First Quarter 2016 Financial Results*, *supra* note 30.

⁵² See Douglas MacMillan & Daisuke Wakabayashi, *Yahoo to Partner With Yelp on Local Search*, THE WALL STREET J. (Feb. 10, 2014), <http://www.wsj.com/articles/SB10001424052702304680904579371263386333816>. See also Figures 3 and 4 above, in which Yelp's reviews are given priority placement.

⁵³ See *Yelp (YELP) Q4 2014 Results - Earnings Call Transcript*, SEEKING ALPHA (Feb. 5, 2015), available at <http://seekingalpha.com/article/2891066-yelp-yelp-q4-2014-results-earnings-call-transcript> ("Over the last couple of years, we've participated in the industry-wide shift away from SEO-driven desktop traffic to a mobile first world. We believe that local is the perfect space for mobile and that Yelp is uniquely positioned to succeed."); *Yelp's CEO Discusses Q2 2012 Results - Earnings Call Transcript*, SEEKING ALPHA (Aug. 2, 2012), available at <http://seekingalpha.com/article/773261-yelps-ceo-discusses-q2-2012-results-earnings-call-transcript> ("the underlying power of the Yelp model, by focusing almost singularly on cultivating rich, authentic local content, we've created a site that is rapidly becoming the de facto local search engine for connecting consumers to create local businesses.").

Yelp, represented by the study's authors, may find fault with Google's search results because they would *prefer* something they perceive as better for Yelp. But in terms of the evidentiary, economic, and legal prerequisites for converting such a preference into a successful antitrust case, the Wu, *et al.* study falls woefully short.