

Defining and Measuring Search Bias:

Some Preliminary Evidence

Joshua D. Wright*

October 20, 2011

* Professor, George Mason University School of Law and Department of Economics. I thank Elyse Dorsey, Stephanie Greco, Whitney Scherck, and Katie Schewietz for excellent research assistance. The International Center for Law & Economics (ICLE) provided financial support for this project. ICLE has received financial support from several companies and individuals, including Google. The ideas expressed here are the authors' and do not necessarily reflect the views of ICLE, its board of directors, advisors, affiliates or supporters.

INTRODUCTION

Search engines produce immense value by identifying, organizing, and presenting the Internet's information in response to users' queries.¹ Search engines efficiently provide better and faster answers to users' questions than alternatives. Recently, critics have taken issue with the various methods search engines use to identify relevant content and rank search results for users. Google, in particular, has been the subject of much of this criticism on the grounds that its organic search results—those generated algorithmically—favor its own products and services at the expense of those of its rivals.

It is widely understood that search engines' algorithms for ranking various web pages naturally differ. Likewise, there is widespread recognition that competition among search engines is vigorous, and that differentiation between engines' ranking functions is not only desirable, but a natural byproduct of competition, necessary to survival, and beneficial to consumers.² Nonetheless, despite widespread recognition of

¹ Yan Chen et al., *A Day without a Search Engine: An Experimental Study of Online and Offline Search* (Nov. 15, 2010), http://yanchen.people.si.umich.edu/papers/VOS_20101115.pdf (finding that the average search time online is only 7 minutes, whereas the average search time offline is 22 minutes); Hal Varian, *Economic Value of Google* (PowerPoint presentation) (on file with author) (estimating that Google provides \$65 billion of value to consumers in time saved). See also KRISTEN PURCELL, PEW INTERNET & AM. LIFE PROJECT, SEARCH AND EMAIL STILL TOP THE LIST OF MOST POPULAR ONLINE ACTIVITIES 2-3 (2011), available at http://pewinternet.org/~media/Files/Reports/2011/PIP_Search-and-Email.pdf (finding that search engine use among all Americans surged from 52% in January 2002, to 72% in May 2011).

² See Danny Sullivan, *Study: Google "Favors" Itself Only 19% of the Time* (Jan. 19, 2011, 5:22 PM), <http://searchengineland.com/survey-google-favors-itself-only-19-of-the-time-61675>; Tom Zeller, Jr., *Gaming the Search Engine, in a Political Season*, N.Y. TIMES (Nov. 6, 2006), <http://www.nytimes.com/2006/11/06/business/media/06link.html> ("And while competition dictates that as search engines get better at this, their results will be similar, they aren't precisely the same. Each

the consumer benefits of such differentiation, complaints from rival search engines have persisted and succeeded in attracting attention from a number of state, federal and international regulatory agencies. Unfortunately, much of this attention has focused on the impact upon individual websites of differences among search engines' algorithmic methods of identifying and ranking relevant content, rather than analyzing these differences from a conventional consumer-welfare driven antitrust analysis. For example, many of these complaints ignore the fact that search engine users self-select into different engines or use multiple engines for different types of searches when considering the competitive implications of search rankings.³

Rather than focus upon competition among search engines in how results are identified and presented to users, critics and complainants craft their arguments around alleged search engine "discrimination" or "bias."⁴ The complainants must have in mind something other than competitive decisions to rank content that differ from the decisions made by rivals; bias in this sense is both necessary to and inherent within any

engine has a slightly different magic formula for indexing the incomprehensibly huge universe of Web pages out there.").

³ Sullivan, *supra* note 2 (noting that consumers are likely searching on a given engine because they prefer that engine's products); *Google's Competition is One Click Away*, GOOGLE OPERATING SYSTEM (May 11, 2009), <http://googlesystem.blogspot.com/2009/05/googles-competition-is-one-click-away.html> (discussing how Yahoo's traffic volume doubled in the half hour during which Google's search results marked all returns as malware and pointing to a 2008 Forrester study finding that 55% of United States Internet users regularly conduct searches on more than one engine).

⁴ *Can Search Discrimination by a Monopolist Violate U.S. Antitrust Laws?*, FAIRSEARCH (July 12, 2011), <http://www.fairsearch.org/wp-content/uploads/2011/07/Can-Search-Discrimination-by-a-Monopolist-Violate-U.S.-Antitrust-Laws1.pdf> (referring to search engine "discrimination") [hereinafter FAIRSEARCH]. We will use the term "bias" throughout without loss of generality to refer to algorithmic differences among search engines that result in relatively favorable ranking for an engine's own content.

useful indexing tool. Yet, critics have generally avoided a precise definition of the allegedly troublesome conduct. Indeed, the term “bias” is used colloquially and is frequently invoked in the search engine debate to encompass a wide array of behavior—generally suggesting a latent malignancy within search engine conduct—with some critics citing mere differences in results across engines as evidence of harmful conduct.⁵

The more useful attempts to define “bias,” however, focus upon differences in organic rankings attributable to the search engine ranking its own content (“own-content bias”); that is, a sufficient condition for own-content bias is that a search engine ranks its own content more prominently than its rivals do. To be even more precise about the nature of the alleged “own-content bias,” it should be clear that this form of bias refers exclusively to organic results, i.e., those results the search engine produces algorithmically, as distinguished from the paid advertisements that might appear at the top, bottom, or right-hand side of a search result page.⁶ Critics at the Senate’s recent

⁵ See, e.g., Adam Raff, *Search, But You May Not Find*, N.Y. TIMES (Dec. 27, 2009), <http://www.nytimes.com/2009/12/28/opinion/28raff.html> (describing bias as any deviation whatsoever from comprehensive, impartial, and relevant results).

⁶ OneBox results are also not necessarily part of organic search, but involve rich text, including data for which Google has paid. “OneBox results are when Google shows information within a special unit, often with images associated with them. OneBox unit often appears to highlight news, shopping, image and other results that are blended into regular listings using Universal Search.” Danny Sullivan, *Meet the Google OneBox, Plus Box, Direct Answers & the 10-Pack*, SEARCH ENGINE LAND (Sept. 28, 2009, 6:12 PM), <http://searchengineland.com/meet-the-google-onebox-plus-box-direct-answers-the-10-pack-26706>. Google refers to OneBox results as “Search Features” within its “Integrated Results.” See *Results Full Page Overview*, GOOGLE, <http://www.google.com/support/websearch/bin/static.py?hl=en&page=guide.cs&guide=1186810&answer=35891&rd=1> (last visited Sept. 1, 2011) (discussing Google’s integrated results); *Search Features*, GOOGLE,

hearing on the “Power of Google” were particularly vociferous on this front, accusing Google of having “cooked”⁷ its algorithm and of “rig[ging] its results, biasing in favor of Google.”⁸

Competition economists and regulatory agencies are familiar with business arrangements which give rise to concerns of own-content bias.⁹ Complaints and economic theories of harm assert that a vertically integrated firm (in this case, Google offers search results as well as products like YouTube and Google Maps) might discriminate against its rivals by “foreclosing” them from access to a critical input. Here, the critical input necessary for rivals’ success is alleged to be prominent placement in Google’s search results. The economics of the potential anticompetitive exclusion of rivals involving vertically integrated firms are well understood in antitrust. The conditions that must be satisfied for these concerns to generate real risk to consumers are also well known. Over a century of antitrust jurisprudence, economic study, and enforcement agency practice have produced a well-understood economic analysis of the competitive effects of a vertically integrated firm’s “discrimination” in

<http://www.google.com/help/features.html> (last visited Sept. 1, 2011). As discussed below, we generally include OneBox results in our analysis unless otherwise specified in order to remain consistent with Edelman & Lockwood, including in cases where it is clear that a rich text result is not an organic result.

⁷ *Google Denies Abusing power of its Search*, SKY NEWS HD (Sept. 22, 2011, 4:29 PM), <http://news.sky.com/home/technology/article/16075171> (quoting Senator Mike Lee).

⁸ Testimony of Jeff Katz, Chief Exec. Officer, Nextag, Inc., *The Power of Google: Serving Consumers or Threatening Competition?*, Before the Senate Comm. on the Judiciary Subcomm. on Antitrust, Competition Policy, and Consumer Rights (September 21, 2011).

⁹ Michael H. Riordan & Steven C. Salop, *Evaluating Vertical Mergers: A Post-Chicago Approach*, 63 ANTITRUST L.J. 513 (1995).

favor of its own products or services, including widespread recognition that such arrangements generally produce significant benefits for consumers. Modern competition policy recognizes that vertical integration and contractual arrangements are generally procompetitive; it also understands that discrimination of this sort may create the potential for competitive harm under some conditions. Sensible competition policy involving vertical integration and contractual arrangements requires one to be sensitive to the potential consumer welfare-enhancing potential of such vertical integration while also taking seriously the possibility that a firm might successfully harm competition itself (and not merely a rival).

In addition to the failure to distinguish procompetitive conduct from anticompetitive behavior, critics' allegations of own-content bias suffer deeper conceptual ambiguities. The perceived issue for Google's rivals is not merely that Google links to a map when responding to search queries, suggesting one might be relevant for the user; indeed, rival search engines frequently respond to similar user queries with their own or other map products. Rather, critics find problematic that Google responds to user queries with a *Google Map*. This is a critical distinction because it concedes that rivals' complaints are not satisfied by the response that consumers are better off with the map; nor do critics pause to consider that perhaps the Google search

user prefers the Google Map to rival products.¹⁰ Thus, critics brazenly take issue with the relationship between Google and the search result even where they concede Google produces more relevant results for consumers.¹¹ Rather than focusing upon consumers, critics argue that the fact that Google is affiliated with the referred search result is itself *prima facie* evidence of competitively harmful bias.¹² On its face, this argument turns conventional antitrust wisdom on its head. Conduct that harms rivals merely because it attracts consumers from rivals is the essence of competition and the logical core of the maxim that antitrust protects “competition, not competitors.”¹³

Critics’ failure to account for the potential consumer benefits from “own-content bias” extends beyond ignoring the fact that users might prefer Google’s products to rivals’. Most critics simply ignore the myriad of procompetitive explanations for vertical integration in the economics literature. This omission by critics, and especially by economist critics, is mystifying given that economists have documented not only a plethora of procompetitive justifications for such integration, but also that such vertical

¹⁰ Sullivan, *supra* note 2 (“If someone’s searching for “maps” on Google, they may be more likely to want Google Maps than Yahoo Maps – and vice versa.”).

¹¹ Joshua D. Wright, *Sacrificing Consumer Welfare in the Search Bias Debate, Part II*, TRUTH ON THE MARKET (June 28, 2011), <http://truthonthemarket.com/2011/06/28/sacrificing-consumer-welfare-in-the-search-bias-debate-part-ii/> (quoting Benjamin Edelman: “If your house is on fire and you forgot the number for the fire department, I’d encourage you to use Google. When it counts, if Google is one percent better for one percent of searches and both options are free, you’d be crazy not to use it. But if everyone makes that decision, we head towards a monopoly and all the problems experience reveals when a company controls too much.”).

¹² See, e.g., Martin Cowen, *Expedia Boss Warns Google/ITA over Bias*, TRAVOLUTION (July 30, 2010), <http://www.travolution.co.uk/articles/2010/07/30/3795/expedia-boss-warns-googleita-over-bias.html>.

¹³ *Brown Shoe Co. v. United States*, 370 U.S. 294, 320 (1962).

relationships are much more likely to be competitively beneficial or benign than to raise serious threats of foreclosure.¹⁴

The critical antitrust question is *always* whether the underlying conduct creates or maintains monopoly power and thus reduces competition and consumer welfare, or is more likely efficient and procompetitive. To be clear, documenting the mere existence of own-content bias itself does little to answer this question. Bias is *not* a sufficient condition for competitive harm as a matter of economics because it can increase, decrease, or have no impact at all upon consumer welfare; neither is bias, without more, sufficient to state a cognizable antitrust claim.¹⁵ Nonetheless, documenting whether and how much of the alleged bias exists in Google's and its rivals' search results can improve our understanding of its competitive implications—that is, whether the evidence of discrimination in favor of one's own content across search engines is more consistent with anticompetitive foreclosure or with competitive differentiation. Critically, in order to generate plausible competitive concerns, search bias must be sufficient in magnitude to foreclose rivals from achieving minimum efficient scale. It follows from this necessary condition that not all evidence of "bias" is relevant to this competitive concern; in particular, Google referring to its own products

¹⁴ Francine LaFontaine & Margaret Slade, *Vertical Integration and Firm Boundaries: The Evidence*, 45 J. ECON. LIT. 629 (2007).

¹⁵ Geoffrey A. Manne & Joshua D. Wright, *If Search Neutrality is the Answer, What's the Question?*, (Int'l Ctr. for Law & Econ. Antitrust & Consumer Prot. Program, White Paper Series, 2011).

and services more prominently than its rivals rank those same services has little to do with critics' complaints unless they implicate general or vertical search.

Despite widespread discussion of search engine bias, virtually no evidence exists indicating that bias abounds—and very little that it exists at all. Edelman & Lockwood recently addressed this dearth of evidence by conducting a small study focused upon own-content bias in 32 search queries; they contend that their results are indicative of systemic and significant bias demanding antitrust intervention.¹⁶ The authors define and measure "bias" as the extent to which a search engine's ranking of its own content differs from how its rivals rank the same content. This approach provides some useful information concerning differences among search engine rankings. However, the study should not be relied upon to support broad sweeping antitrust policy concerns with Google.

The small sample of search queries provides one reason for caution. Perhaps more importantly, the non-random sample of search queries undermines its utility for addressing the critical antitrust policy questions focusing upon the magnitude of search bias, both generally and as it relates to whether the degree and nature of observed bias satisfies the well-known conditions required for competitive foreclosure. Further, evaluating their evidence at face value, Edelman & Lockwood misinterpret its relevance (Edelman & Lockwood in fact find almost no evidence of bias) and, most

¹⁶ Benjamin Edelman & Benjamin Lockwood, *Measuring Bias in "Organic" Web Search* (Jan. 19, 2011), <http://www.benedelman.org/searchbias/>.

problematically, simply assume that own-content bias is inherently suspect from a consumer welfare perspective rather than considering the well-known consumer benefits of vertical integration. Despite these shortcomings, Edelman & Lockwood's study has received considerable attention, both in the press and from Google's critics, who cite it as evidence of harmful and anticompetitive search engine behavior.¹⁷

In the present analysis, as a starting point, we first "replicate" and analyze Edelman & Lockwood's earlier study of a small, non-random sample of search queries in the modern search market. We then extend this methodology to a larger random sample of search queries in order to draw more reliable inferences concerning the answers to crucial questions for the competition policy debate surrounding search engine bias, including: (1) what precisely is search engine bias?; (2) what are its competitive implications?; (3) how common is it?; (4) what explains its existence and relative frequency across search engines?; and, most importantly, (5) does observed search engine bias pose a competitive threat or is it a feature of competition between search engines?

Part I of this paper articulates an antitrust-appropriate framework for analyzing claims of "own-content bias" and delineates its utility and shortcomings as a theory of antitrust harm; it further evaluates Edelman & Lockwood's study, methodology and

¹⁷ FAIRSEARCH, *supra* note 4; MARTIN CAVE & HOWARD WILLIAMS, *The Perils of Dominance: Exploring the Economics of Search in the Information Society*, INITIATIVE FOR A COMPETITIVE ONLINE MARKETPLACE (March 2011); James Temple, *Ben Edelman Says Google Favors Its Own Results*, SFGATE.COM (March 21, 2011), http://articles.sfgate.com/2011-03-21/business/29149701_1_google-competitors-google-s-gmail-search-engines-link.

analysis using this framework. Part II lays out the methodology employed in our own studies. Part III presents the results of our replication of Edelman & Lockwood and analyzes antitrust implications for the search engine bias debate; Part IV does the same for our larger, random sample of search queries. Part V concludes.

I. Defining and Measuring Search Engine “Bias”

A. Defining Search “Bias”

Google critics and search neutrality proponents employ the term “bias” to describe the general conceptual idea of differentiation of organic search results based upon criteria other than “the merits.” For example, some define the relevant bias as any conduct that “involve[s] the manipulation or shaping of search engine results.”¹⁸ Adam Raff of Foundem goes so far as to claim that any deviation from results that are comprehensive, impartial and relevant constitutes bias.¹⁹ The antitrust policy focus upon search results, however, has a narrower scope: a search engine’s treatment of its own content. Google’s general and vertical search competitors often claim that Google purposefully refers to its own content more prominently than that of its rivals.²⁰

¹⁸ Oren Bracha & Frank Pasquale, *Federal Search Commission? Fairness, Access, and Accountability in the Law of Search*, 93 CORNELL L. REV. 1149, 1167 (2008).

¹⁹ Raff, *supra* note 5.

²⁰ Edelman & Lockwood, *supra* note 16; Thomas Catan & Amir Efrati, *Feds to Launch Probe of Google*, WALL STREET J. (June 24, 2011), <http://online.wsj.com/article/SB10001424052702303339904576403603764717680.html> (noting that Expedia, Kayak.com, TripAdvisor, WebMD.com, Yelp.com, Citysearch.com, and Sabre Holdings have all criticized Google for precisely these reasons); *see also* AMIR Efrati, *Rivals Say Google Plays Favorites*, WALL STREET J. (Dec. 12, 2010), <http://online.wsj.com/article/SB10001424052748704058704576015630188568972.html>; *Foundem’s Google Story*, SEARCHNEUTRALITY.ORG (Aug. 18, 2009), <http://www.searchneutrality.org/foundem-google-story>; *Making the Case for Search Neutrality*,

As discussed, the implicit antitrust claim is that Google's discrimination against its rivals results in foreclosure from access to web users, and ultimately in harm to competition. Each of these complaints, however, relies upon a definition of bias that misconstrues the very nature of search engine results and thus the role of competition among search engines in providing those results to consumers. The complaints presume the existence of some intrinsically correct and true list and sequence of results exists and is readily identifiable; however, "bias" is not only inherent in, but also necessary to, any workable indexing system of any size.²¹ Search engines create immense value by serving a set of customers with remarkably heterogeneous preferences.²² Indeed, search engines face downward sloping demand for their services

SEARCHNEUTRALITY.ORG (Oct. 11, 2009), <http://www.searchneutrality.org/search-neutrality>; Steve Lohr, *Antitrust Cry from Microsoft*, N.Y. TIMES, March 31, 2011, at B1, available at <http://www.nytimes.com/2011/03/31/technology/companies/31google.html?pagewanted=all>; Greg Sterling, *EU Antitrust Complaints against Google Grow to Nine*, SEARCH ENGINE LAND (Aug. 2, 2011, 7:44 PM), <http://searchengineland.com/eu-antitrust-complaints-against-google-grow-to-nine-87915>. See also, Wright, *supra* note 11 (quoting Benjamin Edelman: "I don't think it's out of the question given the complexity of what Google has built and its persistence in entering adjacent, ancillary markets. A much simpler approach, if you like things that are simple, would be to disallow Google from entering these adjacent markets. OK, you want to be dominant in search? Stay out of the vertical business, stay out of content.").

²¹ Eric Goldman, *Search Engine Bias and the Demise of Search Engine Utopianism*, 8 YALE J.L. & TECH. 188 (2006); Chris Sherman, *Are Search Engines Biased?*, SEARCH ENGINE WATCH (March 10, 2002), <http://searchenginewatch.com/article/2067657/Are-Search-Engines-Biased> ("[N]o search technology, or for that matter, paper finding tool exists without bias. . . . Given that no finding aid exists without bias, does less of it make a better search engine? . . . [N]ot necessarily.") (quoting Genie Tyburski).

²² JACQUES BUGHINET AL., *The Impact of Internet Search Technologies: Search*, MCKINSEY & CO. (July 2011) (finding that search technology adds approximately \$780 billion annually worldwide, and that \$540 billion of this contributes directly to GDP). See also Chen et al., *supra* note 1; Varian, *supra* note 1. Accordingly, the quest to define search bias and to enforce the elusive and mythical search "neutrality" has thus far proven to be more of a distraction than a useful construct. See, e.g., Manne & Wright, *supra* note 15; Eric Goldman, *Revisiting Search Engine Bias* 9-13 (Santa Clara Univ. Sch. of Law Legal Studies Research Papers Series, Accepted Paper No. 12-11, June 2011), available at <http://ssrn.com/abstract=1860402> ("[T]he term 'search neutrality' implies the existence of 'neutral search

because they are differentiated from one another upon many dimensions. Not only is this differentiation innocuous as a competitive matter, but competition among search engines to satisfy diverse consumer preferences drives this outcome and encourages innovation. Accordingly, a naked identification of bias is simply meaningless for antitrust purposes because it says nothing about its impact upon consumers. Further analysis, at minimum including a determination of its magnitude and whether it in fact implicates anticompetitive foreclosure, is required.

B. Edelman & Lockwood's Study of Search Engine Bias²³

Edelman & Lockwood "investigate . . . [w]hether search engines' algorithmic results favor their own services, and if so, which search engines do most, to what extent, and in what substantive areas."²⁴ Their approach is to measure the difference in how frequently search engines refer to their own content relative to how often its rivals do so. While this approach provides useful descriptive facts about differences among search engines with respect to links to their own content, as discussed, it does little to inform antitrust analysis because the authors begin with the rather odd assertion that competition among differentiated search engines for consumers is a puzzle that creates an air of suspicion around the practice: "it is hard to see why results would vary . . .

engines," but those are entirely mythical."); Abbe Mowshowitz & Akira Kawaguchi, *Measuring Search Engine Bias*, 41 INFO. PROCESSING & MGMT. 1193, 1194 (2005) ("Bias is a relative concept. A search engine is being weighed against its peers, not against an absolute norm derived from features of the universe.").

²³ Edelman & Lockwood, *supra* note 16.

²⁴ *Id.*

across search engines.”²⁵ This assertion completely discounts both the vigorous competitive product differentiation that occurs in nearly all modern product markets as well as the obvious selection effects at work in own content bias (Google users likely prefer Google content). This combination detaches Edelman & Lockwood’s analysis from the consumer welfare perspective, and thus antitrust policy relevance, despite their vigorous claims to the contrary (and the fact that their results actually exhibit very little bias).²⁶

Several other methodological issues undermine the policy relevance of Edelman & Lockwood’s analysis. Edelman & Lockwood hand select 32 search queries and execute searches on Google, Bing, Yahoo, AOL and Ask. Edelman & Lockwood’s hand-selected non-random sample of 32 search queries cannot generate reliable inferences concerning the frequency of bias, a critical ingredient to understanding its potential competitive effects. Indeed, Edelman & Lockwood concede their queries are chosen *precisely because* they are likely to return results including Google content (e.g., email, images, maps, video).²⁷ The 32 search queries are:

²⁵ *Id.* Others have remarked upon the absurdity of this assertion. Danny Sullivan, for example, states “It’s not hard to see why search engine result differ at all. Search engines each use their own “algorithm” to cull through the pages they’ve collected from across the web, to decide which pages to rank first Google has a different algorithm than Bing. In short, Google will have a different opinion than Bing. Opinions in the search world, as with the real world, don’t always agree.” Sullivan, *supra* note 2.

²⁶ Edelman & Lockwood, *supra* note 16, Table 3, Appendix 3.

²⁷ Edelman & Lockwood, *supra* note 16 (“[W]e formed a list of 32 search terms for services commonly provided by search engines. . . .”).

academic article	directions	markets	scholarly journals	voicemail
blog	email	movies	shop	web hosting
books	finance	news	spreadsheet	web publishing
browser	health	photos	stocks	word processor
calendar	images	pictures	translate	
chat	mail	rss reader	translation	
compare prices	maps	satellite images	video	

Edelman & Lockwood analyze the top three organic search results for each query on each engine. They find that 19% of all results across all five search engines refer to content affiliated with one of them.²⁸ Edelman & Lockwood focus upon the first three organic results and report that Google refers to its own content in the first ("top") position about twice as often as Yahoo and Bing refer to Google content in this position. Additionally, they note that Yahoo is more biased than Google when evaluating the first page rather than only the first organic search result.²⁹

Edelman & Lockwood also offer a strained attempt to deal with the possibility of what we've referred to as competitive product differentiation among search engines. They discuss the possibility of "random variation across search engines."³⁰ However,

²⁸ *Id.* ("We preserved and analyzed the first page of results from each search . . . a significant fraction [of results] – 19% – came from pages that were obviously affiliated with one of the five search engines.").

²⁹ On its first page, Yahoo refers to Yahoo content in 37 results, while Bing and Google refer to Yahoo content in just 19 and 15 results, respectively. Meanwhile, Google both refers to its own content in fewer instances and exhibits far less bias in its first page of results: Google refers to its own content in just 32 results; Yahoo refers to Google content in 28 results; and Bing refers to Google content in 26 results.

³⁰ *Id.* This choice of terminology is misleading and obfuscates important and policy relevant economic forces. Search engines do not randomly rank results. They are the product of competition, including systematic and continually scrutinized algorithmic decisions – which are (1) unique to each engine and simply cannot be expected to yield identical results (nor would such an outcome be desirable) and (2)

their evidence undermines claims that Google's own-content bias is significant and systematic relative to its rivals'. In fact, almost zero evidence of statistically significant own-content bias by Google emerges. Edelman & Lockwood examine differences among search engines' references to their own content by "compar[ing] the frequency with which a search engine links to its own pages, relative to the frequency with which other search engines link to that search engine's pages."³¹

Edelman & Lockwood find, in general, Google is no more likely to refer to its own content than other search engines are to refer to that same content. While they do find that both Google and Yahoo are significantly more likely to refer to their own content in their first position than the other engines,³² this is an anomalous result. Across vast majority of their results, Edelman & Lockwood find Google search results are not statistically more likely to refer to Google content than rivals' search results. For example, Edelman & Lockwood find that Google is not more likely to refer to its own content when focusing upon the entire first page or the Top 3 results.³³ In an analysis of 90 common search terms in Google's Keywords tool for "internet software," they find yet again that Google search results are not statistically significantly more likely to refer

reflect search engines' conscious decisions to focus upon different characteristics of search results within their results.

³¹ Edelman & Lockwood, *supra* note 16.

³² Google's odds ratio is 3.1 and is statistically significant at the 2% level, while Yahoo's odds ratio is higher at 3.3 and more statistically significant (at the 1% level). *Id.* at Table 3. An odds ratio of 1 indicates that Google (Yahoo) refers to its own content at the same rate that other engines refer to Google (Yahoo) content. *Id.*

³³ *Id.*

to its own content than its rivals do, while Yahoo is significantly more likely to refer to its own content than other search engines.³⁴

Edelman & Lockwood's same data can be examined to test the likelihood that a search engine will refer to content affiliated with a rival search engine. Rather than exhibiting bias in favor of an engine's own content, it might conceivably be less likely to refer to content affiliated with its rivals. Table 1 reports the likelihood (in odds ratios) that a search engine's content appears in a rival engine's results.

Table 1

	First Result		First Page		
	Google Result	Yahoo Result	Google Result	Yahoo Result	Microsoft Result
<i>Ask</i>	0.4406 (0.2583)	0.5345 (0.3030)	0.6459 (0.1951)	0.4422 * (0.1377)	0.3012 * (0.1977)
<i>Bing</i>	0.3535 * (0.2167)	0.3535 * (0.2167)	0.7930 (0.2195)	0.4654 * (0.1370)	
<i>Google</i>		0.0616 *** (0.0666)		0.3638 *** (0.1154)	0.7226 (0.3408)
<i>Yahoo</i>	0.1975 ** (0.1405)		0.8906 (0.2423)		0.8401 (0.3837)
Chi-squared(3)	6.3900 *	8.0500 **	2.2600	14.6300 ***	3.4400
N	128	128	1224	1224	1224

* Significant at 10%, ** Significant at 5%, *** Significant at 1%
Standard Errors in Parentheses

The first two columns of Table 1 demonstrate that, both Google and Yahoo content are referred to in the first search result less frequently in rivals' search results

³⁴ Google's odds ratio for its Top 1 result, Top 3 results, and First Page are 1.100, 1.207, and 1.084, respectively; and Yahoo's odds ratios for these iterations is 21.118, 2.984, and 2.327 and each is statistically significant at the 1% level. Edelman & Lockwood find that Google, Bing and Yahoo all refer to their own results more frequently than the other engines do in the full first page of results for these searches when rich results are included. This finding, however, merely highlights the importance of analyzing the actual effects of such rankings upon consumers, as such results are not only apparently the industry standard, but also generally perceived as desirable by users.

than in their own. Although Bing does not have enough data for robust analysis of results in the first position in Edelman & Lockwood's original analysis, the next three columns in Table 1 illustrate that all three engines' (Google, Yahoo, and Bing) content appears less often on the first page of rivals' search results than on their own search engine. However, only Yahoo's results differ significantly from 1. As between Google and Bing, the results are notably similar.

Edelman & Lockwood make a limited attempt to consider the possibility that favorable placement of a search engine's own content is a response to user preferences rather than anticompetitive motives. Using click-through data, they find, unsurprisingly, that the first search result tends to receive the most clicks, at an average rate of 72%; while the second and third results receive on average 13% and 8% of clicks, respectively. Furthermore, they identify one search term for which they believe bias plays an important role in driving user traffic. For the search query "email," Google ranks its own Gmail first and Yahoo Mail second; however, Edelman & Lockwood also find that Gmail receives only 29% of clicks while Yahoo Mail receives 54%. Edelman & Lockwood assume that this finding strongly indicates that Google is engaging in conduct that harms users and undermines their search experience. However, from a competition analysis perspective, that inference is not sound. Indeed, the fact that the second-listed Yahoo Mail link received the majority of clicks demonstrates Yahoo was

not competitively foreclosed from access to users. Taken collectively, Edelman & Lockwood are not able to muster evidence of potential competitive foreclosure.³⁵

Claiming that their results collectively identify *prima facie* evidence of inherently anticompetitive search engine bias, Edelman & Lockwood argue that search engines should be subject to more exacting scrutiny and regulatory involvement.³⁶ FairSearch (a compilation of Google rivals) and others have embraced this concept, arguing that Google should be condemned under antitrust laws for manipulating its results in its favor.³⁷

We agree it is important to have an evidence-based discussion surrounding search engine results and their competitive implications; but as we've observed, it 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.³⁸ Edelman & Lockwood's analysis

³⁵ Moreover, any number of other benign reasons could explain this anomalous ranking; for example, users might realize after running this search that they know of a more efficient way of accessing Gmail, or they may simply have clicked on Yahoo Mail first, immediately returned to the search page, and subsequently clicked on Gmail. Sullivan, *supra* note 2. Note additionally that popularity is not always equivalent to relevance. *Id.*

³⁶ Edelman & Lockwood, *supra* note 16 ("[B]y comparing results across multiple search engine[s], we provide prima facie evidence of bias . . . as Google becomes even more dominant, we envision substantially greater investigation of the effect of Google's linking policies, ultimately including deeper outside verification and oversight.").

³⁷ FAIRSEARCH, *supra* note 4.

³⁸ See Danny Sullivan, *The Incredible Stupidity of Investigating Google for Acting Like a Search Engine*, SEARCH ENGINE LAND (Nov. 30, 2010, 7:52 AM), <http://searchengineland.com/the-incredible-stupidity-of-investigating-google-for-acting-like-a-search-engine-57268> ("Google is a search engine. A search engine's job is to point you to destination sites that have the information you are seeking, not to send you to other search engines. Getting upset that Google doesn't point to other search engines is like getting upset that

provides a useful starting point for describing how search engines differ in their referrals to their own content. However, they 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. Further, taken at face value, Edelman & Lockwood’s results actually demonstrate little or no evidence of bias.

II. Replicating and Extending Edelman & Lockwood’s Analysis

Initially, we execute searches for Edelman & Lockwood’s original 32 non-random queries using three different search engines (Google, Bing, and Blekko) to reflect developments in the modern search engine market and in an attempt to produce results relevant to current policy debates.³⁹ We record each organic search result on the first page (up to twelve) as well as whether the result refers to Microsoft- or Google-affiliated content.⁴⁰ To replicate Edelman & Lockwood’s inclusion of Oneboxes and other rich results, we include them in our analysis unless otherwise specified. We record screen shots of all the search results.⁴¹ This initial coding reveals that a total of 97 URLs across the three search engines refer to Google content: Google, Bing and Blekko refer to Google content in 51, 26 and 20 results, respectively. A total of 74 URLs

the New York Times doesn’t simply have headlines followed by a single paragraph of text that says ‘read about this story in the Wall Street Journal.’”).

³⁹ We conducted all queries between June 23, 2011 and July 5, 2011.

⁴⁰ Because Google, Bing, and Blekko do not always report URLs in the same manner, we gave each Google- and Microsoft-related URL a common name to facilitate comparisons. For instance, we coded “maps.google.com/” as “Google Maps” and “office.microsoft.com/en-us/excel” as “Microsoft Office.”

⁴¹ Data available from the author upon request.

reference Microsoft content: Bing, Google, and Blekko refer to Microsoft content in 56, 14 and 4 results, respectively.

Edelman & Lockwood's search queries were recorded in August 2010. Search technology has changed dramatically since then.⁴² Further, Bing now powers Yahoo, and Blekko has had more time to mature and enhance its results. Blekko serves as a helpful "control" engine in this study as it is totally independent of Google and Microsoft, and thus has no incentive to refer to Google or Microsoft content unless it is actually relevant to users. Blekko also provides an interesting comparison because its general approach to search differs significantly from Google and Bing, which have more in common.⁴³ Blekko's goal is to rid its results of spam entirely, and it employs slash tags and user intervention to achieve its objectives.⁴⁴ Thus, if Blekko, Google, and Microsoft results for a particular query each agree that specific content is highly

⁴² For example, Bing has since begun returning results that take account of the user's location and search history; Google introduced Panda – a significant algorithm update affecting 12% of its United States search results; and Ask.com vacated the web crawling market to focus solely upon providing a comprehensive question-and-answer service. Danny Sullivan, *Bing Gets Localized and Personalized*, SEARCH ENGINE LAND (Feb. 10, 2011, 12:00 PM), <http://searchengineland.com/bing-results-get-localized-personalized-64284>; Danny Sullivan, *Google Forecloses on Content Farms with "Panda" Algorithm Update*, SEARCH ENGINE LAND (Feb. 24, 2011, 9:50 PM), <http://searchengineland.com/google-forecloses-on-content-farms-with-farmer-algorithm-update-66071>; Danny Sullivan, *Ask.com to Focus on Q&A Search, End Web Crawling*, SEARCH ENGINE LAND (Nov. 9, 2010, 1:50 PM), <http://searchengineland.com/ask-com-to-focus-on-qa-search-end-web-crawling-55209>.

⁴³ See Danny Sullivan, *Google: Bing Is Cheating, Copying Our Search Results*, SEARCH ENGINE LAND (Feb. 1, 2011, 8:45 AM), <http://searchengineland.com/google-bing-is-cheating-copying-our-results-62914> (discussing how Google intentionally returned irrelevant results for obscure, "long tail" queries when it suspected Bing of copying its results, and noting how, within a few weeks, Bing's results in fact synced with Google's for these queries).

⁴⁴ Aaron Wall, *Rich Skrenta Talks about Blekko Search*, SEARCH NEWZ (Nov. 1, 2010), <http://www.searchnewz.com/topstory/news/sn-2-20101101RichSkrentaTalksAboutblekkoSearch.html>.

relevant to the user query, it lends significant credibility to the notion that the content places well on the merits rather than being attributable to bias or other factors.

It also bears repeating that our purpose in replicating and updating Edelman & Lockwood is to develop a better and more contemporary description of how search engines differ in their treatment of their own content. Yet it should be clear that we do not believe our analysis of the non-random Edelman & Lockwood sample provides dispositive proof of the competitive nature and magnitude of any such bias. Accordingly, to further develop an understanding of own-content bias, we employ the methodology just described to a random sample of 1,000 Google search queries.⁴⁵ Note that even our analysis of a random sample of search queries is one step removed from a direct evaluation of the critical link between competitive differentiation in organic search and impact upon consumers. However, both analyses presented here are useful first steps in documenting and evaluating differentiation in organic search results among search engines and allow those data to be analyzed through the lens of antitrust economic lens to assess the competitive implications of search bias.

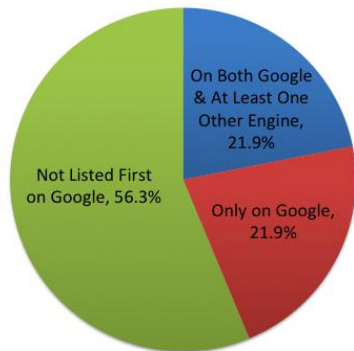
III. Replicating the Edelman & Lockwood Study

A. How Do Search Engines Rank Their Own Content?

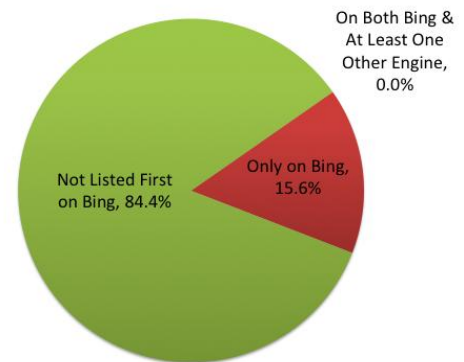
⁴⁵ In August 2006, AOL released a list of 20 million search queries that hundreds of thousands of its users actually ran between March and May of that year. Ellen Nakashima, *AOL Takes Down Site with Users' Search Data*, WASHINGTON POST (Aug. 8, 2006), <http://www.washingtonpost.com/wp-dyn/content/article/2006/08/07/AR2006080701150.html>. AOL published this list on a special website, to allow researchers to study how people search for information online. We randomly selected 1,000 of these queries for our sample. Searches were executed on these queries between July 20, 2011 and August 20, 2011.

Focusing solely upon the first position, Google refers to its own products or services when no other search engine does in 21.9% of queries; in another 21.9% of queries, both Google and at least one other search engine rival (i.e. Bing or Blekko) refer to the same Google content with their first links. The following two charts illustrate the percentage of Google or Bing first position results, respectively, dedicated to own content across search engines.

Search Results List Google Content in the First Position
(N = 32)

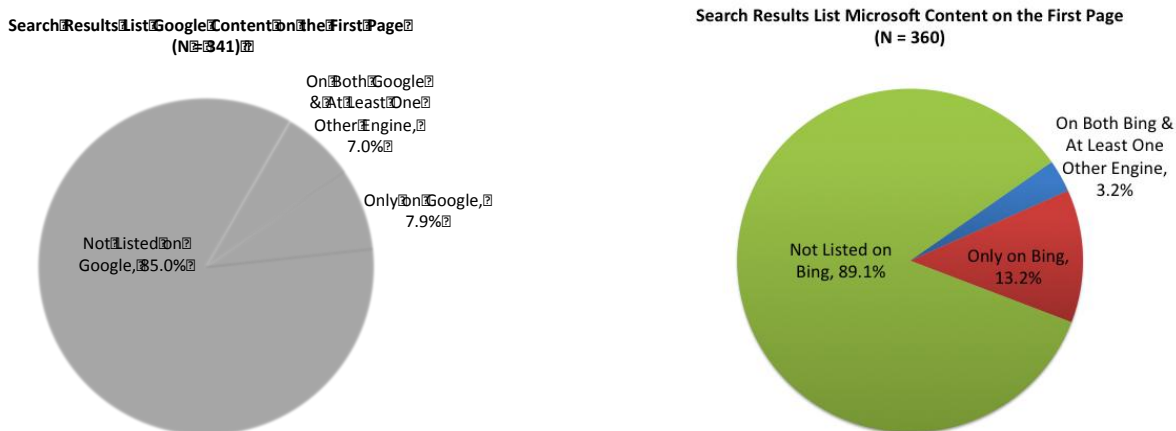


Search Results List Microsoft Content in the First Position
(N = 32)



Restricting focus upon the first position is too narrow. Assuming that all instances in which Google or Bing rank their own content first and rivals do not amounts to bias would be a mistake; such a restrictive definition would include cases in which all three search engines rank the same content prominently — agreeing that it is highly relevant—though not all in the first position. The entire first page of results provides a more informative comparison. We find that Google and at least one other engine return Google content on the first page of results in 7% of the queries. Google

refers to its own content on the first page of results without agreement from either rival search engine in only 7.9% of the queries. Meanwhile, Bing and at least one other engine refer to Microsoft content in 3.2% of the queries. Bing references Microsoft content without agreement from either Google or Blekko in 13.2% of the queries.



This evidence indicates that Google's ranking of its own content differs significantly from its rivals in only 7.9% of queries, and that when Google ranks its own content prominently it is generally perceived as relevant. Further, these results suggest that Bing's organic search results are significantly more biased in favor of Microsoft content than Google's search results are in favor of Google's content.

B. Examining Search Engine "Bias" on Google

Table 2 presents the percentages of queries for which Google's ranking of its own content differs significantly from its rivals' ranking of that same content.

Table 2

Percentage of Google Organic Results with Google Content Not Ranked Similarly by Rival Search Engines

	Google Content Not Mentioned in Corresponding Top 1, 3, 5 or First Page of Results		
	Bing	Bleko	Bing & Bleko
Top 1 N= 14	78.6% 11	57.1% 8	50.0% 7
Top 3 N= 24	37.5% 9	58.3% 14	29.2% 7
Top 5 N= 31	38.7% 12	64.5% 20	35.5% 11
First Page N= 45	51.1% 23	68.9% 31	48.9% 22

Note that percentages below 50 in Table 2 indicate that rival search engines generally perceive the referenced Google content as relevant and independently believe that it should be ranked similarly. When Google ranks its own content highly, at least one rival engine typically agrees with this ranking; for example, when Google places its own content in its Top 3 results, at least one rival agrees with this ranking in over 70% of queries. Bing especially agrees with Google's rankings of Google content within its Top 3 and 5 results, failing to include Google content Google ranks thusly in less than 40% of queries.

Table 3 focuses upon the rare case in which Google ranks its own content within the first page of results and rivals do not refer to the same content at all.

Table 3

Percentage of Google Organic Results with Google Content Not Ranked At All by Rival Search Engines

	Google Content Not Mentioned At All on First Page of Results		
	Bing	Blekkio	Bing & Blekkio
Top 1 N= 14	14.3% 2	42.9% 6	7.1% 1
Top 3 N= 24	16.7% 4	41.7% 10	12.5% 3
Top 5 N= 31	35.5% 11	54.8% 17	32.3% 10
First Page N= 45	51.1% 23	68.9% 31	48.9% 22

Table 3 further reveals a general consensus across search engines as to the relevancy of Google content. Particularly when Google ranks its own content prominently, at least one rival engine agrees that this Google result belongs on the first page; there is only one query for which Google refers to its own content in the first position while no other engine references that same content anywhere on its first page.⁴⁶

We also ran several simple regression models to compare the results from these new data to Edelman & Lockwood's results. Table 4 reports the likelihood Google content will be referred to in a Google search result relative to searches performed on rival engines.

⁴⁶ For the query "blog," Google returns blogger.com in the first position.

Table 4

	First Result		First Page	
	Google Result	Google Result	Google Result	Google Result
<i>Google</i>	4.2000 *** (2.0810)		2.6685 *** (0.5737)	
<i>Bing</i>		0.1330 *** (0.0936)		0.4426 *** (0.1124)
<i>Blekko</i>		0.3600 * (0.2004)		0.3124 *** (0.0860)
Chi-squared(2)		9.1400 **		21.5500 ***
N	96	96	1085	1085

* Significant at 10%, ** Significant at 5%, *** Significant at 1%
Standard Errors in Parentheses

Focusing upon Edelman & Lockwood's small and non-random sample of search queries, the first and third columns report results indicating that Google affiliated content is more likely to appear in a search executed on Google rather than rival engines. Both Bing and Blekko are significantly less likely to refer to Google content in their first result or on their first page. Interestingly, Bing's first result is highly unlikely to include Google content, and much less likely than Blekko.

C. Examining Search Engine "Bias" on Bing

Bing refers to Microsoft content in its search results far more frequently than its rivals reference the same Microsoft content. For example, Bing's top result references Microsoft content for 5 queries, while neither Google nor Blekko ever rank Microsoft content in the first position (see Table 5).

Table 5

Percentage of Bing Organic Results with Microsoft Content Not Ranked Similarly by Rival Search Engines

	Microsoft Content Not Mentioned in Corresponding Top 1, 3, 5 or First Page of Results		
	Google	Blekko	Google & Blekko
Top 1 N= 5	100.0% 5	100.0% 5	100.0% 5
Top 3 N= 8	100.0% 8	100.0% 8	100.0% 8
Top 5 N= 13	69.2% 9	100.0% 13	69.2% 9
First Page N= 54	79.6% 43	92.6% 50	79.6% 43

Table 5 illustrates the significant discrepancies between Bing's treatment of its own Microsoft content relative to Google and Blekko. Neither rival engine refers to Microsoft content Bing ranks within its Top 3 results; Google and Blekko do not include any Microsoft content Bing refers to on the first page of results in nearly 80% of queries.

Moreover, Bing frequently ranks Microsoft content highly even when rival engines do not refer to the same content at all in the first page of results, as Table 6 demonstrates.

Table 6

Percentage of Bing Organic Results with Microsoft Content Not Ranked At All by Rival Search Engines

	Microsoft Content Not Mentioned At All on First Page of Results		
	Google	Blekko	Google & Blekko
Top 1 N= 5	80.0% 4	100.0% 5	80.0% 4
Top 3 N= 8	62.5% 5	87.5% 7	62.5% 5
Top 5 N= 13	69.2% 9	92.3% 12	69.2% 9
First Page N= 54	79.6% 43	92.6% 50	79.6% 43

For example, of the 5 queries for which Bing ranks Microsoft content in its top result, Google refers to only one of these 5 within its first page of results, while Blekko refers to none. Even when comparing results across each engine's full page of results, Google and Blekko only agree with Bing's referral of Microsoft content in 20.4% of queries.

Simple regression analysis corroborates these results. Table 7 shows the likelihood of Microsoft content referred to in a Bing Search or other rivals' search results (reported in odds ratios).

Table 7

	First Page	
	Microsoft Result	Microsoft Result
<i>Bing</i>	7.2354 *** (2.0222)	
<i>Google</i>		0.2324 *** (0.0719)
<i>Blekkio</i>		0.0571 *** (0.0299)
Chi-squared (2)		46.2300 ***
N	1085	1085

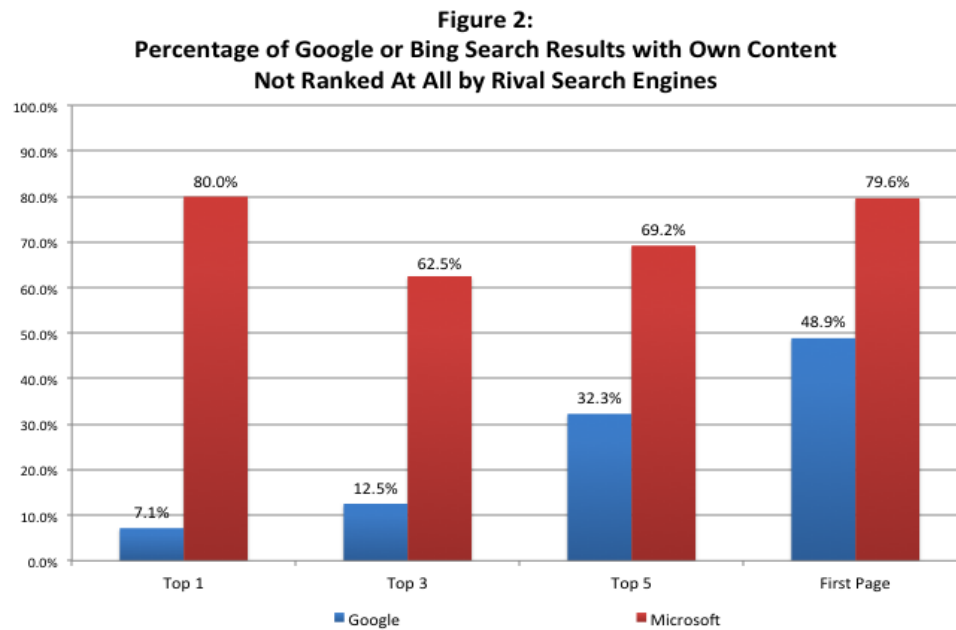
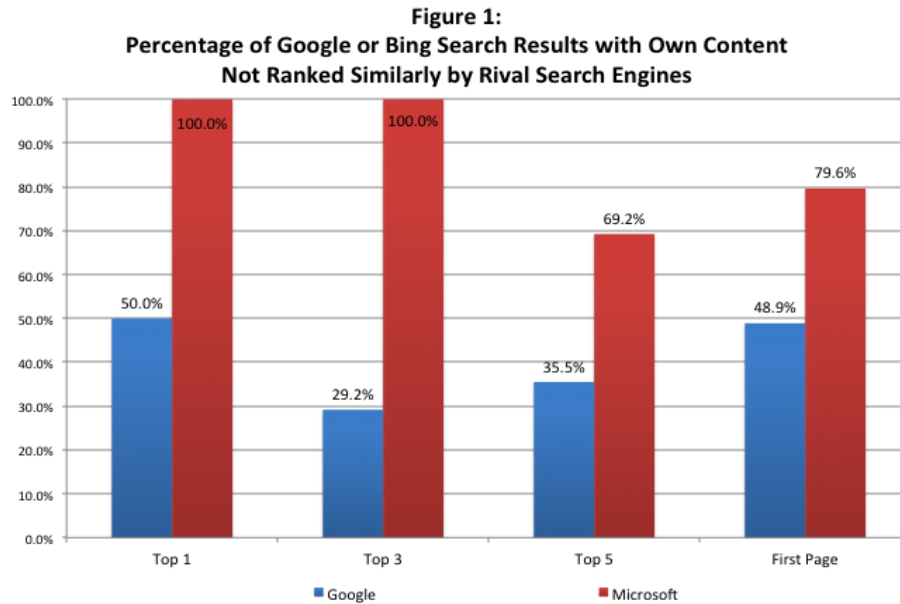
* Significant at 10%, ** Significant at 5%, *** Significant at 1%
Standard Errors in Parentheses

Although there are not enough Bing data to test results in the first position in Edelman & Lockwood's sample, Microsoft content appears as results on the first page of a Bing search about 7 times more often than Microsoft content appears on the first page of rival engines. Also, Google is much more likely to refer to Microsoft content than Blekko, though both refer to significantly less Microsoft content than Bing.

D. A Closer Look at Google v. Bing

By Edelman & Lockwood's measure, Bing results are more biased than Google results. The reason for this result is clear from the data. Bing's own-content referring search results rank that content significantly higher than its rivals do more frequently than Google does. In other words, rivals are more likely to agree with Google's algorithmic assessment that its own content is relevant to user queries. Bing refers to Microsoft content other engines do not rank at all more often than Google refers its own

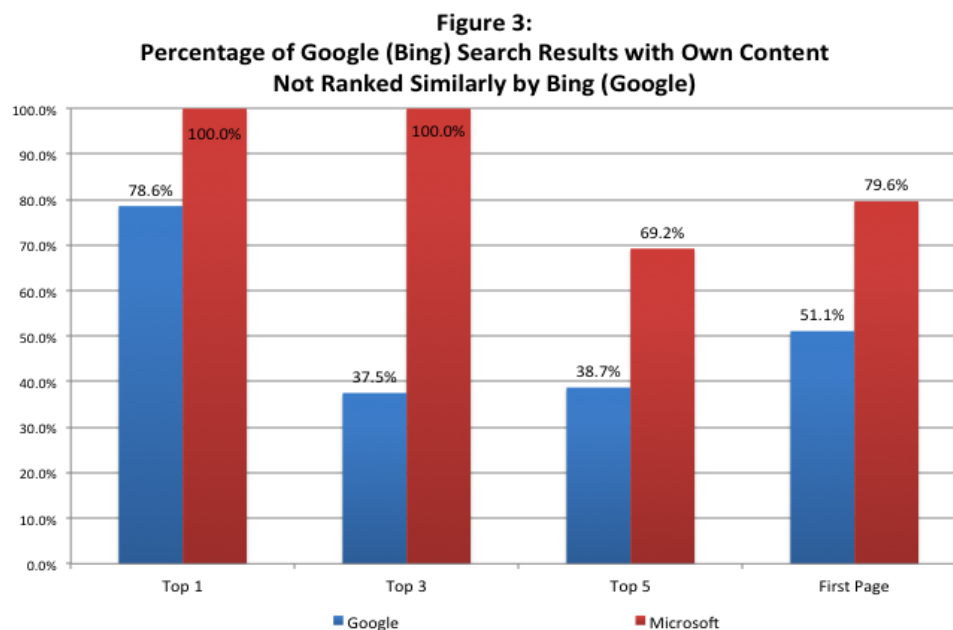
content without any agreement from rivals. Figures 1 and 2 display the same data presented above in order to facilitate direct comparisons between Google and Bing.



As Figures 1 and 2 illustrate, Bing search results for these 32 queries are more frequently "biased" in favor of its own content than are Google's. The bias is greatest

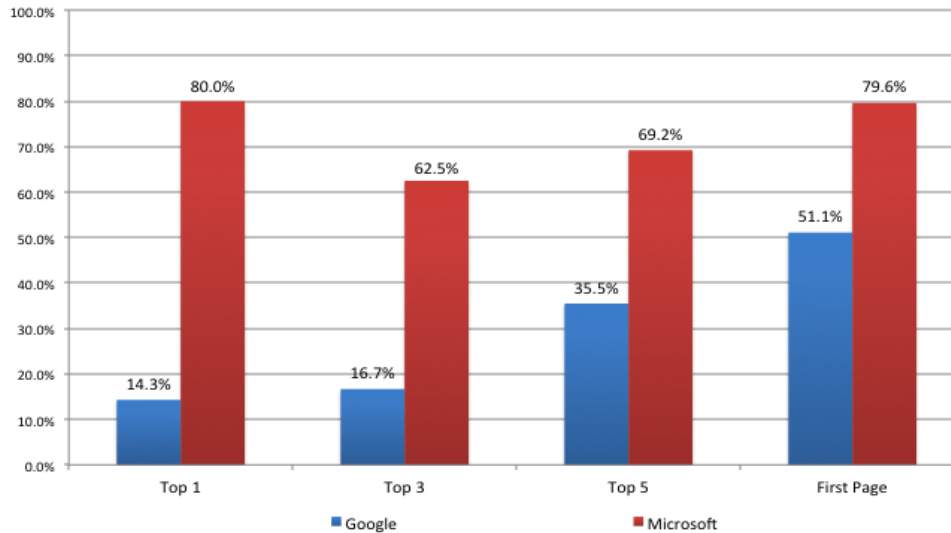
for the Top 1 and Top 3 search results. While it is important to stress Edelman & Lockwood's limited and non-random sample, and to emphasize the danger of making strong inferences about the general nature of magnitude of search bias based upon these data alone, the data indicate that Google's own-content bias is relatively small even in a sample collected precisely to focus upon the queries most likely to generate it.⁴⁷

Figures 3 and 4 present a direct head-to-head comparison of Bing and Google (excluding Blekko from the analysis).



⁴⁷ It is important to recognize that our definition of bias does not imply that either Google or Bing results are "biased" in favor of own content as a result of manual manipulation of organic search results. The definition employed here merely measures the differences among engines in how Google and Microsoft content are displayed.

**Figure 4:
Percentage of Google (Bing) Search Results with Own Content
Not Ranked At All by Bing (Google)**



Notably, Bing ranks Microsoft content more prominently than Google's more frequently than Google ranks its own content more favorably than Microsoft's.⁴⁸

E. Is Google Search Bias Consistent with Anticompetitive Foreclosure?

As we've repeatedly emphasized, while we describe differences among search engines' rankings of their own or affiliated content as "bias," without more these differences do not imply anticompetitive conduct. It is wholly unsurprising and indeed consistent with vigorous competition among engines that differentiation emerges with respect to algorithms. However, it is especially important to note that the theories of anticompetitive foreclosure raised by Google's rivals involve very specific claims about

⁴⁸ Our study finds that Bing exhibits far more "bias" than Edelman & Lockwood identify in their earlier analysis. For example, in Edelman & Lockwood's study, Bing does not refer to Microsoft content at all in its Top 1 or Top 3 results; moreover, Bing refers to Microsoft content within its entire first page 11 times, while Google and Yahoo refer to Microsoft content 8 and 9 times, respectively. Most likely, the significant increase in Bing's "bias" differential is largely a function of Bing's introduction of localized and personalized search results and represents serious competitive efforts on Bing's behalf. Sullivan, *Bing Gets Localized and Personalized*, *supra* note 42.

these differences. One necessary condition of these anticompetitive theories of own-content search bias is obvious: the bias must be sufficient in magnitude to exclude rival search engines from achieving efficient scale.⁴⁹ A corollary of this condition is that the bias must actually be directed toward Google's rivals. The condition is difficult to evaluate with such a small non-random sample of queries designed to generate results with Google content; however, the fact that Google displays less own-content bias than its closest rival even in this sample designed to identify maximum bias, and that such bias is nonetheless relatively infrequent, renders unsurprising our finding that the first condition is also not satisfied in a larger universe of queries, as discussed below.

But are the few instances in which Google ranks its own content as more relevant than its rivals consistent with claims of competitive foreclosure? Are these instances tailored toward the exclusion of search and vertical search rivals? Consider the few queries in the sample for which Google returned Google content within the top three results but neither Bing nor Blekko referenced the same content anywhere on their first page of results. For the query "voicemail," Google refers to both Google Voice and Google Talk; both instances appear unrelated to the grievances of general and vertical search rivals. The query "movie" results in OneBox with the next 3 organic results including movie.com, fandango.com, and yahoo.movies.com. The single instance in Edelman & Lockwood's sample for which Google ranks its own content in the Top 3

⁴⁹ See, e.g., FAIRSEARCH, *supra* note 4 (emphasizing the "importance of scale to competition in search").

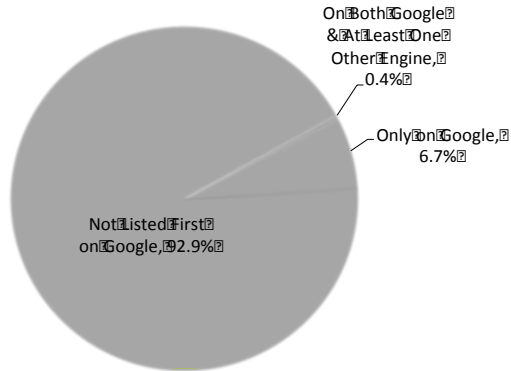
positions but this content is not referred to at all on Bing's first page of results is a link to blogger.com in response to the query "blog." Thus it appears that the small handful of cases in which Google ranks its own content more favorably than its rivals appears to be more consistent with simple and expected competitive product differentiation rather than anticompetitive foreclosure. In Part IV, we extend this analysis to a larger and random sample of actual queries.

IV. Own-Content Bias With a Random Sample of 1,000 Search Queries

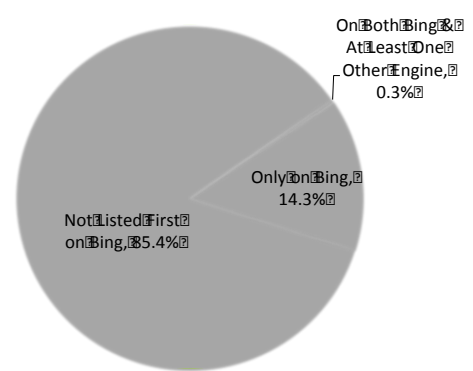
A. How Do Search Engines Rank Their Own Content?

Examining just the first position, we find that Google simply does not refer to its own content in over 90% of queries. Similarly, Bing does not reference Microsoft content in 85.4% of queries. Google refers to its own content in the first position when other search engines do not in only 6.7% of queries; while Bing does so over twice as often, referencing Microsoft content that no other engine references in the first position in 14.3% of queries. The following two charts illustrate the percentage of Google or Bing first position results, respectively, dedicated to own content across search engines.

Search Results List Google Content in the First Position
(N=1,000)

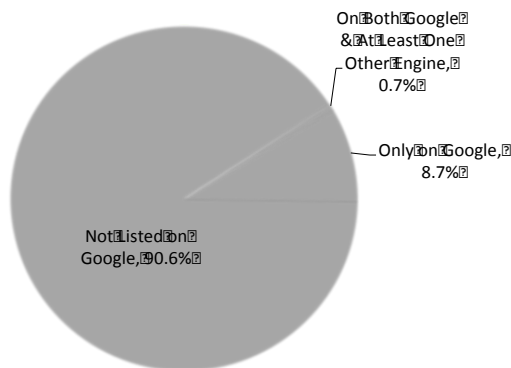


Search Results List Microsoft Content in the First Position
(N=1,000)

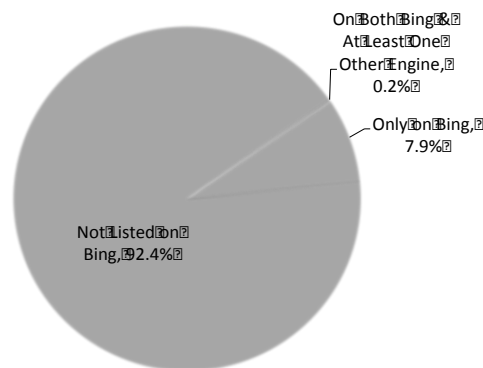


The most striking aspect of these results is the small fraction of queries for which placement of own-content is relevant. Similarly, when we expand consideration to the entire first page of results we find that both Google and Bing do not return their own content in over 90% of queries. Again, the fraction of queries for which placement of own-content is relevant is relatively small. Google returns its own content when no other engine does for only 8.7% of queries, while Bing does so in a comparable 7.9% of queries. Each engine returns own content that at least one other engine also returns on its first page in less than 1% of queries.

Search Results List Google Content on the First Page
(N=10,628)



Search Results List Microsoft Content on the First Page
(N=10,684)



Our data illustrate that levels of own-content bias are similar between Google and Bing when considering the entire first page of results, but that Bing is far more likely than Google to reference its own content in its very first results position.

B. Examining Search Engine “Bias” on Google

We notice two distinct differences between the results of this larger study and our replication of Edelman & Lockwood: (1) Google and Bing refer to their own content in a significantly smaller percentage of cases here than in the non-random sample; and (2) in general, when Google or Bing does rank its own content highly, rival engines are unlikely to similarly rank that same content.

Table 8 reports the percentages of queries for which Google’s ranking of its own content and its rivals’ rankings of that same content differ significantly. When Google refers to its own content within its Top 5 results, at least one other engine similarly ranks this content for only about 5% of queries.

Table 8

Percentage of Google Organic Results with Google Content Not Ranked Similarly by Rival Search Engines

	Google Content Not Mentioned in Corresponding Top 1, 3, 5 or First Page of Results		
	Bing	Bleko	Bing & Bleko
Top 1 N= 71	97.2% 69	94.4% 67	94.4% 67
Top 3 N= 205	95.1% 195	98.0% 201	95.1% 195
Top 5 N= 381	95.3% 363	98.2% 374	95.3% 363
First Page N= 920	95.0% 874	97.5% 897	93.4% 859

Table 9 focuses upon the percentages of queries for which Google references its own content but no other engine references that same content at all. Google refers to its own content within its Top 5 results when no other engine references this content at all for 88.7% of queries.

Table 9

Percentage of Google Organic Results with Google Content Not Ranked At All by Rival Search Engines

	Google Content Not Mentioned At All in First Page of Results		
	Bing	Bleko	Bing & Bleko
Top 1 N= 71	88.7% 63	94.4% 67	88.7% 63
Top 3 N= 205	89.8% 184	97.1% 199	89.3% 183
Top 5 N= 381	89.8% 342	96.9% 369	88.7% 338
First Page N= 920	95.0% 874	97.5% 897	93.4% 859

As presented above for the Edelman & Lockwood smaller sample, Table 10 presents the likelihood that Google content will appear in a Google search, relative to searches conducted on rival engines (reported in odds ratios).

Table 10

Odds Ratios of a Google Result Appearing on Google or Rival Engine's Search Results (All Boxes Included)

	First Result		First Page	
	Google Result	Google Result	Google Result	Google Result
<i>Google</i>	21.7503 *** (8.6598)		16.9897 *** (1.5666)	
<i>Bing</i>		0.3933 *** (0.0233)		0.0718 *** (0.0085)
<i>Blekko</i>		0.0527 *** (0.0272)		0.0471 *** (0.0064)
Chi-squared (2)		59.5600 ***		934.8800 ***
N	2,996	2,996	33,035	33,035

* Significant at 10%, ** Significant at 5%, *** Significant at 1%
Standard Errors in Parentheses

The first and third columns report results indicating that Google affiliated content is more likely to appear in a search executed on Google rather than rival engines. Google is approximately 16 times more likely to refer to its own content on its first page as is any other engine. Bing and Blekko are both significantly less likely to refer to Google content in their first result or on their first page than Google is to refer to Google content within these same parameters. In each iteration, Bing is more likely to refer to Google content than is Blekko, and in the case of the first result, Bing is much more likely to do so.

C. Examining Search Engine “Bias” on Bing

For queries within our larger sample, Bing refers to Microsoft content within its Top 1 and 3 results when no other engine similarly references this content for a slightly smaller percentage of queries than in our Edelman & Lockwood replication. Yet Table 11 illustrates that Bing continues to exhibit a strong tendency to rank Microsoft content more prominently than rival engines. For example, when Bing refers to Microsoft content within its Top 5 results, other engines agree with this ranking for less than 2% of queries; and Bing refers to Microsoft content within its Top 3 results that no other engine does so for 99.2% of queries.

Table 11

Percentage of Bing Organic Results with Microsoft Content Not Ranked Similarly by Rival Search Engines

	Microsoft Content Not Mentioned in Corresponding Top 1, 3, 5 or First Page of Results		
	Google	Blekkio	Google & Blekkio
Top 1 N= 146	98.6% 144	98.6% 144	97.9% 143
Top 3 N= 370	99.5% 368	99.5% 368	99.2% 367
Top 5 N= 558	98.9% 552	99.3% 554	98.4% 549
First Page N= 855	97.9% 837	99.2% 848	97.5% 834

This trend holds when the analysis expands to consider the entire first page of results. As Table 12 presents, when Bing references Microsoft content in its first results position, no other engine references this same content at all on its first page of results

for 97.9% of queries. In fact, when Bing refers to Microsoft content in any of its Top 1, 3, 5 or First Page of results, rival engines do not refer to this content at all for over 97% of queries.

Table 12

Percentage of Bing Organic Results with Microsoft Content Not Ranked At All by Rival Search Engines

	Microsoft Content Not Mentioned At All on First Page of Results		
	Google	Blekko	Google & Blekko
Top 1 N= 146	98.6% 144	98.6% 144	97.9% 143
Top 3 N= 370	99.2% 367	98.9% 366	98.6% 365
Top 5 N= 558	98.7% 551	99.1% 553	98.2% 548
First Page N= 855	97.9% 837	99.2% 848	97.5% 834

Regression analysis further illustrates Bing's propensity to reference Microsoft content that rivals do not. Table 13 reports the likelihood that Microsoft content is referred to in a Bing search as compared to searches on rival engines (again reported in odds ratios).

Table 13

Odds Ratios of Bing/Microsoft Result Appearing on Bing or Rival Engine's Search Results (All One Boxes Included)

	First Result		First Page	
	Microsoft Result	Microsoft Result	Microsoft Result	Microsoft Result
<i>Bing</i>	56.7018 *** (23.7329)		25.7873 *** (3.1011)	
<i>Google</i>		0.0117 *** (0.0084)		0.0354 *** (0.0063)
<i>Blekko</i>		0.0236 *** (0.0120)		0.0418 *** (0.0066)
Chi-squared(2)		91.0700 ***		729.2700 ***
N	2,996	2,996	33,035	33,035

* = Significant at 10%, ** = Significant at 5%, *** = Significant at 1%
Standard Errors in Parentheses

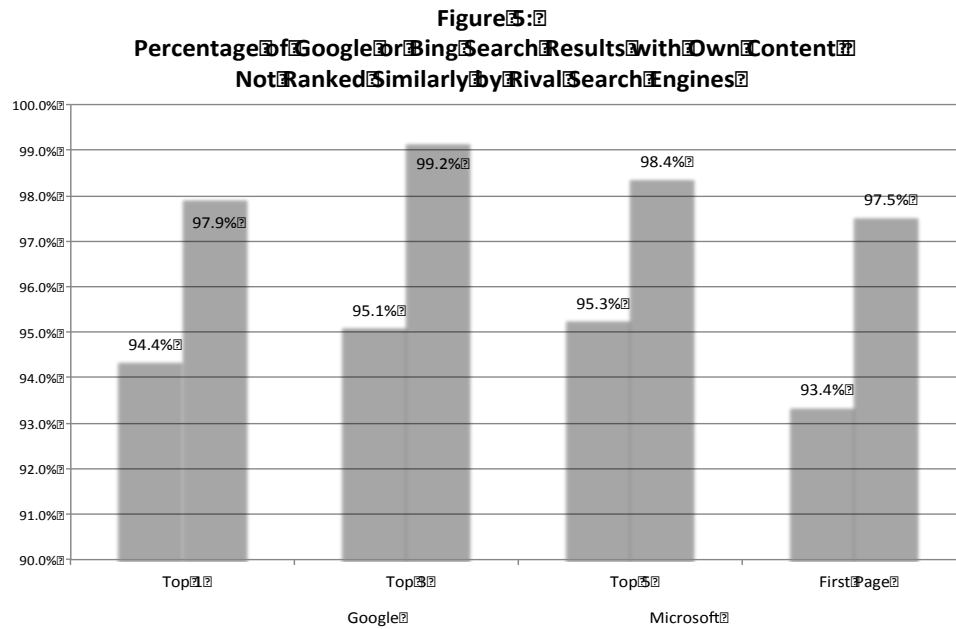
Bing refers to Microsoft content in its first results position about 56 times more often than rival engines refer to Microsoft content in this same position. Across the entire first page, Microsoft content appears on a Bing search about 25 times more often than it does on any other engine. Both Google and Blekko are accordingly significantly less likely to reference Microsoft content. Notice further that, contrary to the findings in the pilot study, Google is slightly less likely to return Microsoft content than is Blekko, both in its first results position and across its entire first page.

D. A Closer Look at Google v. Bing

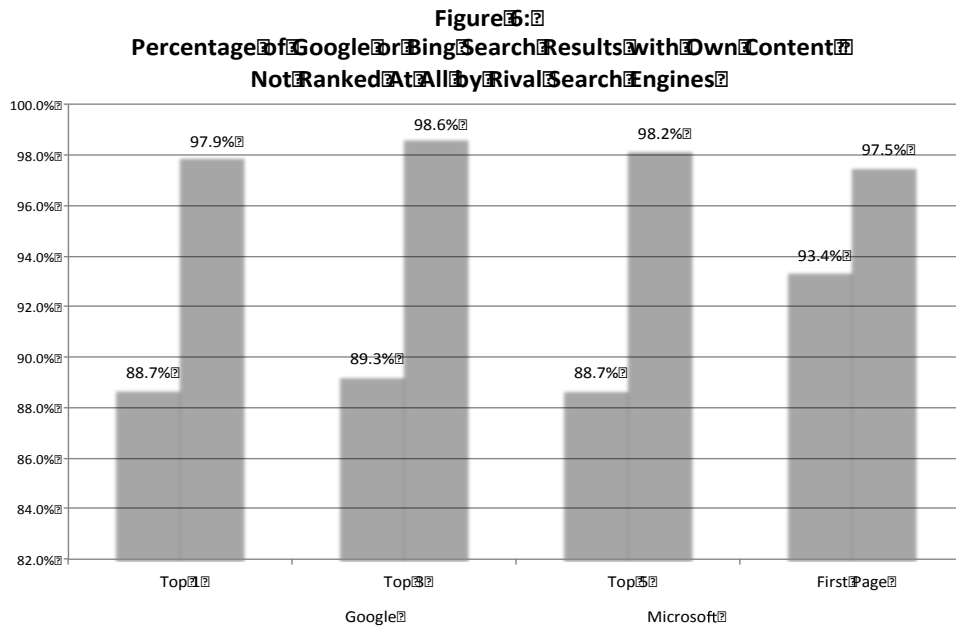
Consistent with our smaller sample, we find again that Bing is more biased than Google using these metrics. In other words, Bing ranks its own content significantly more highly than its rivals do more frequently than Google does, although the discrepancy between the two engines is smaller here than in the pilot study. As noted

above, Bing is over twice as likely to refer to own content in first results position than is Google.⁵⁰

Figures 5 and 6 present the same data reported above to allow for a direct visual comparison of own-content bias between Google and Bing.

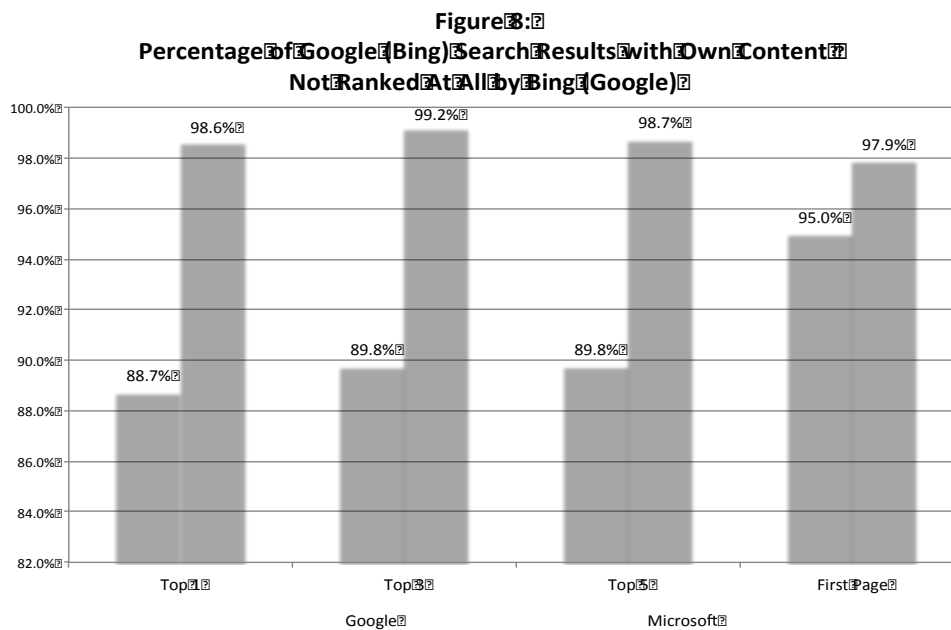
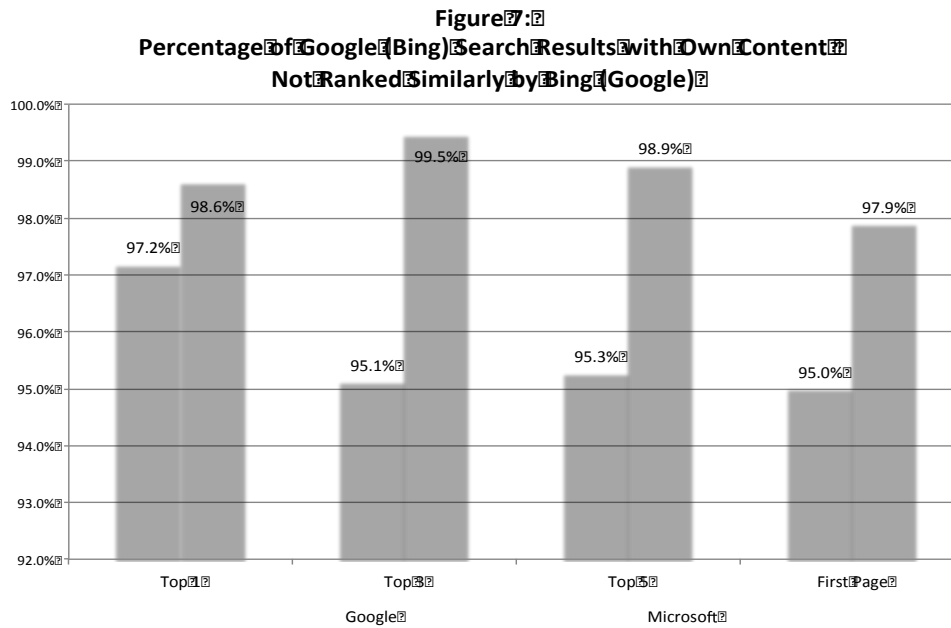


⁵⁰ Google is approximately 21 times more likely to reference own content in position one, while Bing is over 56 times more likely to do so.



The percentages of queries for which both Google and Bing refer to own content that other engines do not similarly reference—or indeed, fail to reference at all—is higher here than in our replication of Edelman & Lockwood’s smaller sample.

As with our analysis of the smaller sample, Figures 7 and 8 again remove Blekko to allow for a direct, head-to-head comparison between Google and Bing.



Consistent with our earlier results, Bing appears to consistently rank Microsoft content higher than Google ranks the same (Microsoft) content more frequently than Google ranks Google content more prominently than Bing ranks the same (Google) content.

E. Is Google Search Bias Consistent with Anticompetitive Foreclosure?

A key antitrust question related to search bias is whether the observed bias, if any, is consistent with the competitive foreclosure arguments raised by Google's critics and rivals. As we have emphasized throughout this paper, vertical foreclosure arguments are premised upon the notion that rivals are excluded with sufficient frequency and intensity as to render their efforts to compete for distribution futile. Our results simply do not support these claims of market conditions conducive to the types of harmful exclusion contemplated by application of the antitrust laws. Rather, the evidence indicates that (1) the absolute level of search engine "bias" is extremely low, and (2) "bias" is not a function of market power, but an effective strategy that has arisen as a result of serious competition and innovation between and by search engines. The first finding would undermine competitive foreclosure arguments on their own terms, that is, even if there were no pro-consumer justifications for the integration of Google content with Google search results. The second finding, even more importantly, reveals that the evolution of consumer preferences for more sophisticated and useful search results has driven rival search engines to satisfy that demand. Both Bing and Google both have shifted toward these results, rendering the complained-of conduct equivalent to satisfying the standard of care in the industry.

A significant lack of search bias emerges in our representative sample of queries. The total percentage of queries for which Google references its own content when rivals

do not is quite low—only about 8%—meaning that Google favors its own content far less often than critics have suggested. This fact is crucial and highly problematic for search engine critics, as their burden in articulating a cognizable antitrust harm includes not only demonstrating that bias exists, but further that it (1) is competitively harmful, and (2) occurs at sufficient levels to actually exclude rivals. As discussed, bias alone is simply not sufficient to demonstrate any *prima facie* anticompetitive harm—as it is far more often procompetitive or competitively neutral than actively harmful. Moreover, given that bias occurs in less than 10% of queries run on Google, anticompetitive exclusion arguments appear unsustainable.

Indeed, we find that theories of vertical foreclosure find virtually zero empirical support in our data. Moreover, it appears that, rather than being a function of monopolistic abuse of power, search bias has emerged as an efficient competitive strategy, allowing search engines to differentiate their products in ways that benefit consumers. We find that when search engines do reference their own content on their search results pages, it is generally unlikely that another engine will reference this same content. However, that both this percentage and the absolute level of own content inclusion is similar across engines indicates that this practice is not derivative of one's misuse of its market power, but an industry standard. In fact, despite conducting a much smaller percentage of total consumer searches, Bing is consistently more biased than Google, illustrating that the benefits search engines enjoy from integrating their

own content into results is not necessarily a function of search engine size or volume of queries. These results are consistent with a business practice at issue is efficient and at significant tension with arguments that such integration is designed to facilitate competitive foreclosure.

Inclusion of own content accordingly appears to be just one dimension upon which search engines have endeavored to satisfy and anticipate heterogeneous and dynamic consumer preferences. Consumers today likely make strategic decisions as to which engine to run their searches on, and certainly expect engines to return far more complex results than were available—and the industry standard—just a few years ago.⁵¹ While the traditional 10 blue links results page is simply not an effective competitive strategy today, it appears that own-content inclusion is.⁵² By developing and offering their own products in search results, engines are better able to directly satisfy consumer desires. For example, when Google or Bing returns its own map in response to a search run for “maps,” consumers can click once and arrive at their desired outcome—rather

⁵¹ For example, over the last few years, search engines have begun “personalizing” search results, tailoring results pages to individual searchers, and allowing users’ preferences to be reflected over time. See, e.g., Danny Sullivan, *Google Now Personalizes Everyone’s Search Results*, SEARCH ENGINE LAND (Dec. 4, 2009, 6:18 PM), <http://searchengineland.com/google-now-personalizes-everyones-search-results-31195> (“For example, let’s say someone else prefers Barnes & Nobles. Over time, Google learns that person likes Barnes & Noble. They begin to see even more Barnes & Nobles listings, rather than Amazon ones.”). See also *supra* note 42.

⁵² See, e.g., Greg Sterling, *Yahoo: We’re Moving from Web of Pages to Web of Objects*, SEARCH ENGINE LAND (May 19, 2009, 4:36 PM), <http://searchengineland.com/yahoo-were-moving-from-web-of-pages-to-web-of-objects-19524> (“The big idea (now familiar) is moving beyond ‘10 blue links’ (popularized as a criticism of search by former Ask CEO Jim Lanzone) to a ‘web of objects.’ The ‘web of objects’ presented by Yahoo is a better representation of the ‘real world’ in search results. In other words: more closely aligning user intent with search results and mapping those to real-world tasks.”).

than having to click to another site, locate the map on its site, and finally engage in their desired activity.

Our evidence reveals very little search engine bias, and no overwhelming or systematic biasing by Google against vertical or general search competitors. To the contrary, our results strongly suggest that own-content bias fosters natural and procompetitive product differentiation. Accordingly, search bias is likely beneficial to consumers—and is clearly not indicative of harm to consumer welfare.

V. Conclusion

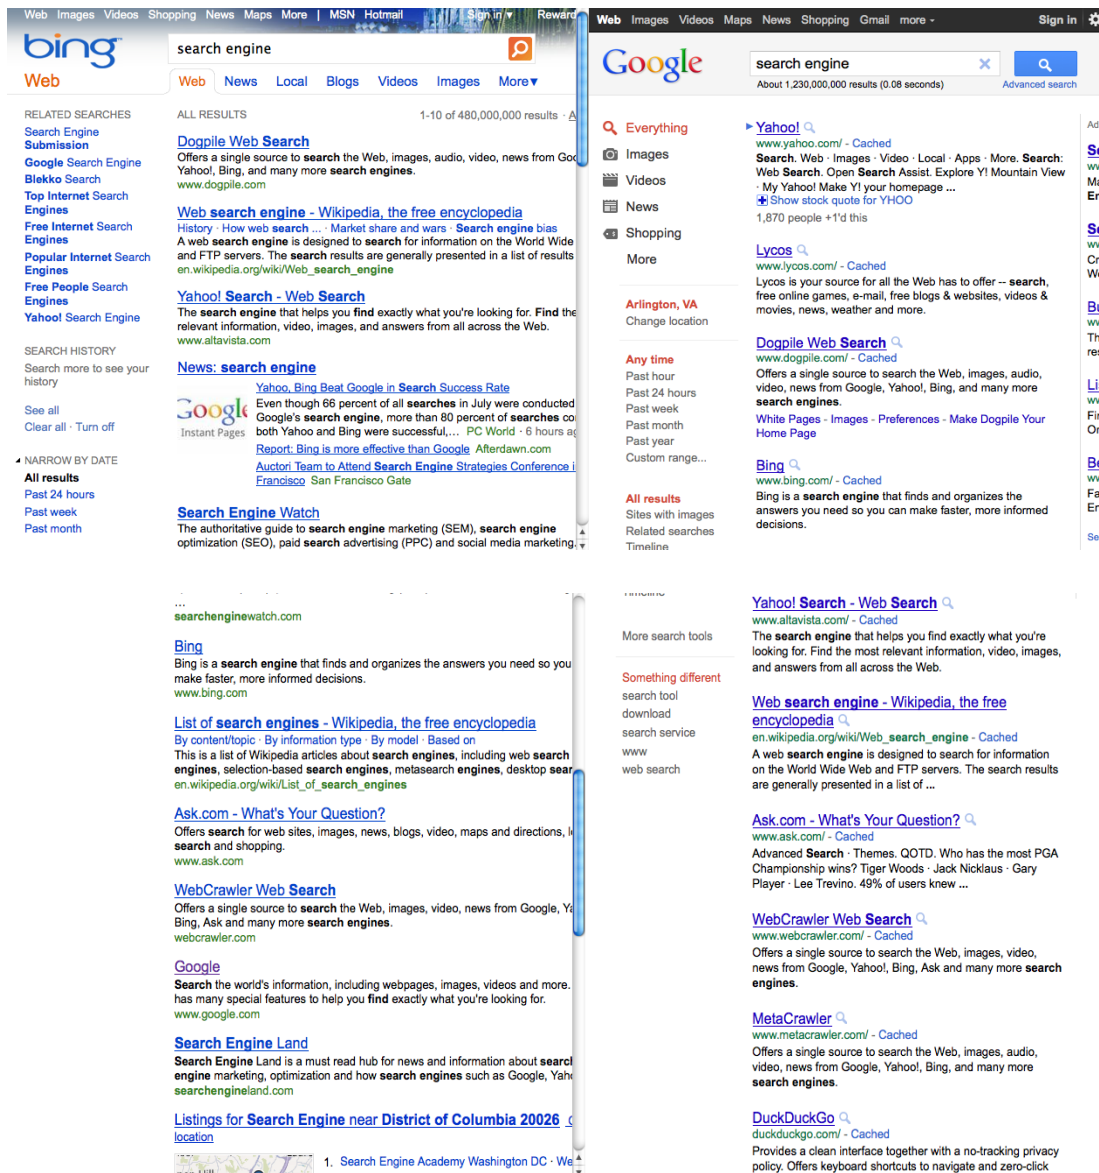
Competition among search engines can be expected to result in different rankings of any number of websites, including those featuring their own or affiliated content, as engines compete for diverse consumers with heterogeneous preferences. Using a small and non-random sample designed to maximize the incidence of search bias, Edelman & Lockwood misinterpret their own findings to conclude that search bias is not only prevalent, but sufficient to warrant antitrust intervention. Our critique, replication and extension of Edelman & Lockwood's analysis suggest these concerns are overstated and misguided. Most importantly, the authors treat "bias" as inherently and deeply suspect, do not engage in any analysis of its actual competitive effects, and simply disregard the well-known procompetitive explanations for vertical integration. This analysis is perverse from an antitrust perspective precisely because it ignores the welfare of consumers in favor of focusing upon individual websites.

Our analysis finds that own-content bias is a relatively infrequent phenomenon. Google references its own content more favorably than rival search engines for only a small fraction of terms, whereas Bing is far more likely to do so. For example, in our replication of Edelman & Lockwood, Google refers to its own content in its first page of results when its rivals do not for only 7.9% of the queries whereas Bing does so nearly twice as often (13.2%). Further, using Edelman & Lockwood's own data, neither Bing nor Google demonstrate much bias when considering Microsoft or Google content, respectively, referred to on the first page of search results. Collectively, rivals do not mention Google (or Microsoft) content any less often than Google (or Bing) mentions its own content. In our analysis of a large, random sample of search queries we find that Bing generally favors Microsoft content more frequently—and far more prominently—than Google favors its own content. Google references own content in its first results position when no other engine does in just 6.7% of queries, while Bing again does so over twice as often (14.3%). Importantly, when search engines appear to more prominently rank their own content in these samples, it is clear both Bing and Google do so. Indeed, Bing appears to be more biased than Google.

Moreover, the remaining bias appears to be much more consistent with natural and procompetitive product differentiation than with anticompetitive foreclosure. These few results did not involve attempts to prominently display Google's own general or vertical search content over that of rivals. Consistent with these data, and as

others have observed, if Google were foreclosing search rivals from prominent display in its organic search results, one would expect the bias to be present for the query "search engine," with Google favoring its own core business at the expense of leading search rivals.⁵³ Instead, a side-by-side comparison of Google and Bing results for that query demonstrates that Google not only is not prominently ranked, but that Google does not reference its own search engine at all; perhaps this is because its algorithm figures that if the user is entering the query from Google he or she already knows how to find it. However, Google's chief search rivals fare quite well on Google compared to Google's rank on Bing, where it appears tenth. While such evidence is anecdotal, the limited evidence available is facially inconsistent with the theories of anticompetitive foreclosure proffered by Google's critics.

⁵³ Sullivan, *supra* note 2 ("Having watched this results set for literally years, I'm borderline believing that rather than favoring itself, Google is deliberately downgrading itself here, as a way to show the world how it doesn't favor itself.").



From an antitrust perspective, as we've explained, differences in own-content references across engines fail to indicate consumer harm; to the contrary, it is quite possible—and indeed likely—that these differences imply the existence of intense competition among engines. For bias to have antitrust relevance, it must be linked to an analysis of its impact upon consumer welfare. In order to generate plausible competitive concerns, search engine bias must be both prevalent and sufficient in

magnitude to deprive rivals of efficient scale, exclude rivals on grounds other than superior quality, and reduce consumer welfare.

Evaluation of our large random sample of user queries significantly improves our understanding of competitive differences among search engines both generally and with respect to ranking their own content. However, all of the data presented here—our large study, our replication of Edelman & Lockwood, and Edelman & Lockwood’s own analysis—simply do not support claims that own-content bias is of the nature, quality, or magnitude to generate plausible antitrust concerns. Indeed, antitrust regulators should proceed with caution when evaluating such claims given the overwhelmingly consistent economic learning concerning the competitive benefits of vertical integration for consumers.⁵⁴ Serious care must be taken so as not to deter vigorous competition between search engines and the natural competitive process between rivals constantly vying to best one another to serve consumers.

⁵⁴ Lafontaine & Slade, *supra* note 14; Cooper et al., *Vertical Antitrust Policy as a Problem of Inference*, FED. TRADE COMM’N (Feb. 18, 2005); Benjamin Klein & Joshua D. Wright, *The Economics of Slotting Contracts*, 50 J. L. & ECON. 421 (2007); Joshua D. Wright, *Sacrificing Consumer Welfare in the Search Bias Debate*, TRUTH ON THE MARKET (April 22, 2011), <http://truthonthemarket.com/2011/04/22/sacrificing-consumer-welfare-in-the-search-bias-debate/>.