



Defining and Measuring Search Bias: Some Preliminary Evidence

November 3, 2011

[Joshua Wright](#)

Summary

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 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 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 (“owncontent 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 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 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 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, at minimum, be sufficient in magnitude to foreclose rivals from achieving minimum efficient scale (otherwise, if it merely represents effective competition that makes life harder for competitors, it is not an antitrust concern at all). 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 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 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 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.

[View Article](#)