Rev Ind Organ (2015) 46:25–57
DOI 10.1007/s11151-014-9434-z

Economics and the FTC's Google Investigation

Michael A. Salinger · Robert J. Levinson

Published online: 8 August 2014

Abstract We explain the issues in the Federal Trade Commission (FTC's) antitrust investigation into whether Google's use of "Universal" search results violated the antitrust laws and assess the role for economics in the FTC's decision to close the investigation. We argue that the presence of the Bureau of Economics infuses the FTC with an economic perspective that helped it recognize that "Universals" were a product innovation that improved search rather than a form of leveraging. Labeling them as "anticompetitive" would have confused protecting competition with protecting competitors.

Keywords Internet search · Antitrust · Market definition · FTC · Universal search · Google · Two-sided markets

1 Introduction

In January 2013, the Federal Trade Commission (FTC) closed its 19-month investigation that focused on whether alleged "search bias" by Google violated US antitrust law.[1] According to the FTC's brief closing statement, the bias allegations were that "Google unfairly preferences its own content on the Google search results page and

[1] The FTC investigation covered some issues besides Google's use of Universals. These issues included "scraping," Google's AdWords API, and standard essential patents. The aspect of the investigation that drew the most attention concerned Google's use of Universals.

M. A. Salinger (✉)
Boston University School of Management, 595 Commonwealth Ave., Boston, MA 02215, USA
e-mail: salinger@bu.edu

R. J. Levinson
Charles River Associates, 1201 F St NW, Suite 700, Washington, DC 20004, USA
e-mail: rlevinson@crai.com

selectively demotes its competitors' content from those results." The closing statement went on to explain that the key finding that convinced the FTC to close its investigation was that Google did not change "its search results primarily to exclude actual or potential competitors." Instead, it concluded, "The totality of the evidence indicates that, in the main, Google adopted the design changes that the Commission investigated to improve the quality of its search results" (Federal Trade Commission 2013a).

The Bureau of Economics is a distinctive feature of the organizational structure of the Federal Trade Commission. The internal organization of the FTC is designed to preserve an independent voice for economists. One widely accepted role for economists at the FTC is to perform the statistical analyses that are needed for both the Commission's antitrust and consumer protection enforcement missions. But, as important as econometrics has become in some cases, the Commission's need for statistical expertise is not sufficient to explain the number of economists employed by the FTC or the prominence of economics in the FTC's organizational structure.[2]

The FTC's Google investigation is an important case for evaluating the role of economists and economics at the FTC because the stakes were high and because, while econometric analysis was not central to the investigation, economics principles were. Assessing the unique role of the Bureau of Economics in such cases is harder than in cases where econometrics plays prominently. In antitrust investigations, the Commission receives separate memos from the Bureau of Economics and the Bureau of Competition (which houses the attorneys who work on antitrust enforcement) as well as separate recommendations from the two Bureau Directors.[3] When the Bureaus disagree and the Commission heeds the advice of the Bureau of Economics, the effect of the Bureau of Economics is clear. But such cases are rare.

As we were not employed by the FTC during the investigation, we cannot know what positions the Bureau of Economics took or what direct influence it had on the final outcome. But focusing on the positions taken by the two Bureaus and, in particular, the differences between them in a particular case can miss the broader impact of the Bureau of Economics. The value of economics at the FTC often manifests itself in concurrence between the Bureaus on an economically sound decision that results from their history of collaboration in antitrust enforcement.

We believe that the FTC's decision to close its investigation into Google's search practices was economically sound. We identify, on the basis of information that was revealed in the Commission's closing statement in this matter, two ways in which economic reasoning appears to have helped the FTC arrive at its decision and a third way in which it might have helped.

One possible role for economics was to provide a broad policy perspective on the appropriateness of government intervention in the design of Internet search sites. The FTC's investigation concerned Google's product design, which is the most important dimension of competition in Internet search. As a broad policy matter, we would expect

[2] One important source of information about the role of the Bureau of Economics is the articles on the FTC in the Review of Industrial Organization's annual Antitrust and Regulatory update (Carlson et al. 2013; Shelanski et al. 2012; Farrell et al. 2011).

[3] In most but not all cases, the Bureau Directors endorse their staffs' recommendations.

most economists to argue that the government should exercise extreme caution about intervening in product design.

A second possible role for economic reasoning was to help distinguish competitive from anticompetitive behavior. The Commission's clear articulation of the point that Google's competitive efforts to improve its search results for Google users would naturally harm some competitors reflects a key economic insight. However obvious the point might seem, antitrust enforcers in the past have often failed to understand it when economic analysis played less of a role in antitrust enforcement than it currently does.

A third possible role for economics was in applying legal standards with economic content. Market definition arises in all antitrust cases (except possibly for price fixing cases). We argue that product definition would also have been an issue if the FTC had brought a case. We argue further that the practice of addressing legal standards with economic content is not merely a requirement of the modern legal system (as some economists both inside and outside the FTC sometimes argue) but is in fact an important mechanism for ensuring that antitrust enforcement is economically sound.

The remainder of this article is organized as follows: A fundamental issue in the investigation was whether Google's search results were biased toward its "Universals." One cannot analyze that issue without understanding exactly what Google's Universals are. Section 2 explains them. Section 3 discusses the nature and history of Internet search. This material provides the essential context for understanding Google's development of Universals. Section 4 then turns to our analysis of the role of economics in the case. It is divided into three sections, each of which is devoted to one of the three possible roles for the Bureau of Economics. Section 5 contains a brief set of conclusions.

2 Universal Search

To quote from the FTC's closing statement:

> Some vertical websites alleged that Google unfairly promoted its own vertical properties through changes in its search results page, such as the introduction of the "Universal Search" box, which prominently displayed Google vertical search results in response to certain types of queries, including shopping and local (Federal Trade Commission 2013a).

Before going to the substance of these concerns, it is necessary to deal with issues of nomenclature. While the term "vertical" seems to imply a relationship to "vertical integration" or "vertical foreclosure," a "vertical website" is a specialized search site. Examples include Travelocity (travel), Orbitz (travel), CNET (electronics shopping),[4] Yelp! (local businesses), NexTag (shopping), and Fandango (movie information).

[4] CNET is not just a shopping site as it also publishes content about the electronics and information technology. But it is a good site for looking for electronics shopping.

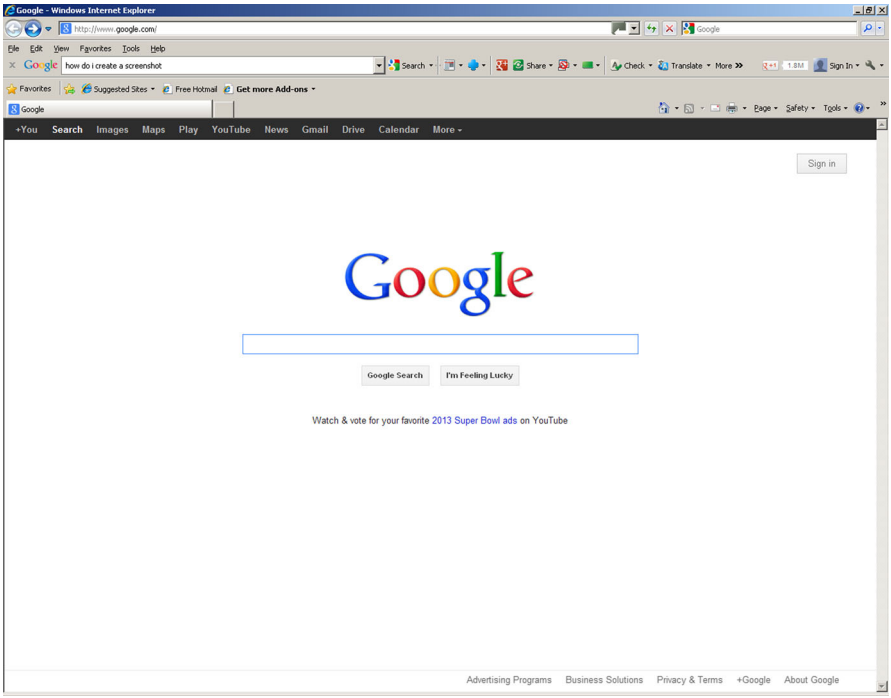

Fig. 1 A screen shot of upper left hand portion of Google home page taken May 2013. Notice the black bar with the words "You Search Images Maps Play YouTube News. . .." Clicking on the appropriate label within the black bar was one way to access Google's thematic results. Entering a query and clicking the "I'm Feeling Lucky" icon took users directly to what would have been the first Web site listed on Google's general SERP

To understand what "Universal Search" is, consider Fig. 1, which is a screen shot of the upper-left hand side of the Google home page as it appeared in May 2013.[5] The black bar near the top includes hyperlinks labeled "You," "Search," "Images," "Maps," "Play," "YouTube," "News," "Gmail," "Drive," "Calendar," and "More." The selection labeled "Search" is in a boldface font because the page that was being viewed was the Google Search screen.[6] This screen was a point of entry to Google's general search engine. Google's search engine is referred to as a "general" search engine because it is capable of providing responses to virtually any conceivable search term issued by the user. The results returned by Google in response to a user's query include, among other things, information

[5] This section of this article makes extensive use of Salinger and Levinson (2013). Screen shots from that time illustrate the Universal search results at the time of the FTC investigation better than do more recent screen shots. Not only do Google results change over time (both because of changes in its algorithms and changes in available content), but they can vary by user (based, for example, on location or search history). Someone else who attempted the searches that we describe may have gotten different results.

[6] The bolder font may not be clear in the screenshot, but it was clear when one used Google.

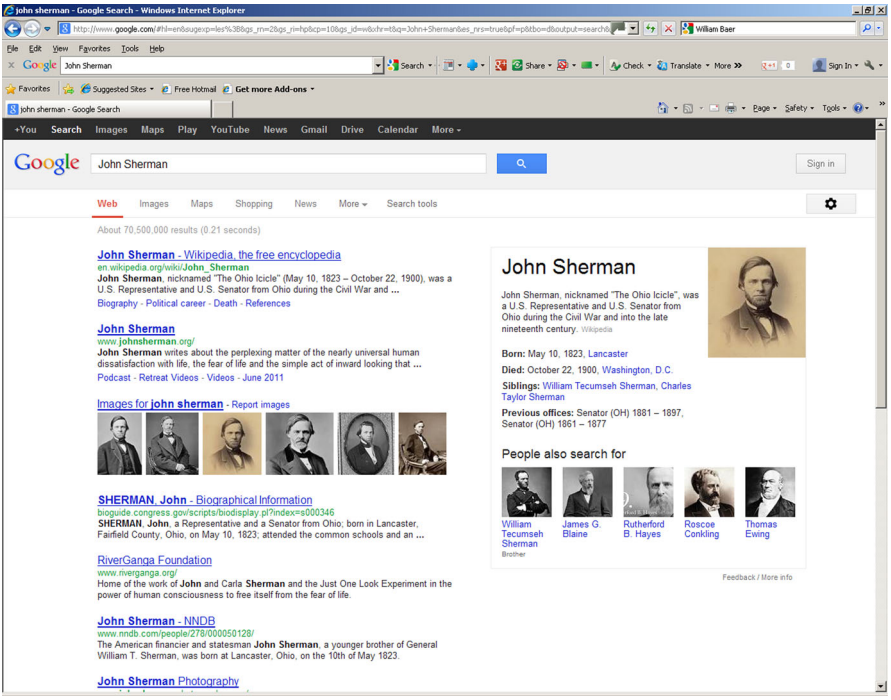

Fig. 2 A screen shot of the upper portion of first Google results page from a query for "John Sherman," May 2013. The third item down, with the title "Images for john Sherman" and showing six pictures, is an example of Google's Images Universal

found and indexed by Google as its software automatically "crawls" the World Wide Web.[7]

Suppose one is interested in pictures of John Sherman, the sponsor of the Sherman Antitrust Act. One way to find such pictures is to type "John Sherman" into the search box on this Search page. Figure 2 displays the results we obtained.

[7] Software programs used by general search engines to crawl the Web are known generically as "spiders," "bots," or "crawlers." Google crawls the web using its "Googlebot." See, e.g., Hayes (n.d.). Although they harvest enormous amounts of data, crawlers such as Googlebot do not access every site on the Web. One reason for this is that only a small fraction of the Web, known as the "surface" or "public" Web, can be accessed by crawlers. The remainder, known as the "deep" or "invisible" Web," includes concealed content and material that is "either in a database or stored in HTML pages many layers deep with complex URL addresses." See, for example, the links "Definition of: Surface Web" (n.d.) and "Definition of: Deep Web" (n.d.) provided in the references to this article.. A second reason is that Web site administrators often can block crawlers' access to their sites by including appropriate directives in their Web site code, either in a special file called "robots.txt" or in meta tags embedded in individual Web pages. Google and most other reputable users of crawlers respect these directives. See Hayes, *Ibid*. See also Google (n.d. a). The "invisible" Web also includes Web content that is generated dynamically as the result of user actions, rather than being stored on static Web pages. Dynamically-generated content cannot be found or indexed by Web crawlers because it does not exist on the Web except in response to user requests.

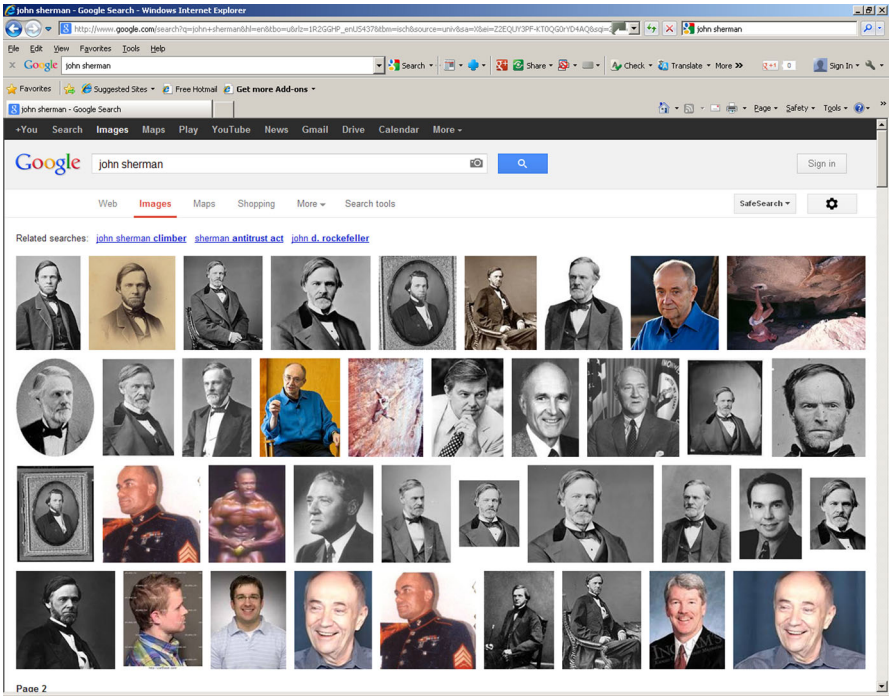

Fig. 3 The results from clicking on the link to the Images Universal in Fig. 2, May 2013

The third item on the left-hand side reads, "Images for john sherman—Report images" and has six images below it. Figure 3 shows the result of clicking on the "Images for john sherman" blue link.

Another way to obtain such images is first to click "Images" in the black menu bar on the main Google search page, then type "John Sherman" into the search box.[8] Figure 4 shows the results from doing so. The key point about Fig. 4 is that it is identical to Fig. 3.

Figures 2, 3 and 4 provide an example of one type of Universal Search result: in this case the "Images Universal." Figure 2 shows the results of a "general search" (or what Google labels a "Web" search in Figs. 2, 3 and 4). Figures 3 and 4 show the results from one of Google's thematic searches.[9] Note that the images in Fig. 2 are the first six images in Figs. 3 and 4. A Universal Search result contains direct links

[8] Note that "Images" is now brighter than the other words in the black bar.

[9] In Figs. 1, 2, 3 and 4, the black rectangle near the top of the page says, ". . . Search Images Maps Play YouTube News . . ..." Each is a clickable "tab" that leads to a page with a search bar (as well as content in the case of Maps, Play, YouTube and News). The same query in these different tabs yields different results because Google uses different algorithms to generate them. As described above, "Search" is Google's general search. Searches in the other tabs are thematic searches. For example, a search in the "Images" yields results based on an image theme, meaning that the results are images. In addition to being based on a different algorithm, a thematic search might be based on a more limited set of crawled sites.

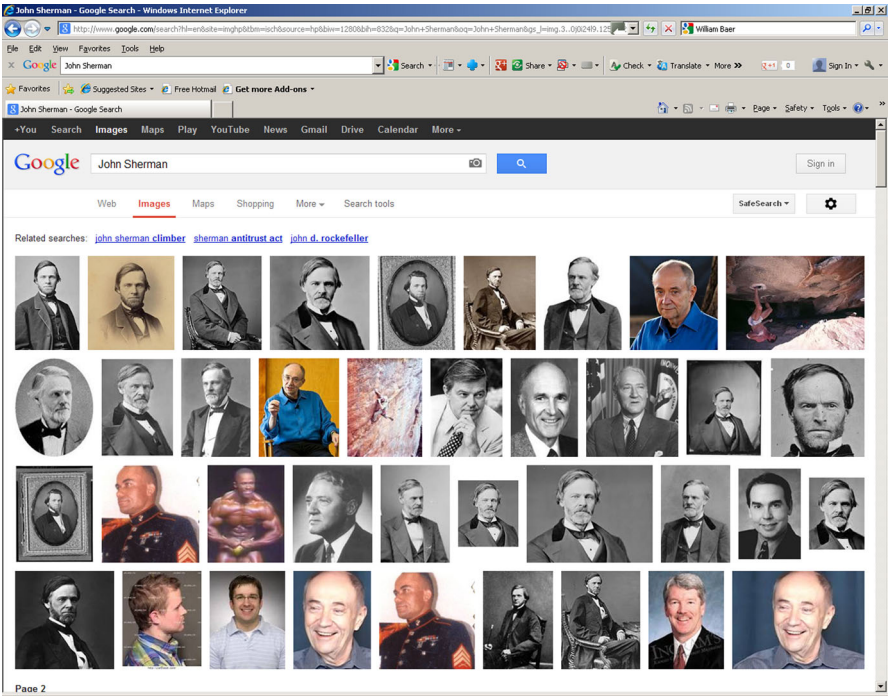

Fig. 4 A screen shot of the the results of a query for John Sherman in Google's thematic Images Search, May 2013. One could access Images Search by navigating to Google and then clicking "Images" in the black bar shown in Fig. 1

to the top results from one of Google's thematic search algorithms in the results of Web Search as well as a link to the more complete set of thematic results. Google's Images Search is an example of what the FTC closing statement refers to as a Google "property."

Unlike the Images Universal, which was not a focus of the Commission's investigation, Google's "Shopping" and "Local" Universals were identified in the FTC closing statement as having been the subject of complaints from competing "vertical" Web sites. Shopping and local "vertical" search sites[10] that valued placement in Google's Web search results had complained that Google's algorithms placed its Shopping and Local Universals above links to their sites. They seem to have argued that Google should be required to treat its Universal Search results as Web pages to be ranked

[10] We are not privy to the identities of all the complaining publishers of "vertical" Web sites, but Foundem and NexTag are examples of shopping sites whose publishers have complained publicly about Google bias. More generally, there are many "vertical" Web sites that provide specialized search capabilities that are tailored to specific user wants. Examples of "vertical" sites that compete with Google's Local Universal, in that they provide links to local businesses, include: Yelp! (providing reviews and links to local restaurants, shopping, entertainment venues and services); OpenTable (providing links and reservations to local restaurants); and Yahoo! Local (listings of local businesses, services and events). Examples of "vertical" sites that compete with Google's Shopping Universal include Amazon.com; Yahoo! Shopping; and Shopping.com.

according to the same general Web search algorithm Google uses to evaluate their sites, and that, had Google done so, the links to their sites would appear above (or appear above more frequently) Google's Shopping or Local Universals. They contended that the effect of this would have been to attract more user traffic to their sites.

3 Search Engines and Search Engine Innovation

The Internet gives people access to a trove of information; but for that access to be useful, people have to be able to locate the information that they want. Exactly how people would be finding the information they want on the Internet in 2014 was not obvious in, say, 1995; and exactly how they will do so in 2023 (and perhaps even 2015), is not completely obvious today.

3.1 Early Internet Search

One of the earliest Internet services was America Online (now named officially after its erstwhile acronym, AOL). AOL tried to create a relatively closed environment in which users were directed to news, shopping, travel, and other sorts of information services that were created and maintained by AOL. While successful with that approach for a while, AOL ultimately had to abandon it because people wanted access to the information available more broadly on the Internet rather than being limited to AOL's offerings.

Yahoo! provided another early approach to locating information on the Internet. At its inception in 1994, Yahoo! was a catalog of Web sites where people could click on different themes (e.g., Sports or Shopping) and then, within each theme, subthemes (e.g., baseball or football within Sports and clothing or electronics within Shopping). Unlike AOL, Yahoo! attempted to help people locate the best information available on the Web. A fundamental problem with that approach, however, was that it required humans to catalog the available information. As early as 1996, when Yahoo! added search capability, the amount of information available on the Internet had grown to the point where Yahoo!'s manual indexing approach was impractical.

The earliest versions of AOL and Yahoo! were not search engines as we use the term today because they could not respond to search queries. They were, however, general search sites (as distinct from general search engines) as they were starting points to look for information on the Internet. There is no such thing as an episode of general search. All (or at least virtually all) searches on the Internet are for specific information. AOL and Yahoo! could be the starting point to search for many different types of information: e.g., news, sports, shopping, and travel. Their approach to being useful for a wide range of searches was to have categories that were devoted to specific types of information. That is, the cataloging approach to helping people find information on the Internet virtually requires a thematic structure that resembles more recent thematic approaches to search.

3.2 First Generation General Search Engines

The first general search engines were Lycos and WebCrawler, both of which were introduced in 1994. AltaVista was launched in 1995, and was initially quite successful. Several other general search engines started in the mid-1990's.[11]

As we use the term, a first-generation general search engine had three defining characteristics: First, unlike the initial versions of AOL and Yahoo!, users could search by entering a query term. Second, it used Web crawlers to access and index (in principle, at least) the entire accessible Web.[12] Third, its response to a query was a list of Web sites that the user who issued the query might find helpful.

This first generation of search engines had advantages and disadvantages over other ways of locating information on the Internet. Because they are automated, Web crawlers can access information far faster and more comprehensively than is possible using a cataloging approach like the one that Yahoo! used when it started. On the other hand, a fundamental challenge for general search engines is how to link queries to the information that is cataloged by the Web crawler.

A purely algorithmic approach to the second step assigns for each query a numerical score that is designed to measure the likely value of each Website to the person who is performing the search.[13] The search engine then sorts the Web sites in descending order of the value of this measure of search "relevance," placing the Web site that receives the top score first, the one that receives the second-highest score second, etc. An example of a very simple algorithm would be to use the number of times that the search term appears on a Web page as the ranking criterion.[14] If one issued a query for "Barack Obama" to a search site using that algorithm, the first page listed would be the Web page containing the name "Barack Obama" the most times, the second site would be the Web page containing the name "Barack Obama" the second most times, etc.

This simple example illustrates four essential points about search and search algorithms: First, algorithms must be based on measurable criteria that can be processed automatically (without human intervention).[15] In the context of our hypothetical algo-

[11] For a discussion of early Internet search sites, see Sullivan (2003).

[12] As noted earlier, the portion of the Web that is accessible to crawlers is known as the "surface" or "public" Web. We say "in principle" because "even large search engines [index] only a portion of the publicly available part" of the Web. See "Web Crawler" (n.d.).

[13] One approach to search would be to have human-generated answers to some queries (perhaps augmented by machine-learning about which answers users clicked on) and then supplement those with results based on Web crawling and algorithms for which the site did not have human-generated answers. Ask Jeeves used this approach when it started in 1998.

[14] The science of assessing the relevance of documents for queries is known as "Information retrieval." Bush (1945) is credited with having introduced the idea of a systematic approach to information retrieval. One of the earliest approaches suggested in the 1950's was based on word overlap. The science had advanced well beyond that by the mid-1990's, although the appearance of query terms in a document continues to be an important consideration. The earliest Web browsers made use of developments up to that time. See Singhal (2001) for a discussion.

[15] That is, there is no human intervention at the time of the search. The design of the algorithm can entail human intervention, which can range in terms of how "heavy-handed" it is. One form of intervention is to augment or diminish the scores given particular sites. A still more heavy-handed approach would be to program directly the response to a particular query (without any reliance on a formula calculated about

rithm, there is some site on the Web that says "Barack Obama" the most times, some site that says "Barack Obama" the second most times, etc. Second, the measurable information that one can incorporate into a computer algorithm is only a proxy for how valuable a user finds a Web site for providing the information that he is looking for. The number of times the query terms appear on a page is simpler and more naïve than what a modern search engine would use as the sole basis for matching Web sites to queries. But even the scores that are generated by modern algorithms (or by the more sophisticated ones that will emerge in the future) are proxies for, rather than direct measures of, the quality of the match between a site and a query.

Third, different people issuing the same query are not necessarily looking for the same information, as each may have different "user intents." Some people who issue the query "Barack Obama" may, for instance, be looking for biographical information, while others may be interested in a recent news story about him. But the algorithm produces a single ranking.[16] No single ranking can be perfect for both users. Fourth, if Web sites benefit from placement on a search engine results page (SERP) and know the underlying algorithm, they can manipulate the design of their pages to improve their placement. So, in our example, a publisher that wants to appear higher on Google's results for "Barack Obama" could game the process by rewriting its content to state more frequently the President's name.

3.3 Google

Google started in 1997. The source of its initial success was its PageRank algorithm,[17] which used data on the importance and number of external links to a Web page as an indicator of that page's "quality." In broadly conceptual terms, the value of a Web site as a response to a query depends both on its relevance and its quality. The earliest search engines based results on relevance without regard to quality. Because PageRank captured aspects of quality as well as relevance, Google generated results that searchers found far more useful than the results generated by AltaVista and the other general search engines that were available at the time.[18]

External links to a page are an indicator of (or proxy for) page quality, but they do not measure quality directly the same way that a yardstick measures length or inches

Footnote 15 continued
each crawled page). Of course, any change in an algorithm designed to modify Google results is arguably human intervention.

16 To be sure, an algorithm might incorporate user-specific information, such as location or search history. But the fact remains that two searchers issuing the same query and that otherwise look identical to Google or any other search engine might be interested in quite different information.

17 More specifically, PageRank is an algorithm that "assigns an 'importance' value to each page on the Web and gives it a rank to determine how useful it is by taking into account the number and quality of other Web pages that contain links to the Web page being ranked by the algorithm." See Google, Inc. (n.d. b).

18 The potential use of links between pages was one fundamental way in which the Internet provided opportunities for information retrieval that had not been available in other applications of computerized information retrieval. Another, which Google's founders were not the first to realize, is that the volume of queries on the Internet is so great that many users issue the same query. As a result, a search engine can track user responses to a query and then use those data to modify its subsequent responses to the same query.

of mercury measure temperature. To be sure, there is a logical connection between links to a page and page quality, as Web page publishers tend to place links to sites they find useful.[19] Still, the logical connection between links to a page and quality is one that is plausibly true on average, but not a direct measure. By the same token, while higher-quality Web pages are, all else held equal, more likely to satisfy any given user's intent when issuing a given query, different users that issue the same search term may seek different sorts of information. A search engine algorithm that is intended to present a sorted list of relevant results in response to individual searches will rely on assumptions as well as user-specific information (which need not be limited to the search term he submitted). Thus, like page quality, the relevance of search results cannot be measured directly.

The key implication of these points is that even though a search algorithm entails sorting scores from numerical calculations on objectively observable data, the relevance of search results is based on a process that relies on subjectively chosen criteria. The only "objective" basis for believing that Google's initial algorithm using PageRank was better than the alternatives available at the time would have been evidence that users preferred Google's search results. Without knowing whether users preferred the results from Google's algorithms to those from other algorithms, there was no objective way to ascertain that Google's results were "better" (in the sense of providing more relevant results).

Notwithstanding Google's initial success, the earliest versions of the Google general search engine were limited in several important respects: First, its responses to queries were limited to links to Web pages it had crawled, which are candidate "blue links."[20] The ability to crawl the Web and link the results to queries was a crucial step in making the information available on the Web useful and accessible, but links are an inherently indirect way of answering questions. When a search engine merely provides links, its role is analogous to a card catalog in a library. It suggests where people might find the answer they are looking for, but does not provide actual answers.

Second, the information that is available on the Internet is not limited to information on Web pages that are reached by a Web crawler. As the Web has developed, a substantial amount of information is dynamically generated, which means that the publisher generates the information by accessing a data base in response to user input. As explained earlier, this sort of information is a part of the "invisible Web" that is not

[19] Assessing quality and incorporating those assessments into its search algorithms has been an important focus of innovation at Google. These assessments are the results of judgments made by Google's developers and managers. For example, Google considers the originality of a Web site's content (in contrast to links to content on other sites) to be an important indicator of quality. As noted in a Google blog entry, "low-quality sites [are] sites which are low-value add for users, copy content from other websites or sites that are just not very useful. . . [while] [h]igh-quality sites [are] sites with original content and information such as research, in-depth reports, thoughtful analysis and so on." See Google, Inc. (2011). While the determination of whether, and to what extent, a Web site's content is "original" may often be empirically observable, the emphasis placed on originality reflects Google's judgments regarding the relationship between a Web site's quality and originality.

[20] Google gives each of its search results a blue title. When the result is a link to a Web site, the title is itself a link to the site (meaning that clicking on the blue title takes the searcher to the Web site).

visible to Web crawlers (or at least first-generation Web crawlers—which can only access the "surface Web").[21]

Third, assuming that the earliest general search engines, including the earliest version of Google, were based on sorting scores from an algorithm that assigns a single value to each page with respect to each match, the results could not capture the diversity of possible motivations behind a search. In a single-valued algorithm, the second-ranked item is the one that would be first if the first-ranked item did not exist. In some cases, the second-ranked item might be very similar to the first item. That link may not be the most useful (i.e., the most relevant) link available to someone who clicks on the second link because he was dissatisfied with the first link.

Fourth, the earliest versions of Google were best suited to evaluating textual content. As an increasing fraction of the material on the Internet became images, videos, and audio files, Google needed algorithms that were well suited to evaluating and helping its users find such content.

Finally, concerns that third-party publishers (that is, operators of Web sites that are independent of Google) can manipulate and so "game" search engine algorithms are real. Web site publishers routinely employ search engine optimization companies (SEOs) to obtain information that can be used to tweak their Web sites in ways that increase their representation and improve their positions in search results. If search algorithms measured consumer utility perfectly, search engine optimization would pose no problems because the changes would improve a Web site's ranking only when they correspondingly improve the quality of the Web site itself. Precisely because the algorithms cannot generate search results pages that measure consumer welfare, however, changes in a Web site's design can simultaneously improve a Web site's ranking by an algorithm and lower its quality.[22] Such changes reduce the quality of the search engine itself, both to its users and its owner.

The result of this gaming is a constant cat-and-mouse game between Web sites and search engines as the former try to increase their visibility and profits by gaming the search provider's algorithms, and the latter adjust their algorithms to thwart such efforts. Indeed, Google's original insight about the "importance" of Web pages (as manifested in its PageRank approach) as an indication of quality is subject to manipulation.[23] Google devotes substantial effort to detecting sites that have artificially engineered external links so as to improve their placement in Google's search results.

[21] See note 7, above, for definitions of the "invisible" or "deep" Web. Additional sources on this topic include, e.g., Bergman (2001) and Sherman and Price (2003). Even if it is technically possible for a search engine to evaluate dynamically generated content, doing so would presumably require a substantial change in how Google crawls the Web and in its algorithms to evaluate the results.

[22] The trade press distinguishes between "White Hat" and "Black Hat" search engine optimizers. "White Hat SEO" improves the quality of a Web page whereas "Black Hat SEO" exploits imperfections in search algorithms that allow a Web site to get more prominent placement without improving (and perhaps even lowering) quality. See "Search Engine Optimization" (n.d.).

[23] The term "Google bomb" refers to a situation in which people have intentionally published Web pages with links so that Google's algorithms generate an embarrassing or humorous result. For example, by creating Web pages that linked the term "miserable failure" to George W. Bush's White House biography page, people were able to "trick" Google's algorithms into returning that page as the top link to a Google query for "miserable failure." See "Google Bomb" (n.d.).

3.4 "Vertical" Search

Sites that specialized in particular types of information became available at about the same time as the early general search engines. Travelocity and Expedia, which are currently two of the top three specialized Web travel sites, launched in 1996.[24] Amazon.com, which has become the starting point for a large number of shopping searches, started in 1995, albeit in the more limited role of an on-line book retailer. MapQuest, an early Internet mapping service that AOL acquired in 2000, launched in 1996. CitySearch, the first online source that was devoted to providing information about local merchants and locally available services, also launched in 1996.

In the late 1990's, the development of thematic search was not limited to specialty sites. Yahoo! was a leader in developing sites that focused on particular types of information. It started Yahoo! News in 1995, which by 1997 had become the most popular on-line source of news (Yahoo! Inc 1997). It launched Yahoo! Finance in 1996. As noted above, Microsoft started Expedia (albeit before it developed MSN, its first collection of Internet sites) in 1996.

Specialty search has one natural advantage over general search. By going to a specialty site (whether it is a stand-alone site or a thematic section within a more general site), a user reveals a great deal about what he is looking for; and that implicit knowledge simplifies the search for information that the user is seeking. A fundamental challenge for general search engines is to ascertain the broad intent of a search, whether it be for news, an item to buy, driving directions, or a trip to plan. To the extent that a specialty site is gathering information from the Internet, it can draw from a smaller set of information than does a general site;[25] and it can tailor the algorithm for ranking how useful a site is likely to be to someone issuing a particular query to the category of the search. For example, the algorithm for a news site is likely to give greater weight to the date (placing positive values on more recent pages) than would a site focusing on images. Finally, developers of a specialty site can custom-design its user interface to elicit specific types of information. For example, travel sites design their pages to solicit from users their origin and destination and when they want to travel.

While some specialty sites did draw information from the Internet, they were not restricted to such information. They licensed some of it from the information originators and produced some of it themselves. For example, the only practical way for a finance site to provide real time (or nearly real time) stock price data is to license it from the relevant stock exchanges. Yahoo!, despite its origins as a site that cataloged information available on the Web, licensed stock price data for Yahoo! Finance rather than referring users to some other site where they could find it. Travel sites get feeds from providers or, in some cases, commit to inventory of items like airplane seats or hotel rooms and then sell their offerings.

[24] The third is Orbitz, which five of the six major airlines launched in 2001.

[25] A specialty site that gathers information by crawling the Web can limit its crawling to sites that provide the class of information its users want. The point is not limited to Web crawlers, however. A specialty site that relies on human cataloguing of sites can limit the sites that it catalogs.

3.5 Google's Thematic Search

Google's first thematic search offering was for Images, which it started in July 2001. Google's apparent motive for introducing Google Images was to provide access to the sort of Web content that was not easily captured by Google's main algorithm at the time, which was focused on textual content.

Google started its news thematic search in the wake of the September 11, 2001, terrorist attacks.[26] At the time, Google's general search algorithm was not designed to identify important breaking news stories or to determine whether a user was seeking that type of information. Google failed miserably for people who tried to find the news about the September 11 attacks by going to Google and entering a query for "World Trade Center" or "World Trade Center attack." Four hours after the attack, the top link for a query for "World Trade Center" was the site for the World Trade Center Association: an association of 300 world trade centers in over 100 countries. Recognizing how badly it was serving its users, Google devised an ad hoc solution by placing in a portion of the page usually reserved for advertisements the heading, "Breaking News: Attacks hit US" along with links to the *Washington Post* and CNN Web sites. As Sullivan (2011) observed, "Google was literally telling people not to try searching."

Google started developing Google News shortly after September 11, 2001. In contrast to Google's general search algorithm, search results in Google News relied on "crawls" just of news sites, and the crawls occurred every hour. Google launched the beta version of Google News in 2002.[27]

Google launched Product Search, originally called "Froogle," in 2002. In addition to using crawled results, Google gave merchants the opportunity to provide direct feeds to Google about their product listings.

Initially, the results from Google's thematic searches were available primarily to users who navigated to the relevant Google thematic search page before entering their query. As a result, Google, like other general search sites, had what one Internet commentator likened to a Swiss Army knife (Sullivan 2001). Like a Swiss Army knife, Google contained many different tools. To use them, however, one had to open each relevant "blade" separately.

One of the problems with the Swiss Army knife approach was that many Google users did not use the thematic tabs that appeared on Google's landing pages. Instead, even when they had queries for which one of Google's thematic search results would have been most appropriate, they entered them into Google's general search bar. If one entered a shopping query into Google's general search bar, Google might return a link to a vertical search site that it had crawled (perhaps along with results for non-shopping Web sites), but it generally would not provide a link to the results from entering the query into the relevant Google thematic search site.

[26] See Kramer (2003) for a description of the start of Google News.

[27] It did not remove the "Beta" label until 2005, but it was used widely and was well-reviewed before that. For example, as noted in Kramer (2003), it won the "Best News" Webby in 2003. The Webbys are the equivalent of Academy Awards for the Internet.

3.6 Universals as a "Second Generation" of General Search

In Section B, we noted several key limitations of what we characterized as "first generation general search engines." In its original incarnation and, to a large extent, until the introduction of Universals,[28] Google's results were limited to answering queries with links to external Web sites. To the extent that the search algorithm is best suited to evaluate textual content, it might fail to generate relevant images, video, or audio files. An algorithm that assigns a single value to serve as the basis for ranking a Web site's potential usefulness to a search cannot inherently value diversity of results. Finally, a single algorithm for all types of searches was at an inherent and considerable disadvantage with respect to thematic search sites. Unlike a query made in a general search bar, one made on a thematic search site inherently provides a user-defined indication of the category of information being sought. In principle, one might argue that Google could address the difficulty of ascertaining intent by trying to teach Google users to start at its own thematic sites. But Google has to design itself to provide valuable results given the way people actually use Google rather than the way Google might like them to use it.

While ascertaining user intent from a query perfectly is generally not possible, probabilistic inferences about intent can be made. Someone who issues the query "Barack Obama" might be looking for a recent news story about Barack Obama. Even though a query for Barack Obama in Google Product search does yield results, the probability that the intent behind a query for "Barack Obama" in Google general search is to purchase a product seems far more remote than the probability that the intent behind such a search is to find recent news.

As described in Sect. 2, the Universals at issue in the FTC investigation were links to the results of Google's thematic search within its general SERPs. Another and more descriptive term used to describe Universals is "blended search," as a Universal entailed blending the result from one or more different algorithms into Google's general search algorithm. Universals represent a probabilistic approach to understanding user intent. When Google's algorithms detect a significant probability that a user's intent is one for which the results from one of its thematic searches would best meet the user's needs, the link to the Universal serves as a type of follow-up question. In the query for Barack Obama, the link to the News Universal that appears asks, in effect, "Are you interested in news about Barack Obama?" Clicking on the Universal is implicitly an affirmative response while not clicking on it may be a negative response.

The introduction of Universals addressed some of the limitations that were inherent in the first generation of general search results. Because a Universal gives users access to a different "thematic" algorithm that reflects a different set of search objectives, its appearance on a Google SERP provides a more diverse set of results than the results that emerge from the ranking produced by a single algorithm. Because some Google Universals focus on non-text content—i.e., images and video—they enrich the type of content to which Google can point users. To the extent that Google thematic

[28] A precursor to Universals at Google was "OneBoxes," which were links to Google thematic results that appeared at the top (or, in some cases, the bottom) of Google's SERP. The introduction of Universals provided for more flexible placement of the links to Google's thematic search sites within its SERP.

search was more likely to rely on content not generated by crawling the Web (such as the merchant-provided information in Google's Shopping Universal), Universals can help users find information that the first generation of general search engines could not locate. Because each of these changes—providing more diverse results that reflect different possibilities of user intent, making images and video content more accessible through a general search, and providing additional classes of content—represent solutions to such fundamental limitations of the first generation of general search engines, Google's introduction of Universals delineated a second generation of general search.

The introduction of Universals represented not only an innovation in Google's algorithms but also in how it presented results. A general search engine could use multiple algorithms and still return "ten blue links." Rather than the ten being those receiving the highest score from a single algorithm, a meta-algorithm could, for each "slot," decide which algorithm it would use, and then pick the top available link from that algorithm.

Consider the search for "John Sherman" that we discussed in the previous section. Suppose that a meta-algorithm assessed a 70% probability that a query for "John Sherman" was a search for biographical information and a 30% probability that it was a search for a photo of John Sherman. The first result would then be the top-rated site from the algorithm that would be best for yielding biographical information. One of the subsequent links—perhaps but not necessarily the second—could be a link to the one photo that received the top ranking in the "Images" algorithm.[29] That is not how Google handled the blending of results from multiple algorithms. Instead, rather than placing a single thematic search result on the SERP, Google returned a group of links, together with a link to the full set of thematic results that are deemed relevant to the query, all at the same position on the page, in each Universal. Moreover, if the Google search engine determined a significant probability that more than one "thematic" search was relevant to the user's query, more than one Universal could appear on the Google search results page.

Google was not unique among general search engines in developing Universal Search. Bing and Yahoo! use them as well. Indeed, many aspects of Bing search resemble Google. If one enters a query into the search bar on the Bing home page, one gets the results of Bing general search. As with Google, the top left hand corner has links to Bing thematic search pages such as "Videos," "Images," "Maps," and "News." The results of some queries into Bing's general search engine contain links to the results of entering the same query on one or more of Bing's thematic search pages, i.e., Universals. Perhaps the similarity is mere imitation; but another and more compelling explanation is that Universals are an obvious and effective approach to improving upon the limitations that were associated with the first generation of general search engines.

[29] This is not the only feasible way to generate diversity in results. For example, using what we have called the first generation of general search algorithms, a search engine would, for each query, compute a single score for ranking and the top ten listings would be those with the top ten scores. A search engine could conceivably place the Web site with the top score first and then generate another set of scores for all remaining Web sites based on an algorithm that is contingent on features of the first listing.

3.7 Further Developments

While Google's introduction of Universals, such as its Products and Local Universals, helped it address some quite fundamental issues about its general search product, Universals did not address all the limitations. Universals were an innovation in the underlying algorithm for identifying sites that a searcher might find useful and in how Google displayed those results. They represented an approach to combining search for different classes of Web content: text, video, and images; and, to the extent that some of Google's thematic search results relied more heavily on direct feeds from Web sites that, for whatever reason, are part of the "invisible Web" that cannot be reached by Web crawlers, they enabled Google to give users access to information that was beyond the results of Web crawling.[30] They did not, however, change the fact that the ultimate outcome of a Google search was a link to an external Web site, which might in turn provide the user with the information he sought. It was not *directly* providing users with the information that they wanted.

An early example of Google directly providing an answer is the result from entering into Google the query, "What is the square root of 2?" Google returned (and still returns) the answer together with a link to Google Calculator: a program that is hosted by Google and that generated the answer. Google has subsequently expanded the extent to which it returns answers created by Google itself rather than links to third-party sources of information. In Sect. 2, we used a search on John Sherman to exemplify the Images Universal. We did not resort to a relatively obscure query merely to pander to an audience that is interested in antitrust. Figure 5 illustrates a newer and different response to a search that might be expected to return images, which resulted from the less obscure query, "Barack Obama."[31]

This newer type of response appears to the right of the "blue links." It includes some basic biographical information that is generated by Google and, above this, a set of photographs with a link to "More images," which is the Images Universal. Unlike the Images Universal that is returned by Google in response to our query "John Sherman" (see Fig. 2), the placement of this new result was not interspersed among the external Web links.[32]

There are many other ways in which Google now provides information directly. Examples of queries for which Google has returned direct answers to us include, "New York weather," "Who wrote *War and Peace,*" and "Who is Barack Obama's wife." Google's responses to these queries require a specific triggering mechanism. That is, the source of the response to these queries is something other than an algorithm generating the top ten Web links that respond to the query, and Google needs a meta-

[30] To the extent that Google licensed some content for its thematic search results, some successful Google searches may not have ended with a referral to the relevant data from the licensor's site, if, for example, such data are available only by entering a database search on that third-party site.

[31] The information about Barack Obama on the right hand side of Fig. 5 is an example of a "Knowledge Graph," which Google introduced in 2012. Google considered it a sufficiently important innovation to list it as a corporate highlight in its 2012 10-K. See Google, Inc. (2013) at p. 4.

[32] The sort of information on the right-hand side of the SERP in Fig. 5 is present in Fig. 2 as well. The key difference between Figs. 2 and 5 for the points we are trying to illustrate is that Fig. 2 includes an Images Universal on the left-hand side of the page where Google's organic search results appeared originally.

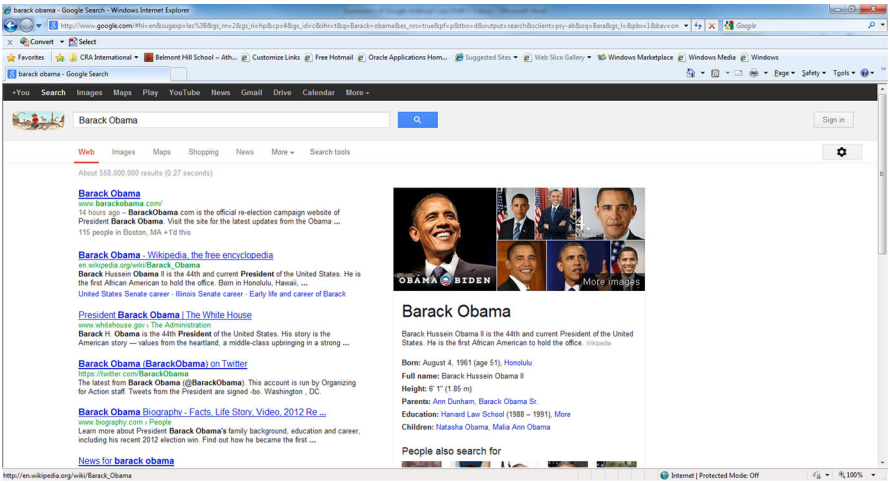

Fig. 5 A screen shot of the upper portion of the first page returned by Google in response to a query for Barack Obama, May 2013. It shows how Google integrated the Images Universal into a subsequent innovation in which Google directly provided information in response to some common queries (Figures 3 and 4 show the same innovation).

algorithm to determine when to use the results from these alternative sources in its SERPs.

We would characterize the direct provision of answers to queries as a third "generation" of general search. In using the term "generation," we do not mean that the second generation supplanted the first and the third supplanted the second, which happens with some products. Instead, the generations we describe, like generations of people, coexist. However, also like generations of people, the second generation was necessary to produce the third.

This last point is important because there can be little doubt that answering questions directly benefits consumers. In order to respond to a query with information, a general search engine needs to identify the likely intent behind a query, ascertain a degree of confidence about the intent, and identify the relevant information from the range of information resources that are available to Google. The introduction of Universals, which required Google to refine its (probabilistic) assessment of the intent behind a search and then provide a link to the best available information for that intent regardless of its form, was an important intermediate step toward the ultimate goal of providing information directly.

4 The FTC's Decision and the Role of Economics

Having explained what Google's Universals are and how their development fit into search innovation, we now turn to the FTC's decision not to bring a case and the role of economics in it.

4.1 Broad Policy Perspective

If the FTC had brought a case against Google, it would have had to articulate exactly what Google did to violate the antitrust laws and what changes Google would have to make to its algorithms to bring it into compliance. In effect, it would have placed the Commission and/or a court in a position of regulating Google's search results.

An important way in which economics contributes to antitrust enforcement at the FTC is by helping the agency approach enforcement activities within the broader context of sound economic policy. This important aspect of the agency's work takes place in tandem with its more narrowly-focused analysis of what constitutes a violation under existing legal standards. It is now a widely accepted principle of economics that the analysis of whether government intervention is appropriate cannot stop with an assessment of whether the conditions for the market outcome to be Pareto efficient are present. Real-world markets rarely satisfy the conditions for such an outcome—and these conditions plainly are not met in the case of Internet search. Rather, one must assess whether intervention by a necessarily imperfect government subject to bias imposed by the political system will in fact improve matters.

To be sure, the relative weights that one gives to market and government imperfections reflects ideology to some extent. Still, an economically sophisticated approach to government intervention in the market place requires at least some consideration of arguments made by Hayek (1945), Schumpeter 1950, and Stigler (1971).

Hayek stressed that information is itself a scarce resource and hypothesized that market signals are, in effect, sufficient statistics for summarizing a potentially large amount of information relevant to the decisions that individuals and businesses make. While failure of the conditions underlying the first welfare theorem implies that market signals are not perfectly efficient, a lesson many take from Hayek is that the efficiency of government intervention depends on whether the government has the information and expertise it needs to improve upon imperfect market outcomes.

Schumpeter stressed the role of what he called "creative destruction" and what modern management scholars call "disruptive technologies" in market competition. Wang Computer, Research in Motion's Blackberry, Polaroid cameras, and AOL are just four examples of once-successful companies that fell prey to the tide of creative destruction. The threat from disruptive technologies means that markets can be more competitive (in a dynamic sense) than a survey of the visible competitive landscape reveals.

Finally, Stigler's theory of economic regulation stressed that influencing regulation is itself an economic activity. He argued that government intervention is biased toward policies that provide large benefits to a small number of entities even if there are large offsetting costs that are dispersed. The principle applies to antitrust, as competitors frequently lobby antitrust authorities to help them accomplish what they failed to achieve in market competition with widespread benefits to consumers. Mindful of this bias, antitrust authorities in the U.S. (and many other jurisdictions) recognize that they should be skeptical of competitor complaints. But there can be an element of Greek tragedy to the phenomenon. Even an agency that is mindful of the trap can fall into it.

While these issues arguably arise at least to some extent in virtually all antitrust cases, all three played particularly prominent roles with respect to the FTC's Google

investigation. Internet search is still in its infancy and it is advancing rapidly. Innovation in search is an extremely important dimension of competition. In 2010, Chairman Jonathan Leibowitz designated the FTC's first chief technologist. While that development may have substantially increased the Commission's competence in dealing with technically advancing industries (and the Internet in particular), one chief technologist and a small supporting staff cannot possibly be sufficient to give the FTC the expertise to intervene in Internet search design without the danger of erecting roadblocks to innovation.

The FTC's interest in Google's use of Universals illustrates the point. As we described in the previous section, Internet search is developing from a service that responds to queries with links to Web sites that might provide the answer to answering questions directly. Google's development of Universals was one intermediate step in that process. If, as seems unlikely, the FTC could have crafted an intervention that was isolated to Google's use of Universals during the time period of the investigation, the intervention would have been ineffectual because subsequent changes in the product would have made the remedies irrelevant. More likely, any remedies would have had implications for future unanticipated developments in search design. It would be naïve to assume that these unforeseeable effects would turn out to be benign.

As we discuss below, the argument that Google's search results were anticompetitively "biased" toward Google Universals implicitly treats Google's Universals as being distinct products from Google's general search results. We question below whether they are. But, as Google moves toward providing answers directly to questions, it needs increasingly to develop (or at least license) content and choose where to place it. Restrictions on how Google places and presents content, including content that may compete with that of third-party Web sites, by definition limits Google's ability to innovate. Such interference runs the risk of preventing Google from introducing improvements to its service that benefit consumers.

This is all the more likely when the restrictions at issue are called for by Google's search engine and thematic search competitors, who would stand to gain if Google innovated less effectively. In particular, it appears that Microsoft urged the FTC to bring a case against Google.[33] It is unlikely that Microsoft did so because it believed that FTC intervention to eliminate alleged bias in Google's search results would improve Google's results for users (which likely would have decreased usage of Bing).

It is not a criticism of the FTC to suggest that, when deciding whether to issue a complaint, the agency needed to consider whether it had the expertise to intervene productively in such a technologically advancing business. If it concluded that it did not, it might have chosen not to bring a case even if it believed that it could prove an antitrust violation.[34]

[33] According to the Wall Street Journal, "The FTC's decision [to close its investigation of Google's search practices] also shows how anti-Google lobbying from rivals like Microsoft Corp. ... had little effect. Microsoft had pressed regulators to bring an antitrust case against Google." See Kendall et al. (2013). Microsoft's efforts to lobby the FTC to bring a case can also be inferred from Microsoft's strong and negative reactions to the FTC's decision to not bring a case against Google. See, e.g., Kaiser (2013).

[34] In making this point, we do not mean to suggest that the FTC could have demonstrated an antitrust violation. We argue below that it could not have.

As the FTC explained in its closing statement:

> Product design is an important dimension of competition and condemning legitimate product improvements risks harming consumers. Reasonable minds may differ as to the best way to design a search results page and the best way to allocate space among organic links, paid advertisements, and other features. And reasonable search algorithms may differ as to how best to rank any given website. Challenging Google's product design decisions in this case would require the Commission—or a court—to second-guess a firm's product design decisions where plausible procompetitive justifications have been offered, and where those justifications are supported by ample evidence.

Commissioner Maureen Ohlhausen's concurring statement made the point even more forcefully:

> Technology industries are notoriously fast-paced, particularly industries involving the Internet. Poor or misguided antitrust enforcement action in such industries can have detrimental and long-lasting effects. This agency has undertaken significant efforts to develop and maintain a nuanced understanding of the technology sector and to incorporate an awareness of the rapidly evolving business environment into its decisions. The decision to close the search preferencing part of this investigation, in my view, is evidence that this agency understands the need to tread carefully in the Internet space (Federal Trade Commission 2013b).

4.2 Protecting Competition vs. Protecting Competitors

Economically sound antitrust enforcement requires distinguishing harm to competition from harm to competitors. All modern antitrust enforcers—economists and lawyers alike—agree on this principle, but how to apply it in practice can be a source of contention. Applying the principle appropriately in this case was arguably more complicated than it usually is because of Google's two-sided business model.

4.2.1 Implications of the Literature on Two-Sided Markets

Google does not charge for search either in the form of subscription fees (per time period) or usage fees (per search). Instead, it is an advertising-supported platform. Its business model[35] is two-sided because it has two distinct sets of customers and because demand by at least one depends on demand by the other.[36] There is no con-

[35] While the term "two-sided market" is more common in the economics literature than "two-sided" business, the latter term is often more accurate. As Evans and Schmalensee (2007) correctly observe, firms with two-sided business models often compete on one or both sides with firms that have one-sided business models. For example, a basic cable network that operates on a two-sided model (with revenues from viewer subscriptions and from advertising) might compete for viewers with pay cable networks that get revenue only from subscription fees and not from advertising.

[36] Demand by advertisers to advertise on Google depends on Google's ability to attract searchers. In some cases, people who search on Google get the information they want from advertising-supported links. To the extent that they do, demand for Google by its search customers might depend in part on its success in

troversy about this point; and, because of it, one would expect a familiarity with the growing economics literature on "two-sided markets" to be useful for understanding competitive behavior by advertising-supported Internet search engines.

While, as we discussed above, the FTC recognizes that it must be skeptical of competitor complaints, customer complaints generally have more credibility. As a result, whether the vertical Web sites that complained about Google should be viewed as competitors or as customers could affect how the staff viewed their complaints.

We understand that Web sites that seek placement (and more prominent placement) in Google's search results argued that Google is a three-sided business—with themselves and others like them counting as the "third side"—and that Web sites that Google lists (or might list) in its organic results are therefore "users" of the Google platform. In part, their argument was that the traffic they get from Google is a promotional service for them. Indeed, their argument that Google "bias" violated the antitrust laws presumed that this service was essential to their business model. But these Web sites did not pay Google for the service that they received, and there is no meaningful sense in which Google has to compete to attract Web sites to be listed in its results.[37]

Google and Web sites that want to appear in Google's search results provide externalities for each other. For any Web site that wants to attract viewers, Google generates a positive externality. Web sites are free to base their business model on continuing to receive this externality, but that business decision on their part neither makes them a Google customer nor creates an obligation for Google to continue to provide that externality. The externalities that Web sites generate for Google can be positive or negative. High quality Web sites that Google helps searchers find provide a positive externality for Google, because users who locate quality information at those sites are likely to use Google again. Poor quality Web sites—particularly those that use Search Engine Optimization to boost their placement on Google—can provide negative externalities for Google to the extent that they crowd out the sources of information that searchers would find useful.

4.2.2 The Meaning of "Bias"

As the Commission explained, it "conducted a comprehensive investigation of the search bias allegations against Google." Before addressing whether alleged bias in Google's results is harmful to competition, one needs to be clear on what one means by the term "bias." In assessing the bias allegations, it is useful to distinguish between specific allegations and what might be termed general bias. The specific allegations included explicit demotion of competing Web sites in Google's algorithm and what Edelman (2010) termed "hard coding." The general allegations were that Google

Footnote 36 continued
attracting advertising customers. But this linkage is not necessary for Google to be a two-sided business. The dependence of demand by advertising customers on Google's success in competing for search customers is sufficient.

[37] As a technical matter, Web pages can choose not to appear in Google search results by denying access to Google's crawler. As an economic matter, though, Google does not have to compete to get Web page publishers to grant them access. Web page publishers benefit from appearing in search results.

triggered its Universals "too frequently" and placed them "too prominently" in its SERPs.

General Allegations The general allegations about Google bias were that Google places its Universals above competing sites of higher quality from the perspective of Google users.

As intuitively simple as the allegation might appear, demonstrating such bias objectively is impossible.[38] Indeed, any suggestion to the contrary reflects a fundamental misunderstanding of the nature of algorithmic search and of search competition.

To begin, some have suggested that, to "level the playing field" for Google's Universals and vertical search sites with similar themes, Google should score its Universals with the algorithm it uses to rank Web pages and use that score to trigger and place its Universals as if they were Web pages. Whatever intuitive appeal this argument might have, it makes no technical sense. Google's thematic results are dynamically generated content, not Web pages. Google does not crawl and index its own Universals and so cannot score their quality as it does third-party Web sites.

Moreover, while Google's search algorithms (and the search algorithms of other search engines) generate numerical scores for Web sites with respect to a search. as we argued above, the algorithmic scores are proxies for the relevance and quality, not direct measures. If the numerical scores were perfect measures of the relevance and quality of a site as a response to a search query, then any deviation of Google's ordering from the algorithmic scores would be bias. But, search engines change algorithms precisely because they conclude that the rankings that emerge from their existing algorithms are imperfect. Such conclusions necessarily reflect judgment rather than an objective measure (which would have been generated by an alternative and also inherently imperfect algorithm).

Internet search engines can make use of live traffic data that provide some evidence of the relevance and quality of search results. Not only can they use the data as a tool for evaluating results, they can automate changes to algorithms based on user responses. (The term "machine learning" refers to these automated updates.) For example, the FTC closing statement mentioned click data. It is intuitively plausible that high click rates on a site indicate that users find the site useful. If Google relies on click rates to guide its decisions about its algorithms and the design of its search pages, then click rates are relevant. But if Google does not rely on click rates, then the FTC or some other complainant could not prove Google bias with click data. Moreover, while Google's reliance on click data would make those data relevant, one would need to understand how Google might use them. If Google considers them to be merely suggestive and just one piece of evidence to consider when determining the relevance and quality of its historical search results, then antitrust authorities and courts should not view them as definitive.

Even if one could measure quality objectively or score Google's Universals as if they were Web pages, using those approaches to measure bias would rest on the

[38] While this point is not an economic point, most economists have some programming experience and an appreciation of what computer algorithms are. Thus, we would expect that the FTC economists would have been more attuned than the FTC attorneys to the difficulties in demonstrating bias objectively.

assumption that the relevant benchmark for unbiasedness is maximizing user welfare. For antitrust purposes, however, maximizing user welfare is not the relevant standard for unbiasedness. Because of its two-sided business model, Google might rationally make trade-offs in its search design between the interests of the people who search on Google and its advertisers without in any way deviating from competitive behavior.

A good analogy is broadcast and basic cable television networks, which also have two-sided business models. Inserting advertisements into television programs makes the programs generally less attractive to viewers.[39] Moreover, a television network might decline to broadcast a show that it knows that viewers would like to watch but that advertisers do not want to sponsor. The decision to air a less popular show that will generate more valuable advertising may be biased relative to a viewer welfare standard; but it does not constitute bias relative to competitive behavior that maximizes profits under competition.

Absent an objective measure of bias, the only practical meaning of bias is based on Google's intent. Google does have techniques for evaluating proposed changes to its algorithms and for assessing the on-going quality of its search results. Those techniques are fundamental to the way that Google competes. To assess the bias claims against Google, one can reasonably ask whether the objective behind those processes is to maximize the quality of search for users and whether Google applied those processes in evaluating Universals.

If the answer to both questions is "Yes," which implies that Google has developed its Universals to improve the quality of its search for users (and to provide a better competitive offering than rival search engines), then the bias allegations have to be dismissed. If the answer to one or both questions were "No," the inquiry would not end. One would still have to investigate what objective other than improving search for users Google was pursuing. A finding that Google adopted changes that it believed lowered the quality of its search results so as to sell more advertising would not be sufficient to establish bias that would be relevant for antitrust purposes, because this finding would be evidence that Google is competing for the business of advertisers. The potential finding that might have suggested a competitive concern would have been that Google reduced search quality in order to drive rivals from the market (presumably by demoting high-quality competitive sites).

That is not, however, what the FTC found. In its closing statement, the FTC stated: "The totality of the evidence indicates that, in the main, Google adopted the design changes that the Commission investigated to improve the quality of its search results, and that any negative impact on actual or potential competitors was incidental to that purpose" (Federal Trade Commission 2013a, p. 2).

Specific Allegations Some Web site publishers have alleged that Google sometimes explicitly demotes specific Web sites or classes of Web sites. Such behavior is not necessarily bias. Because search algorithms are inherently imperfect, Google might judge that some sites that score well (perhaps because their designers "game" the algorithm) are poor quality sites. In such cases, it might identify a generally observable

[39] Even if television users find some advertisements entertaining and informative, the market evidence is that viewers are generally willing to pay a premium for video entertainment without advertising.

feature of sites that its existing algorithms value inappropriately and alter its algorithms in a general way so that the algorithm better captures essential features of quality and relevance. As a practical matter, however, it might simply demote sites that its algorithms overvalue when it is not able to find a more general algorithmic solution. As long as it demotes sites that it views as being low quality or of low relevance, it is not engaging in bias even relative to a searcher welfare standard.

Demoting sites for reasons other than quality or relevance would be bias relative to a searcher utility standard. If the rationale for the demotion is to sell more sponsored links (perhaps by inducing the sponsored sites to buy sponsored links), then the behavior is not bias relative to a standard of (non-exclusionary) profit-maximization. The one rationale for demotion that might raise antitrust concerns would be if Google were demoting the sites because it viewed them as competitive threats.

The other specific allegation of bias concerned "hard coding": For example, Google used to have health content. If one queried for the name of a disease or medical condition, the first listing would be information from Google Health. Edelman (2010) compared Google search results from correctly typing the name of a medical condition with those from appending a comma to the query. At that time, the addition of the comma prevented the algorithm from triggering a result from Google Health. Edelman (2010) interpreted this as evidence of Google bias for its own content.

Google discontinued Google Health (and eliminated the effect of commas on search results), but the syntax of a query can affect whether the SERP includes a Universal. An example is that when we queried Google for "FTC v. Staples", the top link Google returned was "Court Opinions for FTC v. Staples." Clicking on the link took us to Google Scholar and a link to the district court opinion. Querying Google for "FTC vs. Staples" (i.e., including "vs" instead of "v") did not yield the same result. Thus, it appears that the triggering of Google Scholar for court decisions rested on the specific syntax of the query. But that is not bias. If Google's designers judged (with the aid of whatever evidence they use) that the most useful response to "FTC v. Staples" is a link to the court decisions, then providing that link (and placing it first) is not a bias.

Initially, Google's response to queries was restricted to links to web sites that were generated by applying an algorithm to the results of crawled Web sites. Google Calculator, Google Flights, and Google Health all represent qualitatively different responses. Moreover, these responses reflect classes of searches for which Google has or had decided it can provide users with a better response than is available from a Web search. If Google is going to have some responses generated through Web search and others generated in other ways, it has to have a meta-algorithm for deciding which type of result to place where. The fact that the placement of different types of responses rests on the precise syntax of a query may be necessary. And while one might debate whether Google Flights is more convenient for users than Travelocity or Orbitz, application of a search algorithm to them cannot resolve the issue.

4.2.3 Broad Assessment

The FTC's factual finding that "Google adopted the design changes that the Commission investigated to improve the quality of its search results" was sufficient to determine the outcome of the investigation, but it did so only because the FTC under-

stood the distinction between protecting competitors and protecting competition. As it went on to explain, "...[A]ny negative impact on actual or potential competitors was incidental to that purpose;" and "While some of Google's rivals may have lost sales due to an improvement in Google's product, these types of adverse effects on particular competitors from vigorous rivalry are a common byproduct of 'competition on the merits' and the competitive process that the law encourages" (Federal Trade Commission 2013a, p. 2). Antitrust enforcers have not always recognized these basic economic principles, and the FTC's clear articulation of them reflects the impact that an increased role for economics has had on antitrust enforcement.

4.3 Legal Standards with Economic Content

In subsection A, we argued that a possible role for economics at the FTC is to provide a broad perspective on whether government intervention in the market is warranted. When economics plays this role, the effect can be to get the Commission to exercise prosecutorial discretion not to bring a case that it might be able to win in court. The Commission's lack of enforcement of the Robinson-Patman Act is a case in point.[40] But the Commission cannot bring a case just because economic principles dictate that government intervention can improve upon the market outcome. The FTC is a law enforcement agency. To bring a case and prevail in court, it has to demonstrate a violation of one of the statutes that it enforces. While the key sections of the core antitrust statutes have very broad wording, the case law that has evolved from the enforcement of those statutes has established standards for proving a case.

At least part of the explanation for the importance of economics in antitrust enforcement at the FTC is that the legal standards that courts have created in antitrust cases have economic content. Courts have looked to the economics literature to shape these standards and to economists for expert testimony in applying them.

As a general matter, economists play a much more central role in antitrust cases at the FTC than they do in consumer protection cases. The difference is not because consumer protection cases fail to raise important issues of public policy toward business. They do. As a rule, though, the consumer protection attorneys do not need economists when they go to court whereas the antitrust attorneys do.[41]

[40] The Robinson-Patman Act is a 1936 amendment to Sect. 2 of the Clayton Act, which strengthened the Clayton Act's limitations on price discrimination.

[41] A key institutional feature of the FTC is that it enforces both antitrust and consumer protection statutes. The enforcement of these missions has historically been more separate than one might expect, given that they share the common objective of protecting consumers. The lawyers that enforce the competition and consumer protection statutes are organized in different bureaus (the Bureau of Competition and the Bureau of Consumer Protection). This institutional separation is probably not mere historical accident but, instead, reflects the fact that the statutory provisions are distinct and any law enforcement action must allege a violation of a particular statutory provision. As privacy has emerged as a central issue in consumer protection, how a wide array of Internet companies, including Google, collect and use data has been a concern at the FTC. Those concerns are, however, irrelevant for analyzing the allegations about Google's use of Universals. Despite the institutional separation between competition and consumer protection enforcement in

Footnote 41 continued

While, as we explained above, the FTC concluded that Google developed Universals to improve the quality of its search results, it is interesting to consider the economic hurdles that the FTC would have had to surmount to satisfy the legal framework that is needed for an antitrust claim if it had found otherwise.

One hurdle would have been market definition. To allege harm to competition, agencies must identify the relevant antitrust market in which competition is harmed. The suggestion that bias in Google's organic search results might have been an antitrust offense rested on an implicit assertion that Google is dominant in some relevant market. Most likely, the FTC would have alleged a relevant market for search on general search engines. Had it done so, the FTC would have had to defend that definition in court, presumably with expert economic testimony.

Somewhat ironically, economists have expressed reservations about this role.[42] In their book about the *U.S. v. IBM* case, Fisher, McGowan, and Greenwood argued:

> We have seen that market definition is, in a sense, unnecessary for the analysis of competition or monopoly. The constraints on the power of the alleged monopolist are the same regardless of how the market is defined, and, properly handled, different market definitions can only succeed in describing the same phenomena—the constraints—in different but equivalent ways (Fisher et al. 1983, p. 43).

In our experience, many economists at the FTC share this reservation; and, as a result, market definition often plays less of a role in internal deliberations than it might.[43]

Economists are often reluctant to delineate markets because doing so implicitly might be viewed as treating all sources of supply inside the relevant market's boundaries as (undifferentiated) competitive constraints and all sources of supply outside the market as no competitive constraint at all. This dichotomous treatment of supply sources as being entirely in or entirely out of the market is at odds with market reality when, as is often the case, firms face a virtual continuum of competitors that range with respect to the strength of the constraint they impose.

But reservations about drawing exact boundaries around a market cannot and do not obviate the need to conduct a rigorous analysis of competition, and defining the relevant market can provide a useful framework for structuring that analysis. Input from economists on market definition would have been particularly important in the Google investigation because non-standard features of the markets made it impossible (or at least inappropriate) to use techniques that have become standard in other cases.

individual cases, the participation of the Bureau of Economics in both provides an institutional mechanism to harmonize the broad enforcement philosophy that underlies the FTC's pursuit of its two main missions.

[42] Werden (1992) provides an excellent historical discussion of this reluctance. It dates back to the literature on monopolistic competition in the 1930's.

[43] The current DOJ/FTC Horizontal Merger Guidelines indicate that market definition is not always a necessary element in the economic analysis of mergers, noting that the diagnosis of " . . . unilateral price effects based on the value of diverted sales need not rely on market definition or the calculation of market shares and concentration." U.S. Department of Justice and Federal Trade Commission (2010) at § 6.1.

One of the most important contributions of economic reason to antitrust doctrine is the "SSNIP" test for market definition.[44] The SSNIP approach tests whether a hypothetical monopolist of a candidate relevant market could profitably sustain a "*S*mall but *S*ignificant and *N*on-transitory *I*ncrease in *P*rice" relative to actual pre-merger prices. When this criterion is satisfied, suppliers of products outside of the candidate relevant market do not significantly constrain the prices of the products within it.

The FTC could not have used a standard SSNIP test in this case. This approach was inappropriate for two reasons. First, the SSNIP test focuses on price increases as the way in which market power might be exercised. Google does not charge usage or subscription fees for use of its search engine, and no one suggested (as far as we know) that Google sought to use Universals to start charging for search (or, for that matter, that Google had any plans to start charging for search under any plausible circumstances). Second, the SSNIP test was developed for merger review. One longstanding issue in applying the SSNIP test in mergers is whether the benchmark price should be the current price or the competitive price. It asks whether a hypothetical monopolist would charge more than current (and presumably competing) market participants do. This approach runs into trouble when the pre-merger status quo features a dominant supplier that already exercises market power.

All versions of the Horizontal Merger Guidelines since 1982, including the most recent, indicate that the benchmark price used in the SSNIP test is the currently observable price, not the "competitive" one predicted by theory. Whatever the merits of this position for merger analysis, it is inappropriate for monopolization cases (White 2001, 2008). Because an existing monopolist sets price to maximize its profits, it will never do better by increasing price (by a SSNIP or any other amount). Using the SSNIP approach to test whether the scope of the relevant market is limited to an actual monopolist's products will therefore (and wrongly) indicate that the market is broader. This source of false negatives in market definition is frequently referred to as the "*Cellophane* fallacy," after a famous antitrust decision in which a firm with a dominant share of a candidate market for cellophane (du Pont) was found to compete in a broader market for flexible packaging materials.[45]

Another approach to market definition that has received increased emphasis in recent years is to demonstrate anticompetitive effects directly (Salop 2000). While this approach can be used to avoid the pitfalls of the *Cellophane* "trap," it is easiest to apply when the anticompetitive effect of concern is an increase in price—an approach that would have been irrelevant with regard to any concerns the FTC might have raised about Google's search business.

Even though these standard approaches would not have been practical, the FTC would still have had to define a market had it brought a suit; and one would expect the Bureau of Economics to help formulate a conceptually valid approach to doing so. While a standard SSNIP test would not have been suitable, the economic logic behind the SSNIP test (relative to a competitive benchmark) would have been. To

[44] U.S. Department of Justice and Federal Trade Commission (2010) at §4.1.

[45] *United States v. E. I. du Pont de Nemours & Co.,* 351 U.S. 377 (1956).

adapt the SSNIP logic, one first needs to identify the nature of the anticompetitive harm that is creating concern. In this case, it would be a reduction in search quality.[46] Then, to analyze the relevant market, the hypothetical monopolist approach would be conceptually valid. One could ask the question, what set of Web sites acting in concert could profitably reduce the quality of search by a "small but significant" amount relative to a competitive benchmark?

While it is not clear how one could answer the question quantitatively or even how one would measure "a small but significant" reduction in the quality of search, simply asking the question conceptually is useful. Suppose one were to hypothesize a market for "general search," meaning queries on general search engines. For simplicity, suppose further that the only three significant general search engines are Google, Bing, and Yahoo! For the hypothesized market to be a relevant market for antitrust purposes, one would need to consider whether Google, Bing, and Yahoo! would find it jointly profitable to reduce the quality of their search results by a small but significant amount. To answer that question, one would need to assess what alternatives searchers would resort to in the event that they did and then ask how many searches the general search engines would lose as a result. If the loss would be so great as to make the quality reduction unprofitable, then general search would not be a relevant antitrust market.

Once one poses the question of what alternatives searchers have to Google, it should be clear that it depends on the class of search. If general search engines were to degrade the quality of their responses to shopping queries (or perhaps some particular types of shopping queries), one would need to consider the possibility (indeed, likelihood) that searchers would start shopping searches at Amazon.com instead. If general search engines reduced the quality of their responses to searches about baseball scores, Amazon.com would be an unlikely source of baseball score updates. But that is completely irrelevant for considering the incentives of general search engines to degrade the quality of their search results to shopping queries.

The question of whether department stores constitute a relevant market closely parallels the question of whether general search engines do. Department stores sell a variety of goods. The precise mix varies across department stores, but a single department store might sell men's clothes, women's clothes, jewelry, housewares, and furniture among other items. All the types of items they sell are available at stores that specialize in those items. If Macy's offerings of men's clothing are not competitive with the offerings at men's clothing stores, it cannot expect to sell much men's clothing; and the failure of men's clothing stores to offer housewares does not prevent them from competing against Macy's to sell men's clothing.

As with general search engines, the portfolio of offerings at a department store is what, in the technology world, would be called a *feature*. Some people might find the variety of products that are offered by a department store to be a desirable feature perhaps because the consistency of a department store's choices with respect to price, quality, and taste across product classes reduces search costs and avoids the need to identify specialty outlets in each product class that match their tastes. Others might

[46] Since search is a new and rapidly developing product, the historical quality of search would not be the relevant "but-for" benchmark. Rather, the conceptual benchmark would be the quality of search under competitive conditions.

find the department store's broad assortment to be undesirable, perhaps because they find it difficult to locate what they want in a large store. But the feature does not define an antitrust market.

Whatever reluctance economists might have about exact market delineations, there are sound reasons, rooted in decision theory, for market definition to be part of the legal standard for alleging an antitrust claim. If antitrust enforcement were both perfect and costless, then one might argue that it should be sufficient to for the FTC to demonstrate that a defendant had harmed one or more competitors without benefiting any customers. But antitrust enforcement is both imperfect and costly, so it is important to ask whether the market would punish the defendant if such allegations were true.

The allegations in the case were that Google intentionally degraded its responses to shopping-related searches. Evaluating those allegations within a hypothetical market for searches on general search engines ignores the competitive constraint imposed on Google by Amazon.com with respect to shopping searches. That perspective obviously misses a crucial feature of market competition for Google. The legal standard that requires defining a relevant market ensures that the FTC could not overlook this point.

Another legal standard with economic content may have come into play if the FTC had chosen to bring a case is the separate products test. If Google's Universal results are simply a component of Google's SERP, then it is not clear what it would mean for Google's results to be "biased" toward its Universals. Thus, an implicit assumption underlying the bias allegations is that Google's SERPs and Universals are distinct products. According to this perspective, when one navigates to Google and enters a query, the resulting SERP constitutes an entire "unit" of Google search. Any further interaction with Google, whether it be a mouse click or the entry of a follow-up query, constitutes units of an additional, distinct product.

The area of antitrust law in which this issue has arisen is tying doctrine. Justice O'Connor articulated the underlying problem well in her concurring opinion in *Jefferson Parish* v. *Hyde*. As she explained,

> [T]here must be a coherent economic basis for treating the tying and tied products as distinct. All but the simplest products can be broken down into two or more components that are 'tied together' in the final sale. Unless it is to be illegal to sell cars with engines or cameras with lenses, this analysis must be guided by some limiting principle.[47]

In our opinion, the suggestion that Google's Universals and its SERP are separate products lacks any such limiting principle.[48] As Justice O'Connor wrote, the appropriate limiting principle must be "economically coherent," but we are not aware of any analysis in the economics literature for how to conduct a separate products inquiry

[47] Jefferson Parish v. Hyde, 466 U.S. 2 (1984).

[48] In our opinion, an episode of Google search begins when one enters a query into Google and ends when one exits Google with respect to that query (either because one goes to a site identified by Google, gets an answer directly from Google, or abandons Google as a source of useful information). Under our definition, clicks on Universals, clicks on pages other than the first results page, clicks on alternative spelling suggestions, clicks on suggested alternative queries, and even entry of so-called "refinement queries" are all part of a single episode of Google search.

even in the normal case in which people pay for goods. How to do one in the case of Internet search is even less clear.

If the FTC had brought a case against Google alleging search bias and Google had raised the separate products issue, we doubt that the FTC could have succeeded in convincing a court that Google's Universals are separate products. One role that would have been proper for the Bureau of Economics to have played in the investigation would have been to point out both the issue and the challenges as a matter of economics of establishing what the Commission would have had to establish in order to prevail.

5 Conclusions

The FTC's investigation into Google's search practices provides an interesting case to evaluate the role of the Bureau of Economics because it was a high-stakes case that forced the Commission to confront a (relatively) novel set of facts and that raised fundamental issues of government intervention in the market place. In this paper, we have discussed three possible roles for economics in the Google investigation:

One would have been to urge caution about intervening in a rapidly evolving, technologically advanced market subject to disruptive change. Another was to help the FTC recognize that if, as the evidence demonstrated, Google developed Universals to generate better search results for its users, it was behaving competitively even (and, arguably, especially) if competitors were harmed. A third was market and product definition. Even when the immediate goal of such analysis is to satisfy a legal standard, recognizing that the competitive constraints on Google extend far beyond Bing and Yahoo! required economic analysis that provided an essential piece of the reason why, in our opinion, the FTC was correct in deciding to close the investigation.

Acknowledgments Salinger and Levinson were consultants to Google during the FTC's investigation. This paper draws heavily on Salinger and Levinson (2013), for which Google provided financial support. The views expressed in this paper are the authors' alone. We are grateful to Mike Baye, Susan Creighton, Joe Farrell, Franklin Rubinstein, Scott Sher, and Elizabeth Wang for helpful comments.

References

Bergman, M. K. (2001, September 24). The Deep Web: Surfacing hidden value. Retrieved June 3, 2013, from BrightPlanet Web site: http://brightplanet.com/wp-content/uploads/2012/03/12550176481-deepwebwhitepaper1.pdf.

Bush, V. (1945). As we may think. *Atlantic Monthly*, *176*, 101–108.

Carlson, J. A., Dafny, L. S., Freeborn, B. A., Ippolito, P. M., & Wendling, B. W. (2013). Economics at the FTC: Physician acquisitions, standard essential patents, and accuracy of credit reporting. *Review of Industrial Organization*, *43*, 303–326.

Definition of: Surface Web. Retrieved February 3, 2014 from PCMag.com Web site: http://www.pcmag.com/encyclopedia/term/52273/surface-web.

Definition of: Deep Web. Retreived February 3, 2014 from PCMag.com Web site: http://www.pcmag.com/encyclopedia/term/41069/deep-web.

Edelman, B. (2010, November 15). Hard-coding bias in Google 'algorithmic' search results. Retrieved May 28, 2013, from http://www.benedelman.org/hardcoding/.

Evans, D. S., & Schmalensee, R. (2007). The industrial organization of markets with two-sided platforms. *Competition Policy International*, *3*, 150–179.

Farrell, J., Balan, D. J., Brand, K., & Wendling, B. W. (2011). Economics at the FTC: Hospital mergers, authorized generic drugs, and consumer credit markets. *Review of Industrial Organization*, *39*, 271–296.

Federal Trade Commission. (2013a, January 3). Statement of the Federal Trade Commission regarding Google's search practices, In the Matter of Google Inc., FTC File Number 111–0163. Retrieved February 9, 2014, from Federal Trade Commission Web site: http://ftc.gov/os/2013/01/130103googlesearchstmtofcomm.pdf.

Federal Trade Commission. (2013b, January). Statement of Commissioner Maureen K. Ohlhausen in the Matter of Google, Inc., Retrieved March 6, 2014, from Federal Trade Commission Web site: http://www.ftc.gov/sites/default/files/attachments/press-releases/google-agrees-change-its-business-practices-resolve-ftc-competition-concerns-markets-devices-smart/130103googlesearchohlhausenstmt.pdf.

Fisher, F. M., McGowan, J. J., & Greenwood, J. E. (1983). *Folded, spindled and mutilated: Economic analysis and U.S. vs. IBM*. Cambridge: MIT Press.

Google bomb. (n.d.). Retrieved March 6, 2014 from Wikipedia Web site: http://en.wikipedia.org/wiki/Google_bomb.

Google, Inc. (2011, February 24). Finding more high-quality sites in search. Retrieved February 5, 2014, from Google Official Blog Web site: http://googleblog.blogspot.com/2011/02/finding-more-high-quality-sites-in.html.

Google, Inc. (2013). Form 10-K for fiscal year ended December 31, 2012. Retrieved March 6, 2014, from U.S. Securities and Exchange Commission Web site: http://www.sec.gov/Archives/edgar/data/1288776/000119312513028362/d452134d10k.htm#toc.

Google, Inc. (n.d. a). Block or remove pages using a robots.txt file. Retrieved January 31, 2014, from https://support.google.com/webmasters/answer/156449?hl=en.

Google, Inc. (n.d. b). Google fun facts. Retrieved February 5, 2014, from Google Press Center Web site: http://web.archive.org/web/20090424093934/http://www.google.com/press/funfacts.html.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, *35*, 519–530.

Hayes, A. (n.d.). What exactly is a bot like Googlebot: Googlebot and other spiders. Retrieved January 31, 2014, from AFHX Mobile Design and Hosting Web site: http://www.ahfx.net/weblog/39.

Jefferson Parrish Hospital Dist. No. 2 v. Hyde, 466 U.S. 2 (1984).

Kaiser, T. (2013, January 4). FTC, Google officially settle antitrust investigation; Microsoft Cries Foul. Retrieved February 9, 2014, from DailyTech Web site: http://www.dailytech.com/FTC+Google+Officially+Settle+Antitrust+Investigation+Microsoft+Cries+Foul/article29551.htm.

Kendall, B., Efrati, A., Catan, T. and Ovide, S. (2013, January 5). Behind Google's antitrust escape. Retrieved February 9, 2014, from Wall Street Journal Online Web site: http://online.wsj.com/news/articles/SB10001424127887323689604578221971197494496.

Kramer, S. D. (2003, September 25). Google News creator watches a portal quiet critics with 'Best News' Webby. Retrieved June 3, 2013 from USC Annenberg Online Journalism Review Web site: http://www.ojr.org/ojr/kramer/1064449044.php.

Salinger, M. A., & Levinson, R. J. (2013, June). The role for economic analysis in the FTC's Google investigation. Retreived February 12, 2014 from Northwestern University School of Law Searle Center on Law, Regulation and Economic Growth Web site: http://www.law.northwestern.edu/research-faculty/searlecenter/events/internet/documents/Salinger_Economics_of_Google_and_Antitrust_Case_Searle_conference_version.pdf.

Salop, S. C. (2000). The first principles approach to antitrust, Kodak, and antitrust at the Millennium. *Antitrust Law Journal*, *68*, 187–202.

Schumpeter, J. A. (1950). *Capitalism, socialism, and democracy* (3rd ed.). New York: Harper.

Search engine optimization. (n.d.). Retrieved March 6, 2014, from Wikipedia Web site: http://en.wikipedia.org/wiki/Search_engine_optimization.

Shelanski, H. A., Farrell, J., Hanner, D., Metcalf, C. J., Sullivan, M. W., & Wendling, B. W. (2012). Economics at the FTC: Drug and PBM mergers and drip pricing. *Review of Industrial Organization*, *41*, 303–319.

Sherman, C., & Price, G. (2003). The Invisible Web: Uncovering sources search engines can't see. *Library Trends*, *52*, 282–298.

Singhal, A. (2001). Modern information retrieval: A brief overview. *IEEE Data Engineering Bulletin*, *24*, 35–43.

Stigler, G. J. (1971). The theory of economic regulation. *Bell Journal of Economics and Management Science*, *2*, 3–21.

Sullivan, D. (2001, March 4). Being search boxed to death. Retrieved June 3, 2013, from Search Engine Watch Web site: http://searchenginewatch.com/article/2065235/Being-Search-Boxed-To-Death.

Sullivan, D. (2003, May 3). Where are they now? Search engines we've known & loved. Retrieved June 3, 2013, from Search Engine Watch Web site: http://searchenginewatch.com/article/2064954/Where-Are-They-Now-Search-Engines-Weve-Known-Loved.

Sullivan, D. (2011, May 2). Google & the death of Osama bin Laden. Retrieved May 15, 2013, from Search Engine Land Web site: http://searchengineland.com/google-the-death-of-osama-bin-laden-75346.

U.S. Department of Justice and Federal Trade Commission. (2010, August 19). Horizontal merger guidelines. Retrieved February 9, 2014, from U.S. Department of Justice Web site: http://www.justice.gov/atr/public/guidelines/hmg-2010.html#4.

United States v. E. I. du Pont de Nemours & Co., 351 U.S. 377 (1956).

Werden, G. J. (1992). The history of antitrust market delineation. *Marquette Law Review*, *76*, 123–215.

Web Crawler. (n.d.). Retrieved February 4, 2014, from Wikipedia Web site: http://en.wikipedia.org/wiki/Web_crawler.

White, Lawrence J. (2001). Present at the beginning of a new era of antitrust: Reflections on 1982–1983. *Review of Industrial Organization*, *16*, 131–149.

White, Lawrence J. (2008). Market power and market definition in monopolization cases: a paradigm is missing. In W. D. Collins (Ed.), *Issues in competition law and policy* (pp. 913–924). Chicago: American Bar Association.

Yahoo! Inc. (1997, March 25). Yahoo! ranks no. 1 in news. Retreived June 3, 2013, from Yahoo! Investor Relations Web page: https://investor.yahoo.net/releasedetail.cfm?ReleaseID=173352.