Is Google degrading search? Consumer Harm from Universal Search

Michael Luca, Tim Wu, and the Yelp Data Science Team[1]

Abstract

While Google is known primarily as a search engine, it has increasingly developed and promoted its own content as an alternative to results from other websites. By prominently displaying Google content in response to search queries, Google is able to leverage its dominance in search to gain customers for this content. This yields serious concerns if the internal content is inferior to organic search results. To investigate, we implement a randomized controlled trial in which we vary the search results that users are shown - comparing Google's current policy of favorable treatment of Google content to results in which external content is displayed. We find that users are 45% more likely to engage with universal search results (i.e. prominently displayed map results on Google) when the results are organically determined. This suggests that by leveraging dominance in search to promote its internal content, Google is reducing social welfare - leaving consumers with lower quality results and worse matches.

Introduction

"Universal search" refers to a search engine feature through which the search engine operator blends specialized search properties - often from

[1] Harvard Business School and Columbia Law School, respectively. The Yelp data science team consists of Travis Brooks, Sebastien Couvidat, Daniel Frank, and Will Seltzer. Luther Lowe on the public policy team also contributed. The work for this paper was financially supported by Yelp.

proprietary databases - in priority over the results generated by an organic, or algorithmic search. Pioneered and widely used by Google, universal search, by its nature, is a form of self-dealing that excludes competitors in specialized search. As such, over the last decade, the practice has attracted scrutiny from competition law enforcers in the United States, Canada, Europe, and Latin America.

In its defense of universal search, Google has relied on a simple and powerful argument: it has done nothing more than create a better product. "Your search just gets better and better over time" said Eric Schmidt recently, "Google 'Berlin weather' and you'll no longer get ten blue links … you'll get the weather forecast for the next few days at the top result, saving you time and effort."[2] When charged that its "universal search" unfairly excludes competitors, Google argues, in legal terms, that universal search is a pro-competitive innovation that serves the interests of users.[3]

Through a series of online experiments, this paper reveals a more complex reality and sharply questions the assertion that Google's implementation of universal search has been uniformly good for consumers or pro-competitive. Our findings suggest that Google is - in some instances - actually making its

---

[2] Eric Schmidt, *The New Gründergeist*, Google Europe Blog (Oct 13, 2014, 10:28 AM), http://googlepolicyeurope.blogspot.com/2014/10/the-new-grundergeist.html

[3] Communication From the Commission: Guidance on the Commissions' Enforcement Priorities in Applying Article 82 of the EC treaty to Abusive Exclusionary Conduct by Dominant Undertakings, 2009 O.J. (C 45) 7, 12.

overall product worse for users in order to provide favorable treatment to Google content.

On the one hand, there are instances – such as displaying time, or presenting a calculator to arithmetic problems where universal search has been deployed in a manner that likely benefits consumers. After all, to the query 2+2 there is no better or worse answer, but simply the right answer or wrong answer - and the answer can be easier or harder to find. Therefore, a calculator presented via universal search, as we demonstrate below, is preferred by users.

In contrast, consider the question "which doctor should I see for this condition?" Clearly, there is a wide variety of content that one might provide to facilitate a user's search for a doctor. Competing rankings, ratings, and reviews ranging from US News & World Report to ZocDoc to HealthGrades invest heavily in developing such content. In this situation, Google's content may be more or less useful to users than other content. If Google provides favorable treatment to Google content in a world in which ZocDoc is more useful to users, this comes with a real cost to consumers. Our results provide empirical evidence that universal search has in some situations been deployed in a manner that degrades the search product and harms Google's users. Such conduct therefore cannot be described as pro-competitive.

"Local intent"-based searches comprise the largest single category of search, representing roughly one-third[4] of total desktop search volume, and over one-half of mobile search.[5] Today, Google presents users with local search results that are a mixture of its organic results along with a user interface object known as the "Local OneBox."[6] The OneBox typically includes a list of seven business pins populated by exclusively querying Google's proprietary local product, Google+ Local; this set of seven business pins is attached to a Google Map.

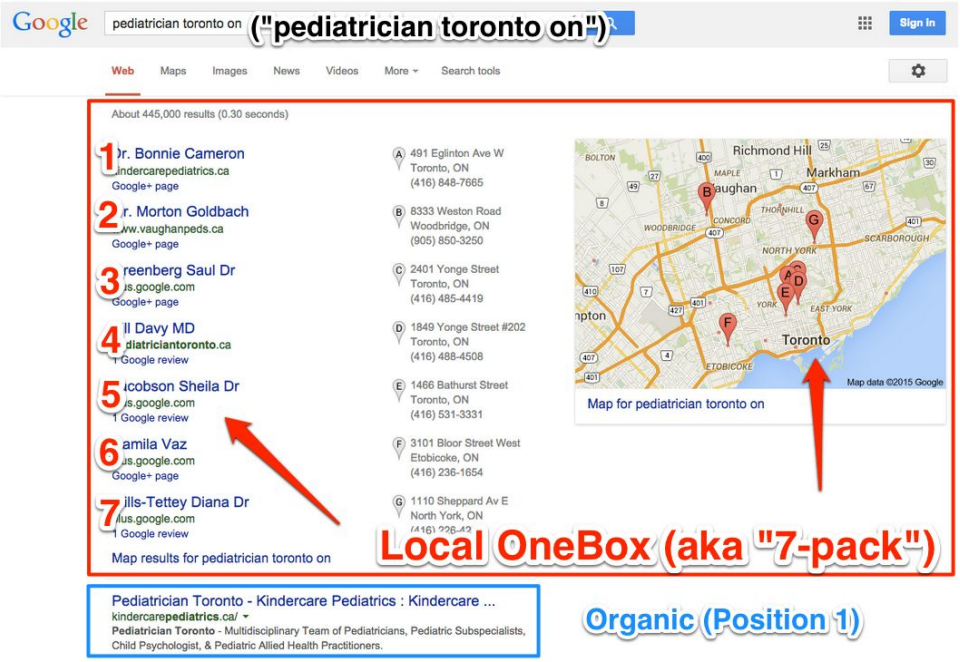

---

[4] Mike Blumenthal, *Ed Parsons: 1 in 3 Searches at Google are Local*, Understanding Google My Business & Local Search (Nov. 14, 2012, 9:26 AM), http://blumenthals.com/blog/2012/11/13/ed-parsons-1-in-3-searches-at-google-are-local/

[5] Greg Sterling, *Research: 56 Pct. of "On the Go" Searches Have Local Intent*, Search Engine Land (May 6, 2014, 2:07 PM), http://searchengineland.com/research-56-percent-go-searches-carry-local-intent-190655

[6] Danny Sullivan, *Meet the Google OnBox, Plus Box, Direct Answers & the 10-Pack*, Search Engine Land (Sept. .28, 2009, 6:12 PM), http://searchengineland.com/meet-the-google-OneBox-plus-box-direct-answers-the-10-pack-26706

However, Google's organic search employs a merit-based algorithm that can easily be used to identify better candidates to populate its local search boxes, enabling the creation of an alternative version of the search engine results page. Based on a sample of 2,690 Internet users, this study demonstrates that users would be more likely to engage with local specialized search results if Google were to replace its proprietary "answers" in universal search with results drawn from the whole web based on the same merit-based algorithm that it uses to populate organic search (as opposed to being exclusively drawn from Google+). The study is conducted by randomly displaying one of two sets of search result screenshots to thousands of internet users: one set of users sees the Google universal search page as it is currently constituted; the second set sees an alternative version of universal search. In the alternative version, a browser plug-in nicknamed "Focus on the User - Local" (FOTUL)[7] has queried third-party review sites and ranked them, using Google's own organic algorithm, according to which site delivers the best or most relevant information for the content in question.

[7] See *Focus On The User - Local,* Github, https://github.com/FocusOnTheUserLocal/FocusOnTheUserLocal. (last visited June 23, 2015) (explaining the Focus On The User - Local as a widget formed in collaboration by local search companies to educate about consumer harm brought about when Google+ is artificially promoted in local search). *See also* Focus On The User, http://focusontheuser.eu/ (last visited June 23, 2015).

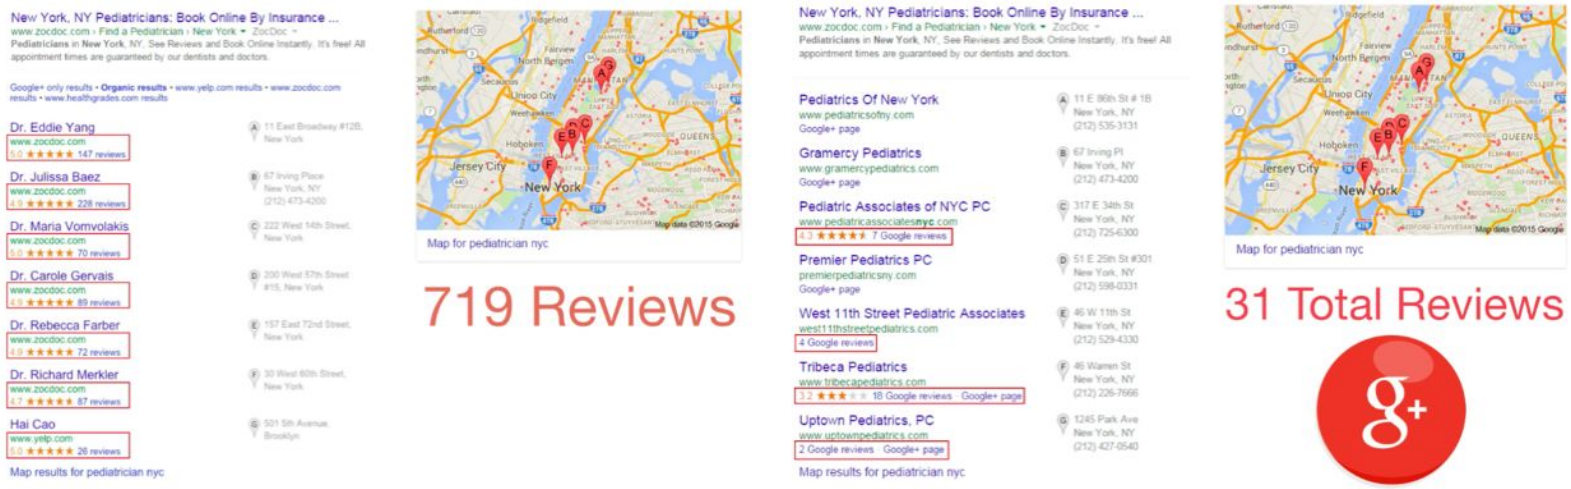

"pediatrician nyc": Focus on the User (FOTUL) results versus Google+

The results demonstrate that consumers vastly prefer the second version of universal search. Stated differently, consumers prefer, in effective, competitive results, as scored by Google's own search engine, than results chosen by Google. This leads to the conclusion that Google is degrading its own search results by excluding its competitors at the expense of its users. The fact that Google's own algorithm would provide better results suggests that Google is making a strategic choice to display their own content, rather than choosing results that consumers would prefer.

In economic terms, we can understand Google's utility as a reduction of search frictions by helping to create a match between users and their preferred goods or services. In a two-sided market analysis (such as Rochet and Tirole), Google is the intermediary or platform that puts together buyers with sellers of desired goods or services. By providing suboptimal results, Google's conduct

eliminates some of the aforementioned welfare gains, hence undermining the value of search. As such, when Google knowingly degrades its search so as to harm its competitors, the impact can be felt by both sides of the market – by the consumers who don't get the results (and products) they value most highly, and the merchants who might otherwise have sold to them.

The demonstration of consumer harm is, we think, an important conclusion standing on its own that should influence any competition law analysis. However, it intersects with several widely-recognized criteria for enforcement action in competition law. First, whatever the general utility of universal search, we have shown that, as implemented in some segments, universal search is harmful both to merchants, consumers and competitors while lacking redeeming qualities. As such, in some implementations it may be categorized as a species of "naked exclusion" – in other words, conduct that excludes competitors without any countervailing benefit.[8]

Alternatively, Google's conduct can be understood as the knowing neglect of a "less restrictive alternative" for achieving legitimate goals.[9] Google's development of universal search, in general, can be accepted as an important innovation that can improve consumer welfare. But it seizes on the fact that, as implemented, Google appears to have chosen to do so in a way that neglects an obvious and clearly more effective alternative, resulting in harm to consumers,

---

[8] See Eric B. Rasmusen, J. Mark Ramseyer, & John S. Wiley, Jr., *Naked Exclusion*, 81 Am. Econ. Rev. 1137 (1991) (describing exclusion as "naked" when the conduct unabashedly excludes rivals without offering any efficiency justification).

[9] See Scott Hemphill, "Less Restrictive Alternatives in Antitrust & Constitutional Law" working paper (2015).

merchants, and its competitors. Important to this conclusion is evidence that Google is sacrificing a higher quality and potentially more profitable product in favor of a more exclusionary option. That fortifies the intuition that the conduct is suspect.

From this paper one thing should be abundantly clear. The easy and widely disseminated argument that Google's universal search always serves users and merchants is demonstrably false. Instead, in the largest category of search (local intent-based), Google appears to be strategically deploying universal search in a way that degrades the product so as to slow and exclude challengers to its dominant search paradigm.

I. Industry Background and the Development of Universal Search

Since its inception, the World Wide Web, now representing more than 4.7 billion linked web pages, has depended for its utility on tools for locating desired information. The tools for finding information on the web have always had an important influence on nearly all aspects of economic activity and innovation on the web and its connection to "Main Street" small businesses. The ability of buyers and sellers to connect can do much to determine which businesses succeed and which fail, what innovations catch on, and which flounder. They also have a preeminent role in influencing the speech environment centered on the web.

The history of tools for finding information on the web is one of technological evolution. In the early days of the world wide web, simple lists of links and directory services, (like the early Yahoo!) were used, in competition with early, primitive search engines. Eventually, in the early 2000s, the "general" search engine (which searches the entire web)[10] became the tool of choice of which Google is the preeminent example. General search, today, remains a dominant tool by which web sites and their users find each other, and buyers and seller are matched.

The evolution of the web's information retrieval tools did not end with general search. Over the 2010s, just as Google's general search replaced

---

[10] More precisely, everything that makes itself available or crawlable with [a standard internet DNS address and uses the HTML protocol] and does not specifically tell the search engine not to make its content searchable.

directory services, its general search began to face its own challengers. The most important challenge has come from *specialized* search; that is, search engines that deliver information based not on searching the entire web, but rather a specific category of information. Prominent examples include the search for books on Amazon, plane tickets on Kayak or Orbitz, hotel reviews on TripAdvisor, or for a local doctor on Yelp or ZocDoc. As it stands, specialized search is not in direct competition to Google; instead, it represents a threat to the general search paradigm (in the sense that Netscape, in combination with other middleware were a paradigm threat to the Windows OS paradigm).[11] Eric Schmidt captured this reality when he recently suggested "our biggest search competitor is Amazon."[12]

Google has, over the decade, aggressively fought the threat from specialized search in the following ways. First, beginning in the 2010s, it made copies of the most successful of the specialized search engines. Many of these are simple clones of competitors, like Kayak or Yelp (yielding, for example, "Google travel" "Google shopping," "Google+ Local," etc.). To improve their popularity, the earlier versions of Google's clones borrowed data from the originals; in particular, early versions of Google local took the reviews collected by Yelp and TripAdvisor and incorporated them into Google's product.

---

[11] The two paradigms are not direct competitors in the market sense, but rather in the Schumpeterian sense (in the sense that Netscape & middleware were not a direct threat to Microsoft Windows, but as a platform). *See* United States v. Microsoft Corp., 253 F.3d 34 (D.C. Cir. 2001).

[12] Schmidt, *The New Gründergeist,* supra note 1.

Overall, with notable exceptions, like Google Maps, the clones have not proved as popular or successful with users as the originals (for example, Google's clones typically suffer from poor search rankings in Google's organic search results). To deal with this challenge Google has leveraged the power of its dominant, general search engine to give advantages to its otherwise unpopular specialized search products. The primary means for doing so is what is called the "universal search" or the "OneBox." Universal search operates by incorporating the general ("ten blue links") and specialized ("OneBox") search paradigms into the same user interface on the search engine results page. Google deploys the OneBox when it detects keywords, or the fact that a natural search returns a specialized competitor as a leading result.

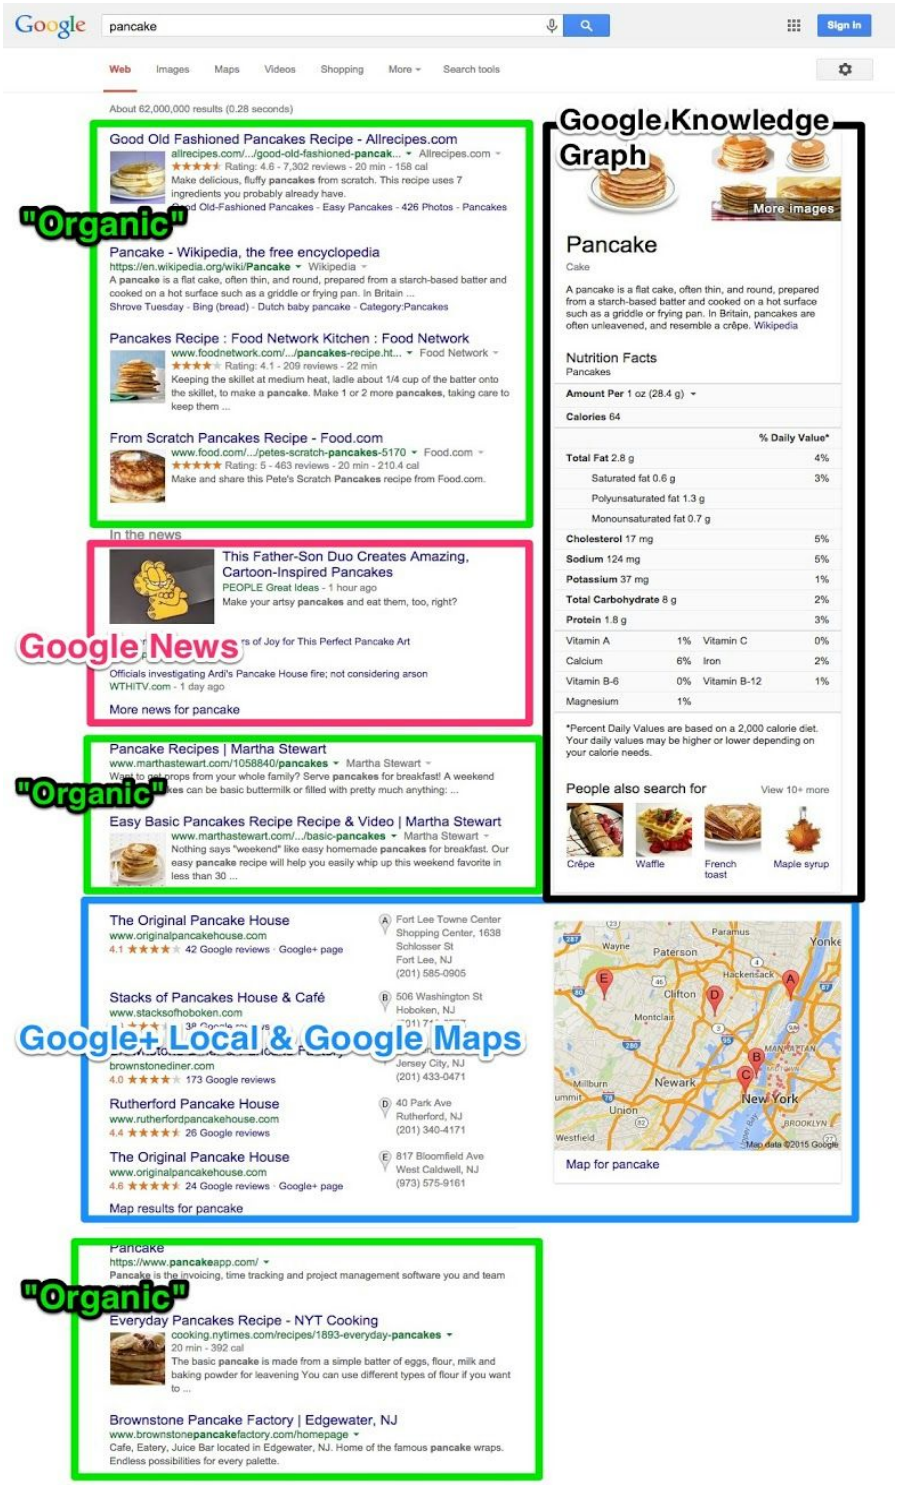

(An example of Google's "Universal search" for the query "pancake" shows a blend of organic results alongside other sources of data: Google+ Local, Google Maps, Knowledge Graph, Google News)

For example, a search for “restaurants San Francisco” might trigger Yelp as the first natural result, causing Google to deploy a Local OneBox.[13] The OneBox, with a few exceptions, is populated by Google’s own clones of its competitors’ specialized search services (e.g., Google+ Local or travel).

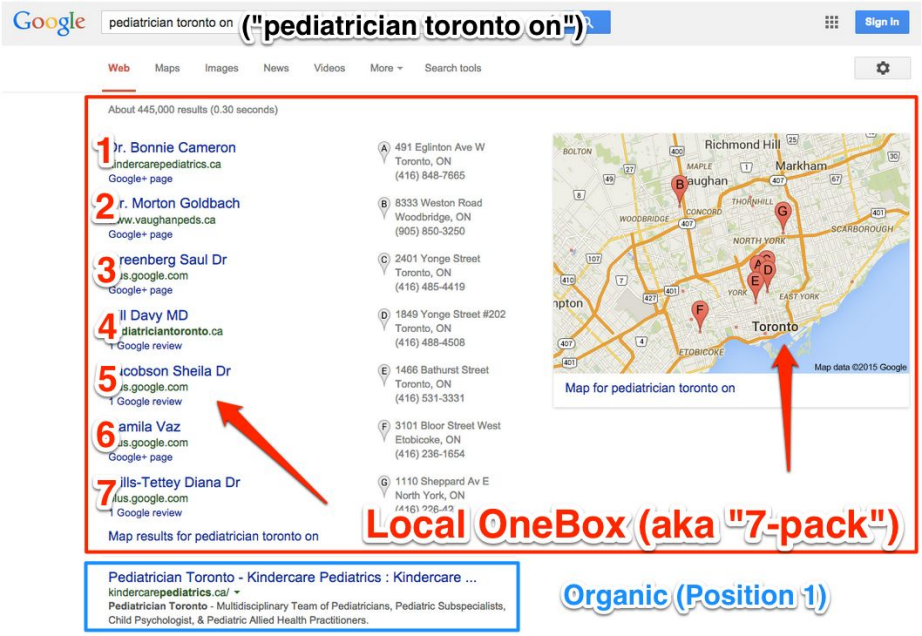

Universal search, in its operation is inherently exclusionary, for it uses the dominant general search engine to divert traffic from Google’s specialized competitors (Expedia, Yelp, etc.) to its own versions of those companies. That fact has, unsurprisingly, led to scrutiny both by Europe and American competition regulators.

Google’s strategy follows a well-known historic pattern. Technological monopolists, facing threats from innovative competitors, often engage in vertical

---

[13] Memorandum, Recommendation to the Federal Trade Commission on Google Antitrust Litigation 24, Page 70 (Aug. 8, 2012), available at http://graphics.wsj.com/google-ftc-report.

self-dealing to protect their monopoly. A prominent example was the defensive self-dealing by the American phone monopolist AT&T, which led to prolonged litigation over the 1970s and 80s, where AT&T was accused of excluding both handset and long-distance competitors, despite federal regulations requiring interconnection.[14] Another historical parallel can be seen in the US and European Microsoft litigation of the 1990s, where Microsoft was accused of using the power of its operating system monopoly to exclude various competitors in favor of its own clones of their products. The most prominent victim of that conduct was Netscape.[15]

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

---

[14] *See* AT&T, 38 F.C.C.2d 213 (1972); AT&T, 64 F.C.C.2d 1 (1977); Communications Satellite Corp., 56 F.C.C.2d 1101 (1975); United States v. AT&T, 461 F. Supp 1314 (D.D.C. 1978); United States v. AT&T Co., 524 F. Supp 1336 (D.D.C. 1981); United States v. AT&T Co., 552 F. Supp 131 (D.D.C. 1982), aff'd sub nom. Maryland v. United States, 460 U.S. 1001 (1983).

[15] *See* United States v. Microsoft Corp., 56 F.3d 1448 (D.C. Cir. 1995); United States v. Microsoft Corp., 147 F.3d 935 (D.C. Cir. 1998); United States v. Microsoft Corp., 253 F.3d 34 (D.C. Cir. 2001); Case COMP/C-3/37.792 Microsoft, 2004 O.J. (C 290).

engineer it any other way."[16] The object of this paper is to test the truth of that proposition.

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

II. Methodology and Summary of Results

A. A/B testing in competition analysis

The A/B testing methodology is a form of statistical testing that is widely used for web development. Pioneered in major part by Google, it also serves as a particularly useful tool for detecting consumer harm in competition analysis, as it creates a direct means for observing the impact of design decisions on consumer behavior. A/B testing presents a random sample of users with one of two variants of a website and then records the user's response to the presentation. The A/B test, also known as a "split" test, is typically used in commercial web development to determine the effectiveness of competing designs for some goal or objective, such as inducing users to click on advertising. But that same method can also serve competition analysis by allowing other interested parties to test the impact on consumers of different designs. As such it allows a direct and realistic test of impacts on consumers.

In a simple implementation to demonstrate the usage of an A/B test in the search context, we tested the effects of presenting users with two different search engine result pages in response to an arithmetic problem ("42 people have given you $54. You decide to use the search engine Google to determine how much total money you have been given. You type "calculator" into a Google search bar and come upon the following page. Where do you click first?") Variant A presented users with merely blue links, followed by "blue link" search results; Variant B presented an interactive calculator on the search engine result

page. When presented with the interactive calculator above the organic links, over 90% of the users clicked on it, suggesting a gain in consumer utility from the presentation of the calculator OneBox. This result, incidentally, is consistent with Google's claims about the efficiency of universal search, but, as we shall see, only tells part of the story.

Instructions: 42 people have given you $54. You decide to use the search engine Google to determine how much total money you have been given. You type "calculator" into a Google search bar and come upon the following page. Where do you click first?

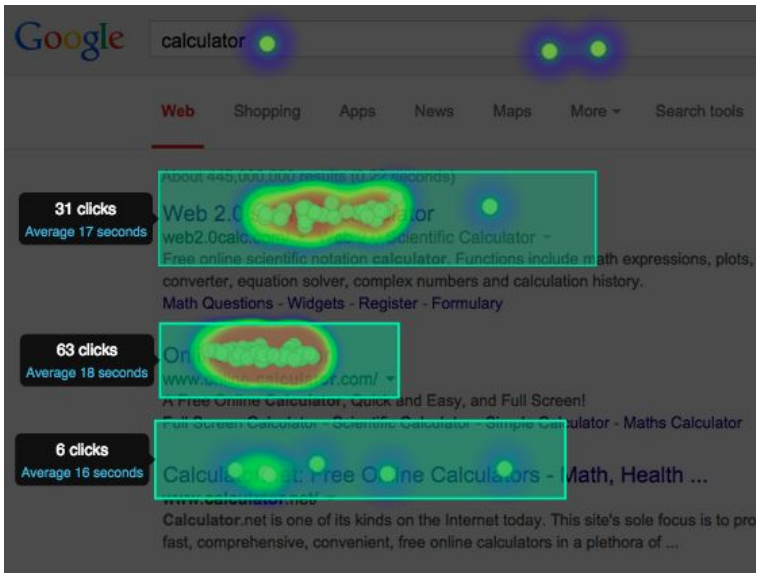

Position 1: **31**
Position 2: **63**
Position 3: **6**

"Answer Box" (Calculator): **96**
Positions 1-3: **8**

Turning to the present case, how might A/B testing measure the effects of a search engine design – here universal search – on consumer welfare? A search engine is a platform or intermediary whose utility comes from matching buyers and sellers at minimum costs, whether measured in time or other efforts. Universal search, therefore, serves the interests of consumers if it helps to facilitate matching between buyers and sellers at minimal cost.

The tests that follow center on universal search as implemented to answer questions that trigger a "local" specialized result, such as "nearby restaurant" or "pediatrician nyc." We first pose a question to users, and follow by randomly showing the user one of two sets of results.[17] The first, the control, is exclusively populated by Google's own products, and reflects the results currently created by Google. The variant, or alternative, populates the universal search results by relying on Google's organic algorithm to determine what might be the best sites for providing the information. This variant, herein called Focus on the User - Local or FOTUL, is described in the detailed methodology section below. Users are randomly presented with either the control or the variant.

We conduct our experiment on UsabilityHub,[18] an online platform that is specifically geared toward testing website designs before a website is introduced to a market; the survey respondents were sourced from MechanicalTurk. This

---

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

[18] See UsabilityHub, http://www.usabilityhub.com (last visited June 24, 2015).

makes an appealing context for this experiment for several reasons. First, it provides a larger subject pool than we would be able to obtain in a typical social science laboratory. Second, this is the exact platform that businesses use to test websites, suggesting the external validity of the results.

There is growing acknowledgment that online labor markets provide a compelling setting in which to investigate real world problems. For examples of field experiments in online labor markets, see Gilchrist, Luca, and Malhotra (forthcoming) and Pallais 2014. Moreover, Horton et al 2011 ran a series of classic economics experiments on MechanicalTurk (a similar, but less specialized, market) and find that the results are qualitatively similar to findings from a physical laboratory. This suggests that we might expect findings to generalize across settings.

Nonetheless, one might still be concerned that results from MechanicalTurk and UsabilityHub do not generalize to actual Google searches. While we do not have access to internal Google data, we do have access to internal Yelp data, which we use to cross-validate a series of experimental findings about Yelp from UsabilityHub. While there were differences between the click-surveys and actual search results, we find that click surveys are both directionally current and provide reasonable estimates of actual user engagement. This provides further evidence of the external validity of our experimental design.

B. Results

Based on click-surveys of 2,690 users, this study demonstrates that Internet users would be approximately 45% more likely to engage with local specialized universal results if the results are drawn from Google's merit-based algorithmic process used to populate organic results, compared with Google's current local search results.

The click-survey revealed that 32% of users would click on Google's current local 7-pack of local results, when prompted with a need that would trigger a local query, and shown current results. By contrast 47% of users clicked on the same 7-pack presentation, in the same scenarios, when the 7-pack was filled with the FOTUL results, using Google's merit based algorithm to rank results from 3rd party review sites.

A nearly 50% increase in CTR is immense in the modern web industry, and is surprising given the absence of any modification to the universal search user interface. We can become more confident in the results, and the user reaction can perhaps be understood, by examining the alternatives that were presented to the users. The FOTUL alternative, since it brings in content from review ecosystems larger than Google's own, fills the local 7-packs with results that have more review content that the Google local results. The presentation of the results makes the review count visible as a clear indicator of the quality and depth of content.

New York, NY Pediatricians: Book Online By Insurance ...
www.zocdoc.com › Find a Pediatrician › New York ▾ ZocDoc ▾
Pediatricians in **New York**, NY, See Reviews and Book Online Instantly. It's free! All appointment times are guaranteed by our dentists and doctors.

Pediatrics Of New York www.pediatricsofny.com Google+ page	(A) 11 E 86th St # 1B New York, NY (212) 535-3131
Gramercy Pediatrics www.gramercypediatrics.com Google+ page	(B) 67 Irving Pl New York, NY (212) 473-4200
Pediatric Associates of NYC PC www.pediatricassociates**nyc**.com 4.3 ★★★★½ 7 Google reviews	(C) 317 E 34th St New York, NY (212) 725-6300
Premier Pediatrics PC premierpediatricsny.com Google+ page	(D) 51 E 25th St #301 New York, NY (212) 598-0331
West 11th Street Pediatric Associates west11thstreetpediatrics.com 4 Google reviews	(E) 46 W 11th St New York, NY (212) 529-4330
Tribeca Pediatrics www.tribecapediatrics.com 3.2 ★★★☆☆ 18 Google reviews · Google+ page	(F) 46 Warren St New York, NY (212) 226-7666
Uptown Pediatrics, PC www.uptownpediatrics.com 2 Google reviews · Google+ page	(G) 1245 Park Ave New York, NY (212) 427-0540

Map results for pediatrician nyc

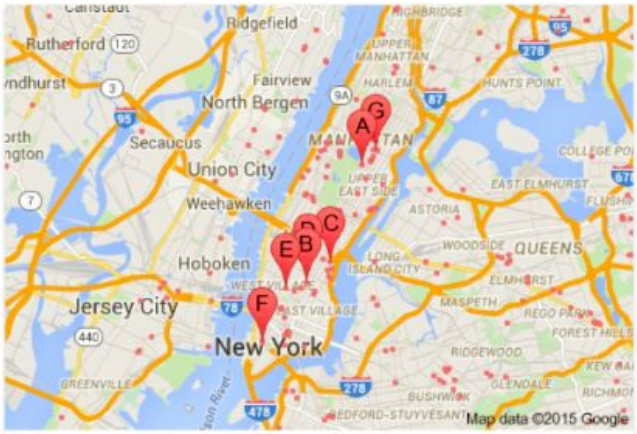

Map for pediatrician nyc

31 Total Reviews

The query for "pediatrician nyc" yields 31 reviews from Google+.

New York, NY Pediatricians: Book Online By Insurance ...
www.zocdoc.com › Find a Pediatrician › New York ▾ ZocDoc ▾
Pediatricians in New York, NY, See Reviews and Book Online Instantly. It's free! All appointment times are guaranteed by our dentists and doctors.

Google+ only results • **Organic results** • www.yelp.com results • www.zocdoc.com results • www.healthgrades.com results

Dr. Eddie Yang
www.zocdoc.com
5.0 ★★★★★ 147 reviews

Dr. Julissa Baez
www.zocdoc.com
4.9 ★★★★★ 228 reviews

Dr. Maria Vomvolakis
www.zocdoc.com
5.0 ★★★★★ 70 reviews

Dr. Carole Gervais
www.zocdoc.com
4.9 ★★★★★ 89 reviews

Dr. Rebecca Farber
www.zocdoc.com
4.9 ★★★★★ 72 reviews

Dr. Richard Merkler
www.zocdoc.com
4.7 ★★★★★ 87 reviews

Hai Cao
www.yelp.com
5.0 ★★★★★ 26 reviews

Map results for pediatrician nyc

(A) 11 East Broadway #12B, New York
(B) 67 Irving Place New York, NY (212) 473-4200
(C) 222 West 14th Street, New York
(D) 200 West 57th Street #15, New York
(E) 157 East 72nd Street, New York
(F) 30 West 60th Street, New York
(G) 501 5th Avenue, Brooklyn

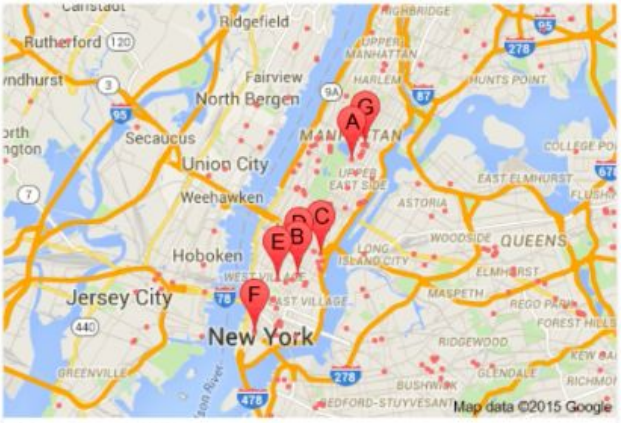

Map for pediatrician nyc

719 Reviews

The same query for "pediatrician nyc" yields 719 reviews when the Focus on the User software is activated.

Users, unsurprisingly, are attracted to this content when attempting to carry out the task set out before them, such as finding a hotel or pediatrician. These results strongly suggest that Google's present implementation of universal search is making consumers worse off than the obvious alternative presented here. Users are engaging at a lower rate with the universal results, which suggests they would prefer the alternative version presented. Stated simply,

when it comes to local search, Google is presenting its users with a degraded version of its search engine.

The suggestion that Google purposely and consciously deployed universal search to maximize damage to its competitors, even at the expense of consumers, is moreover backed by evidence compiled by the Federal Trade Commission Bureau of Competition over the course of its investigation. As the Commission concluded, "Evidence shows that Google sought to increase such "triggering" of Universal Search results not only to provide users with the "right" answer to their queries, but also to drive traffic to Google properties."[19]

We now consider in more detail some of the theories of consumer harm that are supported by these results.

III. Theories of Harm

Beginning in the 2000s, economists have used the phrase "two-sided" markets to describe those environments where buyers and sellers are brought together by a company acting as an intermediary, or platform. As described by Jean-Charles Rochet and Jean Tirole, in such a two-sided market, the economic utility of the platform is a byproduct of the matching of the buyers and sellers on both sides of the market.[20] For example, the utility of a bazaar lies in matching buyers and sellers; a dating application like Tinder creates utility by successfully

---

[19] Memorandum, Recommendation to the Federal Trade Commission on Google Antitrust Litigation 24 (Aug. 8, 2012), available at http://graphics.wsj.com/google-ftc-report.

[20] Jean-Charles Rochet & Jean Tirole, Platform Competition in Two-Sided Markets, J. Eur. Econ. Ass'n 990 (2003).

matching couples; the utility of a travel platform is matching people with flights, and so on.

This analysis provides the first part of a helpful and simple way of understanding the utility of Google search. Web searching comes close to being a paradigmatic example of the two-sided market intermediary. On the one side of the market we have consumers with preexisting preferences for goods and services (European soccer games, vinyasa yoga, New York-style bagels and so on.) On the other side are merchants and website operators who seek consumers (to simplify the analysis, we ignore advertisers, or assume that they are also among those seeking a match). Search creates economic welfare by *matching* consumers with the objects of their preferences. In the language of two-sided markets, the search engine serves as the "platform" or "intermediary," who matches buyers with sellers of desired goods or services.[21] Search, stated more simply, generates utility when someone finds a yoga studio they decide to join, locates where to buy tickets to a soccer game, identifies a good pediatrician for their sick child, and so on. In this manner, by matching buyers and sellers, search directly generates economic welfare. The flipside of understanding the utility of search is that a purposefully degraded search product can therefore do economic harm, particularly if it is widely relied upon. For completeness, we now consider five theories by which degrading of search creates harm to consumers and merchants.

21

Welfare Loss. We have described the utility of a search engine in terms of matching buyers and sellers and the reduction of search costs in that process. Ellison and Ellison[22] (2009) demonstrate that firms can engage in obfuscation to make it harder for customers to acquire information in an attempt to maintain market power. We show that even the platform itself may have the incentive to increase search frictions for their own benefit. Because the resulting quality and volume of matches is worse when search frictions are increased, there is a drop in total welfare. This welfare loss is then split between consumers and merchants.

There are, to be more precise, several species of harm caused. First, some consumers may simply not find what they are looking for in the time they have, and will give up, yielding some number of unconsummated transactions. Second, some consumers will be, in fact, determined enough to eventually find their desired target, but simply suffer greater search costs in the process. A third kind of harm arises when a buyer ends up patronizing a business or other service provider who would *not* have been their first choice, but for the degrading of the search. Consider, for example, a consumer who is misdirected and ends up at a bad restaurant; or the parents who are looking for a top-notch pediatrician, but because of search degradation, patronize a subpar practitioner. The harm caused by such misdirection when it occurs, will vary, but is undeniable in the aggregate. The point is simply that a degraded search engine

[22] *Search, Obfuscation, and Price Elasticities on the Internet*, Econometrica, 2009, Vol. 77, No. 2

will invariably, as compared to its alternative, yield some consumer harm from misdirection.

Search-Advertising Monopoly Maintenance. Our results suggest that Google has chosen a path that excludes its specialized search competitors at the expense of its users. The result is to weaken nascent competitors to the general search paradigm. While it is beyond the scope of this paper to prove that Google's dominance of search advertising allows it to charge supra-competitive prices for search-advertising, assuming it does, the exclusion of competitors likely sustains those elevated prices. Such conduct might be defensible if, as Google claims, its exclusionary conduct was justified by procompetitive benefits for consumers. However, as this paper has shown, Google's implementation is in some areas actually harmful to consumers. As such, the maintenance of inflated prices constitutes a form of economic harm that the competition laws were designed to remedy.

Innovation Harms. The pattern by which Google uses general search to exclude and harm competitors in specialized search poses several threats to innovation, of which two may be highlighted.

It is important to note that most of the successful, pioneer specialized search operators earned their success through investment and innovation. Yelp made deep investments in the successful cultivation of offline communities which encouraged prolific creation of detailed reviews, and also developed important technologies to weed out various forms of fraud and spam. Kayak developed

new ways of presenting flight information and incorporated prediction algorithms that suggest the best time to buy a ticket. TripAdvisor succeeded where others had failed in achieving a critical mass of reviews of hotels and attractions for much of the globe.

Serious harm to future innovation is caused by reducing the incentives of existing and potential verticals to invest in the innovative and disruptive technologies of specialized search. If it is understood by entrepreneurs and innovators that any firm that relies on a specialized search will face an effort by Google to clone its product and be leveraged by the power of a dominant search engine, the shadow cast by Google's search engine becomes long indeed. Just as Microsoft was able to dampen innovation in software that it might incorporate into Windows, so too Google can use the terror of linking specialized searches to general search to discourage investment in specialized search products that might ultimately challenge its dominance.

Second, Google's conduct may create harm by slowing the evolution of search technology. As we have seen, the history of the web has witnessed market-altering improvements in information location technology arising each decade or so (The rise of Yahoo!, Google, etc). To the degree that universal search delays or preserves the general search paradigm over the rise of specialized search alternatives, it may be understood as the potential slowing or blockage of Schumpeterian "competition for the market" that has been a trademark of the high-tech and internet industries for much of the last 40 years.

The tendency may be particularly important in the mobile environment, where general search has been less-well established and presumably weaker than specialized search.

To be sure, the very introduction of universal search, as Google has highlighted, is itself a form of innovation. However, everything, once again, depends on implementation. To the extent that universal search is implemented in a manner that benefits buyers and sellers its value cannot be contested. To the extent it is deployed to damage competitors at the expense of consumers, it represents no real innovation at all.

Speech and Self-Expression. The World Wide Web has been celebrated over the last decade for its widespread promulgation of speech and other forms of self-expression captured in the phrase "user-created content." That phrase reflects in a multitude of forms ranging from blogs, user-created videos, reviews of films, products, or restaurants, and so on. More generally, Google operates in the information industries, where "the cost of monopoly must not be measured in dollars alone, but also in its effect on the economy of ideas and image, the restraint of which can amount to censorship."[23]

Search engines are widely understood as key mediators of the web's speech environment, given that they have a powerful impact on who gets heard, what speech is neglected, and what information generally is reached. It is telling that in censorial regimes, search engines are invariably the targets of strict

---

[23] Tim Wu, The Master Switch: The Rise and Fall of Information Empires 69 (2011).

government controls.[24] As professor Jeffrey Rosen has written, Google has a particularly central role in this process.[25] Search engines like Google have substantial influence to the point that "[u]nderstanding free speech in America has become a matter of understanding the behavior of intermediaries, whether motivated by their own scruples, law, or public pressure."[26]

The decisions made about search, and in particular, the decision to self-deal at the expense of other entities on the web, have implications both for web "speakers" and also listeners, or users. The more that Google directs users to its own content and its own properties, the more that speakers who write reviews, blogs and other materials become invisible to their desired audiences. Similarly, those users who might want to sample a broader, more diverse range of opinions will, unless they undertake more efforts, find themselves with a more constrained range of views. This fits with the general implications of vertical integration in an information industry – that self-dealing tends to yield a more centralized and narrow availability of views as compared to the broader and decentralized presentation of viewpoints that has characterized the web since its infancy.

---

[24] *See* Jack Goldsmith & Tim Wu, Who Controls the Internet? (2006)

[25] *See* Jeffrey Rosen, *The Deciders: The Future of Privacy and Free Speech in the Age of Facebook and Google*, 80 Fordham L. Rev. 1525 (2012).

[26] Tim Wu, Is Filtering Censorship? The Second Free Speech Tradition, Governance Studies at The Brookings Institute (Dec. 27, 2010), available at http://www.brookings.edu/~/media/research/files/papers/2010/12/27-censorship-wu/1227_censorship_wu.pdf; *see also* Jeffrey Rosen, *Google's Gatekeepers*, N.Y. TIMES MAG. (Nov. 28, 2008), available at http://www.nytimes.com/2008/11/30/magazine/30google-t.html).

IV. Detailed Methodology

The testing procedure described herein consisted of three separate efforts. Firstly, the FOTUL software, which generated an alternative set of local search results, using Google's own algorithm. Secondly, a click-survey to test user reaction to FOTUL, compared to user reaction to Google's local results. Finally, an attempt to understand the degree to which the click-survey results are indicative of real-world user behavior.

A. The FOTUL Software

Software was developed in the form of a Chrome Browser Extension called Focus on the User – Local (referred to herein as FOTUL) designed to detect Local OneBoxes. Upon detecting the presence of a Local OneBox in Google search results, FOTUL conducted a separate search in the background including major third party local review websites separated by "OR" statements (e.g. "site:yelp.com OR site:zocdoc.com OR site:tripadvisor.com OR …") The first ten pages of results are analyzed for links to business pages appearing on popular third party review websites such as Yelp and TripAdvisor. Google's universe of local review content on the subdomain plus.google.com is also included.

FOTUL extracted and ranked these results according to a combination of:

(a) Google's organic ranking

(b) average star rating

(c) quantity of reviews

The extracted, ordered lists then replaced the default contents of the Local OneBox, with a new set of results that Google's algorithm selected as relevant, and that had high review content. These results could come from local review websites like Yelp, TripAdvisor and ZocDoc or from Google's own ecosystem of local reviews on plus.google.com.

In this way FOTUL was able to present the user with a live comparison of the Google's existing local results and an alternative method of presenting local results, based on the algorithm that powers Google organic search. Because Google search result pages, local content, and domains change often, this methodology may not work in all locales or at all times, however at the time of this study, this work did produce a comparison of the content as described.

B. Click-survey Testing of FOTUL

A total of five distinct varieties of Local OneBoxes were identified. While FOTUL was disabled, screenshots were taken of the search results to establish a control group. FOTUL was then enabled and used to generate experimental variations within each of the five varieties of Google Local OneBoxes.

Click-surveys were presented to a sample of 2,690 Internet users, sourced using MechanicalTurk. Respondents were prompted with a hypothetical question based on contextual information derived from each illustration and asked to

answer by clicking on the illustration.

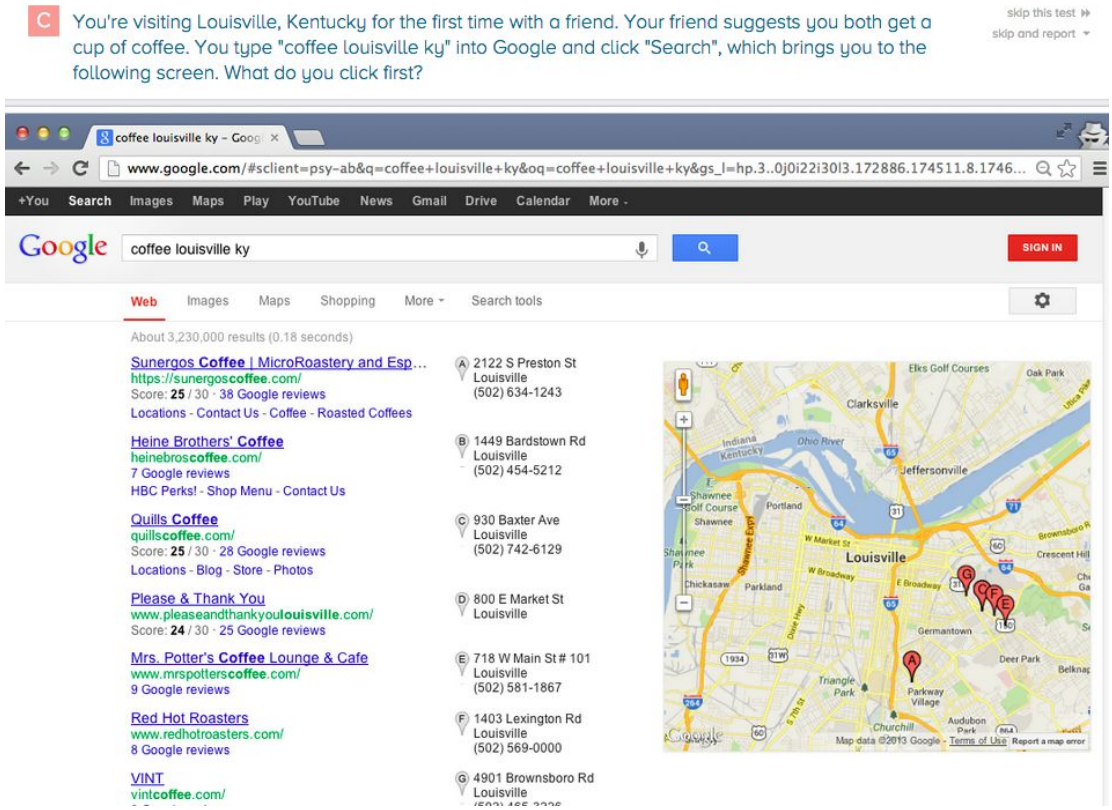

The relative performance, as measured by the clickthrough rate (CTR) of each Local OneBox was analyzed and compared.

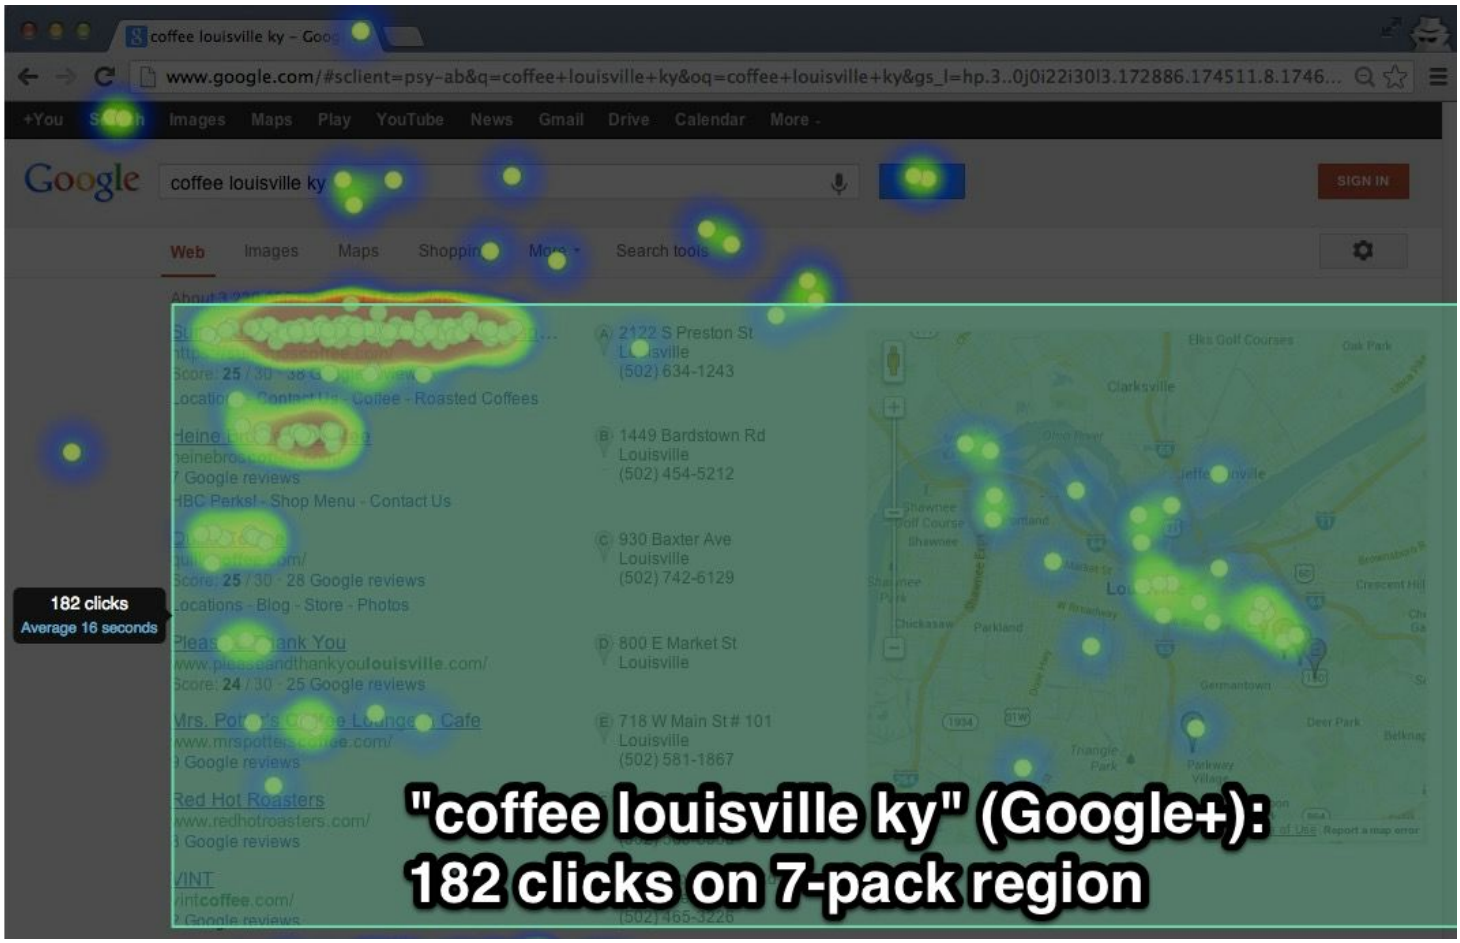

An example from the series of tests run shows 182 clicks on 7-pack region

powered by Google+ (ie status quo Google).

An example from the series of tests run shows 241 clicks on 7-pack region variant powered by FOTUL.

The following calculations were used to determine relative performance of the Experimental Group against the Control Group

$$\frac{OneBoxClick}{(TotalClicks - OtherClicks)} = PerformanceScore$$

$$\frac{PerformanceScoreExperimental}{PerformanceScoreControl} = PerformanceScoreRatio$$

	CTR on OneBox + Map	N
Google+ Local	32%	1,121
FOTUL powered	47%	1,442

Comparing the experimental groups gives a Performance Ratio of 1.44 or a 44% increase in engagement when users were shown FOTUL powered results. In order to determine the statistical validity of observed values of the Performance Score Ratio, 95% confidence bounds were computed using the methodology described in [http://www.jstor.org/stable/2531405] producing a bound of [1.303, 1.594].

C. Click-Text Survey Validation

Click-test surveys on sites such as UsabilityHub are a commonly used tool in the industry, and it is widely understood that such tools provide useful

guidance, but are not as accurate a predictor of user behavior as direct A/B testing on a live platform with the real user community in question.

Factors such as intent of the user in the moment, demographics of the studied population, and the context of the live site all create differences between the results of usability testing sites and live A/B tests. Of course, in this particular case, only Google is in a position to do a real, live A/B test on its users, so a relevant question becomes “to what extent are these test results reflective of reality?”

We gain confidence in the validity of using a click-survey to draw conclusions about actual internet behavior by examining A/B tests run on Yelp internet traffic, and comparing these with click-surveys designed to mimic similar experimental conditions. This comparison shows that changes in behavior in a click-survey do indicate the directionality of changes in actual internet behavior.

Specifically, Yelp internet traffic was exposed to a control condition involving a hidden set of search filters versus an experimental condition involving an exposed set of search filters (see image below). A click-survey was also created which reproduced these conditions and was administered to a set of

7,062 users in the same manner as the FOTUL study described above.

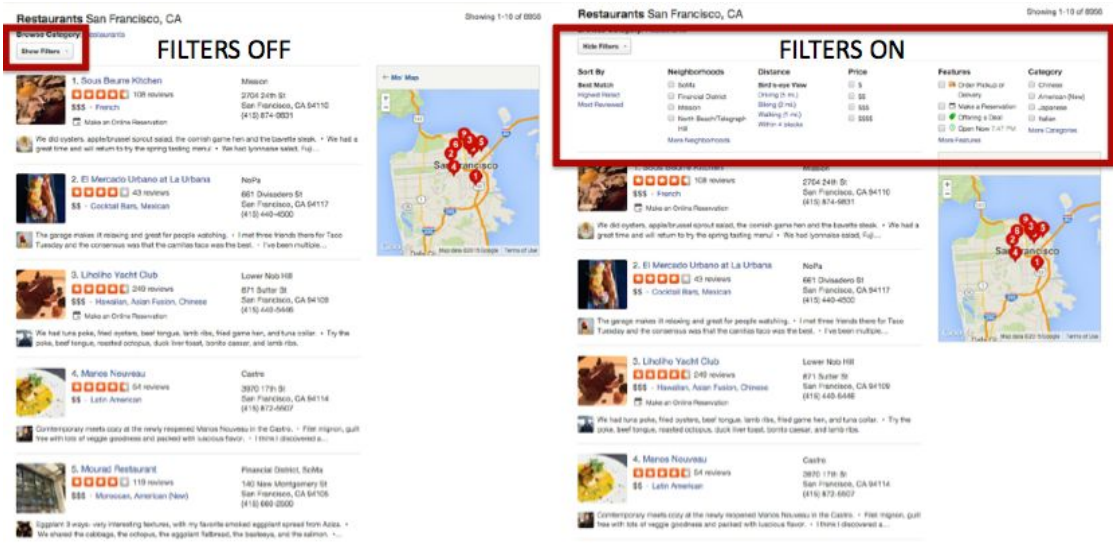

The image above is an example of survey performed: for the “restaurants near San Francisco, CA” query, users were shown the search results with filters hidden (left panel) or visible (right panel). For both Yelp traffic and the click-survey, user engagement with the map was measured and is summarized below.

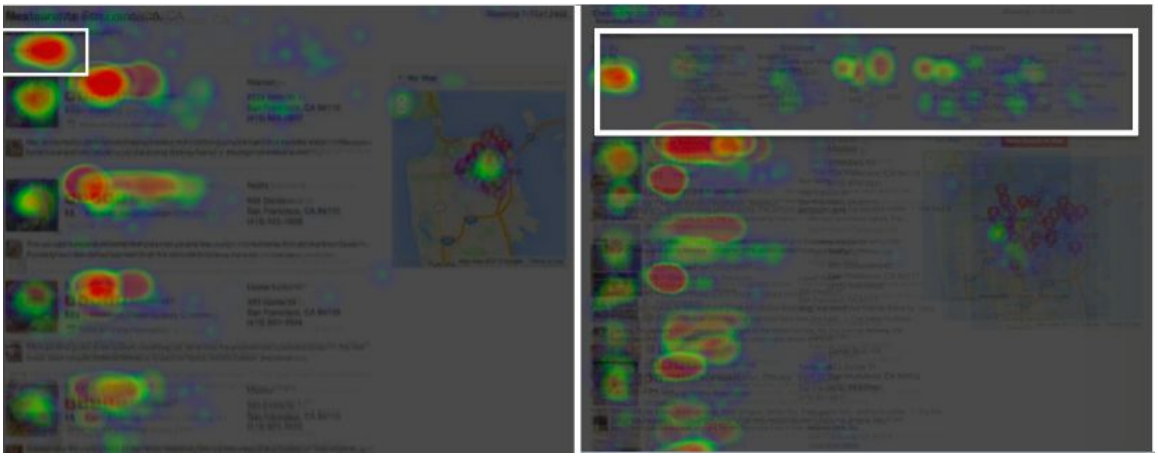

Map engagement	Yelp	N	UsabilityHub	N
Filters shown	X%	57131	0.12*X%	3531

Filters hidden	Y%	530	0.127*Y%	3531
CTR Ratio	1.22		1.29	

The actual CTR on Yelp is hidden because the data is proprietary, however note that while Click-test study CTR was much smaller than that of the Yelp live traffic, the change due to the experimental condition remained quite similar: 1.22 vs 1.29. This agreement of the magnitude and direction of the user response to the experiment across different experimental techniques suggests that the click-test technique can deliver relevant insights to real-world user behavior.

Differences in the context of a click-survey as compared to actual internet usage clearly affect user behavior. This comparison does not aim to establish that the behavior of click-surveys is identical to the behavior of actual internet users, but instead shows that while such differences obviously exist, the directionality of the difference is preserved and, for example, positive changes in engagement in a click-survey are mirrored by positive changes in actual internet traffic. Thus, an observed click-survey study change can be used as evidence of an actual internet traffic change in the same direction.

Click-surveys are used here as a substitute for a live A/B test on Google with FOTUL as an experimental condition compared to a control condition using current search results, because Google does not provide these data, which could definitively show the superiority or inferiority of the current local search algorithm. The raison d’etre of click-surveys is to provide insight into user behavior when

live A/B testing data is not available, and thus using them for this purpose, while perhaps novel in this context, aligns with industry practices.

Indeed, in the absence of internal data from sites such as Google, click-surveys appear to be a useful and directionally accurate way to estimate user reactions to proposed changes. We suggest more generally that the low cost and high speed of A/B click surveys may generally serve as a useful and accurate tool for competition enforcement authorities in their investigations of anticompetitive conduct as it provides a dynamic measurement tool in these fast-moving markets.

V. Legal Criteria of Actionable Exclusion

A demonstration of consumer harm is relevant to nearly any competition law theory one might care to invoke. We therefore think the conclusion that Google is knowingly degrading its search at the expense of consumers stands on its own as an crucial fact. However, we here develop three ways of thinking about the information which might prove relevant to the question as to whether an enforcement authority ought take action.

A. Naked Exclusion, Neglect of Less Restrictive Alternatives & Sacrifice of Product Quality

In any competition case centered on exclusionary conduct, authorities face the following hard question: When should conduct that excludes competitors be excused by virtue of the fact that it may also be efficient or

beneficial for consumers? How can "pro-competitive" efficiencies be weighed against anti-competitive exclusion? The problem shows up with great frequency in cases arising under Section 2 of the Sherman Act or Para 102 of the Treaty on the Functioning of the European Union.[27]

American judges have been particularly sensitive to this challenge, and have warned of the dangers of the "false positive" – that is to say, condemnation of practices that might, on balance, be good for consumers or the economy. As the American Supreme Court has stressed, to condemn exclusionary conduct that is pro-competitive might "chill the very conduct that the antitrust laws are designed to protect."[28] The European Commission is also wary of this concern, and allows the justification of conduct "leading to foreclosure of competitors on the ground of efficiencies that are sufficient to guarantee that no net harm to consumers is likely to arise."[29]

As a response to this concern, courts and commentators have isolated scenarios where action in cases of exclusion is most clearly warranted. We

---

[27] *See* Sherman Antitrust Act, 15 U.S.C. §2 (2000). See *also* Consolidated Version of the Treaty of the Functioning of the European Union, art. 102, Oct. 26, 2012, 2012 O.J. (C 326) 47.

[28] Matsushita Elec. Indus. Co. v. Zenith Radio Corp., 475 U.S. 574, 575 (1986). The challenge is augmented in software or technologies cases where exclusion is often accomplished by product design, for competition authorities have sometimes been inclined to defer to the designer. In the American Microsoft case, for example, the court pointed out that "In a competitive market, firms routinely innovate in the hope of appealing to consumers, sometimes in the process making their products incompatible for those of rivals; the imposition of liability when a monopolist does the same thing will inevitably deter a certain amount of innovation." United States v. Microsoft Corp., 253 F.3d 34, 65 (D.C. Cir. 2001).

[29] Communication from the Commission, supra note 4.

consider three here: "naked exclusion," "neglect of less restrictive alternatives" and "sacrifice of product quality."

The first is the category of "naked exclusion," where the exclusion of competitors is simply not justified by any real efficiencies or benefits for consumers (even if they may be claimed.). As Susan Creighton puts it, such conduct is that which "is likely … to have only anticompetitive effects."[30] Among many classic examples is the *Lorain Journal*[31] case, where a newspaper refused to print advertisements from companies that patronized its rivals, or *Allied Tube*[32], where the makers of a form of steel pipe conspired to prevent plastic pipes from being accepted by a standards body.

The second category is one where a dominant firm forgoes an obvious, less restrictive alternative course of conduct that would be equally, or more effective in serving the pro-competitive goals articulated. The reliance on a less-restrictive alternative analysis, as commentators point out, serves several purposes. For one thing, it serves as an aid to the balancing of procompetitive efficiencies with the harms of exclusion. As Herbert Hovenkamp has written, given complex balancing, "first and foremost, the antitrust decision-maker must

---

[30] *See* Susan A. Creighton, Dir. Bureau of Competition, Fed. Trade Comm'n, Remarks at the Charles River Associates 9th Annual Conference: Current Topics in Antitrust Economics and Competition Policy (Feb. 8, 2005), available at https://www.ftc.gov/sites/default/files/documents/public_statements/cheap-exclusion/050425cheapexclusion.pdf.

[31] Lorain Journal Co. v. United States, 342 U.S. 143 (1951).

[32] Allied Tube and Conduit Corp. v. Indian Head, Inc., 486 U.S. 492 (1998).

look for less restrictive alternatives."[33] Second, as Scott Hemphill points out, the use of less restrictive alternatives may also serve as important tool for "smoking-out" an illicit motive.[34]

In European practice, an "indispensability" analysis mirrors the American "less restrictive alternatives" jurisprudence. As the Commission has written, a firm that offers a pro-competitive efficiency must show that "its conduct is indispensable to the realization of those efficiencies: there must be no less anticompetitive alternatives to the conduct that are capable of producing the same efficiencies."[35]

The indispensability or less-restrictive alternatives analysis aims to identify the following situation. Consider a dominant firm who claims that its exclusionary practices serves a laudable, pro-competitive goal. It faces two equally effective strategies for doing so: strategy A and B, but strategy A is more exclusionary of its competitors. In that context, the choice of A is suspicious on its face. But matters become much worse if strategy A is, moreover, not equally, but actually *less* effective in achieving the stated pro-competitive goals; in that case suspicion may be replaced with outright incredulity. The choice of a less effective but more exclusionary alternative strongly suggests both the intent and

---

[33] Herbert Hovenkamp & Areeda Phillip, Antitrust Law: An Analysis of Antitrust Principles and Their Applications, § 2111c (2nd ed. 2005).
[34] C. Scott Hemphill, Less Restrictive Alternatives in Antitrust and Constitutional Law (working paper 2015) , available at http://scholarship.law.berkeley.edu/cgi/viewcontent.cgi?article=1031&context=law_econ.
[35] Communication from the Commission, supra note 4.

implementation of an anticompetitive and consumer harming course of conduct that should attract serious scrutiny.

Third, enforcers may consider, as evidence of actionable exclusion, the "sacrifice" of profit or product quality so as to damage competitors. When a dominant firm degrades its own products so as to damage competitors, it is often reasonably inferred that the goal of such conduct is not, in fact, innovation, but the maintenance of dominance. Here, Google is sacrificing quality and profits on its search platform to exclude rivals to its local product. Such sacrifices are often linked to a change in an existing course of dealing.

The idea of a sacrifice in product quality or of profit has an important pedigree in competition law. In the European Commission case British Midland v. Aer Lingus[36], for example, regulators were greatly concerned that Aer Lingus sacrificed a better product – the issuance of "interlined" airline tickets - for the evident purpose of slowing a rising competitor (British Midland). At issue was an existing practice whereby the airlines, who shared a route, issued a ticket that would be honored by either airline. The Commission noted that the practice was both preferred by consumers and revenue generating for all involved. When Aer Lingus began to refuse to interline with British Midland, it concluded that the denial of interlining was intended to damage competition and would harm consumers. "A significant number of passengers consider the possibility to change tickets and organize complex journeys on a single ticket as necessary"

[36] Commission Decision No. 92/213/EEC, 1992 O.J. (L 96) 34 (EC).

wrote the Commission, and "a refusal to interline will have the effect of diverting many of these passengers away from the new entrant airline. In this respect, a refusal to interline affects in particular the well-informed business travellers who require fully flexible tickets and who make a disproportionately large contribution to the revenue of the new entrant."[37]

The basic premise of the sacrifice analysis is a recognition that that even though exclusionary conduct can be ambiguous in its effect on consumers, where the exclusion creates a demonstrably lower quality product it becomes inherently suspicious. An American take on such conduct can be found in the *Aspen Skiing* case, where the dominant skiing company in the Aspen area purposely destroyed a previously popular "all mountain" pass for no reason other than doing damage to its competitor.[38] The strategy has the consequence of both harming consumers and damaging a competitor, and as such, bears some resemblance to the conduct at issue here.

B. As Applied

Our analysis and results suggest that Google's implementation of universal search in the local category sets of warning signs, whether it is considered through an instance of naked exclusion, neglect of a less restrictive alternative, or a sacrifice of product quality. That should lead enforcers to treat the implementation of universal search with suspicion.

---

[37] *Id.* at 40.
[38] Aspen Skiing Co. v. Aspen Highlands Skiing Corp., 472 U.S. 585 (1985).

Naked exclusion is an appropriate conclusion if considering universal search not as a whole, but specific implementations. This paper demonstrates that, in local search, the effect of Google's OneBox on consumers is clearly negative. There may be other examples, discoverable through further testing. It may be appropriate then, to consider implementations particularly those serving up non-fact-based information of universal search as harmful to both consumers and competitors as simply the exercise of naked exclusion.

Alternatively, one might choose to view universal search more broadly, and credit the instances where Universal Search does benefit consumers (such as its more rapid responses to inquiries for raw factual data). In that case Google's implementation of universal search more readily fits within the category of forgoing more effective, less restrictive alternative in favor of a strategy that both does more to exclude competitors and less to serve users. As such, the test gives strong reason to suggest that Google is opportunistically deploying universal search to prevent any threats to the general search engine, at the expense of both competition and consumers.

Third, it should not be difficult to conclude that Google is sacrificing product quality in the pursuit of the exclusion of its competitors. As demonstrated above, consumers actually preferred a universal search that includes both Google and its competitors, just as travelers preferred the interlined airline tickets. As such, just as in the Aer Lingus case, there is "no persuasive and

legitimate business justification for its conduct."[39] Nor, as the Commission has recognized, can Google's desire to grow its own specialized search product competitors serve as a justification for the conduct. The "desire to avoid loss of market share … [does] not make this a legitimate response to new entry."[40]

The statements tend to support the theory that Universal Search was at minimum a two-edged instrument. It has been, at times, deployed in a way that benefits consumers. But everything also points to the use of the same instrument, in at least some areas, to exclude competitors at the expense of consumers, a finding that must be taken seriously.

---

[39] *Id.* at 41.
[40] *Id.*