

Does Google Content Degrade Google Search? Experimental Evidence

Michael Luca
Sebastian Couvidat
William Seltzer

Timothy Wu
Daniel Frank

Working Paper 16-035



Does Google Content Degrade Google Search? Experimental Evidence

Michael Luca

Harvard Business School

Timothy Wu

Columbia Law School

Sebastian Couvidat

Yelp

Daniel Frank

Yelp

William Seltzer

Yelp

Working Paper 16-035

Copyright © 2015 by Michael Luca, Timothy Wu, Sebastian Couvidat, Daniel Frank, and William Seltzer

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Does Google content degrade Google search? Experimental evidence^{*}

Michael Luca[†], Timothy Wu[‡], Sebastian Couvidat[§], Daniel Frank^{**}, William
Seltzer^{††}

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 use its dominance in search to gain customers for this content. This may reduce consumer welfare if the internal content is inferior to organic search results. In this paper, we provide a legal and empirical analysis of this practice in the domain of online reviews. We first identify the conditions under which universal search would be considered anticompetitive. We then empirically investigate the impact of this practice on consumer welfare. To investigate, we implement a randomized controlled trial in which we vary the search results that subjects 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 roughly 40% more likely to engage with universal search results (which receive favored placement) when the results are organically determined relative to when they contain only Google content. To shed further light on the underlying mechanisms, we show that users are more likely to engage with the OneBox when there are more reviews, holding content constant. This suggests that Google is reducing consumer welfare by excluding reviews from other platforms in the OneBox.

^{*} The authors thank Susan Athey, Ben Edelman, Joshua Gans, Luther Lowe, and participants at the Antitrust Enforcement Symposium for helpful comments. All remaining errors are our own. This project received financial support from Yelp. Wu received compensation for his time and travel on this paper. Luca has a consulting arrangement with Yelp, but did not receive compensation for this paper. Brooks, Frank, and Seltzer are members of the Yelp data science team.

[†] Harvard Business School, corresponding author: mluca@hbs.edu

[‡] Columbia Law School, corresponding author: wu@pobox.com

[§] Yelp

^{**} Yelp

^{††} Yelp

1. Introduction

Google is the world's dominant search engine, accounting for roughly two-thirds of US searches and more than 90% of searches in parts of the world, far exceeding competitors such as Bing and Yahoo (comScore qSearch™ Analysis 2014). With a market capitalization of roughly \$350 billion, Google is also among the most valuable companies in the world. In 2006, "Google" even became a verb in the Oxford English Dictionary, meaning to search for something on the web (Google Operating System Blog 2006).

In an effort to expand its downstream product offering, Google has begun to develop its own content over time, such as its own price results for shopping and its own reviews for local businesses. In these situations, Google is acting both as a search engine and a content provider. Google shopping competes with Amazon; Google reviews compete with TripAdvisor. To use its search dominance to promote this content, Google has developed a feature called "universal search," through which it blends specialized search properties - often from proprietary databases - in priority over the results generated by an organic, or algorithmic search. Universal search intentionally excludes content competitors and only shows Google's content. For example, Google places and reviews may substitute for content from other specialized search platforms such as TripAdvisor and Yelp.

Should this raise concerns from competition authorities? From a legal perspective, Google's prioritization of its own content raises concerns about

consumer welfare and illegal uses of monopoly power if the internal content is inferior to organic search results. However, this behavior may be acceptable if the internal content is simply a better product. The impact of universal search on consumer welfare is then an empirical question with important implications from an antitrust perspective. In this paper, we develop a full legal analysis exploring antitrust implications of universal search.

Does Google's favorable placement of its own content results in harm to consumers? Because it is difficult to observe consumer preferences from observational data, the Federal Trade Commission and other regulatory bodies often struggle to identify this behavior despite widespread allegations of illegal exercises of monopoly power.

We propose the use of randomized controlled trials to identify whether Google's favorable treatment of Google content increases or decreases consumer welfare. Specifically, we implement an experiment 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. Through a series of online experiments, we provide a case study of a situation in which Google is systemically making its overall product worse for users in order to provide favorable treatment to Google content. To populate the treatment group, we use content selected by Google's own organic algorithm.

We begin with a simple thought experiment. Suppose you are planning a trip to Louisville, Kentucky and are searching for a coffee shop through Google. Clearly, there is a wide variety of content that might facilitate this search.

Competing ratings and reviews ranging from Yelp to TripAdvisor to Food & Wine invest heavily in developing such content. In this situation, Google's content may be more or less useful to users than other content. If Google provides favorable placement to Google content in a world in which Food & Wine is – on average - more useful, then this creates harm to consumers.

“Local intent”-based searches (including areas such as coffees shops, doctors, and mechanics) comprise the largest single category of search behavior. Google currently 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.” 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. 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.

Exploiting these institutional features of Google, we construct two sets of results for searches for coffee shops in different markets. In the control group, we display Google's actual results. In this situation, the OneBox contains Google content. In the treatment group, we display the exact same organic content below the OneBox. However, instead of filling the OneBox with only Google content, we fill the OneBox with listings and reviews from the content providers that Google's algorithm organically isolates (Yelp, TripAdvisor, etc.). We construct these alternative results through a browser plug-in called “Focus on the User - Local”

(FOTU), which has queried third-party review sites and ranked them, using Google's own organic algorithm, according to which site delivers the most relevant information for the content in question.

If Google's placement of its own content in the OneBox makes users better off, then we should expect users to engage with the OneBox more often with Google content than with organic content offered by Google's algorithm. In contrast, if users are less likely to engage with the OneBox when it is filled with Google content, this suggests that Google is creating harm by tying its own content through the OneBox.

We then task subjects on UsabilityHub with searching for a local coffee shop, and compare the performance across these two sets of results. We find that users would be more likely to engage with local specialized search results if Google were to replace its proprietary results in universal search with results drawn from the web based on the same merit-based algorithm that it uses to populate organic search (as opposed to being exclusively drawn from Google+).

The results demonstrate that consumers prefer the second version of universal search. Stated differently, consumers prefer results scored by Google's own organic search engine to the content currently developed for Google+ Local. 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. The easy and widely disseminated

argument that Google's universal search always serves users and merchants is demonstrably false. 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.

The demonstration of consumer harm is 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 (Rasmusen et al. 1991)

Alternatively, Google's conduct can be understood as the knowing neglect of a "less restrictive alternative" for achieving legitimate goals (Hemphill 2015). 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, 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.

Overall, our findings contribute to a growing literature on the economics of online platforms. Early work showed that the growth of the Internet and online search created consumer welfare gains ranging from lower prices of life insurance (Brown and Goolsbee 2002) to increased variety of books (Brynjolfsson, Hu, and Smith 2003). Specialized search platforms such as Yelp, TripAdvisor, and Amazon provide vast amounts of user-generated content about goods and services. Online marketplaces such as TaskRabbit, Airbnb, and Uber create markets where none existed. These platforms have influenced markets ranging from restaurants and hotels to books to labor markets, through lower search costs and increased information flows.⁷

Our results suggest that some of the welfare gains are lost due to Google's market power and the practice of tying Google content to Google search results. Our findings also relate to the literature on market power in platforms. Hagiu and Julien (2011) develop a theoretical model of search diversion. Edelman and Lai (2015) empirically show that the introduction of Google's preferential placement of its flight search tool crowded out visits to organic pages. Edelman (2015) provides the first legal analysis of Google's practice of tying. We contribute to this literature by providing further legal analysis as well as empirical evidence of consumer harm from a randomized controlled trial.

In addition to our contribution to the legal analysis of antitrust cases in this domain, our experimental approach provides a new tool for regulators who are

⁷ See Chevalier and Mayzlin 2006, Luca 2011, Ghose et al 2012, and Einav et al. 2015 for evidence on the impact of online reviews and welfare gains from online marketplaces.

seeking to better understand consumer preferences and harm in the online world. To the best of our knowledge, our paper is the first to incorporate the use of randomized controlled trials into an antitrust analysis.

2. Background on Online Search and OneBox

The Internet, with more than 4.7 billion linked pages, contains more information than any one user could read in a lifetime. Because this content is extremely decentralized, users frequently begin their quests for online information through *search engines* such as Google, Bing, and Yahoo. Users enter a word, number, or phrase, and the search engine will return a large list of webpages that the user may want to click on. For example, a Google search for “restaurants near me” yields roughly 443 million results.

Because there is so much information to sift through, 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 determines 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 central task of search engines is to choose the order in which to list results. The history of search engines 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) became the tool of choice; 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 sellers are matched.

General search engines now use an algorithm that determines which results are most likely to be clicked and orders them based on the percent of people who clicked on that result, along with other factors. This approach relies on the *revealed preference* of users. Developing high quality search results has been the core objective of search engines.

The evolution of the web’s information retrieval tools did not end with general search. From around 2000—2014, just as Google’s general search replaced 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). Eric Schmidt captured this reality when he recently suggested “our biggest search competitor is Amazon.”

Google has, over the decade, 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 essentially clones of competitors, like Kayak or Yelp (yielding, for example, “Google travel” “Google shopping,” “Google+ Local,” etc.). Earlier versions of Google’s clones also used data from the originals; for example, early versions of Google local took the reviews collected by Yelp and TripAdvisor and incorporated them into Google’s product.

Some of these products – such as Google Maps – have gained traction and become dominant over time, in part due to excellent engineering and superior product design. Others have not proved as popular with users as the originals. Many of these clones suffer from poor search rankings in Google’s organic search results.

To improve the popularity of its specialized search features, Google has used the power of its dominant general search engine. The primary means for doing so is what is called the “universal search” or the “OneBox.” 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 that indicate that the user is performing a specialized search. The OneBox is one of several Google products that has given rise to antitrust concerns, and is the focus of this paper.

A search for “restaurants San Francisco” might trigger Yelp as the first natural result. Google deploys a Local OneBox above the natural results. The OneBox, with a few exceptions, is populated by Google’s own versions of its competitors’ specialized search services (e.g., Google+ Local or travel). Figures 1 and 2 provide sample screenshots of universal search results.

3. Legal Assessments of Universal Search

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 led to scrutiny both by European and American competition regulators.

Google’s strategy follows a well-known historic pattern. Technological monopolists, facing threats from innovative competitors, often engage in vertical 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. 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 versions of their products. The most prominent victim of that conduct was Netscape.

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

Based on the facts available to them, the economic experts employed by Google and other academics have tended to assert, generally, that universal search serves consumer interests, a point this paper does not contest. For example, Michael A. Salinger and Robert J. Levinson, Google’s economic consultants, describe universal search as an important innovation, a point with which we take no disagreement. “[T]here can be little doubt that answering questions directly benefits consumers... 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.” Similarly, Robert Bork and Greg Sidak, in *What does the Chicago School Teach about Internet Search*, argue that universal search is “a product improvement that consumers value.” They add, using a one-monopoly profit theorem, that Google has no particular reason to want to extend its market power into specialized search.

The assertions made by Salinger, Levinson, Bork and Sidak do not really address the degradation described in this paper. No one truly disagrees that universal search, in concept, can be an important innovation that can serve consumers. The more challenging question arises not from the deployment of universal search to tell people the weather, but its intentional degradation for exclusionary in areas where Google faces the most serious long-term competition.

In this respect, the most in-depth analysis of exactly the dilemma presented by universal search was presented by James D. Ratliff & Daniel L. Rubinfeld. Ratliff & Rubinfeld describe the possibility that third-party content might be better, or more relevant, than Google's own affiliated content, presented by universal. In that case, the authors suggest, "Google faces a tradeoff. If Google listed the less relevant Google-affiliated website more prominently, Google would benefit from the greater 'free publicity' that website would receive. Conversely, choosing the less relevant Google-affiliated website to display more prominently would, by assumption, lower the relevance of the organic search results. This effect, if non-negligible, could cause the user to have a poorer search experience compared with one in which the more relevant link was listed more prominently."

Having captured the search degradation potential presented by universal search, however, the authors, working mainly by assumption, assert that the degradation will not happen. In short, they theorize that any degradation caused by preferential treatment of Google's own properties would be noticed by users,

causing a switch to competing search engines, general or specialized. Hence, the authors conclude, Google will always be incentivized to do what is best for users, and should, by their assumptions, always choose third-party sites when that is good for users. As they state, Google “would be more likely to more prominently display the more-relevant website (not affiliated with Google) the greater the relevance advantage of the non-Google website and the lower the beneficial effect of the enhanced free publicity on Google’s affiliated business.”

As detailed below, the underlying experiment directly contradicts these assertions, suggesting that that the authors have come close to assuming away the problem. Among various implicit assumptions, they assume that users have perfect knowledge of the quality of the search they are using, and would suffer relatively little switching costs in abandoning Google for some constellation of alternatives. They essentially assume that whenever Google chooses its own properties it does so because its own properties are better than those provided by third parties.

That latter assertion is contradicted by Google’s PageRank itself, which consistently ranks Google’s products below those offered by third parties. More generally the fact, demonstrated here, that Google does exactly what Ratliff and Rubinfeld describe — offer a degraded search product that ignores more relevant alternatives, strongly suggests that something is wrong with the Ratliff and Rubinfeld analysis. Among other things, Ratliff and Rubin fail to take into consideration Google’s interest in the exclusion and weakening of long-term, paradigmatic competitors. That additional factor, we theorize, may justify, for the

firm, the degradation of the search product in some areas, where it faces dangerous nascent competitors. The premise is that Google can afford to degrade its search in some areas while retaining top quality searches in others, without notably losing users or market share.

Support for this theory comes from Google's pattern in its deployment of non-affiliated content. In areas where Google does not face a serious specialized search competitor, like general knowledge questions, or health questions, Google relies on the best third party sites as revealed by PageRank (sites like Wikipedia or the Mayo Clinic). It is only where it faces a specialized competitor that Google engages in degradation of the search as described here. That suggests the factor missing from Ratliff and Rubin's analysis is the interests of Google in weakening its nascent competitors, which may be a real interest, but not a pro-competitive interest.

Randy Picker makes this point more explicit when he notes that a media monopolist that depends on advertising as opposed to cash payments need always degrade the product (by adding ads) to make its profit. As Picker points out, the important question, from a competition perspective, is what form the degradation takes. As he writes:

"For the zero-cash-price media monopolist, the exercise of market power is just about product degradation. ... [Y]ou take the product that you would otherwise sell to consumers and make it worse in consumer eyes by adding advertising. You do that to make money."

“We should expect media monopolists to degrade their products and the only issue is precisely how they do that. Degrading the product to make money today through ads is ... a legitimate use of market power. But degrading the quality of the product to maintain that monopoly or to extend it to a new setting is an illegitimate use of market power.

Much of the other academic writing in this area concerns the question of market power in search — a question important to a full competition law analysis but one beyond the scope of this paper (see, for example, Denegri-Knott et al. (2006)). Without dwelling on the topic, we make just one point. The ability of Google to profitably degrade its search in some areas, as suggested by the results here, has natural empirical relevance to the firm’s possession of market power. Indeed, it might be taken as direct evidence of market power,⁸ though we do not fully develop the point here.

The ability to offer a degraded product without losing consumers is relevant, for example, to the argument that Google does not have market power because of an absence of switching costs; Professors Aaron Edlin and Robert Harris, among others, make the argument that switching costs are “trivial” in search markets, and suggest that this greatly constrains Google’s power (Edlin and Harris 2013). In purely competitive market, without switching costs, it would seem to be impossible for Google to sustain a degraded search without immediately losing its users to competitors. The apparent ability of Google to maintain a degraded search product in at least some areas may suggest users

⁸ As in the Toys R’ Us case in the Seventh Circuit Court of Appeals.

will not immediately switch even if the search is selectively degraded, suggesting that there may be something wrong with the Edlin and Harris theory.

A final assessment of the universal search product came from the Federal Trade Commission, both in its statement closing its investigation, and the remarks of Chairman Jon Leibovitz. The FTC, in its closing of the case, stated, based on the evidence it had, that “the documents, testimony and quantitative evidence the Commission examined are largely consistent with the conclusion that Google likely benefited consumers by prominently displaying its vertical content on its search results page.” In his statement, based on the evidence then available to the Commission, Leibovitz wrote “We close that investigation, finding that the evidence does not support a claim that Google’s prominent display of its own content on its general search page was undertaken without legitimate justification ... Google’s primary reason for changing the look and feel of its search results to highlight its own products was to improve the user experience.” While the full scope of the evidence reviewed is not public, it is evident that the Commission did not have the benefit of randomized controlled trials when it offered that conclusion in 2013. The goal of this paper is to test, whether in fact the user experience has been improved. As we stressed, in some areas, the results suggest the opposite.

3.1 Theories of Harm

Google creates economic welfare by reducing search frictions and matching users with the objects of their preferences. Users come to Google with preexisting preferences for goods and services (European soccer games,

vinyasa yoga, New York-style bagels and so on.) They ultimately are matched with 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).

Google serves as the “platform” or “intermediary” that matches buyers with sellers of desired goods or services. 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.

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 (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 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. Yelp 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.

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 used 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.

The very introduction of universal search, as Google has highlighted, is itself a form of innovation. However, when considering welfare implications, 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-generated 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” (Wu, 2011).

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 government controls (Goldsmith and Wu 2006). As professor Jeffrey Rosen has written, Google has a particularly central role in this process (Rosen 2012). 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” (Wu 2010; see also Rosen 2008).

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.

4. Experimental Design

4.1 Focus on the User Plug-In

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

FOTU essentially constructed an alternative method of presenting local results, based on the algorithm that powers Google organic search. Note that maintaining a plug-in such as FOTU in an ongoing basis would be challenging because Google search result pages, local content, and domains change often. In fact, the content has changed even since this version of this plug-in was created, preventing this from being adopted by users. However, for the purpose of this research, we extracted these ordered lists for our key search, which we incorporate into our treatment group below.

4.2 Treatments

We create three experimental conditions, described below.

4.2.1 Control

To create a control condition, we captured a screenshot of results for searches for the phrase “coffee Louisville KY”. Throughout the experiment, this was used as a control display. Figure 3 shows the screenshot.

4.2.2 FOTU

This treatment is identical to the control except for the content of the OneBox. Instead of containing Google content, the OneBox presents content from FOTU. By comparing the control to this treatment, we can explore consumer preferences across the two sets of content.

4.2.3 No Reviews

This treatment is identical to the control except for the content of the OneBox. The OneBox contains the same Google content as the control, except we removed the rating and review counts next to the businesses (which are contained in the control). This will allow us to identify the mediating role of reviews in the comparison between Google content and FOTU.

4.2.4 Experimental Tasks

With treatment groups in hand, we ran two sets of experiments. In the first, we randomized subjects into one of the three conditions above. Users were then asked where they would click if conducting the specified search. One possible concern with this is that clicks may not necessarily measure preferences. To

confirm our interpretation, we run a survey in which we show two sets of results and ask them which they would use.

4.3 Subjects and Experimental Platform

We run our experiment on UsabilityHub, which is a testing platform used by companies that are interested in testing different designs of their webpages. Our subject pool consists of 2,170 users from Amazon Mechanical Turk, which is an online labor market and common platform used for social science experiments. UsabilityHub provides a useful platform for experimentally testing different search features, since it is the same platform that multiple consumer Internet companies actually use for this type of task.

5. Empirical Results

5.1 User Preferences for FOTU

Figure 4 presents the main results. In the control condition, 48% of users clicked inside of the OneBox. In contrast, 66% of users clicked inside of the OneBox when it was filled with FOTU content. In other words, users are roughly 40% more likely to engage with the OneBox when it is filled with FOTU content/using Google's own algorithm.

One might be concerned that user clicks may not necessarily reveal their preferences. To address this concern, we run a simple second survey experiment in which we show a separate set of users both conditions, and ask which they prefer. Per figure 5, roughly 70-80% of users prefer FOTU in a head-

to-head comparison, reinforcing our interpretation that clicks are a good measure of user preferences.

5.2 The Role of Reviews

Our main result shows that users prefer FOTU content, but does not explain the mechanism. Our hypothesis is that users prefer FOTU content because it contains more reviews than the control (464 relative to 117). By construction, FOTU must *always* have more reviews, since it pulls data from reviews platforms and Google reviews, whereas the Google OneBox intentionally excludes other review platforms.

To test this hypothesis, we run a second set of experiments in which we show the identical content, but vary the number of reviews shown next to each business (117 relative to 0). In figure 4, we show that only 33% of users click on the OneBox when it contains no reviews, relative to 48% when it contains 117 reviews – even though all other content is identical. In other words, users are roughly 46% more likely to engage with the OneBox when it contains reviews relative to when it doesn't, holding all else constant. To support our interpretation, we again run a head-to-head comparison in which a different set of subjects is asked to choose between the two sets of content. Figure 6 shows that users overwhelmingly prefer the content that contains reviews relative to the content that does not.

Overall, these findings show that users prefer FOTU content to Google content, and that this is mediated by the fact that FOTU contains more reviews.

By intentionally excluding reviews from other platforms, Google is hence degrading the quality of its product.

5.3 Cross-Validating Subject Responses

While our experimental approach is new to antitrust analysis, randomized controlled trials form the backbone of much of the social sciences and are becoming an increasingly important policy tool. Experiments are also an important approach used when designing online platforms – in fact, this is one frequent use of UsabilityHub.

Because this is a new approach to this context, we cross-validate our results by comparing behavior on UsabilityHub to known behavior in online platforms. Optimally, we would compare our results to results from experiments done within Google. However, we do not have access to their data since it is proprietary. Instead, we ran a separate experiment within Yelp and the exact same experiment on UsabilityHub.

Specifically, we ran an experiment in which a control group was shown a Yelp search page without any filters and a treatment group was shown a Yelp search page with filters – as shown in Figure 7. We then measured the percent of users who engaged with the map on the page. We ran this both on UsabilityHub and in vivo on Yelp. We found that users were 22% more likely to engage with the map in the treatment condition on Yelp, and 29% more likely to engage with the treatment condition on UsabilityHub. The qualitative consistency of results across these platforms provides further support for the validity of our experimental approach.

6 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 a crucial fact. Here, we here develop three ways of thinking about whether an enforcement authority should take action.

6.1 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 (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).

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.” *Matsushita Elec. Indus. Co. v. Zenith Radio Corp.*, 475 U.S. 574, 575 (1986).⁹ 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.” *Commission Decision* No. 92/213/EEC, 1992 O.J. (L 96) 34 (EC).

As a response to this concern, courts and commentators have isolated scenarios where action in cases of exclusion is most clearly warranted. We 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 defines the term, it is such conduct is that which “is likely ... to have only anticompetitive effects.” (Creighton 2005). Among many classic examples is the *Lorain Journal* case, where a newspaper refused to print advertisements from companies that patronized its rivals, or *Allied Tube*¹⁰, where the makers of a form of steel pipe conspired to prevent plastic pipes from being accepted by a standards body.

⁹ 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).

Lorain Journal Co. v. United States, 342 U.S. 143 (1951); *Allied Tube and Conduit Corp. v. Indian Head, Inc.*, 486 U.S. 492 (1998).

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 pro-competitive efficiencies with the harms of exclusion. As Herbert Hovenkamp has written, given complex balancing, “first and foremost, the antitrust decision-maker must look for less restrictive alternatives.” (Hovenkamp and Areeda 2005). Second, as Scott Hemphill points out, the use of less restrictive alternatives may also serve as important tool for “smoking-out” an illicit motive (Hemphill 2015).

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.” *Commission Decision Id.*

The indispensability or less-restrictive alternatives analysis aims to identify the following situation. Consider a dominant firm who claims that its exclusionary practices serve 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 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*¹¹, 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). *Commission Decision Id.* 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” 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.” *Commission Decision Id.* at 40.

The basic premise of the sacrifice analysis is recognition 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.¹² The strategy has the consequence of both harming consumers and damaging a competitor, and as such, bearing resemblance to the conduct at issue here.

6.2 As Applied

Our analysis and results suggest that Google’s implementation of universal search in the local category sets off 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.

¹² Among various American examples of such conduct are *Allied Tube*, discussed above, and *Aspen Skiing Co. v. Aspen Highlands Skiing Corp.*, 472 U.S. 585 (1985), in which a dominant skiing company in the Aspen area purposely destroys a previously popular “all mountain” pass for no reason other than doing damage to its competitor.

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." *Commission Decision Id.* at 41. 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.” *Commission Decision Id.*

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.

6.3 Remedies

Among the greatest challenges facing any antitrust analysis is the question of remedies. During the various investigations of Google, the challenge of finding a remedy has been repeatedly stressed.

Marina Lao asserts, for example, that “the remedy for any incidental exclusionary effects of search ‘bias’ is that it is likely to do more harm than good.” (Lao 2013). In the following section we suggest that randomized controlled trials performed for this paper present a clear means for designing a remedy where Google can be shown to be degrading search for exclusionary purposes, and one that is based entirely on Google’s own operating algorithm. In other words, contrary to the assertions of Lao and others, there is no need for any third party assessment of Google’s search.

The remedy depends on the availability of Google’s PageRank algorithm, which provides a well-recognized gold-standard for assessing relevance. PageRank with its origins as an academic paper written by Google founders,

provides the industry's leading indicator of what is most relevant, and therefore most serving of user welfare, for a given search inquiry. The remedy would, in areas that Google has demonstrated exclusionary activity, enjoin Google to use its PageRank algorithm to populate its universal search results.

A few examples might make clear how this would work in practice, using “local” search as the subject, an area where Google has been widely accused of exclusionary practices. Following a finding that Google, in fact, has market power, and uses that market power in an exclusionary manner through universal search, the firm might be enjoined as follows. For those searches which trigger its local “OneBox” result, the firm must henceforth use the results of its own PageRank algorithm to populate that OneBox, as opposed to arbitrarily populating the OneBox with its own properties.

This remedy would come very close to doing what, at least according to some analysts, Google is already incentivized to do anyhow. Ratliff and Rubinfeld have argued that Google has natural reasons to include third-party content when it is better than that which Google itself can provide.¹³ If that is the case, which is true, then the harm to Google of listening to PageRank cannot be great, and indeed the remedy may make the product better, even if it makes Google more vulnerable to competition on the merits. This is, to be clear, the exact opposite of “do[ing] more harm than good.” Instead, the remedy, if adopted, could be expected to create a product better for consumers, at the expense of universal search as a barrier to entry.

¹³ See argument, above.

References

Armstrong, Mark, John Vickers and Jidong Zhou, 2009. "Prominence and Consumer Search," *RAND Journal of Economics*, Vol. 40, 209-233.

Armstrong, Mark and Julian Wright, 2007, "Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts," *Economic Theory*, Vol. 2, 353-380.

Brown, Jeffrey, and Austan Goolsbee, 2002. "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, Vol. 110, No. 3, 481-507.

Brynjolfsson, Erik, Jeffrey Hu, and Michael Smith, 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, Vol. 49, No. 11, 1580-96.

Chevalier, Judith A., and Dina Mayzlin, 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews," *Journal of Marketing Research* Vol. 43, No. 3, 345-354.

"comScore Releases August 2014 U.S. Search Engine Rankings." *comScore qSearch™ Analysis*, September 17, 2014. <http://www.comscore.com>

Creighton, Susan, Dir. Bureau of Competition, Fed. Trade Commission, Remarks at the Charles River Associates 9th Annual Conference: Current Topics in Antitrust Economics and Competition Policy. February 8, 2005. See more at https://www.ftc.gov/sites/default/files/documents/public_statements/cheap-exclusion/050425cheapexclusion.pdf.

Denegri-Knott, Janice, Detlev Zwick, and Jonathan E. Schroeder, 2006. "Mapping consumer power: An integrative framework for marketing and consumer research." *European Journal of Marketing*, Vol. 40., No. 9/10, 950-971.

De Corniere, Alexandre and Greg Taylor, 2015. "Quality Provision in the Presence of a Biased Intermediary," *working paper*.

Edelman, Benjamin, 2015. "Leveraging Market Power through Tying and Bundling: Does Google Behave Anti-Competitively?" *Journal of Competition Law & Economics*.

Edlin and Harris, 2012. "The Role of Switching Costs in Antitrust Analysis: A Comparison of Microsoft and Google," *Yale Journal of Law and Technology*. Vol. 15, 169.

Einav, Liran, Chiara Faronato, and Jonathan Levin. 2015. "Peer-to-Peer Markets." *Invited by the Annual Review in Economics*.

Ellison, Glenn and Sarah Ellison, 2009. "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, Vol. 77, No. 2.

Ellison, Glenn and Alex Wolitzky, 2012. "A Search Cost Model of Obfuscation," *RAND Journal of Economics*, Vol. 43, 417-441.

Ghose, Anindya, Panagiotis Ipeirotis, Beibei Li. 2012. "Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowd-Sourced Content." *Marketing Science* Vol. 31, No 3, 493-520.

Gilchrist, Duncan, Michael Luca, and Deepak Malhotra, forthcoming. "When $3+1>4$: Gift Structure and Reciprocity in the Field," *Management Science*.

Goldsmith, Jack L., and Tim Wu, 2006. *Who controls the Internet?: Illusions of a Borderless World*, Vol. 89. New York: Oxford University Press.

"Google is Officially an English Word." *Google Operating System Blog*, July 6, 2006. <http://googlesystem.blogspot.com/2006/07/google-is-officially-english-word.html>

Hagiu, Anrei and Bruno Jullien, 2011. "Why do Intermediaries Divert Search?" *RAND Journal of Economics*, Vol. 42, 337-362.

Hemphill, Scott, 2015. "Less Restrictive Alternatives in Antitrust & Constitutional Law," *working paper*.

Hovenkamp, Herbet and Phillip Arreda. 2005. *Antitrust Law: An Analysis of Antitrust Principles and Their Applications*, § 2111c, 2nd ed

Hyman, David and David Franklyn, forthcoming. "Search Bias and the Limits of Antitrust: An Empirical Perspective on Architecture and Labeling Remedies," *Arizona Law Review*.

Lao, Marina. 2013. "'Neutral' Search As A Basis for Antitrust Action?" *Harvard Journal of Law and Technology Occasional Paper Series*.

Luca, Michael, 2011. "Reputation, Reviews, and Revenue: The Case of Yelp.com," *working paper*.

Rasmusen, Eric B., J. Mark Ramseyer, & John S. Wiley, Jr., 1991. "Naked Exclusion," *The American Economic Review*, Vol. 8, No. 5, 1137-1145.

Rochet, J.C. and J. Tirole, 2003, "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, Vol. 1, 990-1029.

Rochet, J.C. and J. Tirole, 2006, "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, Vol. 37, 645-66.

Rosen, Jeffrey, "Google's Gatekeepers," *The New York Times*, November 28, 2008.

Rosen, Jeffrey, 2011. "The Deciders: The Future of Privacy and Free Speech in the Age of Facebook and Google," *Fordham Law Review*, Vol. 80, 1525.

Taylor, Greg, 2013. "Search Quality and Revenue Cannibalisation by Competing Search Engines," *Journal of Economics and Management Strategy*, Vol. 22, No. 3, 445-67.

Wu, Tim. "Is Filtering Censorship? The Second Free Speech Tradition, Governance Studies at The Brookings Institute." December 27, 2012. Available at http://www.brookings.edu/~media/research/files/papers/2010/12/27-censorship-wu/1227_censorship_wu.pdf.

Wu, Tim, 2011. *The Master Switch: The Rise and Fall of Information Empires*, Vintage Books.

Figures

Figure 1: Google Search Results for “Pancake”




Figure 2: Google Search Results for “pediatrician toronto on”

The screenshot shows Google search results for the query "pediatrician toronto on". The search bar includes the text "(\"pediatrician toronto on\")". Below the search bar, the "Web" tab is selected. The results are categorized into a "Local OneBox (aka '7-pack')", which is highlighted with a red border, and organic search results below it.

Local OneBox (aka "7-pack"):

- 1** Dr. Bonnie Cameron
kindercarepediatrics.ca
Google+ page
- 2** Dr. Morton Goldbach
www.vaughanpeds.ca
Google+ page
- 3** Greenberg Saul Dr
js.google.com
Google+ page
- 4** Dr. Davy MD
diatriciantoronto.ca
Google review
- 5** Jacobson Sheila Dr
js.google.com
Google review
- 6** Amila Vaz
js.google.com
Google+ page
- 7** Mills-Tetley Diana Dr
lus.google.com
Google review

Map results for pediatrician toronto on

Map: A map of the Greater Toronto Area showing the locations of the pediatricians listed in the Local OneBox. Red pins are placed on the map, labeled A through G, corresponding to the locations of the pediatricians. A red arrow points from the text "Local OneBox (aka '7-pack')\" to the map.

Organic (Position 1):

Pediatrician Toronto - Kindercare Pediatrics : Kindercare ...
kindercarepediatrics.ca/

Pediatrician Toronto - Multidisciplinary Team of Pediatricians, Pediatric Subspecialists, Child Psychologist, & Pediatric Allied Health Practitioners.

Figure 3: Google Search Results for “coffee Louisville ky”

C 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?

skip this test →
skip and report ▼

The screenshot shows a Google search results page for the query "coffee louisville ky". The browser address bar shows the URL: www.google.com/#client=psy-ab&q=coffee+louisville+ky&oq=coffee+louisville+ky&gs_l=hp.3..0j0i22i30i3.172886.174511.8.1746.... The search bar contains "coffee louisville ky". Below the search bar, there are tabs for "Web", "Images", "Maps", "Shopping", "More", and "Search tools". The "Web" tab is selected. The results show "About 3,230,000 results (0.18 seconds)".

On the left side, there are several search results:

- Sunergos Coffee | MicroRoastery and Esp...**
<https://sunergoscoffee.com/>
Score: **25** / 30 · 38 Google reviews
Locations · Contact Us · Coffee · Roasted Coffees
- Heine Brothers' Coffee**
heinebroscoffee.com/
7 Google reviews
HBC Perks! · Shop Menu · Contact Us
- Quills Coffee**
quillscoffee.com/
Score: **25** / 30 · 28 Google reviews
Locations · Blog · Store · Photos
- Please & Thank You**
www.pleaseandthankyoulouisville.com/
Score: **24** / 30 · 25 Google reviews
- Mrs. Potter's Coffee Lounge & Cafe**
www.mrspotterscoffee.com/
9 Google reviews
- Red Hot Roasters**
www.redhotroasters.com/
8 Google reviews
- VINT**
vintcoffee.com/

On the right side, there is a list of locations with their addresses and phone numbers:

- A** 2122 S Preston St
Louisville
(502) 634-1243
- B** 1449 Bardstown Rd
Louisville
(502) 454-5212
- C** 930 Baxter Ave
Louisville
(502) 742-6129
- D** 800 E Market St
Louisville
- E** 718 W Main St # 101
Louisville
(502) 581-1867
- F** 1403 Lexington Rd
Louisville
(502) 569-0000
- G** 4901 Brownsboro Rd
Louisville
(502) 455-2922

Below the list of locations is a map of Louisville, Kentucky, showing the locations marked with red pins labeled A through G. The map includes major roads like I-65, I-75, and I-275, and landmarks like the Ohio River and the Louisville Slugger Museum & Factory.

Figure 4: Percent of Users Who Click on the OneBox




Figure 5: Percent of Users Who Prefer FOTU to OneBox





Figure 6: Percent of Users Who Prefer More Reviews



[illegible]