Product Quality and Entering through Tying:
Experimental Evidence*

Hyunjin Kim† and Michael Luca‡

November 2018

Abstract

Dominant platform businesses often develop products in adjacent markets to complement their core business. One common approach used to gain traction in these adjacent markets has been to pursue a tying strategy. For example, Microsoft pre-installed Internet Explorer into Windows, and Apple set Apple Maps as the iOS default. Policymakers have raised concerns that dominant platforms may be leveraging their market power to gain traction for lower quality products when they use a tying strategy. In this paper, we empirically explore this question by examining Google's decision to tie its new reviews product to its search engine. We experimentally vary the reviews displayed above Google's organic search results to show either exclusively Google reviews (Google's current tying strategy), or reviews from multiple platforms determined to be the best-performing by Google's own organic search algorithm. We find that users prefer the version that does not exclude competitor reviews. Furthermore, looking at observational data on user traffic to Yelp from search engines, we find that Google's exclusion of downstream competitors may have been effective. The share of Yelp's traffic coming from Google has declined over this period, relative to traffic from Bing and Yahoo (which do not exclude other companies' reviews), and Google reviews has grown quicker than Yelp and TripAdvisor during the period in which they excluded these (and other) reviews providers. Overall, these results shed light on platform strategy: tying has the potential to facilitate entry in complementary markets, even when the tied product is of worse quality than competitors.

1. Introduction

Multisided platforms ranging from Google to Amazon to Airbnb play an increasingly central role in organizing economic transactions in the digital age. In recent years, many dominant platforms have developed new products to complement their core business, using a tying strategy to enter these adjacent markets. For example, Microsoft entered into browsers by pre-installing Internet Explorer in Windows, Apple entered into maps by setting Apple Maps as the iOS default, and Google developed a suite of products tied to its search engine.

A large body of research has explored the theoretical motivations and implications of tying for the platform owner, as well as its competitors and customers (e.g. Gans and Stern (2003), Carlton, Gans, and Waldman (2010), Eisenmann et al (2011)). A key question raised by the growing use of tying is the extent to which tying strategies enable successful entry without a higher quality product – often considered a necessary condition to entry in markets like platforms with sheltered incumbents (Henderson and Clark (1990), Bresnahan (1999), Evans and Schmalensee (2001), Eisenmann et al (2011)). From a managerial and policy perspective, it is important to understand whether tying strategies are being used to gain traction with lower quality products.

Academics and policymakers alike have explored the potential of dominant platforms entering adjacent markets without a product improvement. However, it is empirically challenging to ascertain whether dominant platforms are offering better or worse products (European Commission (2017), Greene (2015), Wu et al (2015)). Moreover, companies using tying strategies have little incentive to provide their own data. Google, the world's largest search engine and one of the world's largest companies, has come under significant scrutiny for entering many markets spanning news, shopping, flights, and reviews by tying these products to the search engine. Google has argued that these are all product improvements: as Eric Schmidt recently described these tying efforts, "Your search just gets better and better."[2] But in principle, tying may enable an entry path even with a worse product than competing products in the market. Theories in economics and strategy show that tying can provide economies of scope, enable aggressive pricing tactics, and ultimately foreclose competitors' access to overlapping consumers – all of which can confer substantial advantages (Whinston (1990), Nalebuff (2004), Carlton, Gans, and Waldman (2010), Eisenmann et al (2011)). Understanding the extent to which tying is used to launch lower quality products can thus provide important insights on platform strategy and market entry, as well as concerns about rising market concentration among dominant platforms.

We explore this question by empirically investigating Google's decision to tie its reviews product to the search engine. In 2010, Google launched an interface for listing and reviewing local businesses, later called

---

[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

"Google Local".[3] At the time, Google had over 70% market share in search, but was a relative newcomer to the reviews market, where businesses like TripAdvisor and Yelp had amassed millions of reviews. Google faced a strategic decision: should it take advantage of its dominant search position and tie reviews to its search engine, or should it leave its reviews platform to compete with other reviews platforms in organic search results? Google ultimately developed a "Onebox" that sits on top of any organic results and excludes reviews from other platforms. In contrast, Bing (which was in a less dominant position in the general search engine market) partnered with Yelp and TripAdvisor.

Did Google's decision to exclude competitors result in a higher or lower quality product? To explore, we construct and compare two choices in Google's feasible set: results with the Onebox showing only Google reviews, as Google ultimately chose, and results with the Onebox showing the most highly-ranked reviews from multiple platforms as assessed by Google's own search algorithm. We conduct an experiment on UsabilityHub, a platform that provides website testing tools for companies, and recruit 15,166 users from Mechanical Turk, an online labor market where companies can hire freelancers for short online tasks. We ask users how they would search for a local pizza restaurant if they were in a new city and randomly assign them to see different versions of Google search results varying the reviews shown in the OneBox. We then observe where on the screen users click, as a measure of whether they find the Onebox results to be more useful compared to incurring further search costs by scrolling down the page. We validate this outcome with a stated preference measure, by asking users to directly choose which search results they prefer when shown options side-by-side. We then explore mechanisms driving user behavior through robustness tests.

We find that users prefer the Onebox that includes competitor reviews, rather than only Google reviews: they are 5% less likely to click on the Onebox when Google reviews are tied to the search engine -- instead choosing to scroll down to organic search results. The data suggest that the preference is driven by both lower quantity and lower perceived quality of Google reviews.

Looking at observational data on user traffic to Yelp from general search engines, we find the share of Yelp's traffic coming from Google has declined over this period, relative to traffic from Bing and Yahoo (which do not exclude other companies' reviews). Moreover, Google reviews has grown quicker than Yelp and TripAdvisor during the period in which they excluded these (and other) reviews providers. Overall, the evidence suggests that Google's decision to exclude competitor reviews may have helped them gain traction in this complementary market, despite consumers' preference for an easily implementable alternative.

Our analysis contributes to the empirical literature on tying, platform strategy, and market entry. Consistent with theory, our results suggest that platforms may at times be using tying as a strategy to enter an adjacent market even with a lower quality product than existing firms. In the case of Google, tying may have

---

[3] While Google had initially allowed users to write a review in 2007, they made significant investments in local reviews as a standalone product in 2010, with an overhaul of Google Places (information on local businesses), launch of Google Hotpot (local search results for "places" on Google powered by reviews and ratings). See https://techcrunch.com/2011/07/21/google-places-stops-stealing-reviews/ and https://techcrunch.com/2010/04/20/google-places/

enabled an entry path to the reviews market by excluding existing competitors. Many dominant platforms have claimed that tying their own downstream products provides a product improvement for its consumers, and this claim may be true in many markets. However, this claim should be evaluated on a product by product basis, rather than to assume that successful entry necessarily implies a product improvement.

2. Why might tying enable entry without product improvements?

Platform-based markets often rely on complementary applications to solve a chicken-and-egg problem: without an existing base of users, no complementors are interested in providing complements, and without complementary applications, no consumers are interested in adopting the platform (Rochet and Tirole (2003), Hagiu and Spulber (2013)). However, after a platform has taken off, it may have various reasons to offer complementary products directly, such as motivating better quality products by introducing competition, or capturing additional value by imitating successful complementors (Gawer and Cusumano (2002), Gawer and Henderson (2007), Jiang et al (2011), Zhu and Liu (2018)).

Prior literature in strategy suggests many reasons why successful entry may be difficult. Platform markets often involve network effects, as the value of a platform for any given user depends on the number of other users to interact with on the platform (Katz and Shapiro (1985)). With strong network effects, users generally converge on a smaller number of established platforms. In markets with these barriers, research has suggested that entrants are more likely to succeed if they offer genuine improvements in performance (Henderson and Clark (1990), Bresnahan (1999), Evans and Schmalensee (2001)).

An increasingly common entry strategy used by dominant platforms has been tying, which involves entering an adjacent market by integrating one's new product with its primary product (see Carlton, Gans, and Waldman (2010) for a review). Economic theories on tying suggest that tying can provide many benefits, such as economies of scope and better price discrimination (Whinston (1990), Nalebuff (2004)). Eisenmann et al (2011) also theorize that tying may enable an entry path by foreclosing competitors' access to overlapping consumers if products are comparable.

In principle, one might expect that market pressure would lead dominant platforms to tie and exclude competitor products only if this leaves users at least as well off. For example, if tying an inferior product to its search engine drove users to switch to Bing, then Google may not have the incentive to tie, for fear of losing share in its primary market. However, there is increasing recognition that the Internet does not eliminate all search frictions. For example, users appear to display widespread inertia in product choice, with few users switching between browsers or other pre-installed programs.[4] Tying may enable platforms to take advantage of this inertia by nudging users to use tied products that are presented as the default. Furthermore,

---

[4] https://www.theguardian.com/technology/2013/dec/01/default-settings-change-phones-computers

a platform may have the incentive to increase search frictions for its own benefit. Ellison and Ellison (2009) demonstrate that firms can engage in deliberate obfuscation to make it more difficult for consumers to search and acquire information, in order to maintain market power. In spite of potential user costs from lower product quality, if these costs do not lead to sufficiently lower demand, dominant platforms could profit from this decision, through the revenue opportunity from keeping users on its pages, as well as debilitating its rivals in adjacent markets.

Understanding whether tying is facilitating entry with lower product quality can provide important insights on platform strategy and market entry, as well as antitrust concerns that dominant platforms might be harming consumers by strategically nudging them toward lower quality products.

3. Methodology

We explore this question by empirically investigating Google's decision to tie its online reviews to the search engine as it entered the reviews market, by developing a "Onebox" that sat on top of any organic results and excluded competitor reviews. To understand the extent to which Google's tied product provides lower or higher quality, we ran an online experiment on 15,166 users looking at 100 US cities, leveraging UsabilityHub and Mechanical Turk (MTurk).

3.1 Using online labor markets to test platform design changes

The ideal way to run our experiment would be to partner with the platform of interest and run a test directly on their users. However, in many situations, companies may not have the incentive to run these experiments or reveal results even after experimentation. Instead, we leverage the user-testing tool UsabilityHub.

UsabilityHub is a user-testing platform that allows companies to create web design tests to source feedback from users before launching. It offers different tests to measure the effectiveness of a webpage design, such as click tests that ask users where they would click on a website and preference tests that show users two designs side-by-side and ask them which one they prefer. These user tests have been a part of product development at many companies, including Google, Amazon, and eBay.[5] UsabilityHub outsources these tasks to designers on their website or workers on MTurk—an online labor market where employers can hire for short online tasks. Prior studies have explored various uses of MTurk, including surveys, mechanism testing, and pilot design (Kuziemko et al (2015)). As tools like UsabilityHub offer design tests used by real

---

[5] https://usabilityhub.com/customers

companies, it can offer a way to gain insights in instances where the implications of a platform design choice may be important, but the relevant company does not have an incentive to pursue experimentation or allow publication of results.

However, one concern that such tools introduce is that of external validity. Prior studies have suggested that results from online experiments using MTurk are qualitatively generalizable in the context they explore (Horton et al (2011), Pallais (2014)). Horton et al (2011) replicated three classic lab experiments using MTurk, and in each of the experiments, found no significant qualitative differences between the lab and MTurk results. These findings suggest that although point estimates are difficult to generalize between any two settings, online platforms like MTurk might provide a relevant setting to test for directional insights, which we leverage in our paper to test hypotheses about platform choices.

3.2 Experimental Design

Jobs were posted on MTurk with a description stating that the task would take 1 to 3 minutes and pay $0.10, which was comparable to the average length and wage of MTurk tasks (Kuziemko et al (2015)). Limiting applicants to those who spoke English, we recruited a total of 15,166 users. Participants were randomly assigned to a city-treatment condition in a single step using UsabilityHub, which assigned each user to one of three conditions through an independent draw. We selected the hundred largest US cities by population to run our experiment.

All participants were asked to imagine that they had just performed a search on Google, with the instructions, "You are planning a meal with friends at a pizza restaurant in [city name]. You decide to conduct research online. You type "pizza [city, state]" into Google and come upon the following page. What's your first click?" The city and state named depended on which city and state each participant had been randomly assigned out of the full set of one hundred. After confirming that they read the instructions, participants were directed to a screen of a Google search result. We then recorded where on the search page they clicked.

Each participant was randomly assigned to one of three conditions, and we verified that they were generally balanced on demographic traits (**Appendix Table 1**). In the *Google + Information* condition, participants were shown Google's actual search results screen, which showed reviews exclusively from Google highlighted in the Onebox above organic search results, along with information snippets on business hours, location, and price that Google had recently added by the time of the experiment (**Figure 1**). In the *Google Only* condition, participants were shown the same screenshot as *Google + Information* where Google reviews were tied to the search results, but with the information snippets in the Onebox removed. In the *Google + Competitor* condition, participants were shown Google's search results screen modified by a plugin that populated the OneBox with reviews from multiple platforms determined to be the best by Google's organic algorithm. Other than the review information presented in the OneBox, all three conditions provided

identical screenshots of Google search results. The *Google + Competitor* condition represented a counterfactual strategy for the *Google Only* condition, which represented Google's use of tying.[6] Comparing user behavior across these two conditions enables us to explore the extent to which Google's tied product provides lower or higher quality for users compared to showing reviews from multiple platforms.

Across the 100 cities, the number of reviews shown in the *Google Only* and the *Google + Competitor* conditions differed. The *Google + Competitor* screenshots on average showed three times the total number of reviews than the *Google Only* screenshots, with a median result in *Google Only* showing 96.75 reviews compared to 309 reviews in a median result in *Google + Competitor* (**Table 1**). This difference implies that Google discarded about two-thirds of the reviews in the process of excluding competitors by not using its organic search algorithm.

We recorded where on the screen participants clicked. Clicks were coded into one of three categories: those on the OneBox displayed at the top of the screen ("Onebox"), those on any of the organic search results displayed below the OneBox ("Organic"), or all other clicks elsewhere on the screen ("Other"). These clicks provide a measure of users preference for whether they found the Onebox results at the top of the search page to be more useful compared to incurring further search costs by scrolling down the page to organic search results. Once participants finished clicking on the screen, they were presented with an easy addition problem to ensure that they were actively completing the task. For robustness, we ran a separate test to validate user clicks as a measure of user preference by asking users to directly choose between two screenshots when shown options side-by-side.

One city (Corpus Christi, TX) was dropped due to an error that showed an incorrect image for one condition. This resulted in a sample of 99 city experiments across 15,014 participants.

4. Results

Pooling the entire set of city-level search experiments, we find that Google's tying strategy significantly reduced users' probability of clicking in the OneBox by 5%. **Table 2** shows user behavior across all conditions. 42.3% of participants in the *Google Only* condition clicked in the OneBox, compared to 44.5% of participants in the *Google + Competitor* condition. This 2.2 percentage point difference represents a decrease of 5% in clicking on the OneBox when Google ties its own reviews. The higher number of clicks on the OneBox in the *Google + Competitor* condition suggests that users prefer the Onebox with results drawn from

---

[6] Our treatments draw on a Chrome plugin called "Focus on the User," which was constructed before Google modified their Onebox to include the information snippets. To create a clean counterfactual, we strip out these information snippets to create the main "Google" condition.

Google's merit-based organic algorithm (*Google + Competitor* condition) rather than systematic favoring of Google's own reviews.[7]

The difference between the *Google + Information* condition and the *Google + Competitor* condition is statistically insignificant. Given that the *Google + Information* condition provides restaurant address, hours, and prices in addition to the number of reviews, the results suggest that users also respond to genuine product improvements. Furthermore, providing reviews from multiple platforms has approximately a comparable effect size as providing useful information about business traits. As a robustness check, we show two screenshots in the *Google Only* and *Google + Information* condition side-by-side and ask users which one they prefer – and find that 86% of users choose the *Google + Information* condition where information snippets are shown (**Appendix Table 2**). These results suggest that the information snippets that Google eventually added improved the search interface. While these improvements are independent of Google's decision to tie its reviews (they could add information snippets regardless of whether they exclude competitors), this does highlight other design choices – such as adding complementary information – that also improve user experience.

While our results showed that Google had fewer reviews for this search relative to competitors, it is possible that the reviews were of superior perceived quality. To explore, we run two additional experiments where we randomly assign participants to view one of two screenshots, again observing where on the screen they click (**Appendix Figure 1**). In the first additional experiment, we show the same screenshot of the *Google + Information* search result in San Francisco across both conditions, but remove the last digit of the review count in one condition (exploring the effect of the quantity of reviews). In the second additional experiment, we show the *Google + Competitor* result for San Francisco across both conditions, but modify the text "Yelp" and "TripAdvisor" to "Google" in one condition, to imply that these review results are from Google (exploring participants' perceived quality of reviews across platforms). We find that users prefer competitor reviews not only for the quantity, but also for the perceived quality of reviews across platforms (**Appendix Table 2**). Reducing the number of reviews reduces clicks on the Onebox by 7% (4 percentage points), though this difference is not statistically significant. Branding Yelp and TripAdvisor reviews as Google's reduces clicks on the Onebox by 20% (11 percentage points), which is statistically significant at 5%.

Taken together, these results suggest that Google reviews provided both a lower quantity and a lower perceived quality of reviews compared to its competitors, as measured by user preferences. We also find that Google's decision to add information (such as hours and address) increased users' preferences for the product.

---

[7] While most participants clicked on either the OneBox or the Organic results, a small percentage clicked elsewhere on the page. However, there is little variation in the likelihood of clicking outside of the main search areas across the main conditions, suggesting that this is not driving the results. As a robustness check, we reanalyze our main results excluding this set of participants and find qualitatively consistent results.

5. Discussion

Our findings suggest that Google's strategy ultimately may have directed users to less preferred reviews, suggesting that platforms may at times be using tying as a strategy to enter an adjacent market despite having a lower quality product than existing alternatives. Our design has several limitations that suggest directions for future research. First, our experiment is not run on Google's platform, and field evidence would help to shed light on magnitudes and boundary conditions. Second, our results focus on a specific search for a set of cities. There is likely to be considerable heterogeneity across dimensions such as time, search topic, and region – which would shed light on the conditions under which this strategy is more or less likely to lead to a lower quality product.

At the time of Google's entry, competitors TripAdvisor and Yelp had already acquired 30 million and 15 million reviews, respectively. One year afterward, Marissa Mayer (the head of Google's local product) noted that they had acquired 3 million total reviews since launching – already one-fifth of Yelp's number of reviews, which had a six-year head start. She also announced that Google was growing at more than 1 million reviews per month.[8] In comparison, Yelp was growing at 8% of that rate, at approximately 80,000 reviews per month.[9]

This rate of growth in Google reviews appears to have been accompanied by a decrease in traffic to Yelp from Google. **Figure 2** shows that the share of Desktop traffic to Yelp from Google decreased rapidly beginning in 2012, dropping from approximately 85% to 68% by mid-2016. Over this time period, Google's overall share of search does not appear to be decreasing, remaining around 65% (**Appendix Figure 2**). These data alone are not conclusive, given that Google made a number of other changes during this timeframe, such as buying Zagat (which in itself highlights the challenge of launching a review platform) and providing tangible incentives for Google users to write reviews. However, they suggest that Google's entry into the reviews market may have been successful in reducing the traffic to its competitors by strategically pushing their links further down the search page and lowering the probability that consumers will discover them. Furthermore, because of network effects, competitor products have the potential to worsen as a result of being excluded, since lower readership and usage can reduce the number of reviews being written. This possibility may further raise the possibility of deterring entrants if they anticipate Google's actions.

More broadly, our results relate to an emerging discussion about dominant online platforms, which have been drawing increasing scrutiny from policymakers. This scrutiny has come as a part of concerns about rising industry concentration – that large firms may be accounting for higher shares of industry profits relative to their investment. On the one hand, higher profits may simply be a result of superstar firms with better ability to innovate and offer better products. On the other hand, concentration may also arise from

---

[8] http://blog.kelseygroup.com/index.php/2011/03/12/googles-marissa-mayer-at-sxsw-hotpot-heating-up/
[9] http://www.theatlantic.com/technology/archive/2011/07/infographic-the-incredible-six-year-history-of-yelp-reviews/242072/

anticompetitive forces where incumbent firms are increasingly able to prevent rivals from entering and growing, even when they can offer better quality products to consumers. While there has been growing evidence that industry concentration may be on the rise, there has been less insight on whether this consolidation is resulting in better quality products (Autor et al 2017, de Loecker and Eeckhout 2017). Our findings suggest early evidence that dominant platforms may, at times, be degrading products for strategic purposes, such as excluding competitors in adjacent markets that they are looking to enter or grow in.

References

Autor, David, David Dorn, Lawrence Katz, Christina Patterson, and John Van Reenen (2017). The Fall of the Labor Share and the Rise of Superstar Firms. No. 23396. National Bureau of Economic Research, Inc.

Bresnahan TF (1999). "New modes of competition" in Competition, Innovation and the Microsoft Monopoly. Eisenach JA, Lenard TM (eds). Kluwer: Norwell, MA; 155–208.

Carlton, D., J. Gans, and M. Waldman (2010). Why tie a product consumers don't use? American Economic Journal: Microeconomics 2(August 2010): 85-105.

De Loecker, Jan, and Jan Eeckhout (2017). The Rise of Market Power and the Macroeconomic Implications. National Bureau of Economic Research.

Edelman, Benjamin (2015). "Does Google Leverage Market Power Through Tying and Bundling?" *Journal of Competition Law & Economics* 11(2): 365–400.

Eisenmann TR, Parker G, Van Alstyne M (2011). "Platform Envelopment." Strategic Management Journal 32, no. 12 (December 2011): 1270–1285.

Ellison, Glenn, and Sara Ellison (2009). Search, Obfuscation, and Price Elasticities on the Internet. Econometrica 77(2): 427-452.

European Commission (2017). Press Release, June 27, 2017. http://europa.eu/rapid/press-release_IP-17-1784_en.htm.

Evans D, Schmalensee R. 2001. Some economic aspects of antitrust analysis in dynamically competitive industries. NBER Working paper no. 8268. National Bureau of Economic Research: Cambridge, MA.

Gans J.S., & Stern S. (2003). The product market and the market for "ideas": Commercialization strategies for technology entrepreneurs. *Research Policy,* 32(2), 333–350.

Gawer A., & Cusumano M.A. (2002). *Platform leadership: How Intel, Microsoft, and Cisco drive industry innovation.* Cambridge, MA: Harvard Business School Press.

Gawer, A., & Henderson, R. (2007). Platform owner entry and innovation in complementary markets: Evidence from Intel. *Journal of Economics & Management Strategy,* 16(1), 1–34.

Greene, Hillary (2015). "Muzzling Antitrust: Information Products, Innovation, and Free Speech." 95 Boston University Law Review 35.

Hagiu, A., & Spulber, D. (2013). First-party content and coordination in two-sided markets. *Management Science,* 59(4), 933–949.

Henderson R, Clark K. 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. Administrative Science Quarterly 35: 9–30.

Horton, J., D. Rand, and R. Zeckhauser (2011). The online laboratory: conducting experiments in a real labor market. Experimental Economics, Springer, 14(3): 399-425.

Jiang B., Jerath K., & Srinivasan K. (2011). Firm strategies in the "mid tail" of platform-based retailing. *Marketing Science,* 30(5), 757–775.

Katz, ML, Shapiro C (1985). Network externalities, competition, and compatibility. American Economic Review 75: 424-440.

Kuziemko, Ilyana, Michael Norton, Emmanuel Saez, and Stefanie Stantcheva. (2015). "How Elastic are Preferences for Redistribution: Evidence from Randomized Survey Experiments." American Economic Review 105 (4): 1478-1508

Nalebuff B. (2004). Bundling as an entry deterrent. *Quarterly Journal of Economics* **119**: 159–187.

Pallais A. (2014). Inefficient Hiring in Entry-Level Labor Markets. American Economic Review 104 (11): 3565-3599.
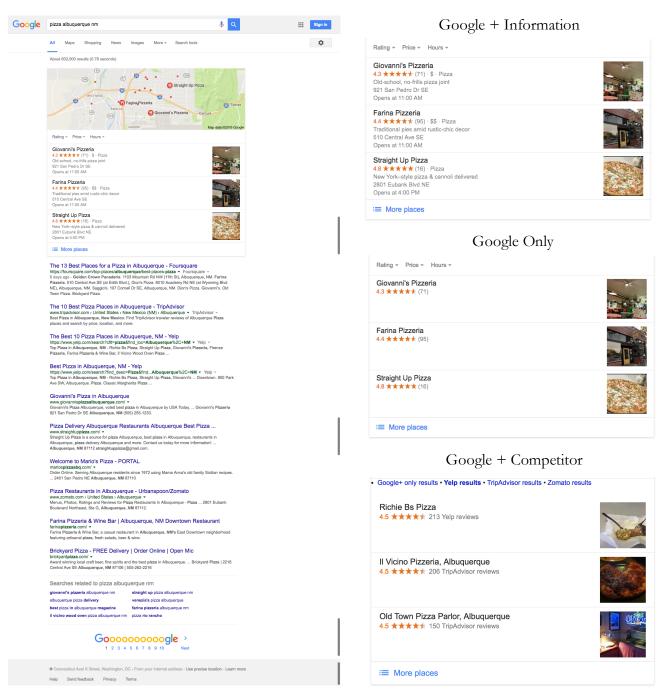
Rochet, Jean-Charles, and Jean Tirole (2003). Platform Competition in Two-sided Markets. Journal of the European Economic Association 1(4): 990-1029.

Whinston M. (1990). Tying, foreclosure and exclusion. *American Economic Review* **80**: 837–859.

Wu, Tim Sebastien Couvidat, Daniel Frank, William Seltzer (2015). Does Google Content Degrade Google Search? Working Paper.
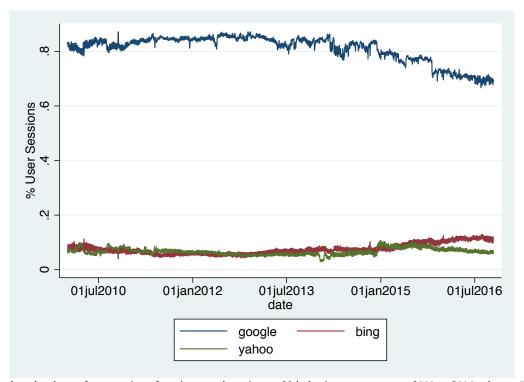
Zhu F, Liu Q (2018). Competing with complementors: An empirical look at Amazon.com. *Strat Mgmt J*. 39:2618–2642. https://doi.org/10.1002/smj.2932

Figure 1   Experimental Conditions

At left is the full screen shown to participants in the Google + Information condition: the unaltered Google search results, including information snippets for each OneBox entry. From top to bottom on the right are the OneBox for each of our three conditions, which only varied the Onebox.

Figure 2   Breakdown of Yelp Desktop Traffic by Search Engine

This figure plots the share of user sessions from key search engines to Yelp business pages across 2010 to 2016, where a "session" is defined (by Google Analytics) as a collection of interactions from a specific user during a defined period of time. A session is defined to end after 30 minutes of inactivity.

Table 1   The Distribution of Reviews across Cities

	Google Only	Google + Competitor
Total Reviews	335.85	1063.79
	(31.79)	(102.07)
Min Reviews	58.28	160.77
	(6.42)	(17.97)
Median Reviews	96.75	308.98
	(8.75)	(31.06)
Max Reviews	180.83	594.05
	(18.32)	(67.49)
Number of Yelp Results (out of 3 shown in OneBox)	0	1.85
		(0.10)
Number of TA Results (out of 3 shown in OneBox)	0	0.87
		(0.093)
Observations	99	99

This table compares the distribution of reviews shown in screenshots of Google Only and Google + Competitor conditions. Means and standard errors of the mean in parentheses are displayed.

Table 2   **User Clicks by Experimental Condition**

	Google Only	Google + Competitor	Google + Information
OneBox	0.42	0.44**	0.45***
	(0.01)	(0.01)	(0.01)
Organic Search Results	0.55	0.52	0.52
	(0.01)	(0.01)	(0.01)
Other	0.03	0.04	0.03
	(0.00)	(0.00)	(0.00)
Observations	5011	5006	4997

This table shows the percentage of participants that clicked within the OneBox, Organic Search Results, and elsewhere on the page (Other) across all three experimental conditions. The standard error of the mean is presented in parentheses. Stars indicate significant statistical difference from the Google Only condition.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$