Does Advertising Serve as a Signal? Evidence from a Field Experiment in Mobile Search*

Navdeep S. Sahni
Assoc. Prof. of Marketing
Stanford GSB

Harikesh S. Nair
Prof. of Marketing
Stanford GSB

This draft: Dec 14, 2016. Previous drafts: June 3, Feb 1, 2016

Abstract

In a field experiment, we demonstrate that advertising can serve as a signal that enhances consumers' evaluations of advertised goods. We implement the experiment on a mobile search platform that provides listings and reviews for an archetypal experience good, restaurants. In collaboration with the platform, we randomize about 200,000 users in 13 Asian cities into exposure of ads for about 600+ local restaurants. Within the exposure condition, we randomly vary the disclosure to the consumer of whether a restaurant's listing is a paid-ad. This enables isolating the effect on outcomes of a user knowing that a listing is sponsored – a pure signaling effect. We find that this disclosure increases calls to the restaurant by 77%, holding fixed all other attributes of the ad. The disclosure effect is higher when the consumer uses the platform away from his typical city of search, when the uncertainty about restaurant quality is larger, and for restaurants that have received fewer ratings in the past. On the supply side, newer, higher rated and more popular restaurants are found to advertise more on the platform. Taken together, we interpret these results as consistent with a signaling equilibrium in which ads serve as implicit signals that enhance the appeal of the advertised restaurants. Both consumers and firms seem to benefit from the signaling. Consumers shift choices towards restaurants that are better rated (at baseline) in the disclosure condition compared to the no disclosure condition, and advertisers gain from the improved outcomes induced by disclosure. The results also imply that search-platforms would gain from clear sponsorship disclosure, and hold implications for platform design.

Keywords: Informative advertising, signaling, field-experiments, restaurants, mobile, paid-search, platforms.

*The views discussed here represent that of the authors and not of Stanford University or Zomato. We thank Daisy Dai, Dafna Goor, Wes Hartmann, Randall Lewis, Puneet Manchanda, Sridhar Moorthy, Thomas Otter, John Roberts, Imran Rasul, Devesh Raval, Catherine Tucker, Kosuke Uteke, Birger Wernerfelt; and Dan Ackerberg, Matt Gentzkow and Sanjog Misra in particular for helpful comments and suggestions. Comments from participants at the 2015 Marketing Science, 2016 SICS, QME conferences; seminars at Goethe-Univ. Frankfurt, HBS, HEC-Paris, MIT-Sloan, Stanford, Univ. of Colorado, Univ. of Mannheim, and Univ. of Michigan are also appreciated. We thank Shuo Xie and Charles Zhang for their excellent research assistance. The usual disclaimer applies. Please contact Sahni (`navdeep.sahni@stanford.edu`) or Nair (`harikesh.nair@stanford.edu`) for correspondence. The authors declare they have no commercial relationship with Zomato, and the research is conducted with no restrictions on the publishability of the findings from the field experiment.

1 Introduction

Despite its prominent influence on how social scientists think about the role of advertising, Nelson's (1970, 1974) celebrated idea that advertising can serve as a signal of product quality has proven difficult to test empirically. Consequently, more than 40 years since it was originally articulated, credible empirical evidence in favor of the signaling view of advertising has remained rare. Understanding whether advertising actually plays a signaling role and how this role materializes is important to assess the welfare consequences of advertising. If advertising can serve as a signal, it can improve the efficiency of markets with search frictions by helping buyers and sellers communicate. It also has implications for firms targeting their ads. If ads convey demand-enhancing information about products beyond informing users of their existence and product attributes, ads could be targeted to users already aware of the product, when there is considerable uncertainty about the products' quality. This paper describes a field experiment implemented in collaboration with a large restaurant search platform that enables a test of the "signaling hypothesis". We find results consistent with signaling.

Nelson's work postulates an indirectly informative view of advertising, suggesting that one role played by advertising is to signal to market participants that the advertising firm is of high quality. Nelson's suggestion is formalized in several well-known canonical models that followed (e.g., Kihlstrom and Riordan 1984 and Milgrom and Roberts 1986). In these models, consumers are *ex ante* uncertain about the quality of a good, which is revealed to them upon consumption. Firms with higher quality benefit more from advertising, for example, when high quality firms obtain more repeat purchase after consumption. A separating equilibrium is achieved when (1) the gain from the repeat purchase relative to the cost of advertising is higher for high quality firms at optimally chosen prices, and (2) lower quality firms do not gain from mimicking the strategies of the high quality firms. In the equilibrium, the *act* of advertising itself conveys information about quality. Because direct claims about quality in an ad cannot be verified prior to purchase of an experience good, the indirect way by which advertising reveals quality is more relevant. Therefore, costly advertising serves as a credible signal.

Subsequently, several empirical studies have built on Nelson's ideas and attempted to test the signaling hypothesis. One group of studies on the "supply-side" have investigated the predicted equilibrium associations amongst the key components of the model – advertising, prices and quality. Another set of studies on the "demand-side" have used micro-data on consumer exposure to advertising to explore patterns suggestive of signaling. Researchers using either strategy have faced significant challenges in establishing a signaling role for advertising.

A supply-side empirical test of the theory requires (a) a measure of quality that is being signaled by the firm, and (b), a way to match its association with firms' observed behavior in a manner that is falsifiable by the theory. Both these steps are difficult. The main challenge in step (a) is in obtaining a measure of quality as conceptualized in the theory, which is a construct that is observed to the firm but is unobserved to the consumer. This is non-trivial when the consumers' actual information sets are unobserved. Even if a researcher obtains a measure of quality unavailable to consumers, it is hard to rule out that it is uncorrelated with some component of consumers' unobserved information-sets. By implication, any covariation between such a quality metric and advertising could reflect the mediation of such omitted variables. Thus, correlations between publicly observed metrics of quality − like peer ratings, consumer reports or time spent at the firm − with advertising actions, present a weaker test of the theory (e.g., see Archibald et al. 1983; Kwota 1984; Caves and Greene 1996; Thomas et al. 1998; Kirmani and Rao 2000; Horstmann and MacDonald 2003; and Horstmann and Moorthy 2003 for further discussion).

The challenge in step (b) is that observed patterns of firm-level advertising can often be explained by reasons other than those postulated by signaling. For instance, the canonical model predicts that ad-intensity starts high for new goods and falls for established goods as information about unobserved quality diffuses in the market. A subset of studies focus on whether observed data are consistent with these life-cycle predictions (e.g., Tellis and Fornell 1988, Horstmann and MacDonald 2003). However, such patterns could also be produced by changing consumers' awareness about the product, changes in competitive intensity due to entry, changes over time in the costs of advertising in media-markets, all of which have to be ruled out to establish the empirical relevance of the signaling mechanism. Additionally, different models of signaling can predict different patterns of movement over time in prices and advertising depending on underlying assumptions imposed, and these underlying assumptions may be hard to test.[1]

Testing whether advertising signals quality on the "demand-side" by directly exploring consumer-level response to ads is also difficult. The main challenge is to disentangle the signaling effect from other effects of advertising. This concern is severe in data where the ad-creative or all relevant aspects of ad-content that impacts on consumer behavior are not observed. Hence, it is usually impossible to say whether consumer response to advertising is due some aspect of the ad's message that is unobserved by the econometrician, or due to the ad reminding the consumer of the product, or due to the consumer's knowledge that the firm has advertised. Only the latter is a pure signaling effect. A separate concern is the endogeneity of advertising exposure due to the targeting of advertising by firms or from user self selection into viewing ads. A final difficulty with consumer-level analysis involves issues of statistical power arising from the large noise-to-signal ratio of ad-effects at the individual level, requiring large datasets that are difficult to collect.[2]

[1] For example, Horstmann and MacDonald (1994, 2003) predict different life-cycle patterns for prices and advertising in a signaling model in which consumers learn only partially about product quality with experience; and Hertzendorf (1993) predicts advertising will be persistent over time when consumers observe ads with noise.

[2] Ackerberg's (2001) analysis outlines how one may distinguish between informative versus persuasive effects of advertising using such consumer-level data, by examining whether experienced consumers respond to advertising (which is consistent with persuasive effects). However, Ackerberg does not distinguish between the direct versus indirect channels of

We design a field-experiment that attempts to address these difficulties in a more direct way compared to the previous literature. Our experiment is implemented in collaboration with Zomato, a worldwide restaurant-search portal. Zomato provides an online platform for consumers to search and browse through information on restaurants in local markets. The experiment is implemented on the platform's mobile app, and introduces search advertising on Zomato's mobile platform. The experiment randomizes users into conditions in which they see search ads for local restaurants. The conditions are similar on all dimensions except the manner in which advertising is disclosed. Users in the first see the advertiser's listing without any disclosure, whereas users in the second see the listing with an indication disclosing that the listing is an ad. Therefore, users in both conditions are exposed to the same ads for the same restaurants with and without disclosure that the ads are paid for by the advertiser. In the data, we observe the user-level browsing behavior, and the restaurants they call.[3] This design enables assessing the effect of a consumer's knowledge that the firm is advertising, separately from the ads' effect on the awareness of the existence of the firm, and from the effect of the content of the ads. Thus, it facilitates a demand-side test of the signaling effect.

The empirical setting has advantages as a field laboratory to assess signaling. First, signaling is most relevant in markets for experience-goods, in which consumers are information-constrained, quality sensitive, have uncertainty about the product prior to consumption, and show repeat purchase for firms revealed to have high quality from their visitation. Restaurants are examples of such goods. Second, search ads on restaurant search platforms are matched to user intent and served in response to users who are searching for information about the goods being advertised, which reduces the chance the ads are annoying and will be skipped. This makes detection of the signaling effect more likely. Third, the mobile application environment has the advantage that one obtains a persistent user identifier defined within a closed system that logs engagement with the ads as well as demonstrated interest in the product. This makes randomization and behavior-tracking at the individual user-level possible. Finally, a large number of users (around 200,000) and advertisers (around 600) enables exploring differential patterns of response along dimensions of consumer- and restaurant-heterogeneity that facilitate additional tests of the signaling theory, and serve as consistency checks on the effect.

The experimental design incorporates features to assess signaling while minimizing experimental interference. One feature derives from the fact that signaling is an equilibrium phenomenon, requiring measurement of the response to advertising disclosure without disturbing equilibrium beliefs that con-

informative advertising on account of the reasons above. Related consumer-level studies include Anand and Shachar (2011); Goeree (2008); Terui et al. (2011) who estimate models where advertising has an informative role by affecting consumer's information sets; Shum (2004) and Erdem et al. (2001) who study the impact of advertising on consumer price sensitivity; and Homer (1995) and Kirmani (1990, 1997) who report on relationships between consumer perceptions of advertising and quality in a lab setting.

[3]The call action we observe occurs after a consumer visits the restaurant's page. As detailed later in the paper, calls represent a consumer's intent to order food from a restaurant.

sumers and restaurants hold about advertising on Zomato. Showing an ad for a *random* restaurant in a local market to a consumer would be problematic if that restaurant would not have advertised in equilibrium.[4] To avoid this, the experimental design here shows ads to users only for restaurants that *choose* to advertise on the platform. By manipulating disclosure for these restaurants, we avoid this experimental interference and measure the impact on a relevant subpopulation for which the effects are meaningful. Another aspect of experimental design is to ensure that the conditions do not induce "Hawthorne effects" or "randomization bias" that confound the signaling effect. In particular, the treatment condition of showing an advertised listing without disclosure to consumers that it is paid for, is harder to conceptualize in some kinds of media like mainstream print or TV. The presence of such an ad would seem odd and may evoke a corresponding consumer response. This condition is more natural as part of an experiment on a search platform, where ads without disclosure look exactly like "organic" listings. Further, we implement our experiment in a manner such that the ads introduced by our experiments are consistent with the user's expectations from search. More details on this are discussed later.

Our main results report response to disclosure based on users' first advertising exposures.[5] We find that disclosing to a consumer that a listing on the platform's search results is an ad increases calls to the restaurant by 77% relative to no disclosure. This represents the causal effect of disclosure, because it keeps all aspects of the ad, including its content and position on the listings page fixed, and holds consumer type fixed in the comparison on account of the randomization. This effect is also large relative to the impact of other informational attributes. An average advertiser in our data obtains roughly the same benefit by disclosing sponsorship as a two decile increase in the number of its ratings on the platform. Analyzing calls conditional on visiting a restaurant's page, we find that conditional call-rate is highest for the subset of users who visit the restaurant's page on the platform while seeing that the listing is a paid ad, and 76% higher compared to showing the listing in the same position but without disclosure. Further, we find that exposure to the listing mainly drives visits to the restaurant's page, whereas disclosing that the listing is an ad helps convert the page-visits into calls to the restaurants. These results suggest that the ad-disclosure implicitly conveys a cue to exposed users, which increases the restaurant's appeal amongst them, and improves their evaluation of the advertised product.

Exploring further, we test whether the effect of the disclosure is higher for subpopulations of users that have more uncertainty about the restaurants they are searching for, as the signaling hypothesis

[4] As an extreme example, suppose we pick a non-advertising restaurant and show ads about it on the app, and it turns out the restaurant is closed on the day the ads are served because the owner is on vacation. Consumers who call the restaurant on seeing the ad and find no answer may change their views about advertisements on Zomato and lose trust in the platform. These consumers may also know *a priori* from other sources (word-of-mouth, its low ratings) that the restaurant is of low quality. Seeing this restaurant advertised may cause the user to revise his beliefs downward about the quality of restaurants advertising on Zomato. If the restaurant owner eventually hears complaints from consumers that his restaurant was closed though it was advertised on Zomato, he may change his views about advertising on the platform.

[5] We do this to avoid bias from within-user feedback effects that may be problematic in longitudinal analysis (described in more detail later in the paper).

would suggest. Using subsets of users that search in cities different from the cities where they usually search as a way to operationalize the uncertainty, we find that users searching in a different city visit listed restaurant's pages at a significantly higher rate if they are advertised with disclosure. We also see that that restaurants that have been rated fewer times on the platform − presumably, those about which consumers *a priori* have more uncertainty − benefit more from the ad-disclosure. These effects are consistent with predictions from signaling theory. Exploring whether consumers are made better or worse off under disclosure, we find that consumers' choices shift systematically towards restaurants that are better rated (at baseline) in the disclosure condition compared to the no disclosure condition. This suggests that consumers overall benefit from the signaling.

Finally, exploring patterns of covariation on the supply-side, we find that restaurants with higher appeal to consumers (with better ratings), with higher prices and which are newer (presumably ones that consumers have more uncertainty about) are more likely to advertise on Zomato. We also find that restaurants that chose to advertise on Zomato during the experiment are likely to receive higher ratings in next the two years, compared to those that did not advertise. This is consistent with these restaurants having higher unobserved quality that is revealed over time through actual consumer experiences. These supply-side patterns, similar to those reported in some of the past literature, are broadly consistent with a signaling equilibrium the demand-side results suggest. Estimating the economic benefits to ad-disclosure, we find that the platform gains roughly about USD2.7 million in incremental annual ad-revenue from disclosing the ads. We broadly conclude that the signaling benefits of advertising on the platform is significant.

This paper is related broadly to an empirical literature on online advertising (e.g., Manchanda et al. 2006; Yang and Ghose 2010; Chan et al. 2011; Yao and Mela 2011; Rutz and Bucklin, 2011; Johnson 2013; Lewis and Reiley 2014), and more pointedly, to a burgeoning literature that leverages experimental or quasi-experimental variation on search platforms to address issues related to search advertising (e.g., Goldfarb and Tucker 2011; Nosko et al. 2015; Sahni 2015; Narayanan and Kalyanam 2015; Ursu 2016). To our knowledge though, none have focused explicitly on the role of signaling or disclosure. Outside of advertising, our analysis is closest to the empirical literature in education that has tested for the signaling role of education (for e.g., Tyler et al. 2000 in particular for a quasi-experiment; and Weiss 1995 for a survey of empirical work). The difficulties in that literature in distinguishing between human capital versus signaling ("sheepskin") explanations for education have parallels to the difficulties here in distinguishing between direct and indirect informative effects of advertising. Finally, this study is also related to recent empirical papers that have investigated signaling via round-numbered asking prices by sellers on `eBay.com` (Backus et al. 2016) and via posted interest rates by loan seekers on `Prosper.com` (Zhang and Liu 2012; Kawai et al. 2014).

The design and data have two shortcomings. First, the design might understate the size of the signaling effect. If the search engine's algorithm is working well, a searching consumer will believe that organic listings are a good match for his taste. Hence, what is measured as the difference in outcomes between an ad with disclosure versus without, is the sum total of a (potentially negative) "not-organic" effect and a (positive signaling-derived) "is-ad" effect. The not-organic effect reduces the estimated effect. So our test of signaling is conservative, but gives a possible underestimate of the signaling effect. Second, our main outcome variable − calls to the restaurant − represents only a proxy for restaurant demand. We do not have access to actual demand/expenditures. We present later in the paper supporting data suggesting that this is a reasonable proxy in the experimental markets we consider; nevertheless, in the absence of actual demand data, this remains a limitation.

The rest of the paper is outlined as follows. The next section briefly describes our empirical strategy. The following section describes the Zomato platform and details of our field experiment. The sections after describe market level advertising patterns, the main results from the experiment, assessment of heterogeneity, robustness and aggregate effects. The last section concludes.

2 Empirical Strategy

To motivate the empirical analysis, consider the behavior of a consumer i who is contemplating buying from a seller of an archetypical experience good like a restaurant, a product repair or other personalized service vendor. Let r index the provider, q_r the provider's appeal to the consumer and $b_i(q_r)$ represent the consumer's ex-ante belief about q_r. For the purpose of this discussion, assume that i is aware of the existence of r. Prior to a possible purchase, suppose the consumer is exposed to information $\mathcal{I}(x_r, a)$ about the seller. $\mathcal{I}(.)$ includes a vector of attributes denoted by x_r which we will refer to as "content," that informs the consumer about the seller's appeal. In the context of a search platform, x_r includes attributes like the seller's average rating, reviews by other consumers and the position on the page of listings. $\mathcal{I}(.)$ is also indexed by a binary variable a that denotes whether the information is part of an advertisement or not. When $a = 1$, the consumer realizes that that information is paid for by the vendor (i.e., "paid listing"); when $a = 0$, the information is presented by the platform (i.e., "organic listing"). When $a = 1$, advertising plays a signaling role. On receiving the information, the consumer's prior beliefs are updated to a posterior $b_i(q_r|\mathcal{I}(x_r, a))$, which drives a purchase or pre-purchase action, $y_i[b_i(q_r|\mathcal{I}(x_r, a))]$. The causal effect of a on y represent a signaling effect of advertising.

As noted, empirical researchers face two main difficulties in assessing the signaling effect of advertising from field data on the demand side. The first is related to self-selection: the set of consumers who get exposed to advertising are typically different from the set of consumers who do not, either because ads are targeted to users who are more likely to respond to them, or because users who search for a particular service and subsequently get exposed to the ad are more likely to prefer the vendor than users who do not. Thus, the comparison of the behavior of consumers who are exposed to the information to those who are not,

$$\triangle_{1ik} = y_i \left[b_i \left(q_r | \mathcal{I} \left(x_r, a = 1 \right) \right) \right] - y_k \left[b_k \left(q_r | \emptyset \right) \right]$$

does not deliver a causal effect valid as a test of the signaling hypothesis. A second difficulty raises because information in an ad typically arrives as a bundle of x_r and $(a = 1)$, which makes it difficult to separate the effect of the content of the ad from the fact that the seller is advertising. Because of this problem, even within-comparisons that may be conceivable in observational data, viz.,

$$\triangle_{2i} = y_i \left[b_i \left(q_r | \mathcal{I} \left(x_r, a = 1 \right) \right) \right] - y_i \left[b_i \left(q_r | \emptyset \right) \right]$$

representing the incremental effect of seeing the ad compared to not seeing the ad do not deliver the signaling effect of advertising, because they do not hold content constant.

The Empirical Strategy in this Paper We compare the behavior of a *similar* set of individuals exposed to the *same* content either as part of a paid advertisement by the seller or not. That is, we construct the comparison,

$$\triangle_i = y_i \left[b_i \left(q_r | \mathcal{I} \left(x_r, a = 1 \right) \right) \right] - y_i \left[b_i \left(q_r | \mathcal{I} \left(x_r, a = 0 \right) \right) \right]$$

representing the causal effect on buying behavior of the consumers' knowledge that the content is paid for by the seller. This contrast is implemented by randomizing the same content to a similar sets of users with and without revealing the content is paid for by the advertising restaurant. By randomizing a across users, we are able to estimate an average treatment effect of the ad across all participating restaurants on the platform,

$$\triangle = \mathbb{E}_r \triangle_r = \mathbb{E}_r \left[\mathbb{E}_i y_i \left\{ b_i \left(q_r | \mathcal{I} \left(x_r, a = 1 \right) \right) \right\} - \mathbb{E}_i y_i \left\{ b_i \left(q_r | \mathcal{I} \left(x_r, a = 0 \right) \right) \right\} \right]$$

to assess the effect of signaling. Further, we also assess heterogeneity in these treatment effects to test for patterns of differential take-up that are predicted by the theory. In the next two sections, we describe the empirical setting, the field-experiment and the control and treatment conditions in more detail.

3 Application Setting and Field Experiment

3.1 Zomato.com

Pursuant to the acquisition of `urbanspoon.com` in 2015, the Zomato platform hosts searchable listings on about 1.4 million restaurants in 22 countries, counting approximately 90 million visits each month across its website and mobile applications. As comparison, Yelp, the market leader for online listings of local businesses (not just restaurants), is present in 32 countries and is visited by approximately 142 million users monthly (Yelp.com 2015). Compared to competing restaurant platforms, Zomato is differentiated by having a strong presence in South Asia and the Middle East, in large cities traditionally underserved by online restaurant search platforms, and by having a more comprehensive and reliable database about restaurant attributes than traditional crowdsourced content platforms.[6] In 2014, 30 million unique users used Zomato every month to search for restaurants.

The Zomato platform is accessible via an internet website or via a mobile application available on Android or Apple iOS smartphones. The website was launched in July 2008, the Android app in Feb 2010 and the iOS app in May 2011. On accessing the platform, users can search for restaurants by inputing a set of text-based keywords (for example, some combination of the restaurant name, location, cuisine or other attribute), or by searching by pre-established categories (for example, a list of recommended restaurants in the users' location that are open for service at the time of search). A variety of filters based on geographic location, cuisine, and intention (as defined as "home-delivery, dine-out or night-life") can be applied as part of the search as desired. In response to the search, a list of restaurants that are determined by the platform to be relevant to the search criteria are displayed to the user on a search results page. Following the online search literature, we refer to these as "organic" listings. If an advertiser or a set of advertisers have contracted with Zomato to show ads for the search criteria and filters used, a set of advertised listings are also displayed on the search results page. We refer to these as "paid" listings or simply "ads". The user can click on any of the displayed listings and subsequently browse a set of pages containing additional information specific to the listed restaurant. At any point of time, thousands of restaurants advertise on the Zomato website. Ads were served on the website from its launch days, but were officially launched on the mobile apps only in November 2014. Our experiment (described in more detail below) is implemented in Aug-Sept 2014 on the Android version of the Zomato app. Thus, in the pre-experiment period, users are exposed to ads on the Zomato website, but see no ads on their mobile apps.

App Search Experience To understand how the experiment works, we describe a user's pre-experiment search experience on the Zomato Android mobile app in more detail. Figure (1) shows a series of snapshots of a search session on the app. Applying a search criteria takes the individual to a search results page that displays listings that satisfy the user's criteria. The search results are sorted by the search

[6] From TechCrunch (2015): "Zomato started in 2008 as a supercharged portal for restaurant search that went beyond basic names and addresses. Zomato staff would visit venues, collecting menus and photos that would be scanned and input into Zomato's larger database (think Google Maps' roving cars but for restaurants) which in turn would be used to power searches not only for certain restaurants but places where consumers could go for very specific dishes, for example. This filled a niche: smaller and independent venues are not always up to date with their online presence (many don't even have websites today) and this provided a way to find them on the web. It also helped differentiate Zomato from the likes of Yelp and others that looped in crowdsourced information, which can be hard to verify as not being biased and more generally keep up to date."

engine's measure of "popularity" of a restaurant, unless the user specifies an alternate sorting criterion. Each listing on the search results page presents the name of the restaurant, its cuisine, its location, a flag for whether or not the menu is available, the number of photos available, along with the number and average value (on a five point scale) of ratings given to the restaurant by past Zomato users (see Figure (1)). Clicking on a listing takes the user to an "info" page that provides more information about the restaurant as shown in Figure (2). This page has tabs that allow the user to view the restaurants' menu, its photos, its location on a map and to add reviews and photos if desired. It also shows the restaurants' rating information and allows the user to browse through reviews provided by other users on the platform. Users interested in the restaurant can click on a button on this page to call the restaurant directly from their mobile phone. All user actions are tracked on the app.

Advertising on Zomato website As mentioned, there is no mobile advertising on Zomato prior to the experiment, but restaurants actively advertise on the Zomato.com website. To advertise, restaurants contract with Zomato to buy ads for a specific set of search criteria. A typical contract outlines the location specified by the user that initiates the search, the category of the search, the day and the position on which the advertised listing will be displayed. To specify location, Zomato divides each city it operates in into a set of non-overlapping (approximately) 5 miles $\times$ 5 miles zones and allows restaurants to buy ads at any level of aggregation over these zones. To specify search intent, Zomato specifies three search categories viz., "home-delivery," "dine-out," or "night-life" as described previously. The contracting advertiser can specify the criteria for his ads narrowly (e.g., "shows ad at position X for any user who searches for home-delivery with location specified as Y on Fridays for the next 2 weeks"), or broadly ("show ad at position X for any user who searches with location specified as $Y + Z$ on Fridays for the next 2 weeks"), depending on its needs. When a search is initiated on the Zomato website that satisfies a desired criteria, the contracted ad is shown. If more than one advertiser has contracted for that search criteria, Zomato uses its own proprietary algorithms to vary across users the positions at which these advertisers will be shown that day. Advertisers can negotiate separately to be exclusively featured at a given position for a given search criteria to avoid this. All ads are local.

3.2 Field Experiment

The field experiment adds search ads into the Zomato mobile platform. The collaboration with Zomato was motivated by the firm's desire to assess the viability of advertising on its mobile platform through pre-launch "A/B" testing. The experiment starts in July 2014, when a new update of the android app was launched with the experiments encoded in it and made available on Google Play app store. The 13 cities in our data are: (India) Delhi, Kolkata, Mumbai, Bangalore, Pune, Hyderabad, Chennai, Lucknow,

Figure 1: Search Flow for User on Android Mobile App

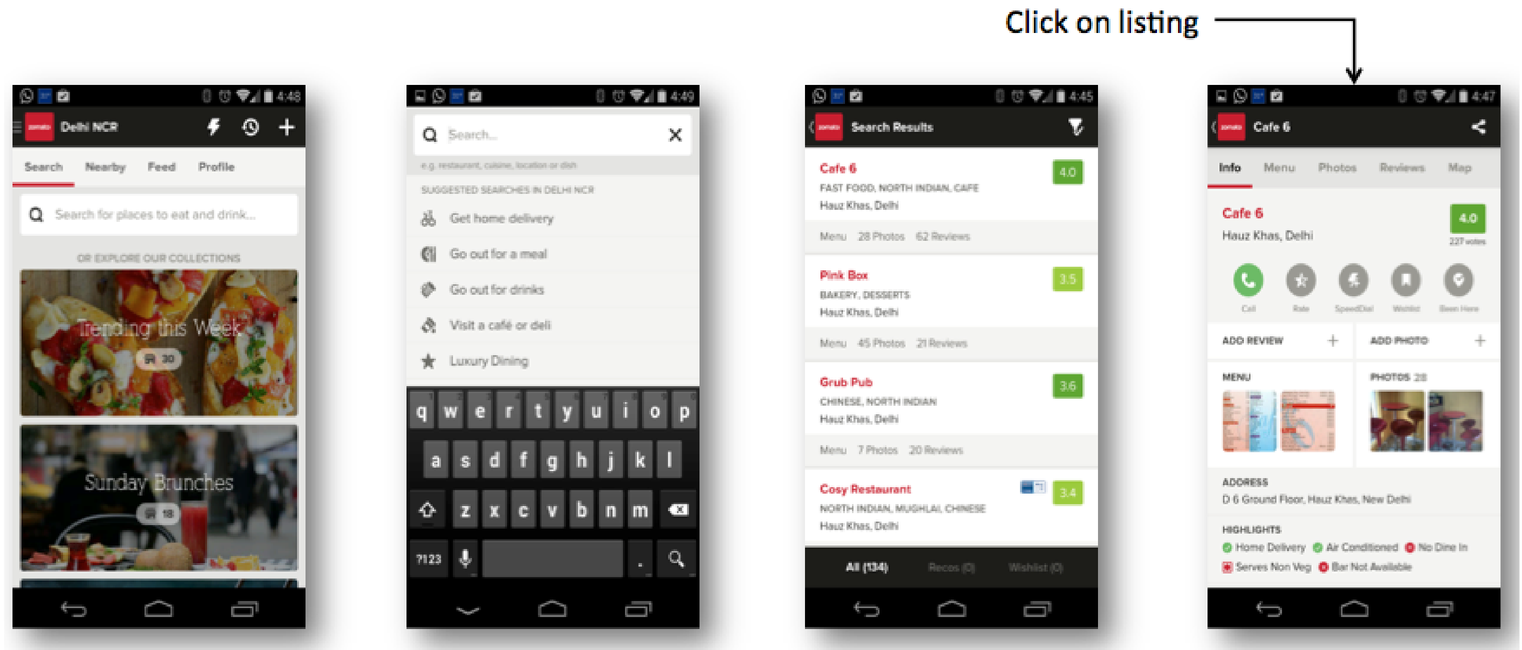

Notes: The figure shows a snapshot of the search experience of a user on the Zomato Android app. The user searches for restaurants by inputing a set of text-based keywords (including for example, some combination of the restaurant name, location, cuisine or other attribute), or by searching by pre-established categories (like requesting a list of recommended restaurants in the users location that are open for service at the time of search), while applying a variety of search filters (including limiting the search by geographic location, cuisine, and intention as defined as "home-delivery, dine-out or night-life"). In response to the search, a list of restaurants that are determined by the platform to be most relevant to the search criteria are displayed on a search results page. The user can click on any of the displayed listings and is led to an info page that contains information specific to the listed restaurant including users reviews, the restaurant's menu, photos and a map.

Figure 2: Information Screen for User Upon Clicking on a Restaurant Listing on Android Mobile App

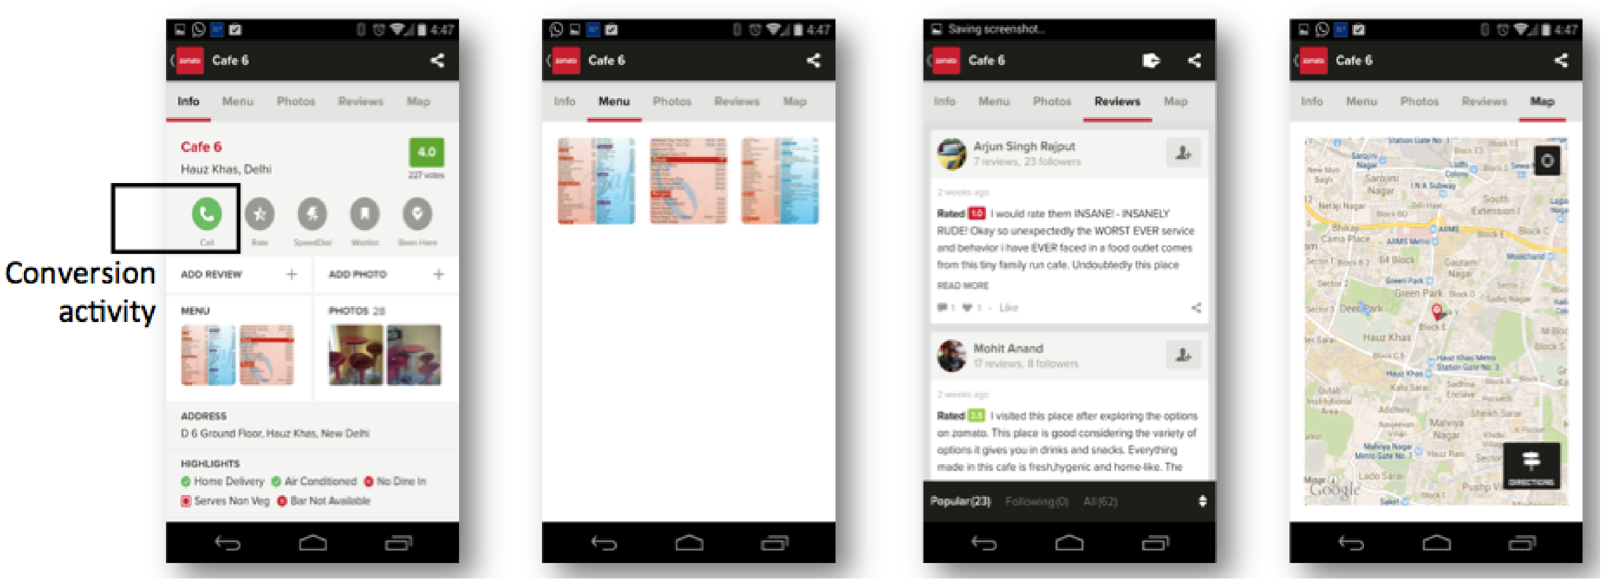

Notes: The figure shows a snapshot of the information screen a user sees when he clicks on the search listing of a restaurant on the Zomato Android app. The user is led to a screen from which he can browse several tabs to view the menu, read the reviews, see a map of the restaurant, or add his own review. The user can also call the restaurant from within the app, which we track as an activity key to conversion.

Ahmedabad; (UAE) Sharjah, Abu Dhabi; (Philippines) Manila; (Indonesia) Jakarta.[7] Our raw data starts on August 9, 2014 and ends on Sept 26th, 2014, when the next update of the app was launched. Following this, advertisements were formally launched on the Zomato mobile platform in November 2014.[8]

Users and Randomization Any user who downloads the updated app with the experiment in it becomes part of the experiment, and is allocated to one of several conditions (described below). Randomization is induced at the user level and is persistent across all sessions by that user. Every user is assigned a unique id, and all activities of the user on the app subsequent to allocation to one of the experimental conditions are tracked.

Experimental Conditions Users are randomized into three experimental conditions:

1. Control Condition: Users in the control condition are shown no advertising and are exposed to only organic listings. These users experience no difference between the experimental and pre-experimental regimes, and serve as a baseline.

2. Treatment Condition A ("*Ad with no disclosure*"): Users in condition A are shown the same organic listings as those in the control condition, along with additional paid listings, but *without* any indication that these paid listings are ads.

3. Treatment Condition B ("*Ad with disclosure*"): Users in condition B are shown the same organic listings as those in the control condition, along with exactly the same additional paid listings as those in condition A, but *with* an indication that the paid listings are ads. Apart from the ad-indication, all other aspects of the ad, including the identity of the advertising restaurant, the content of the listing, and its position in the search results is the same between conditions A and B.

The contrast between conditions A and B helps estimate the signaling effects of advertising. We include the control condition because it benchmarks the firm's no-advertising regime, comparisons to which have business value.

Figure (3) shows an example of the three experimental conditions. The left panel shows the control condition, in which users are shown no advertising. The middle panel shows condition A, wherein users are shown additional paid listings (here for the restaurant "Smoke House Deli"), but without any indication that these paid listings are ads. The right panel shows condition B, in which users are shown exactly the

[7] These comprise large cities in what Zomato refers to as "Full-stack" regions. The firm defines these as "a) large markets b) growing very fast c) [where] Zomato is the strongest player in its space." See `http://blog.zomato.com/post/131277554406/shifting-focus-to-what-matters-and-what-works`

[8] For an announcement of the introduction of the mobile ads, see "Zomato Introduces Ads On The Mobile App; Doubles Consumer Reach For Merchants," `Nov 26, 2014, http://www.thestartupjournal.com/zomato-mobile-app-ads`

same additional paid listings as those in condition A, but with a yellow label indicating that the paid listings are ads. Apart from the label, all other aspects of the ad is the same between conditions A and B.[9]

Choice of Advertisers Included in Experiment The inclusion of advertisers reflects the logic explained in the introduction section. We show ads for restaurants that choose to advertise in equilibrium. When a user searches on the Zomato *mobile app,* his search criterion (the set of filters applied) is directed to an algorithm that reveals the restaurants whose ads the person would have seen if he was on the *website* rather than the mobile app. These restaurants are then advertised on the mobile app's search listings page for users in conditions A and B. This way, we include in the experiment a set of advertisers that are interested in advertising on the platform in response to the search criteria applied by the user, and also mirror on the app the profile of advertisers on the website. The advertisers enjoy the potential benefits of the added mobile exposure for free during the duration of the experiment.[10]

The ads are all served on the search results page shown to a user. A search results page consists of 20 listings. The page may show up to three ads, placed in slots among the organic links as shown in the figures above. The order and position of the slots is decided by Zomato's algorithm (i.e., we do not randomize over these; discussed below). In a typical search in the data, the first ad appears after four organic listings. Clicking on a restaurant's ad takes a user to its info page. Ads are displayed only on the search results pages, and not on the restaurant info pages. There is no change in a restaurant's info pages across the experimental conditions, regardless of whether or not the restaurant is advertised.

To keep the scale of the experiment manageable and to reduce experimental interference, we do not show mobile ads for all possible search criteria that a user inputs on the app. In particular, ads in our experiment appear in search results only for broadly defined search criteria. Specifically, when the search is based on (a) location and/or (a) the three search categories "home-delivery," "dine-out," or "night-life". If a consumer includes a cuisine in the search filter (or any other narrow factor apart from location and search category) he does not see any advertising in the search results, by experiment design. Since advertising is sold on the basis of location and these three category filters, this ensures that the search criteria for which ads are shown are aligned with those actually desired by advertisers (i.e., if an ad is shown as part of the experiment, there will for sure exist an advertiser that desired to advertise to that search). Also, by showing ads based on only these broadly defined searches, we reduce the chance the

[9] The yellow label with the white text denoting "Ad" is chosen to match the disclosure practice of popular search platforms like Google and Yelp. Within condition B, we also sub-randomize users into alternative ways of ad-disclosure used in the industry, so as to explore robustness and alternative mechanisms. The experimental condition and the sub-conditions remain the same for a user across the time period of the experiment (there is no re-randomization over time). This paper focuses on the contrast in outcomes between disclosure versus no-disclosure and aggregates these sub-conditions. In a companion paper that explores these alternatives in more detail (Sahni and Nair 2016), we report no statistically significant difference in outcomes across these sub-conditions.

[10] Advertisers also cannot track if calls to the restaurant originate from the Zomato website or from the app.

advertised restaurants are unrelated to the users' search intent. For example, we avoid a situation where a user searches for Chinese restaurants and sees a series of restaurants in the search results that serve Chinese cuisine but ads that may not satisfy this search criterion.

Consider an individual who applied a search criterion for which a restaurant r's ad is to be placed. If the individual is in the control condition, he may see r's link *once*, as a part of the organic listings. If the individual is in condition A, he may see r's link at least once if r does not appear in the organic listings, or twice if r does appear in the organic listings. Finally, if the individual is in condition B, he may see r's link once as an ad, and again if r appears in the organic listings.[11]

Order and Positions On the mobile app, the order and position of advertised restaurants in both the paid and organic listings is determined by the platform on the basis of its own proprietary algorithms and its contractual arrangements. While ad-position is not under our control nor is randomized, what should be noted is the experiment ensures that if a restaurant's ad is shown in a given position in response to a particular search criteria in condition A, its ad will be shown in response that *same* search criteria in the *same* position in condition B as well. This implies that position is held fixed in comparisons between conditions A and B. Further, while the sequence of organic listings are determined by the platform on the basis of its own proprietary algorithms, the experiment ensures that the sequence of organic listings shown for a given search criteria are the *same* in the control, A & B. This facilitates interpreting the difference between the conditions as driven by the manipulations we induced to paid listings.

To fix ideas, suppose the advertising returned in Figure (3) − "Smoke House Deli" − has contracted for its ads to appear in response to searches from the "Hauz Khas Village" location in New Delhi City on August 9, 2014, and Zomato displays its ad in condition A as shown. Then, all searches in condition B from the "Hauz Khas Village" location in New Delhi City on August 9, 2014 with the "Smoke House Deli" ad served will also feature it on the same position. Further, the same sequence of organic listings that appeared in response to the search in conditions A and B appear in the control condition as well.

Outcome Measures We analyze outcome measures for which signaling theory has a clear prediction. Consider the user actions we track on the app:

- *Calls to the Restaurant*: If advertising serves as a positive signal, we expect disclosing it to enhance demand. Therefore, we expect calls to increase (weakly: not decrease) in the disclosure condition relative to the condition without disclosure. As discussed below, calls act as proxies for demand

[11]Situations in which the advertiser also appears in the organic results on the same search results page as its ad are rare in the data. Specifically, an advertiser and its organic listing appears on the first page together in 5.2% of all searches in our data. Because of randomization, this number is on average the same across the three experimental conditions ($p =$ 0.50). The findings in the paper are not sensitive to omitting searches in which an advertiser appears twice in conditions A and B.

in our context. On the other hand, if users are averse to advertising, they might be less likely to choose an option if they realize it is being advertised. In this case, calls to advertisers may decrease.

- *Information-Acquisition (Search)-related outcomes*: These include other actions a user can take to learn more about a restaurant after reaching its info page on the app, including clicking on its map, browsing through reviews and viewing the menu and photos. Signaling theory does not have a clear prediction for such search-related outcome measures, they may decrease or increase because of ad disclosure. The argument for search activity to decrease is straightforward. Information provided by the ad signal is a free substitute for information a user may gather through effortful search. Therefore, in a world in which a user is informed by the ad signal, the need for search may go down. On the other hand, the ad signal may get some users interested in a restaurant, driving them to explore the restaurant, increasing search activity.

- *Page-visits*: This refers to the action of visiting a restaurant's info page. This action is special because the user cannot acquire more information or call a restaurant without visiting its info page. A user with a low prior who wants to search more or call the restaurant on seeing the ad signal is more likely to visit the restaurant's info page. A user with a high prior who would have visited the info page or called the restaurant anyway, may show no impact on page-visits in response to the signal. Net-net, we predict page-visits to increase (weakly: not decrease) because of disclosure of an ad.

Based on the above rationale, we use calls and page-visits as dependent measures. We use calls instead of actual demand because we are unable to track actual restaurant visits or spending. Both Zomato and restaurant advertisers regard calls to be sales-leads that are a good proxy for sales in the markets we consider. To check this, we analyze historical data on calls made to restaurants. We report this analysis in detail in Online-Appendix A, and provide a short summary here. We use a sample of 1,033 recorded calls made by users to 28 restaurants that advertised on Zomato in Oct-Nov 2010 and content-code them manually. 69.5% the calls involve the caller placing an order for home delivery. 8.5% involve reserving a table at the restaurant. Other calls related to purchase involve those placed to "takeout" food (1.8%) or to arrange for catering (1.3%). The remaining 18.9% are unrelated to purchase. Specifically, 13.4% of these ask for information without expressed purchase intent (e.g., details on the buffet, inquiries about whether the restaurant is open on a holiday). The others include marketing calls, wrong number dials etc. Overall, we find that 81.1% of the calls involve purchase intent from the caller. These data suggest that calls proxy for demand.[12]

[12]The fact that orders comprise a large proportion of the calls is not surprising given that online ordering was not available on Zomato until 2015. Also, there is no well established technology-enabled facilitator of reservations (like OpenTable in the US) in these markets. Quoting an official announcement by Zomato, "A few months ago, we made an important move in India – launching our online ordering service. And we've really been kicking our competition's ass in this

To start the analysis, we begin by exploring patterns of advertising and pricing by the restaurants in the markets Zomato serves in. Then we describe the results from the analysis of the data from the field experiment. We discuss these in sequence below.

4 Supply-side Patterns

To explore these, we look at the set of 142,934 restaurants in our data which received at least one review as of baseline.[13] The database includes for these restaurants the average rating across reviewing users, the number of reviews, the estimate average cost for two people to have a meal at the restaurant (a price-index), as well as the date on which the restaurant is entered into the Zomato database. Unfortunately, a within-restaurant time-series of these variables are not available. We use the values for these variable reported just prior to the beginning of the experiment and compare these in the cross-section. We therefore caution the reader that these comparisons are suggestive and may be picking up unobserved differences in restaurant characteristics.

In the figures below we describe how the probability that a restaurant advertises on Zomato during our study is related to the average rating, number of ratings it receives, cost for two and days since it was added to the database. Since these variables are market-specific and hard compare across regions, we compute the decile of each restaurant's value on each of these variables within the zone in which it is located, and plot the probability of advertising within these deciles. Roughly speaking, we interpret the average rating as a proxy for the appeal of the restaurant, the number of ratings as proxy for its popularity and/or a measure of uncertainty around the average, and the days-since-added as a proxy for the age of the restaurant since its market entry.[14]

Figure (4) shows the probability of advertising during our study against the within-zone decile of a restaurant's average rating at baseline. Restaurants with higher appeal to consumers are seen to advertise more. Figure (5) shows the probability of advertising against the within-zone decile of a restaurant's number of ratings at baseline. We see a positive dependence between the probability of advertising

business with less than 0.1% of their marketing budgets. Why we were able to win? Because we have millions of users who already use us for ordering food over the phone. Now, they have started placing the same orders online using our app. And there is a lot of growth still left; 92% of our users who use Zomato to search for restaurants that deliver haven't even started ordering online on Zomato as yet. Our ticket sizes are more than double our competitors' – because our users are not using us for the discounts. They are using us for the convenience, and a product they already love Zomato for. Our users hold tremendous potential for transaction-based businesses. Getting into transactions was always the natural next step for our business. Online ordering is a natural and logical alternative for our users who, up until now, used to call restaurants to place their orders for delivery. Table reservations fit into Zomato as easily as online ordering did. The time has come for us to focus deeply on transactions in countries where it matters." See, `http://blog.zomato.com/post/131277554406/shifting-focus-to-what-matters-and-what-works`,dated Oct 16, 2015.

[13] We obtain access to a snapshot of the restaurant profile database at baseline. Out of 210,302 restaurants in the database, we retained 142,934 restaurants that had at least one rating provided. Out of these, 15,976 were missing the date added, 31 missing the number of total ratings, 3,522 missing the price index. These restaurants are dropped when creating the respective plots.

[14] Restaurants are added in Zomato's database by a field team that regularly surveys the market and updates the database. Therefore, within a market, days since a restaurant is added is likely to be correlated with when the restaurant started, albeit noisily.

Figure 3: Experimental Conditions

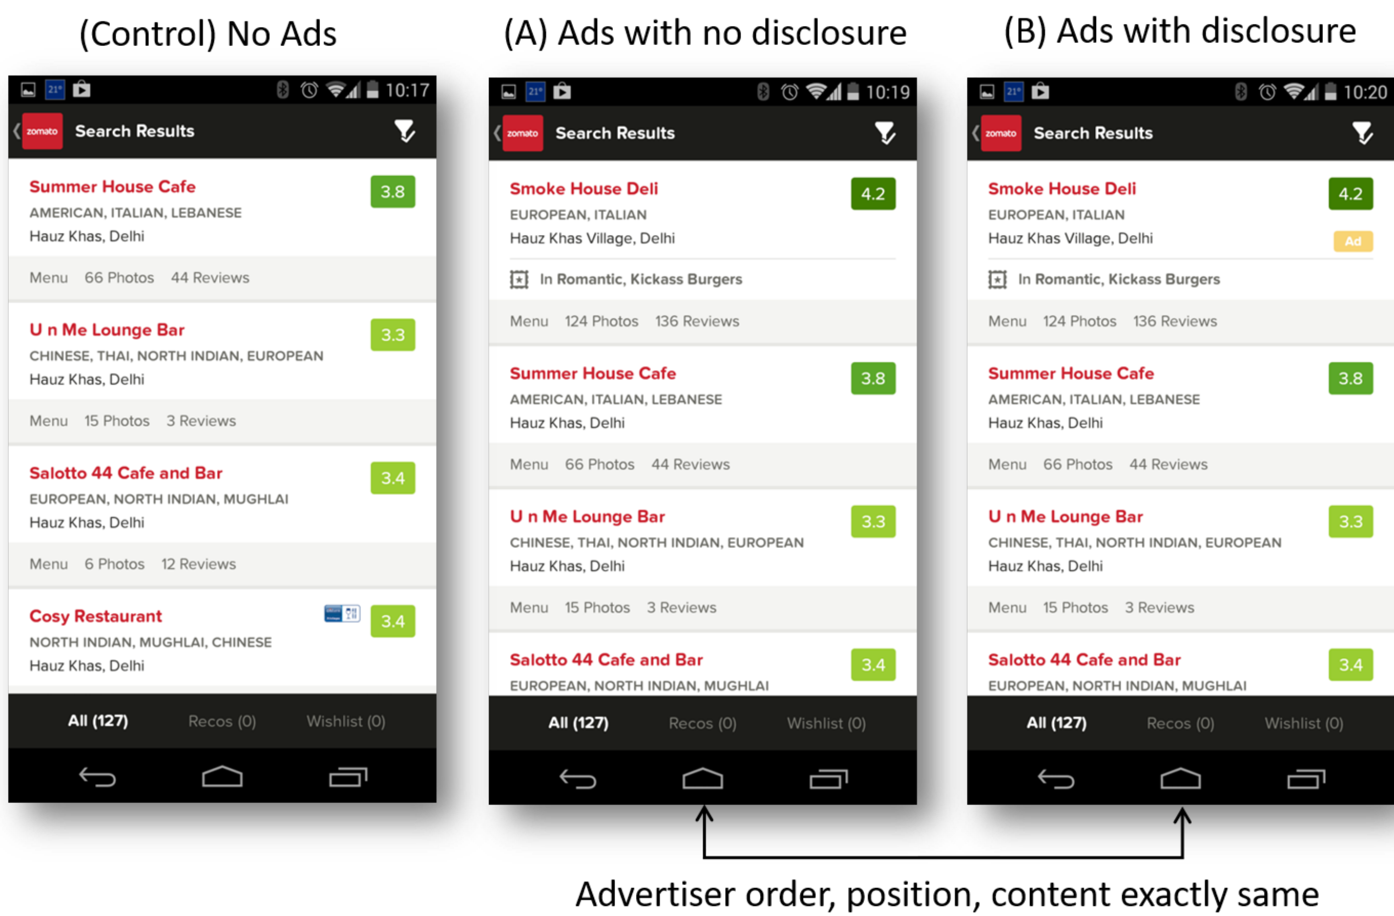

Notes: The figure shows a snapshot of the three experimental conditions. In the control condition, users are shown no advertising and are exposed to only organic listings. In condition A, users are shown additional paid listings (here for the restaurant “Smoke House Deli”), but without any indication that these paid listings are ads. In condition B, users are shown exactly the same additional paid listings as those in condition A, but with an indication –here, the yellow label – indicating that the paid listings are ads. Apart from the indication, all other aspects of the ad, including the identity of the advertising restaurant, the content of the listing, and its position in the search results is the same between conditions A and B. In a typical search in the data, the first ad appears after four organic listings, which is likely to be below the page-fold (i.e., a user has to scroll down after arriving at a search-result page to see an ad).

and the number of ratings, though this correlation is harder to interpret as number of ratings could be proxying for restaurant popularity or age, or for the uncertainty consumers have about the restaurant's appeal. Figure (6) shows the probability of advertising against the within-zone decile of a restaurant's estimated price-index. Restaurants with higher prices are seen to advertise more. Figure (7) shows the probability of advertising against the within-zone decile of the date on which the restaurant is added to the Zomato database. Newer restaurants are seen to advertise more. Finally, Figure (8) shows the probability of advertising against the within-zone decile of the date on which the restaurant is added to the Zomato database, split by within-zone buckets of the average rating. Looking at the figure, newer restaurants, presumably ones that consumers have more uncertainty about, are seen to advertise more; but, holding age fixed, the restaurants with higher appeal also tend to advertise more.

These patterns seem broadly consistent with a signaling equilibrium in which restaurants of higher-than-average "quality" advertise more to signal their appeal to consumers. The higher advertising propensity of newer restaurants, especially those of higher appeal, and its gradual reduction over time can be explained with canonical signaling setups in which firms use ads as signal early in their life-cycle and scale back on dissipative advertising as information diffuses in the market and uncertainty about the restaurant's appeal reduces. The fact that higher priced restaurants tend to advertise more is also consistent with signaling equilibria in product markets where marginal costs increase with quality.[15] While we do not have access to cost data, it seems that most components of marginal costs including food, labor and service, tend to be higher for better restaurants, so this seems the case that applies in practice.

We can also use the ratings information to assess whether the advertising restaurants have higher unobserved appeal. Specifically, under signaling the restaurants that advertise would have aspects of high quality that are unobserved to consumers. This quality would be revealed over time through consumption experiences. If we view ratings as summarizing these experiences, we expect ratings in the future (as well as the improvement in those ratings relative to baseline) to be higher for restaurants that advertised during the experiment relative to those that did not. To check whether this is the case, we pick at random 200 restaurants who advertise during our experimental time period, and 200 restaurants that do not. We

[15] As Milgrom and Roberts (1986) note, whether or not high quality firms use high prices and high advertising to signal quality depend on the relative marginal costs between high and low quality firms. When marginal costs increase with quality, equilibria exist where the high quality firms set high prices and uses high advertising outlays to signal its quality. To see this, assume that there are two firms of high and low quality, such that the high quality firm has higher marginal costs. Assume that advertising is dissipative, and that the higher quality firm obtains more repeat business once consumers visit the firm and experience the good. From a position of equal initial prices, assume that the high quality firm increases its prices a bit. Then, the number of current consumers it loses should be the same as that of the low quality firm, but the "pain" of this current loss should be lower for the high quality firm as its margin per lost consumer is lower. However, if these lost consumers had actually visited the firm, they would have revisited and repurchased at a higher rate for the high quality firm, because their visit reveals its quality is higher. So, the "pain" associated with lost repeats is likely higher for the high quality firm. Since the price increase effects on current and repeats go in opposite directions, its possible that the high quality firm cannot use high-prices alone to signal quality, and will also need to use advertising to signal quality. In the possible equilibria in this scenario, the high quality firm sets its prices above its full-information level and chooses a high level of advertising to separate from the low quality firm. If on the other hand, marginal costs decrease with quality, low prices will signal high quality in this model, and there is no need to use advertising as a signal.

Table 1: Ratings Two-years After Experiment

Dependent Variable: Rating in 2016		
Restaurant advertised in 2014 (0/1)	.411 (.094)	.176 (.076)
Rating in 2014		.241 (.061)
Intercept	3.066 (.080)	2.479 (.208)
N	283	283

Notes: The table shows coefficients and standard errors in parentheses from regressions comparing ratings of restaurants in 2016 depending on whether they advertised in 2014. There are 283 randomly picked restaurants in the data that were found in 2016 on Zomato. The first column regresses the ratings in 2016 on whether the restaurant advertised in 2014. The coefficient is positive and significant suggesting that the restaurants that advertised in 2014 have a higher rating in 2016. The second specification controls for ratings in 2014 to check whether the change in ratings from 2014 to 2016 is predicted by whether the restaurant advertised. Both the coefficients in the second column are positive and statistically significant. This finding suggests that restaurants with higher ratings in 2014 tend to have higher rating in 2016. It also suggests that among restaurants with the same rating in 2014, those that advertised had a higher rating in 2016.

search for their ratings on Zomato in September 2016 (two years after our experiment). We found 137 that advertised and 146 that had not advertised in 2014 (the probability of finding a restaurant is the same whether or not the restaurant had advertised; $p = 0.32$). Table 1 reports regressions comparing the 2016 ratings of restaurants depending on whether they advertised in 2014. Looking at the first column, we see the restaurants that advertised in 2014 have a higher rating in 2016. The second column reports the same regression but controlling for 2014 ratings. Among restaurants with the same rating in 2014, we see those that advertised also have a higher rating in 2016, i.e., the fact that a restaurant advertised is able to predict a better consumption experience for consumers beyond what is predicted by its contemporaneous ratings.[16]

Finally, note the positive correlation between advertising and ratings we observe on the platform may make equilibria with ad-signaling easier to sustain. Because of this, consumers may form casual associations between restaurant quality and the propensity to advertise on Zomato. What is required for a signaling equilibrium to obtain is that these associations are not violated by actual experiences when those consumers visit the advertised restaurants. If restaurants realize that consumers respond to advertising, and those of higher "quality" advertise more anticipating they will obtain more repeat purchase, this will be the case, and the casual associations will be reinforced. Thus, we do not have to assume consumers are endowed with unrealistic levels of rationality and are able to "compute the equilibrium" in their minds in order for advertising to have indirect informational value in this setting.

Overall, these descriptive facts on the "supply-side" are broadly consistent with signaling theory. However, we note that they are only suggestive, as alternative models of signaling can imply different life-cycle predictions, and because tests of signaling in across-restaurant comparisons are confounded with unobservable product and market-specific differences as cautioned in the introduction. For this reason,

[16]Thanks to Matt Gentzkow for suggesting this test.

Figure 4: What Types of Restaurants are Advertising? Average Ratings

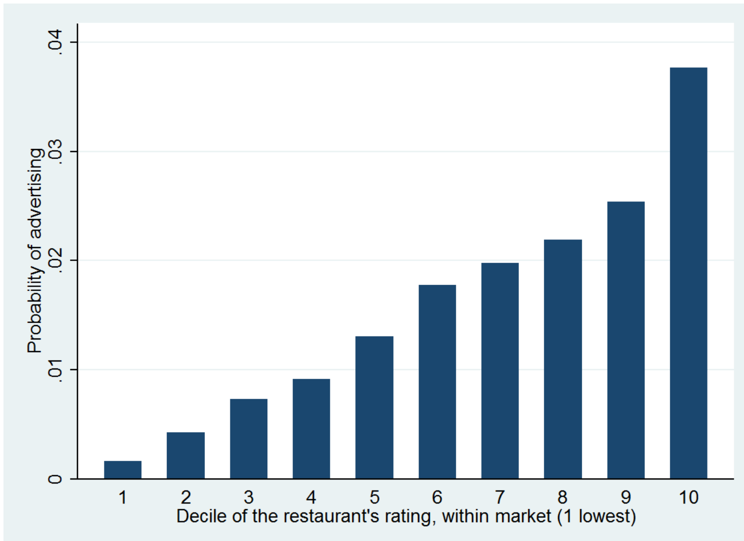

Notes: The figure shows the probability of advertising against the within-zone decile of a restaurant's average rating at baseline. A zone is a 5mi X 5mi area within a city on the basis of which ads are sold and geo-targeted. Restaurants with higher appeal to consumers are seen to advertise more.

Table 2: Split of Users by Experimental Condition

Condition	*Description*	*Num. of Users*
Control	No ad	44,233
A	Ad with no disclosure	44,637
B	Ad with disclosure	177,105

Note: 265,975 consumers who are in the experiment.

we base a test for signaling primarily on more carefully constructed comparisons on the demand-side.[17]

5 Field-Experiment Results

5.1 Analysis Based on First Sessions

The main results reported below utilize data only on users' first exposures to experimental ads. We do this to address the following econometric issue. Consider an individual in condition A or B who

[17]Casual empiricism suggests that the best restaurants do not advertise heavily. How do we reconcile that with the descriptive facts reported here? What is relevant to signaling is the advertising behavior of those restaurants when they were new to the market and consumers did not know they were good. Our data (Figure 8) suggest that better restaurants indeed tend to advertise more on the platform when they are new. Other theory has pointed out that "counter-signaling" equilibria that obtains with multidimensional signals and large heterogeneity across firms can also explain this phenomena (Orzach et al. 2002). In these equilibria, even noisy signals are sufficient to separate the "best" firms from the "low" quality firms. The "medium" quality firms then use advertising to separate themselves from the "low" quality firms, while the "best" firms avoid advertising to separate from those with "medium" quality. Note that even in these models, advertising works as a signaling device, serving to separate firms of differing quality, which is what we are testing here.

Figure 5: What Types of Restaurants are Advertising? Number of Ratings

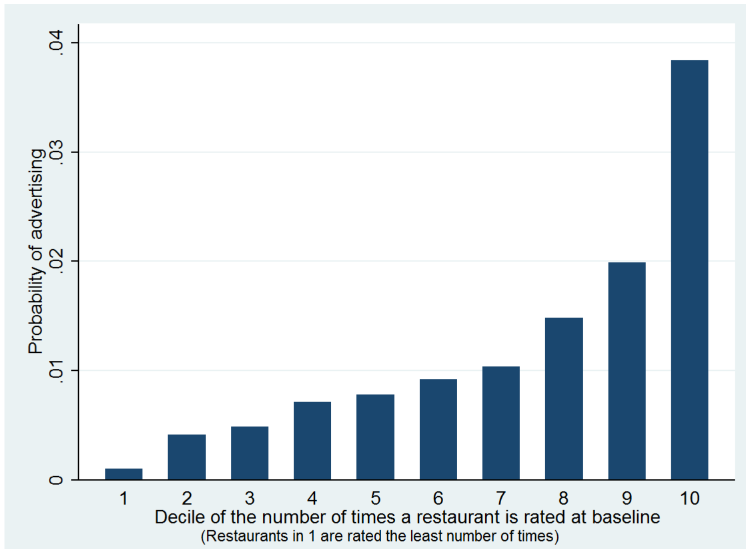

Notes: The figure shows the probability of advertising against the within-zone decile of a restaurant's number of ratings at baseline. A zone is a 5mi X 5mi area within a city on the basis of which ads are sold and geo-targeted. The number of ratings could be proxying for restaurant popularity or age, or for the uncertainty consumers have about the restaurant's appeal.

Figure 6: What Types of Restaurants are Advertising? Prices

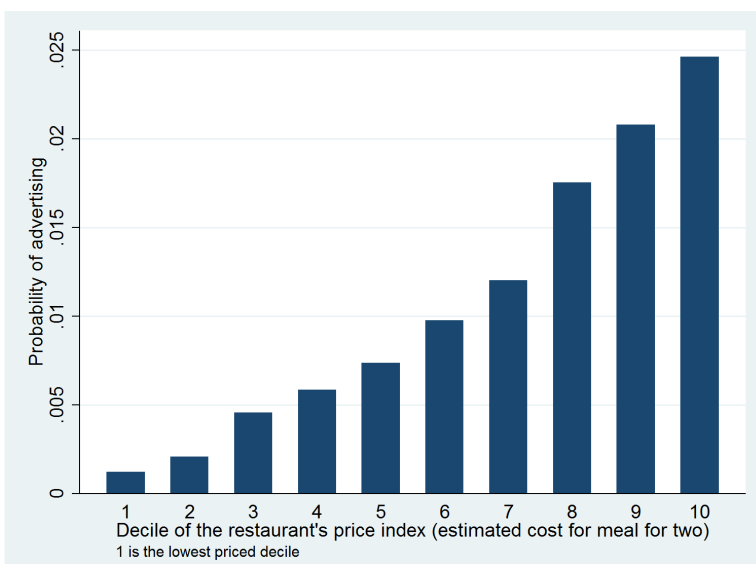

Notes: The figure shows the probability of advertising against the within-zone decile of a restaurant's price-index, estimated by Zomato as the cost for a couple to have a meal at the restaurant. A zone is a 5mi X 5mi area within a city on the basis of which ads are sold and geo-targeted. Restaurants with higher prices are seen to advertise more.

Table 3: Tests of Randomization in the Experimental Design

User Characteristic	*Variable*	*Mean*	*Std. dev*	*Mean in Control*	*Mean in* A	*Mean in* B	*p-value of test*
							H_0: Equal means across conditions
	Number of reviews	0.76	5.26	0.76	0.77	0.76	0.9
Past engagement with Zomato	Number of Ratings	1.68	8.3	1.7	1.68	1.67	0.81
	Days since signed up	285.24	282.98	285.76	286.16	284.89	0.62
	Number of Restaurant Clicks	2.76	11.92	2.71	2.71	2.78	0.33
	Number of menus seen	1.35	6.76	1.31	1.33	1.36	0.28
Search activity prior to the experiment	Number of calls made	0.14	0.95	0.14	0.14	0.14	0.94
	Number of Searches	1.35	4.17	1.36	1.34	1.34	0.56
	Number of Sessions	0.48	1.4	0.48	0.48	0.48	0.98
	Avg. Rating of the restaurants	3.49	0.7	3.5	3.5	3.49	0.17
Characteristics of restaurants visited in the past	Avg. Price (indexed) of restaurants	14335.2	70909.9	13985.9	14568.8	14364	0.82
	Popularity in terms of number of reviews	336.8	414.72	334.83	340.72	336.3	0.53

Note: Table tests whether the baseline characteristics of users in the three experimental conditions are statistically significantly different from each other, in order to assess whether randomization is induced correctly.

Figure 7: What Types of Restaurants are Advertising? Age

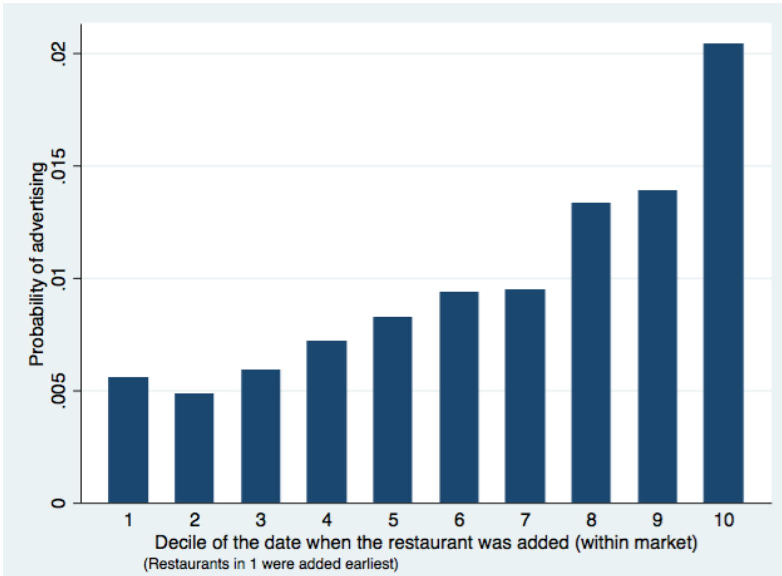

Notes: The figure shows the probability of advertising against the within-zone decile of a the date on which the restaurant is added to the Zomato database, which serves as a proxy for age. A zone is a 5mi X 5mi area within a city on the basis of which ads are sold and geo-targeted. Newer restaurants are seen to advertise more.

searches on the platform and gets exposed to an ad. All of the user's subsequent search behavior on the platform − which drives his propensity to be exposed to additional ads − could be influenced by this initial ad-exposure. For example, if advertising in condition B is more effective compared to A, then users in condition B might end up browsing fewer pages and getting exposed to fewer (and possibly different) ads after the first one. Therefore, comparing users conditional on subsequent searches or subsequent ads seen is subject to selection, and does not provide a valid comparison. To mitigate this, we base our tests on the effects of the first exposure to experimental ads on consumer-decisions.

Next, in order to measure the causal effect of advertising, we need to define a proper counterfactual comparison to the behavior of users in conditions A and B. The relevant counterfactual is what the user's behavior would have been if he had done the same search, but had been in the control condition and seen no ads. This requires finding searches in the control condition that match the search in conditions A or B that caused the user to be exposed to a treatment ad. We implement the following steps to accomplish this.

- We start by defining a *session* as constituting all actions starting with a users' opening of the app on his phone up-to the beginning of a continuous period of inactivity that is longer than 3 hours. A session could comprise one or several searches, page visits and calls and represents a user trying to find a restaurant for a particular consumption occasion.

Figure 8: Profile of Advertising by Average Rating and Age

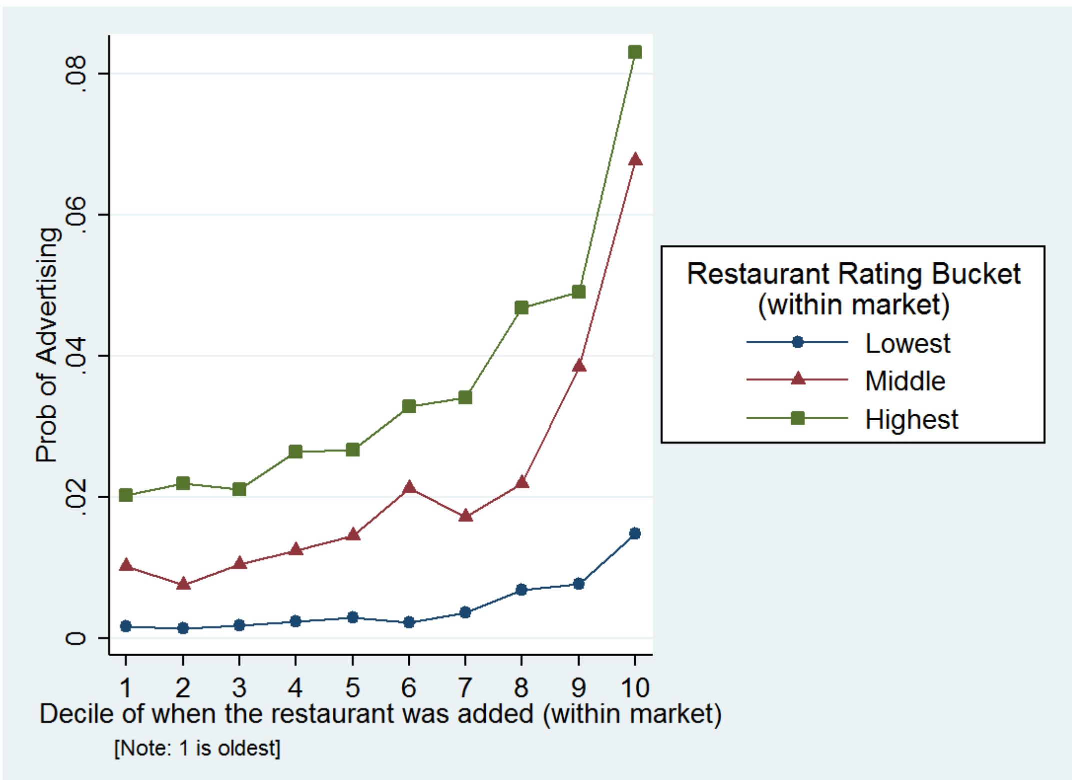

Notes: The figure shows the probability of advertising against the within-zone decile of a the date on which the restaurant is added to the Zomato database, split by within-zone buckets of the average rating. Newer restaurants, presumably ones that consumers have more uncertainty about, advertise more, but within them, the ones with higher appeal also tend to advertise more.

- For each individual (across all three conditions), we examine the sequence of his searches in the data, and determine the first search for which ads would be served. For this search, we determine the restaurants who advertised, using the search-to-advertiser mapping.[18] We restrict our analysis to the users' actions related to the first restaurant that is advertised on this search. This is the first advertiser that a user sees (or could have seen). If multiple restaurants advertised at this position (which happens if Zomato decides to show multiple restaurants given its contractual agreements), we consider the user's actions related to each such restaurant. Further, if the advertising restaurant is a part of a chain, we consider the user's actions related to all restaurants of that chain in the area.

This analysis plan minimizes the bias from feedback effects by basing inference on the response to the first ad exposure for each user. The above procedure results in a dataset in which each observation in conditions A & B is an individual $\times$ the restaurants that are advertised at the first ad-slot in a users' search results page. In the control condition, each observation is an individual $\times$ the first restaurant or restaurants that would have advertised in response to his search if the user had been treated. The dataset is unbalanced (different restaurants r for each user i). Each observation tracks as dependent variables indicators of whether the user visits the advertiser's page and/or calls the advertiser during the session when the exposure occurs.

A more detailed description of the data along with checks of balance is presented in Appendix A.

5.1.1 Data

We observe 265,975 users who download the update, and at least at one instance, apply a search for which ads are served. These users are spread across 321 zones, and exposed to 622 advertisers across conditions. The distribution of users across the three experimental conditions is presented in Table (2). There are roughly four times more users in condition B than A or the control, because B is comprised of additional sub-conditions as noted before.

Table (3) presents tests of mean equivalence in user characteristics across the conditions. We report tests on a variety of pre-experimental variables including users' past engagement with Zomato, search activity prior to the experiment, and the characteristics of restaurants visited in the past. Looking at the p-values reported in the last column, we see that the null of equal means across the three conditions is not rejected for any of the variables, showing that randomization is induced properly.

[18] For any search, depending on the day when the search happens and the filters applied, we know whether any ads would be served, and which restaurants would be advertised.

Table 4: Effect of Ad-disclosure on Call and Page-Visit

Condition	*Description*	Call		Page-visit	
		Mean	*Std. Err.*	*Mean*	*Std. Err.*
A	Ad with no disclosure	0.031%	0.006%	1.10%	0.03%
B	Ad with disclosure	0.055%	0.004%	1.18%	0.02%
	p-value of test, H_0: Equal means	0.002		0.084	
	Effect-size	3.09		1.73	

Notes: There are 366,330 observations, 73,714 in A; and 292,616 in B corresponding to 44,637 unique users in A; and 177,105 unique users in B. Means represent the average of call and page-visit indicators across all users' first-sessions in each condition. p-values computed by running a regression of the call/page-visit indicators for users in A and B on an indicator for condition B, with standard-errors clustered at the user-level. Effect-size computed as the t-statistic on the indicator in the regression. Call probability is 77% higher when disclosing relative to the no-disclosure condition. The visit rate with and without disclosure are statistically indistinguishable. Disclosure that a listing is a paid ad thus drives the incremental conversion.

5.1.2 Main Effects: Visit and Call Rates

Table (4) reports the change in page visit probability and the call probability in response to disclosure and represents the main results of the paper.

Looking at Table (4), we see that disclosure has a positive and statistically significant effect on the probability of calling the advertised restaurant. This is a large effect, producing a 77% increase in B relative to A (from 0.031% to 0.055%). This represents the causal effect of disclosure on calls and is consistent with the signaling prediction.

To obtain a sense of the magnitude of the disclosure effect, we compare it to the call-probabilities the same restaurants would obtain in a world without disclosure, in response to a change in their characteristics. To do this, we estimate how much the call probability of the advertisers in the top ad-slot would change if their characteristics changed. As we do not randomize these, we cannot estimate a causal effect of these on calls. Nevertheless, as a back-of-the envelope estimate, we regress an indicator of whether any of the advertisers are called by users in group A, on the within-zone deciles of the attributes we reported on in the previous section (average rating, number of ratings received, cost for two and days since it was added to the database). Appendix B presents the regression results. We find the disclosure effect is comparable to a two decile increase in the number of ratings received (statistically significantly estimated). So, roughly speaking, an average advertiser who is listed on the top ad-slot in response to a search, obtains the same conversion benefit by disclosing sponsorship as a two decile increase in the number of ratings it has on the platform.

We can use back of the envelope arithmetic to obtain a sense of the size of the effect of ad disclosure for the platform as a whole. Based on conversations with the firm, we obtain that one call is worth roughly INR. 400 (USD $6.2) to the platform. Using the results from Table (4), the incremental improvement in call probability from the disclosure is 0.024% (0.055−0.031). This yields an expected incremental benefit of roughly \$1.48 per 1,000 impressions from disclosure alone ($1000 \times \frac{.024}{100} \times \$6.2$). At the time of the experiment, 30M monthly consumers used the platform on average. Assuming conservatively 5 ad-exposures per month to these users, we obtain a rough annual value of USD\$2.7 million from the ad-disclosure to the platform [30M(users per month)×12(months)×5(exposures per month)×$\frac{\$1.48}{1000}$]. In 2015, given reported average monthly users of 90M, this value scales to thrice as much, to about USD\$8M.

Looking at the same table, we also see the page-visit probability does not decrease when the advertising is disclosed (B relative to A). This is also consistent with the signaling prediction.[19] While disclosure does drive more people to the advertiser's page, an increase is not statistically significant at the 95% confidence level.

What could explain the fact that calls increase more significantly than page-visits? This can occur if many of the users who get affected by the disclosure would have visited the restaurant's info page anyway, but would not have called in the absence of the disclosure. This is likely to happen if the advertising restaurants have higher-than-average observable quality (which makes it worthwhile to users to visit its info page and acquire more information); but in the absence of the signal, there remains enough uncertainty that many users continue to explore other options and do not stop. The stylized facts we reported in section 4 about advertising restaurants suggest that this may well be the case on the platform. There, we reported that restaurants that advertise have higher-than-average ratings. But, they are also newer, implying possibly higher uncertainty about them.

To summarize these results, disclosure has a measurable positive effect on demand that is consistent with the signaling prediction. Our interpretation is that disclosure conveys a cue to exposed users that increases the restaurant's appeal and improves their evaluation of the advertised product. Additional statistical robustness on these results, including exact p-values computed via simulation is reported in Section 6.

Exploring Further: Conditional Call Rate and Continued Search Finally, we report two additional comparisons as a way to better unpack user behavior between the two conditions. Table (5) compares the call probability conditional on a page-visit between A and B. This comparison focuses on the subset of individuals who chose to visit an advertising restaurant's info page in each of the conditions, and then compares the proportion of the subset that called the restaurant. Since each subset within a condition is a selected subsample, this comparison does not hold fixed user profiles across conditions and does not estimate a clean causal effect. Nevertheless, it helps explore further the pattern found in the visit and call tables above, and to check how the conversion of the set of individuals that disclosure attracts is different from those who visit under no disclosure. Table (5) shows that the conversion from

[19] We can reject the null-hypothesis that disclosure *decreases* page visits using a single tailed-test (p-value = 0.04).

Table 5: Call Probability Conditional on Visit, Split by Experimental Condition

Condition	*Description*	*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*
A	Ad with no disclosure	2.33%	0.50%	0.005
B	Ad with disclosure	4.11%	0.30%	

Notes: The table reports on 4,214 users who visited the page of an advertising restaurant across the conditions. There are 814 observations in A; 3,452 observations in B; corresponding to 808 unique such users in A; and 3,406 unique such users in B. Means represent the average of call indicators conditional on visit, across all users' first-sessions in each condition. p-values computed by running a regression of the call/page-visit indicators for users in A and B on an indicator for condition B, with standard-errors clustered at the user-level. The conditional call probability is 76% higher when the ad is placed with disclosure relative to the no-disclosure condition. Disclosure that a listing is a paid ad thus seems to make the restaurant's appeal stronger to exposed users.

a page-visit to a call is highest for the subset of users who visit the restaurant's info page while seeing that the listing is a paid ad, and higher compared to showing the listing in the same position but without disclosure (a 76% improvement for B compared to A).

Table (6) shows the average probability of a user continuing to search during the session, i.e., visiting another restaurant's info page after visiting the advertiser's info page, across experimental conditions. As a first observation, note that the probability of a user continuing to search after visiting a restaurant's page is high: 78.5% in condition B. This shows this is a competitive setting with significant search, and a large majority of the individuals who consider a restaurant continue to explore more options. This allows disclosure to affect consumer decisions. Comparing the estimates across conditions, we see that users in B continue to search at roughly the same (or lesser) rate after seeing the advertised restaurants' page. Yet, Table (4) documented that these users call the advertised restaurant at a higher rate during the session than users in A. This is consistent with the fact that users in condition B see the ad, visit the advertised restaurant's page, continue browsing to explore other options, and then call the advertised restaurant after some deliberate processing. Users do not seem to be drawn into calling a restaurant simply because they are induced to click on a listing by a catchy ad-label. Rather, advertising seems to work by exposure by providing a cue to consumers who choose to call after some thoughtful consideration.[20]

These patterns augment our main findings from the comparison of unconditional means across the experimental conditions. They suggest that the role of the disclosure does not seem to be in increasing user attention. Rather, conveying that a listing is a paid ad makes it more likely that the users who click on the listings convert the page-visit to a call, suggesting that the disclosure made the restaurant's appeal stronger to exposed users.

[20]To underscore this point, note that after exposure to an advertising restaurant's listing, users in condition A browsed past 53.6 listings on average before calling the advertised restaurant in that session; the corresponding number for those in condition B is 69.2.

Table 6: Probability of Continuing Search After a Page-visit, Split by Experimental Condition

Condition	*Description*	*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*
A	Ad with no disclosure	78.50%	1.4%	0.18
B	Ad with disclosure	76.36%	0.7%	

Notes: The table reports on 4,214 users who visited the page of an advertising restaurant across the conditions. There are 814 observations in A; 3,452 observations in B; corresponding to 808 unique such users in A; and 3,406 unique such users in B. Means represent the estimated probability of visiting another restaurant's info page after visiting the advertiser's page. p-values computed by running a regression of the call/page-visit indicators for users in A and B on an indicator for condition B, with standard-errors clustered at the user-level. The probability of a user continuing search after visiting an advertiser's page goes down moving from condition A to B, though the difference is not statistically significant.

The Signaling Effect We infer the estimated effect of ad-disclosure – the consumer seeing that a restaurant is advertised – as evidence supporting the signaling effect. When we say that this is a test of signaling, we mean that the disclosure to the consumer causes the consumer to change his beliefs about the restaurant's appeal, which is revealed by his changed propensity to call the restaurant after some thoughtful deliberation. Since we are holding ratings and other characteristics fixed in our comparisons, the effect we measure represents the informational value of advertising over and above the information provided to consumers by ratings and other characteristics. A restaurant's "appeal" could include low prices, high vertical quality (better ingredients, unspoiled meat), and/or a better match value with the consumer's idiosyncratic tastes. We cannot sort between these different types of beliefs. However, since we control for all other aspects of advertising, we assert that the channel by which disclosure affects a consumer's decisions is via a change in beliefs. We believe that a mere indication disclosing the ad is unlikely to directly affect intrinsic restaurant preferences. While we cannot rule out this possibility emphatically, note that preference changing effects like persuasive or complementary effects of advertising are typically understood as operating through the content of the ad (for example, through the positive feelings evoked by the ad-copy or through positive associations kindled in memory through reminders of existence or enjoyment embedded in the ad). Such effects may well coexist with the signaling mechanism, but may operate through ad-content and position. To the extent that we have controlled for all aspects of ad-content, awareness, positions and existence, we believe what we have isolated in the response to disclosure is a cleaner signaling effect relative to the past literature.

Comparing to the No-Ads Condition While the comparison with the control (no ads) is not directly relevant for our analysis, we conclude this section by reporting on the comparison for completeness. An advertiser placed in an ad-slot in condition A may have an organic listing shown at a much lower position in the control condition (or not shown at all). Going from the control condition to A thus involves an increase in listing position for the advertiser, which might affect outcomes positively due to increased awareness

Table 7: Comparing the Ad with no disclosure to the No-Ad Condition

Condition	*Description*	Call		Page-visit	
		Mean	*Std. Err.*	*Mean*	*Std. Err.*
Control	No ad	0.030%	0.006%	0.83%	0.03%
A	Ad with no disclosure	0.031%	0.006%	1.10%	0.03%
	p-value of test, H_0: Equal means	0.904		0.000	
	Effect-size	0.12		5.39	

Notes: There are 146,809 observations, 73,095 in the control; 73,714 in A; corresponding to 44,233 unique users in the control and 44,637 unique users in A. Means represent the average of call and page-visit indicators across all users' first-sessions in each condition. p-values computed by running a regression of the call/page-visit indicators for users in the control and A condition on an indicator for condition A, with standard-errors clustered at the user-level. Effect-size computed as the t-statistic on the indicator in the regression. Visit probability is 33% higher when the ad is placed, relative to the no-ad condition. The call rate between the conditions is statistically indistinguishable.

and/or exposure. However, going from the control condition to A might also bring the advertiser's listing closer in placement to more relevant listings for that search as determined by the platform's algorithm. This can increase substitution, potentially mitigating any increase in the outcomes. Looking at Table (7) which compares A to the control, we see that though the page visit probability goes up, calls are statistically indistinguishable. Thus, it appears that placing the ad in the ad-slot does not automatically translate into conversion in our setting.[21]

5.1.3 Exploring the Mechanism and Heterogeneity in More Detail

We explore differences in response to ad-disclosure amongst users and restaurants along dimensions of heterogeneity on which signaling theory has a testable prediction. The main comparative static is that signal is likely to be more useful when uncertainty about the quality of the good is higher. As operationalizations of this idea, we test whether the effect of the disclosure is higher for subpopulations of users that have more uncertainty about the appeal of the restaurants they are searching for, and for subpopulations of restaurants about which users are *ex ante* more likely to be uncertain.

Behavior of Users Searching in an Outside City To describe this test precisely, we revisit the notation in section (2) where user $i's$ posterior belief about restaurant r's appeal is represented by $b_i\left(q_r|\mathcal{I}\left(x_r,a\right)\right)$, where $\mathcal{I}\left(x_r,a\right)$ is the information bundle the consumer is exposed to, which includes content x_r and an indicator a of disclosure. Let σ_i represent the variance of the consumer's prior beliefs such that users who have larger σ_i are ex ante "more uncertain" about restaurant appeal. Theory predicts that users with higher prior uncertainty are likely to rely more on the signal. Hence, holding all other

[21] This seems to parallel findings in the search literature. In the context of Google search-ads, Narayanan and Kalyanam (2015) use a regression-discontinuity based argument to document that shifting an ad to the top position (making it more likely to be seen) causes an increase in clicks, but this increase in clicks does not convert to a detectable increase in sales.

characteristics fixed, the disclosure should have more effect for such users. Motivated by this, we test in the data whether,

$$\mathbb{E}_i\mathbb{E}_r\left(y_i\left[b_i\left(q_r|\mathcal{I}\left(x_r,a=1;\sigma_i\right)\right)\right]-y_i\left[b_i\left(q_r|\mathcal{I}\left(x_r,a=0\right);\sigma_i\right)\right]\right)$$
$$>\mathbb{E}_j\mathbb{E}_r\left(y_j\left[b_j\left(q_r|\mathcal{I}\left(x_r,a=1;\sigma_j\right)\right)\right]-y_j\left[b_j\left(q_r|\mathcal{I}\left(x_r,a=0\right);\sigma_j\right)\right]\right)$$

across sets of users i and j who are expected to be *ex ante* identical, except that one set is more uncertain of the restaurant's appeal than the other, i.e., $\sigma_i > \sigma_j, i \neq j$.

To operationalize the uncertainty, we look for users in each experimental condition who search for restaurants in a city which is different from the city where they predominantly searched in the past. Our assumption is that these users are more uncertain about the appeal of the restaurants they are searching for compared to users searching in their home cities. We maintain the assumption that the reason these users travel is not influenced by which experimental condition they are allocated to.

We use the historical data prior to the first session for each individual to identify the city in which the individual conducts most of his searches.[22] We find that 5,031 users search for restaurants in a city that is different from the one they usually conduct their searches. We test whether the ad-disclosure has a differential effect for these user-sessions, compared to user-sessions in their usual cities. Unfortunately, we do not have enough power to explore a similar comparison for the call variable.[23]

Table (8) shows the means and tests for this analysis. Looking across the rows in the left column, we see that users searching in a different city from their usual one behave differently and visit the listed restaurant's pages at a significantly higher rate if they are advertised with disclosure. This behavior contrasts with the rest of the sample, whose behavior looks similar to that reported for the overall group in Table (4): a statistically insignificant difference in the visit propensity with disclosure relative to no disclosure. For another perspective, looking across the columns in the row corresponding to condition B, we see that the mean chance of visit to the advertiser's page is different for the two subsamples (p-value < 0.01). The shift is insignificant in the row corresponding to condition A. These patterns show that a consumer knowing that a restaurant is advertised is more important for individuals in a different city where the uncertainty about the restaurants is likely to be higher. These findings are consistent with a signaling role for advertising, and serve to augment the main effects reported in the previous section.

Behavior of Users with Respect to Advertiser Characteristics The next cut of the data now investigates characteristics of restaurants that users visit/call under ad-disclosure. Under the signaling mechanism, we would expect users to place more weight on the signal-value incorporated into the ad-

[22] Specifically, we use the data from before the individual conducts a search for which experimental ads are shown. Because of limited availability of past data, we are able to do this for 83,245 (of the 265,975) users in the experiment.

[23] Out of the 5,031 users, 860 are in the control condition, 843 in A, and 3,328 in B. We observe 118 visits across conditions that facilitates a test of differences in rates on this variable, but only 4 calls totally.

Table 8: Visit Probability Split by Experimental Condition and By User Uncertainty

		User sessions in a different city from hometown			Remaining User sessions		
Condition	*Description*	*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*	*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*
A	Ad with no disclosure	0.71%	0.22%	0.000	1.11%	0.09%	0.194
B	Ad with disclosure	1.73%	0.18%		1.17%	0.05%	

Notes: The table reports on two subsets of our sample of observations. We observe 1,409 user $\times$advertiser observations in condition A, and 5,489 in B. The right panel reports on the rest of the users. The table reports the differences in visit probability, separately by these two groups of users, across all relevant first-sessions in each condition. p-values computed by running a regression of the info page-visit indicators for users in A and B on an indicator for condition B, with standard-errors clustered at the user-level. Users searching in a different city from their usual one are seen to behave differently: disclosure has a strong effect for users searching in a different city, but there is no statistically distinguishable effect for the rest.

disclosure when they have higher prior uncertainty about the advertising restaurant's appeal. Thus, if we compare across restaurants, we should see that those restaurants about which consumers *a priori* have more uncertainty benefit more from the ad-disclosure, all other things held equal. As an operationalization of this idea, we check whether newer restaurants and those rated by fewer users on Zomato benefit more from ad-disclosure controlling for other observable restaurant attributes. Implicitly, we are treating fewer ratings and newer entry as proxies for higher uncertainty, though these variables could be proxying for other things like popularity or quality, so the test is only suggestive and is weaker than the ones presented previously. Other than the number of ratings, we also observe other characteristics about restaurants including the average rating and price-index presented on the Zomato listings. These characteristics shift the baseline chance of the individual choosing a restaurant. Theory does not have clear predictions about how the value of the signal will change with such characteristics. Pick for instance, the average rating of a restaurant. If the rating is high enough, a user may not need an additional signal to buy from it. On the other hand, if the rating is low enough, a strong signal may not be enough for an individual to buy. Similar arguments hold for other characteristics that shift the baseline.

To hold observed characteristics fixed, we analyze the results in a regression set-up. Let $I_{i,B}$ be an indicator of whether user i is in condition B, and let $d_{r,\text{Num-rating}}$, $d_{r,\text{Ave-rating}}$, $d_{r,\text{Price-Index}}$ and $d_{r,\text{Date-Added}}$ denote the within-zone deciles of restaurant r on the respective variables. Represent by Y_{ir} an indicator for whether user i calls restaurant r. Stacking across all user, first-session, advertiser combinations in conditions A and B, we run the following regression,

$$\begin{aligned} Y_{ir} = \alpha + \beta I_{i,B} + \lambda_1 d_{r,\text{Num-rating}} + \lambda_2 d_{r,\text{Ave-rating}} + \lambda_3 d_{r,\text{Price-Index}} + \lambda_4 d_{r,\text{Date-Added}} \\ + I_{i,B} \times [\gamma_1 d_{r,\text{Num-rating}} + \gamma_2 d_{r,\text{Ave-rating}} + \gamma_3 d_{r,\text{Price-Index}} + \gamma_4 d_{r,\text{Date-Added}}] + \epsilon_{ir} \end{aligned}$$

We also report the same regression using an indicator for page visit as the dependent variable. Recall that in the earlier analysis, we did not find a significant effect of ad-disclosure on page-visits. Therefore, *a priori*, we expect to find systematic heterogeneity in the effect of ad-disclosure on calls but not visits. In both, the main interest is in the interactions of the number of ratings and age variables with the condition B dummy (γ_1 and γ_4). These pick up the comparative static on prior uncertainty and prior appeal and are expected to be negative under the signaling hypothesis. We include interactions with the average rating and the price index so as to describe the heterogeneity in treatment response.

Table (9) presents the results. Looking at the table, we see that restaurants of higher rating benefit more from ad disclosure in terms of both page visits and calls. The interaction of the condition B dummy with the age of the restaurant is not statistically significant. We see that restaurants that have fewer ratings benefit more from the ad-disclosure in terms of their call probability (coefficient on $\mathcal{I}_{i,B} \times d_{r,\text{Num-rating}}$ in the call regression). These interactions are consistent with a signaling mechanism. The effects are also quantitatively significant. Holding everything else the same, a restaurant that is

Table 9: Heterogeneity in Responsiveness to Disclosure Across Restaurant Types

Dependent Variable:	Page Visit		Call	
Variable	*Parameter*	t-stat	*Parameter*	t-stat
$\mathcal{I}_{i,B} \times d_{r,\text{Num-rating}}$	-0.03%	-0.83	**-0.02%**	**-2.51**
$\mathcal{I}_{i,B} \times d_{r,\text{Date-Added}}$	-0.01%	-0.34	0.00%	-1.42
$\mathcal{I}_{i,B} \times d_{r,\text{Ave-rating}}$	**0.05%**	**1.73**	**0.02%**	**3.45**
$\mathcal{I}_{i,B} \times d_{r,\text{Price-Index}}$	-0.01%	-0.27	0.00%	-1.17
$d_{r,\text{Num-rating}}$	0.02%	0.80	**0.01%**	**2.25**
$d_{r,\text{Date-Added}}$	0.00%	0.18	0.00%	0.98
$d_{r,\text{Ave-rating}}$	**0.08%**	**3.15**	-0.01%	-1.50
$d_{r,\text{Price-Index}}$	**0.10%**	**4.53**	0.00%	0.15
$\mathcal{I}_{i,B}$	0.01%	0.04	0.06%	1.33
Intercept	**-0.50%**	**-1.99**	-0.03%	-0.86
Num. Observations		277,894		277,894
Num. Individuals (clusters)		196,552		196,552
R^2		0.0012		0.0001

Notes: The table reports on 196,552 users in conditions A and B in the experiment who are exposed to ads for restaurants on which we have restaurant-attribute information. The table reports on the heterogeneity in responsiveness to ad-disclosure by running a regression of visit/call indicators across all relevant user-first-sessions in conditions A and B, on a dummy for whether the session belonged to condition B ($\mathcal{I}_{i,B}$), and on interactions of this dummy with $d_{r,\text{Num-rating}}$, $d_{r,\text{Ave-rating}}$, $d_{r,\text{Price-Index}}$ and $d_{r,\text{Date-Added}}$, which represent the within-zone deciles of a restaurant r on the respective restaurant characteristics. Standard errors are clustered at the user level. All coefficients multiplied by 100 to express as %. Restaurants that are newer and with fewer ratings (about which users have more uncertainty) are expected to benefit more from the signal. The table suggests these patterns hold.

rated fewer -- by one decile at baseline, is likely to experience an incremental effect of ad-disclosure of +0.02%. This increase is equivalent to 60% of the average baseline call rate of 0.03%.

Overall Changes in Consumer Choice Under Disclosure We now assess whether consumers are made better or worse off under disclosure. A formal analysis of consumer welfare would require taking a stance on a particular representation of utility. Instead, we ask whether consumers are more likely to call better or worse restaurants in the disclosure condition, where we judge "better" or "worse" on the basis of the average baseline rating of the called restaurants. To be clear, this assessment is not picked up in our analysis above. While we documented there that users call *advertised* restaurants at a higher rate in the disclosure condition, it could be that these calls are coming at the expense of calls to better, non-advertised restaurants. If so, that suggests users may be worse off under disclosure.

To assess these, we compare calls to all restaurants (including those that did not advertise) between individuals in condition A and B and examine how ad-disclosure changes users' call choices. As an added test, if ad-disclosure works as a signal and makes the consumers better informed, we again expect the users in condition B to pick options that might appear "risky", *a priori*.

As before, we use a regression setup for this analysis. Let $I_{i,B}$ be an indicator of whether user i is in condition B. Let $d_{r,\text{Num-rating}}$, $d_{r,\text{Ave-rating}}$, $d_{r,\text{Price-Index}}$ and $d_{r,\text{Date-Added}}$ denote the within-zone deciles of restaurant r on the respective variables. Represent by Y_{ir} an indicator for whether user i calls restaurant r. Stacking across all user, first-session, and all restaurant combinations in conditions A and B, we run the following regression,

$$Y_{ir} \ = \ \alpha_r + \beta I_{i,B} + I_{i,B} \times [\gamma_1 d_{r,\text{Num-rating}} + \gamma_2 d_{r,\text{Ave-rating}} + \gamma_3 d_{r,\text{Price-Index}} + \gamma_4 d_{r,\text{Date-Added}}] + \epsilon_{ir}$$

Note that we include a fixed effect for each restaurant, therefore, the main effects of the restaurant characteristics are not included. Each row in the regression is a user-restaurant combination, which, with roughly 221K users and 140K restaurants would amount to roughly about 31B observations. To make the regression manageable, we consider 10,834 restaurants which received at least one call from any user in conditions A or B. Therefore, we have about 2.4 billion observations (221K users in conditions A or B $\times$ 11K restaurants), and we cluster the standard errors at the user-level, which is the unit of randomization.

The estimates from this regression are presented in Table (10). For ease of interpretation, the coefficients are scaled by the average probability of a call in the 2.4B observations (i.e., the coefficients are divided by 9.13×10^{-6}). We run separate regressions with and without including $d_{r,\text{Date-Added}}$ because this measure is missing for about a quarter of the observations, and is consequential for the precision of the estimates. Column 1 shows that the coefficient corresponding to $\mathcal{I}_{i,B} \times d_{r,\text{Num-rating}}$ is negative and $\mathcal{I}_{i,B} \times d_{r,\text{Ave-Rating}}$ is positive. The coefficient corresponding to $\mathcal{I}_{i,B} \times d_{r,\text{Price-Index}}$ is negative but statistically insignificant.[24] When we drop the observations with missing data (column 2), the estimates becomes imprecise, but remain of the same sign. Looking at Table (10), the effect of ad-disclosure decreases by 0.0191 of the average call probability, (a change of 2% to the baseline) if a restaurant moves up by one decile in the number of ratings it receives. Put another way, moving a restaurant from the lowest to the highest decile would produce an increase in call rate that is about 20% of the baseline effect. The effect of average ratings is comparable.

These estimates suggest that consumers are more likely to call restaurants that have higher ratings and have received fewer ratings in the past, when they are in experimental condition B compared to when they are in condition A. This suggests that users' choices are shifting systematically toward options that

[24]Note that the numbers in the table are very small in magnitude compared to the numbers in the previous tables because of the regression setup. This regression considers all combinations of users and restaurants (not just advertisers), including restaurants that (1) may be less popular than the advertisers as noted in section 4, and (2) less likely to be relevant to the user compared to the advertiser whose ad is shown to the user. Both these reasons decrease the baseline propensity of a user calling a restaurant.

Table 10: Heterogeneity in User Choice across all restaurants to Ad-Disclosure

Dependent Variable:	(1) Call		(2) Call	
Variable	*Parameter*	t-stat	*Parameter*	t-stat
$\mathcal{I}_{i,B} \times d_{r,\text{Num-rating}}$	**-0.0191**	**-2.05**	0.0110	-0.92
$\mathcal{I}_{i,B} \times d_{r,\text{Ave-rating}}$	**0.0187**	**2.22**	0.0141	1.48
$\mathcal{I}_{i,B} \times d_{r,\text{Price-Index}}$	-0.0139	-1.84	-0.0125	-1.48
$\mathcal{I}_{i,B} \times d_{r,\text{Date-Added}}$			-0.0008	-0.11
$\mathcal{I}_{i,B}$	0.1117	1.86	0.0931	1.08
Intercept	**1.00**	**51.10**	**1.0011**	**46.96**
Fixed effect for each restaurant		Yes		Yes
Num. Individuals (clusters)		221,742		221,742
Num. Restaurants		10,843		8,266
Num. Observations		2.4B		1.8B
R^2		0.00		0.00

Notes: The Table reports on 221,742 users in conditions A and B in the experiment. The table presents how the characteristics of restaurants called by individuals changes due to ad-disclosure by running a regression of call indicators across all combinations of users and restaurants (that got at least one call from any individual in conditions A or B), on a dummy for whether the user belonged to condition B ($\mathcal{I}_{i,B}$), and on interactions of this dummy with $d_{r,\text{Num-rating}}$, $d_{r,\text{Ave-rating}}$, $d_{r,\text{Price-Index}}$ and $d_{r,\text{Date-Added}}$, which represent the within-zone deciles of a restaurant r on the respective restaurant characteristics. Standard errors are clusters at the user level.

are better rated but, presumably, perceived as risky without the information conveyed by advertising. Net-net, these are again consistent with the signaling mechanism, and suggest that consumers tend to go to better restaurants under disclosure.

6 Robustness

To close the paper, we discuss some statistical considerations pertaining to our main results; discuss some robustness to alternative explanations for our findings; and report on sensitivity analysis related to our results.

Statistical Considerations We use as our main test statistic the difference in mean call rates between the two conditions.[25] Though we have a large number of observations in each condition and standard results based on the Central Limit Theorem apply, one may worry that the normal approximation to the distribution of the test statistic under the null may be poor given the small number of realized calls in both

[25]With binary outcomes and independent samples, the t and χ^2 tests are equivalent, and yield the same p-value.

conditions.[26] To assess this, we report on a simulation that computes an exact p-value that does not rely on the normal approximation. We simulate the sampling distribution of the test statistic assuming the null that disclosure has no effect is true. To do this, we block-sample at the user-level with replacement two datasets from the no-disclosure condition, one with $n = 44,637$ users and the other with $n = 177,105$ users, mirroring our setup. We then take the difference in means of the call indicators between the two datasets, and repeat this procedure 10,000 times to obtain 10,000 such mean differences. We plot the empirical CDF of these in Figure (9), representing the empirical distribution of our test statistic under the null. Denote this empirical CDF as $\hat{\mathcal{F}}_d(d)$. The observed value of the test statistic in the data is 0.024% (0.055–0.031, see Table 4). This is represented as vertical red lines on the plot. The chance of seeing a value more extreme than 0.024% under the null, $\hat{\mathcal{F}}_d(-0.024) + 1 - \hat{\mathcal{F}}_d(0.024)$, is $\frac{11}{10000} = 0.0011$. This can be interpreted as an exact p-value based on the empirical CDF, and not the normal approximation. This is very similar to the 0.002 value reported in Table 4. Visually inspecting Figure (9) also shows that an observed difference of 0.024% would be highly unusual under the null. Overall, we conclude that there is a robust statistically significant difference in the call probability between the two conditions.[27]

Gelman and Carlin (2014) caution that the risk of obtaining a statistically significant estimate of the wrong sign, as well as the risk of obtaining an overstated statistically significant effect, are high if a study is under-powered. To assuage this concern, they suggest ex-post verification of type-S error rates and exaggeration ratios. They define the type-S error rate as the probability in many replications, that the replicated estimate has the incorrect sign, if it is statistically significantly different from zero. They define the exaggeration ratio as the average in many replications, of the absolute value of the replicated estimate divided by the *a priori* expected effect size, if it is statistically significantly different from zero. We simulate these by bootstrapping. We block-sample at the user-level with replacement two datasets, one with $n = 44,637$ users from the no disclosure condition and the other with $n = 177,105$ users from the disclosure condition. We then run a regression of the calls in these data on an indicator for condition B, clustering at the user level, and store the associated coefficient and p-value. We repeat this procedure 10,000 times to obtain 10,000 such coefficients and p-values.

In the 10,000 replications, we found 18 of the coefficients are negative (i.e., wrong sign). But all of these 18 cases are ones which have a p-value > 0.05, i.e., not a statistically significant replication. Thus, the implied probability of a statistically significant wrong sign = type-S error rate = 0.

To assess the possibility of exaggeration, we need to pick an *a priori* expected effect size. We are not aware of reliable past studies of the effect of search-ad disclosure on demand for restaurants which could

[26] The binomial distribution with parameters (N, p) is skewed when the success probability, p is close to 0 or 1, and the normal approximation may be poor unless the number of trials, N, is very large (Blythe and Still 1983; Samuels and Lu, 1992).

[27] The two-tailed p-value of the test is also low enough to alleviate a concern related to multiple hypothesis testing. For example, above, we conducted two tests comparing the disclosure and no-disclosure condition. Adjusting the p-value for multiple hypothesis testing using a conservative Bonferroni correction yields $p = .004$ (.002×2).

Figure 9: Assessing p-values by Simulation

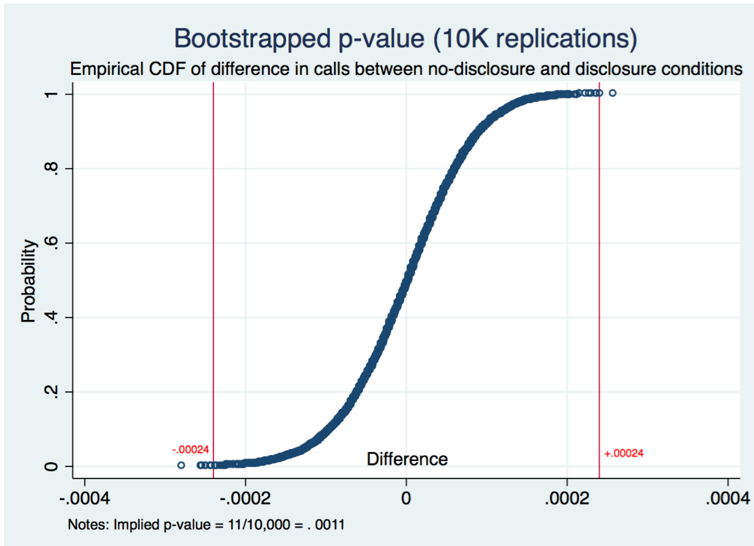

Notes: The figure shows the sampling distribution of the test statistic assuming the null that disclosure has no effect is true. The observed value of the test statistic in the data is 0.024% (0.055–0.031, see Table 4). This is represented as vertical red lines on the plot. The chance of seeing a value more extreme than 0.024% under the null is $\frac{11}{10000} = 0.0011$, which can be interpreted as an exact p-value based on the empirical CDF, and not the normal approximation.

be the basis of forming an *a priori* expected effect size. Hence, we compute the ratio of the statistically significant coefficients to the estimate we obtain (i.e., 0.00024) and present a histogram in Figure (10) representing the distribution of the ratios across the 10,000 replications. The mean (exaggeration) ratio is 1.207 (a very low value per Gelman and Carlin's paper).

Finally, we assess whether the fact that we do not detect disclosure causing a statistically significant change in page-visits reflects low power. We assess power *ex-post* by simulation.[28] We find that if disclosure increases page-visits by 77% (as observed for calls) we would have detected it almost surely (power close to 100%). If disclosure increases page-visits by only 20%, we find that again we would have detected it almost surely (power $> 99\%$). If disclosure increases page-visits by only 10%, we find we would have detected it in about 70% of the cases. Broadly we conclude that if disclosure increases page-visits, we can confidently state that the effect is likely to be smaller than the effect on calls.

Do Calls Change without a Change in Demand? Calls represent a proxy for demand. One possibility for an erroneous conclusion may arise if calls increase in the disclosure condition relative to

[28]We compute the probability of rejecting the null that disclosure has no effect, assuming the alternative (that disclosure has an effect) is true. We block-sample at the user-level with replacement two datasets, one with $n = 44,637$ users and the other with $n = 177,105$ users from the disclosure condition. We then run a regression of the calls in these data on an indicator for whether the observation belongs to the second dataset (representing condition B), clustering at the user level, and store the associated p-value. We repeat this procedure 10,000 times to obtain 10,000 such p-values. We compute power as the proportion of the p-values < 0.05.

Figure 10: Assessing Possibility of Exaggeration Error by Simulation

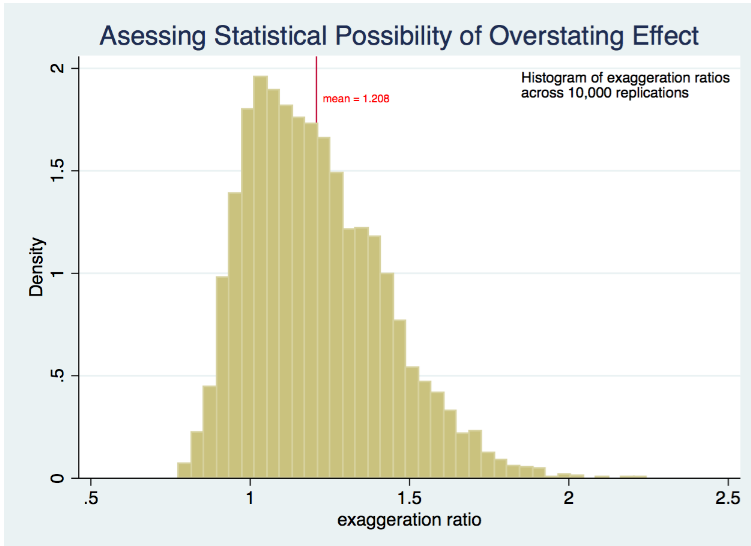

Notes: The figure shows the distribution across 10,000 replications of a measure of the possibility that the disclosure effect is overstated. In each replication in which disclosure is estimated to be statistically significantly different from zero, the absolute value of the estimated disclosure effect is divided by 0.0024 to obtain a ratio. The distribution of these ratios for the across the 10,000 replications is plotted above. The mean ratio is 1.207 (a low value according to Gelman and Carlin 2014).

no disclosure, but actual demand does not change. For example, knowing that a restaurant advertised may cause a consumer to be concerned about the restaurant's seating capacity and call the restaurant to make reservations. Such a call may not occur in the condition without ad-disclosure, causing calls to increase while demand remains unchanged. To address these kinds of considerations, note first that such an explanation implicitly presumes a change in the consumer's belief caused by his knowing that the restaurant advertised. For example, the above example assumes that knowing a restaurant is advertised changes the consumer's beliefs about market demand for the restaurant (causing her to be concerned about capacity). If that's the case, this is consistent with a signaling effect. Second, the analysis below addresses this more directly , by checking whether the effect of ad disclosure exists in the subset of cases where the search category chosen by the consumer is "home-delivery". In this case, we know for sure the user is looking to order food for delivery, wherein a call is likely to match closely with demand.[29] Table (11) shows the means across the three conditions split by the home delivery search filter. The table shows that disclosure produces incremental conversion in both categories, and the pattern reported earlier continues to hold even for calls with the home-delivery filter.

[29]Unlike situations in which a consumer is dining at the restaurant, making a phone call is the primary channel to order food for delivery. Though we find in our analysis of call recordings that a large proportion of calls actually involve requests for delivery, consumers do seem to always add-on this filter to their search. If a user added the "home-delivery" filter to his search, we know it involves delivery for sure, so we're being conservative.

Table 11: Call Probability Split by Experimental Condition and by “Home-Delivery” Search Filter

		Home delivery Searches				Non home delivery Searches			
Condition	*Description*	*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*		*Mean*	*Std. Err.*	*p-value of test, H_0: Equal means*	
				Between control & A	*Between* A *&* B			*Between control &* A	*Between* A *&* B
Control	No ad	0.043%	0.011%	0.82		0.018%	0.007%	0.62	
A	Ad with no disclosure	0.039%	0.010%		0.026	0.024%	0.008%		0.03
B	Ad with disclosure	0.067%	0.06%			0.044%	0.005%		

Notes: The table splits the 265,975 users into two groups, those that searched for home delivery (126,665 users; 212,602 user-advertiser observations) and the rest (139,310 users; 226,823 user-advertiser observations). It reports the call probabilities, separately by these two groups of users, across all relevant first-sessions in each condition. p-values computed by running a regression of the call indicators for users on an indicator for the relevant contrasted condition, with standard-errors clustered at the user-level. We expect calls corresponding to home deliveries to represent orders directly, and not other user motivations for calling. We see that ad-disclosure (condition B) produces incremental calls in both types of searches.

Is Ad-Disclosure Simply Catching User's Attention? We explore whether the reason we see increased user response in condition B is because the disclosure catches user's attention and makes the listing more salient. First, note that if disclosure was catching user attention, we would expect it to have a larger impact on page-visits than call. However, we found that ad-disclosure does not affect page visits significantly but increases calls, as reported in the main results. This result is inconsistent with the view that disclosure works by simply catching users' attention on the platform. Further, we describe a test that is incorporated into the experimental design to assess this directly. The test corresponds to one of the sub-conditions considered in more detail in a companion paper; so we only describe the set-up and the results here, pointing the reader to Sahni and Nair (2016) for exact details.

To implement the test, we utilize the fact that condition "B" with the disclosure comprises two sub-conditions "B1" and "B2" such that an advertised listing in B2 is more prominent to the user than in B1. In particular, condition C2 adds a bold outline over the paid listings. B2 is more likely to catch a user's attention, so if salience is the channel by which disclosure operates, we should see improvements in user outcomes for the advertised restaurants in B2 relative to B1. We find that the improvement in visits and calls in conditions B1 with disclosure and no outline look similar to the patterns reported in Table (4) previously. We find that there is no statistically significant difference between B1 and B2 in terms of both visits and calls, suggesting that salience is not the channel by which disclosure operates in this setting.

Are the Effects Driven by Adverse User Reactions to No-Disclosure? We now check whether there is support for alternative phenomena that explain the difference in outcomes between conditions A and B by positing negative user reactions to condition A (Figure 11). For instance, one possibility is that users experience some "cognitive dissonance" (Festinger 1957) in condition A because the listing shown in the ad-slot (i.e., the advertiser) is not a great match to their search and tastes. This may reduce their trust in the search-engine, and cause them to respond unfavorably to listings in condition A. Dissonance may be lower in condition B because users see there the listing is a paid-ad, and do not blame the search-engine for any reduced match. Alternatively, consumers in condition A may lose trust in the platform because they see a similar listing for the advertised restaurant twice − once at the top (i.e., the advertisement without disclosure), and then in the search feed as an organic listings if they scroll down.

A few aspects of the setting suggest that such consumer reactions are not the driving reasons behind the reported responsiveness to disclosure. First, the experiment design allows advertising to appear only in searches that are broadly specified as described previously. This feature ensures that the advertised options seen in the experiment are not very different from the user's expectation from the search results, limiting the potential for such cognitive dissonance. Second, if dissonance exists, it should be higher for advertisers who get placed on the top ad-slot but have bad reviews. So, the disclosure should help "mute"

Figure 11: Possibility of Adverse Reaction to Ads: No Disclosure Condition

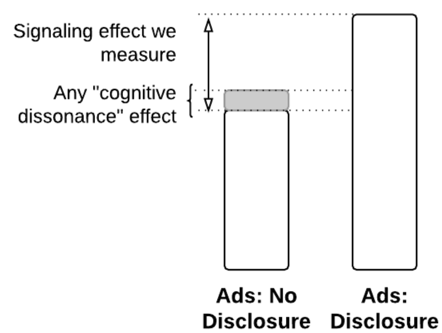

Notes: The figure demonstrates the possibility that the difference between the disclosure and non disclosure conditions may be driven by dampening induced in the non disclosure condition. For instance, it could be could be that users experience some "cognitive dissonance" in condition A if the advertised listing on the top ad-slot is not the best match to their search and tastes.

the dissonance and help restaurants with bad reviews more when they get placed on the ad-slot. This is not what we find − ad-disclosure helps advertisers with better reviews more. Third, if restaurants realize that advertising will help "mute" the dissonance in this manner, those with bad reviews will advertise more on the platform, because they incorporate the fact that the advertising will mitigate their bad reviews. This is not the pattern in the data − restaurants with better reviews tend to advertise more on the platform.

Finally, a reduction in user trust of the platform should manifest itself in reduced calls of all restaurants, including the advertisers and non-advertisers. We do not find this to be the case. This can be seen in Table 10; the coefficient of indicator of condition B is insignificant. (If we run the same regression leaving out the interaction terms, we find the coefficient for the indicator of condition B remains statistically indistinguishable from zero).

Boundary Conditions and Generalizability All things held equal, the signaling value of advertising is likely to be high when consumer uncertainty about product quality is high; when repeat-purchase is more likely; and when the provision of advertising to consumers is costly. Accordingly, the extent to which these factors are relevant moderate the size of the signaling role of advertising. In the paid-search context, consumer uncertainty about the appeal of the restaurant is dependent on the quality of the search-engine's recommendation algorithm and the information content of ratings. The effect we are measuring in our context should be interpreted as the effect of advertising over and beyond the value of ratings and the order of recommended organic listings. In a world where the search-engine's organic listings algorithm is close to perfect and ratings convey all relevant information about restaurant appeal, the signaling role of advertising will be diminished. When restaurant demand becomes more sticky, repeat purchase is more likely, and hence, the value to a high quality firm of advertising and signaling quality increases. Hence, advertising is more likely to play a signaling role in markets with strong repeat purchase propensity.[30] Finally, the credibility of the ad-signal is higher when the cost of delivering an ad-impression is high. Obtaining the top ad-slot in response to a localized directed search on a popular search engine like Zomato is costly. This cost is ultimately driven by the nature of competition in the ad-market and the contractual mechanism by which ads are sold on the platform. Accordingly, as these factors change, the magnitude of the measured signaling effect will also vary.

7 Conclusions

A field experiment to assess the signaling role of advertising is presented. The experiment randomizes users into treatment and control conditions that enable us to measure the causal effect of disclosure to a consumer that the firm has advertised, separately from the content of the search ads served by the platform to the user, thereby pinning down signaling using a "demand-side" strategy. The effect of disclosure on conversion to the advertised restaurants is found to be large, and translates to about USD$2.7 million in incremental ad-revenue to the platform. Separately from testing for signaling effects, this finding also holds importance for platform design and monetization for search platforms because it implies that disclosure of the sponsored nature of advertising is beneficial. Recent reports from online publishers that clearly disclosing the sponsorship status of non-annoying ads improves outcomes, is consistent with this finding (e.g., Moses 2016).[31]

[30] As an empirical test of this, we tried to see whether we could isolate a set of restaurants in our data that were situated closer to tourist spots and less likely to have repeat customers. Unfortunately, we were unable to pin down enough restaurants to reliably run this test.

[31] The article reports that Slate "found that on its more explicitly labeled ads, the click-through rates were three times higher than the previous units (though Slate wouldn't disclose the CTR). The publisher also contends that average time spent on the new units doubled, to 4 minutes, 15 seconds."

References

1. Ackerberg, D. A. (2001), "Empirically Distinguishing Informative and Prestige Effects of Advertising," Rand Journal of Economics, 32, 316-33.

2. Anand, B. and R. Shachar. (2011). "Advertising, the Matchmaker," RAND Journal of Economics 42(2) 205-245.

3. Archibald, R. B., Haulman, C. H., and C. E. Moody, Jr. (1983), "Quality, Price, Advertising and Published Quality Ratings," Journal of Consumer Research, 9, 347-56.

4. Backus, M., T. Blake. and S. Tadelis, (2016), "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining," working paper, Univ. of Berkeley.

5. Blake, T., Nosko, C. and S. Tadelis. (2015), "Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment," Econometrica, 83(1), p. 155-174.

6. Blyth, C.R., and H.A. Still. (1983). "Binomial Confidence Intervals," Journal of the American Statistical Association, 78(381), 108-16.

7. Caves, R. E. and D. P. Greene (1996), "Brands' Quality Levels, Prices and Advertising Outlays: Empirical Evidence on Signals and Information Costs," International Journal of Industrial Organization, 14, 29-52.

8. Chan, T.Y., Wu, C., and Xie, Y. (2011). "Measuring the Lifetime Value of Customers Acquired from Google Search Advertising," Marketing Science, 30 (5), 837-850.

9. Erdem, T., M. Keane and B. Sun (2008), "Advertising and Consumer Price Sensitivity in Experience Goods Markets," Quantitative Marketing and Economics, 6 (2), 139-176.

10. Festinger, L. (1957). "A Theory of Cognitive Dissonance." California: Stanford University Press.

11. Gelman, A. and J. Carlin (2014), "Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors," Perspectives on Psychological Science, 9(6), 641- 651.

12. Goldfarb, A. and Tucker, C. (2011). "Search Engine Advertising: Channel Substitution when Pricing Ads to Context," Management Science, 57(3), 458-470.

13. Goeree, M.S. (2008). Limited Information and Advertising in the US Personal Computer Industry," Econometrica, 76(5), 1017-1074.

14. Hertzendorf, M. (1993), "I'm Not a High-Quality Firm - But I Play One on TV," Rand Journal of Economics, 24, 236-47.

15. Holland, P. W. (1986). "Statistics and Causal Inference." Journal of the American Statistical Association 81(396): 945-960.

16. Homer, P. M. (1995), "Ad Size as an Indicator of Perceived Advertising Costs and Effort: The Effects on Memory and Perceptions," Journal of Advertising, 24, 1-12.

17. Horstmann, I. J. and G. M. MacDonald (1994), "When Is Advertising a Signal of Quality?," Journal of Economics and Management Strategy, 3, 561-84.

18. Horstmann, I. J. and G. M. MacDonald (2003), "Is Advertising A Signal of Product Quality? Evidence from the Compact Disc Player Market, 1983-92," International Journal of Industrial Organization, 21, 317-45.

19. Horstmann, I. J. and S. Moorthy, (2003). "Advertising Spending and Quality for Services: The Role of Capacity," Quantitative Marketing and Economics, 1, pp. 337-65.

20. Johnson, G. (2013). "The Impact of Privacy Policy on the Auction Market for Online Display Advertising," Simon School Working Paper No. FR 13-26.

21. Kawai, K., K. Onishi and K. Uteke, (2014), "Signaling in Online Credit Markets," working paper, Yale University.

22. Kihlstrom, R. E. and M. H. Riordan (1984), "Advertising as a Signal," Journal of Political Economy, 92, 427-50.

23. Kirmani, A. (1990), "The Effect of Perceived Advertising Costs on Brand Perceptions," Journal of Consumer Research, 17, 160-71.

24. Kirmani, A. (1997), "Advertising Repitition as a Signal of Quality: If It's Advertised So Often, Something Must Be Wrong," Journal of Advertising, 26, 77-86.

25. Kirmani, A. and A. R. Rao (2000), "No Pain, No Gain: A Critical Review of the Literature on Signaling Unobservable Product Quality," Journal of Marketing, 64, 66-79.

26. Kwoka, Jr., J. E. (1984), "Advertising and the Price and Quality of Optometric Services," American Economic Review, 74, 211-16.

27. Lewis, R. and Reiley, D. (2014). "Online Ads and Offline Sales: Measuring the Effects of Retail Advertising via a Controlled Experiment on Yahoo!," Quantitative Marketing and Economics, 12, 235-266.

28. Moses, L. (2016), "Slate now relies on native ads for nearly 50 percent of its revenue," `http://digiday.com/publishers/slate-now-relies-native-ads-nearly-50-percent-revenue/`, accessed Dec 14, 2016.

29. Manchanda, P., Dube, J.P., Goh, K.Y., and Chintagunta, P.K. (2006). "The Effect of Banner Advertising on Internet Purchasing," Journal of Marketing Research, pp. 98-108.

30. Milgrom, P. and J. Roberts (1986), "Price and Advertising Signals of Product Quality," Journal of Political Economy, 94, 796-821.

31. Narayanan, S. and Kalyanam, K. (2015), "Position Effects in Search Advertising and their Moderators: A Regression Discontinuity Approach," Marketing Science, 34(3), 388-407.

32. Nelson, P. (1970), "Information and Consumer Behavior," Journal of Political Economy, 78, 311-29.

33. Nelson, P. (1974), "Advertising as Information," Journal of Political Economy, 82, 729-54.

34. Orzach, R., Overgaard, P.B., Tauman, Y. (2002). "Modest Advertising Signals Strength," Rand Journal of Economics, 33(2), 340-358.

35. Rutz, O. J and Bucklin, R.E. (2011). "From Generic to Branded: A Model of Spillover in Paid Search Advertising," Journal of Marketing Research, 48 (1), 87-102.

36. Sahni, N. (2015), "Effect of Temporal Spacing between Advertising Exposures: Evidence from Online Field Experiments," Quantitative Marketing and Economics, 13(3), pp. 203-247.

37. Sahni, N. and Nair, H. (2016), "Native Advertising, Sponsorship Disclosure and Consumer Deception: Evidence from Mobile Search-Ad Experiments ," working paper, Stanford GSB.

38. Samuels, N. L., and T.-F. Lu. (1992). "Sample Size Requirements for the Back-of-the-Envelope Binomial Confidence Interval," The American Statistician, 46(3), 228-231.

39. Shum, M. (2004), "Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Markets," Journal of Economics and Management Strategy, 13, 241-71.

40. TechCrunch (2015), "Restaurant Discovery Site Zomato Eats Up $60M To Fuel Expansion Into New Verticals," http://techcrunch.com/2015/09/07/restaurant-discovery-site-zomato-eats-up-60m-to-fuel-expansion-into-new-verticals/, accessed Dec 8, 2015.

41. Tellis, G., and C. Fornell (1988), "Advertising and Quality Over the Product Life Cycle: A Contingency Theory," Journal of Marketing Research, 15 (1), 64-71.

42. Terui, N., M. Ban, G.M. Allenby. (2011), "The Effect of Media Advertising on Brand Consideration and Choice," Marketing Science, 30(1), pp. 74-91.

43. Thomas, L., S. Shane and K. Weigelt (1998), "An Empirical Examination of Advertising as a Signal of Product Quality," Journal of Economic Behavior & Organization, 37 (4), 415-430.

44. Tyler, J. H., Murnane, R. J., and Willett, J. B. (2000). "Estimating the Labor Market Signaling Value of the GED," Quarterly Journal of Economics, 115(2), 431-468.

45. Ursu, R. (2016). "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions," working paper, Stern School of Business, NYU.

46. Weiss, A. (1995), "Human Capital Versus Signaling Explanations of Wages," Journal of Economic Perspectives, 9(4):133-54.

47. Yao, S. and Mela, C.F. (2011). "A Dynamic Model of Sponsored Search Advertising," Marketing Science, 30 (3), 447-468.

48. Yang, S. and A. Ghose, (2010). "Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?," Marketing Science, 29 (4), 602-623.

49. Yelp.com (2015), "Why Yelp? Yelp for Business Owners," https://biz.yelp.com/support/what_is_yelp, accessed Dec 8, 2015.

50. Yelp Fact Sheet (2015), "An Introduction to Yelp Metrics as of September 30, 2015," http://www.yelp.com/factsheet, accessed Dec 8, 2015.

51. Zhang, J. and Liu, P. (2012), "Rational Herding in Microloan Markets," Management Science , Vol. 58 (5), pp. 892 -912.

Appendix A: Creating the First-Session Dataset

This appendix describes how the first-sessions database used for analysis in the paper is constructed. For each user i (in any of the three conditions), we find the first session on the mobile app in which i is in a search filter that would have caused Zomato to serve him an ad. This session could have one or several such searches. Call the filter of the first search f. On that day, Zomato could have advertised several restaurants in the top position to users searching under filter f. Collect the identities of these restaurants in a vector r. We record all of i's activity related to the restaurants in r during that session and store these in a "first sessions file". The "first sessions file" constructed in this manner serves as the dataset for the results reported in section (5.1.2).

Essentially, what we are doing in this procedure is to match outcomes between the conditions on the filters f. When filter f is seen in a search in the control condition, all we can say is that Zomato would have shown ads for one of the restaurants in r if that search occurred in A or B. Hence, we compare the users' behavior for all the restaurants in r across the conditions. Figure (12) shows a schematic with an example detailing how the data are constructed for two users in the control and treated cells.

It could be that calls and page-visits are highly correlated within individual. Checking this in the data, we find that the intra-class correlation (class being an individual) for outcome measures is very low:

Figure 12: Building the First-Sessions Database

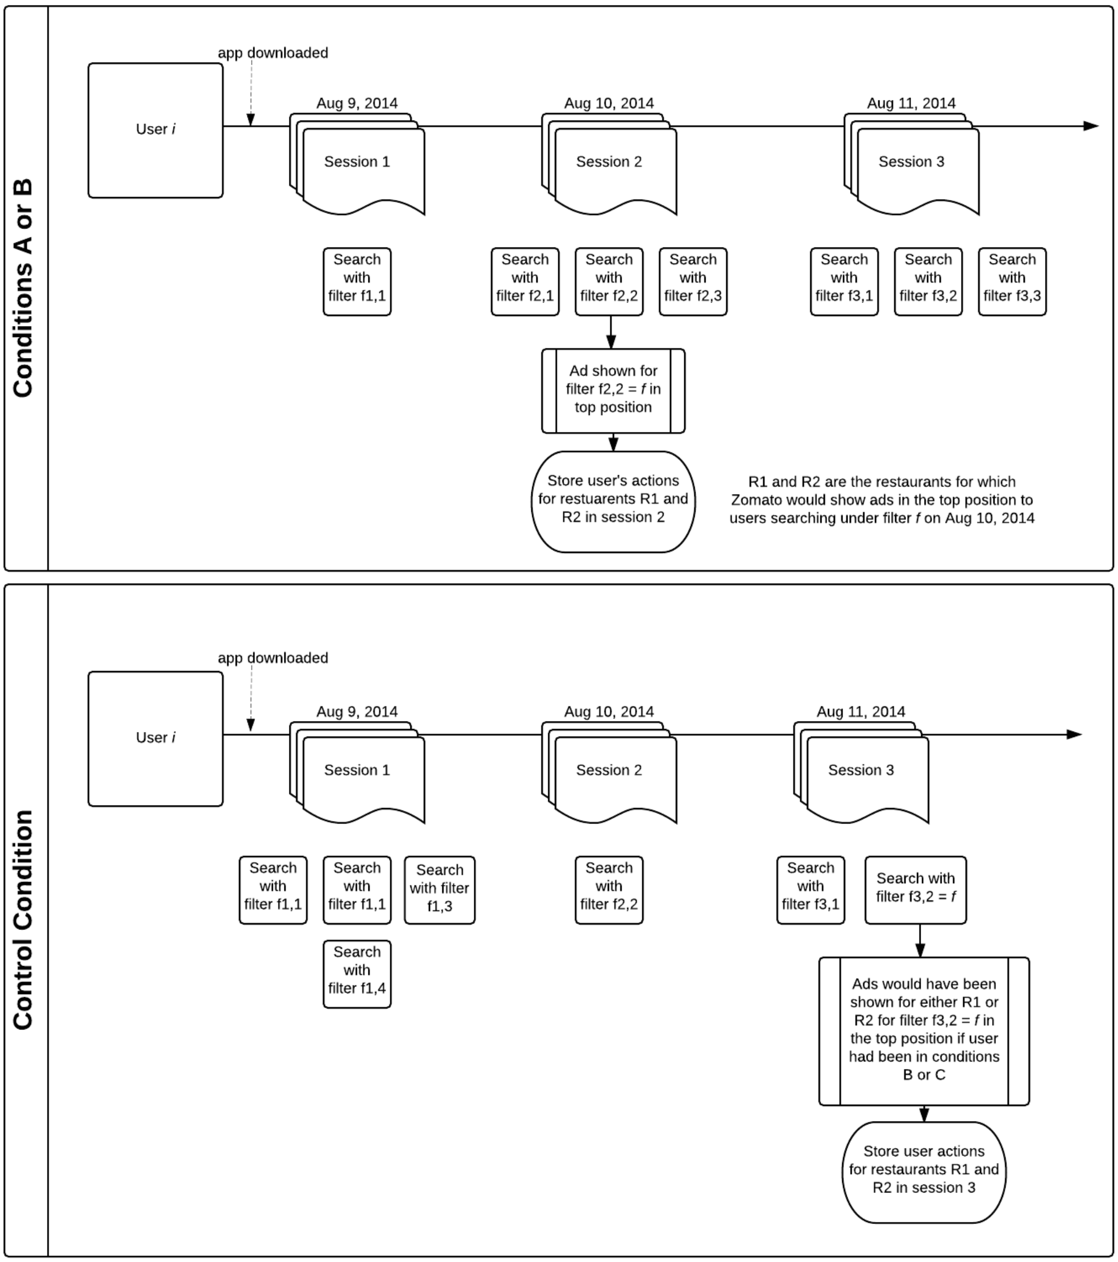

Notes: The figure shows how the first sessions database is built up. The sequence of sessions for a generic user i who downloaded the app on Aug 9, 2014 in conditions A or B is shown in the top panel. The user's session 1 features one search, and sessions 2 and 3 features three searches. The user's second search in session 2 has a filter f that would have caused Zomato to show him an ad if it were on the website. We identify this as the first search for that user that is in a filter that generates an ad. For a generic user in the control condition (bottom panel), we check if there is any search that matches filer f, and if so, we find the first such search, here search two in session 3. If this user were in conditions A or B, he would have been shown ads for restaurants R1 or R2. Thus, the users' behavior with respect to R1 and R2 in the control condition with *no* ads represent a counterfactual to his behavior with respect to R1 and R2 in conditions A and B *with* ads. We store the actions for R1 and R2 for the top-panel user in his session 2, and for the bottom-panel user in his session 3 to build up the first sessions database.

Figure 13: Quantile-Plot of the Estimated p-values

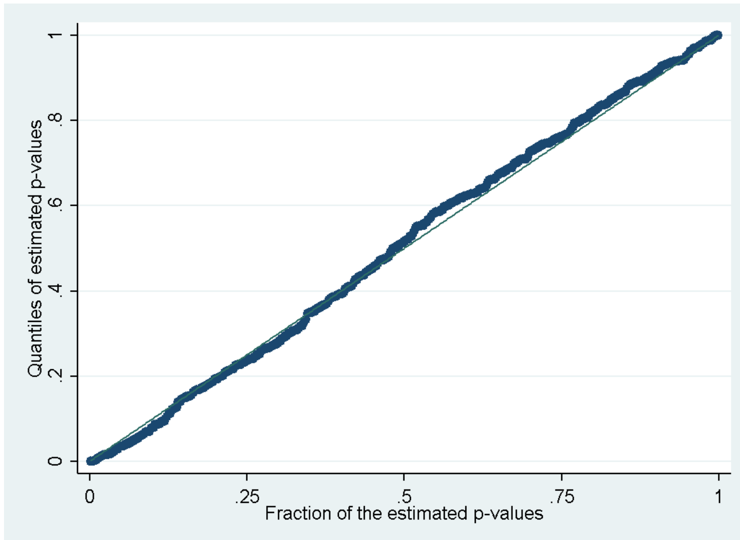

Notes: The figure plots the estimated p-values against a uniform CDF (45 degree line). For each advertising restaurant we estimate a p-value from testing the following null hypothesis: the proportion of individuals that could have been exposed to the ad are equal across the three experimental conditions. The figure shows that the p-values are uniformly distributed between 0 and 1, which is what we expect if the advertisers are not systematically different across the experimental conditions.

0.02 for page visits and -0.0004 for calls. Nevertheless, to allow for arbitrary correlation within individual, we allow for observations to be clustered by individual and estimate standard errors accordingly.

Because of randomization, we expect the identity of the advertisers to be the same across the first sessions of individuals in the three experimental conditions. We check whether the data support this. For each advertising restaurant, we compute the proportion of users who could have been exposed to it in each condition. Then, for each advertising restaurant, we test the null hypothesis that these proportions are equal across the three conditions. This gives us a p-value corresponding to each advertising restaurant. Figure (13) plots the p-values against the uniform CDF. A visual inspection shows that the distribution of p-values is close to uniform. A Kolmogorov-Smirnov test is unable to reject the null hypothesis that the p-values are uniformly distributed (p-value $= 0.49$), as expected. We conduct a similar exercise to check if the first-sessions are equally distributed over the days of the experiment (August 9 to September 26, 2014) across the three conditions. A similar test as above supports this assertion (A Kolmogorov-Smirnov test yields a p-value $= 0.19$).

Appendix B: Comparison to the Effect of Num(Ratings)

Let $d_{r,\text{Num-rating}}$, $d_{r,\text{Ave-rating}}$, $d_{r,\text{Price-Index}}$ and $d_{r,\text{Date-Added}}$ denote the within-zone deciles of advertiser r on the respective variables and let Y_{ir} an indicator for whether user i calls r. Stacking across all user, first-session, and advertiser combinations in the top position in group A, we estimate the following regression (t-stats clustered at the advertiser in parentheses),

$$Y_{ir} = \underset{(-.84)}{-.0003044} \underset{(-1.28)}{-.0000774} d_{r,\text{Ave-rating}} \underset{(2.32)}{+.0001202} d_{r,\text{Num-rating}} \underset{(.13)}{+4.55\text{e-}06} d_{r,\text{Price-Index}} \underset{(1.02)}{+.0000278} d_{r,\text{Date-Added}}$$

At these estimates, an increase in the total number of ratings by one decile is statistically significantly associated with a .012% increase in the call probability. We estimated the disclosure effect for an advertiser as a .024% increase in the call probability. So, roughly speaking, providing disclosure induces the equivalent of a two decile increase in total ratings for a firm in a world without the disclosure.

Online Appendix A: What do Calls to Restaurants Signify? (Not for Publication)

We study the effect of advertising on users making calls to restaurants. This appendix examines data on actual calls made by consumers to document the extent to which it represents purchase intent, and serve as a proxy for demand. The collaborating restaurant search platform and restaurant advertisers regard calls to be sales-leads that are a proxy for sales in the markets we consider.

Data gathering methodology In order to collect the data, calls originating from the platform to a sample of advertising restaurants are recorded. The specific procedure is as follows. Each restaurant whose calls are studied is allotted a unique phone number owned by the platform. If a user visits the restaurant's page during the studied time period, the platform displays the phone number allotted to the restaurant, instead of the restaurant's actual phone number. When a call is made to the allotted number, the call is then redirected to the restaurant's actual number. The redirected call is recorded and stored as an mp3 file. Before redirecting, the caller is told that the call is recorded. Monitoring of customer calls to customer-service phone numbers in this manner is common in the industry.[32]

Description of the data The data are collected in Oct-Nov of 2010 for restaurants in New Delhi and Mumbai. We analyze a sample of 1,033 calls made to 28 advertising restaurants by manually listening to the recorded mp3's of the conversations that took place between the customer and the restaurant's phone attendant.[33] Out of the 1,033 calls, 55 are inaudible to us or to one of the parties on the call. The reports below pertain to the content of the remaining 978 calls. Table 12 shows the distribution of the calls based on the caller's objective.

Calls related to orders. A majority (69.5%) of calls are made to order food for delivery. 8.5% of calls are made to reserve a table at the restaurant. Other calls involving takeout orders, catering orders, and book arrangements for a party comprise another 3.1%. In total, 81.2% of the calls are related to making a purchase (combining delivery, takeout, table reservation, party and catering orders).

Calls asking for information. 13.4% of calls comprise enquiries about the restaurant. Some ask about the location of the restaurant or whether the restaurant is open on that day. Others enquire about details on dishes served at a lunch buffet or the ingredients used in a specific dish. About a third (43, specifically) of the calls asking for information occurred for one restaurant around the day it was featured in a local newspaper.

Other calls. 3.8% of calls are made by another business marketing their product to the restaurant. An example would be a web aggregator asking for details about the restaurant. The remaining categories have only a few entries. 6 calls involve follow up on orders made before the call. A wrong number is dialed in 4 cases. We categorize 6 calls unrelated to food as part of the "other" category. One example is a call from a customer who left his credit card at the restaurant the day before.

Transcripts of a sample of calls We sample a few calls from the main categories and transcribe them to illustrate the content of the calls in various categories.[34] Information that could identify the caller or the restaurant has been masked. Several of the calls took place in English, some are translated from Hindi.

Home delivery

[32] For example, see http://www.nytimes.com/2005/01/11/business/your-call-and-rants-on-hold-will-be-monitored.html?_r=0.

[33] If a restaurant gets multiple calls from the same phone number during one day, we count that as one call. This is consistent with the dependent measure (whether or not there is a call) used in the analysis.

[34] The calls were picked randomly, but oversampled from smaller categories to illustrate examples.

Table 12: Content Across Calls

Call type	*Count*	*Fraction of total audible calls*
Delivery order	680	69.5%
Takeout order	18	1.8%
Table reservation	83	8.5%
Reserve arrangement for a party	9	0.9%
Catering order	4	0.4%
Asking for information (no purchase order)	131	13.4%
Follow up on a previously made order	6	0.6%
Marketing call	37	3.8%
Wrong number dialed	4	0.4%
Other calls	6	0.6%
Total clearly audible calls	978	100%
Call inaudible	55	
Total calls	1,033	

Notes: The table shows the distribution of 1,033 calls across the various objectives of the call. The calls are recorded with permission. We do the content classification manually by listening to the recordings. 81.2% of the calls are made to make a purchase (delivery order + takeout + table reservation + reserve a party + catering order). The majority of orders are for home delivery. 13.4% of the calls are enquiries about the restaurant's food/facilities.

Call 1
Caller: Hello
Restaurant: Hello xx
Caller: Just wanted to give a home delivery order
Restaurant: Your name?
Caller: xx
Restaurant: What is your address?
Caller: xx
Restaurant: Your phone number?
Caller: xx
Restaurant: What would you like to order?
Caller: One assorted veg starters. And one veg fried rice. That's all.
Restaurant: The starter is dry. Is that ok?
Caller: That's ok
Restaurant: And one veg fried rice, that's it?
Caller: Yes

Call 2
Restaurant: Hello
Caller: Hello
Restaurant: Yes madam
Caller: is this xx?
Restaurant: Yes
Caller: I want to make an order
Restaurant: Yes
Caller: One veg hakka noodles
Restaurant: (repeating) One veg hakka noodles
Caller: One veg fried rice
Restaurant: (repeating)
Caller: One chilly chicken with gravy
Restaurant: (repeating)
Caller: One veg manchurian with gravy
Restaurant: (repeating)
Caller: One crispy veg dry
Restaurant: (repeating)
Caller: Order in the name of Mr xx.
Restaurant: What is the address?

Caller: xx
Restaurant: And your contact number?
Caller: xx
Caller: What's the total amount?
Restaurant: Rs 750
Caller: Ok. Thank you
Restaurant: Ok

Call 3
Caller: Hello
Restaurant: Hi Good evening this is xx
Caller: I want to place an order
Restaurant: Your number?
Caller: xx
Restaurant: What is your order?
Caller: Please make it well
Restaurant: Yes sure
Caller: I want one veg hakka noodles
Restaurant: Yes
Caller: and one chicken in garlic sauce
Restaurant: Ok, anything else sir?
Caller: No that's it.
Restaurant: you'll have to pay Rs. 390. The food will reach you in 45 mins or one hour maximum
Caller: thanks
Restaurant: you're welcome

Call 4
Restaurant: Good evening xx
Caller: Wanted to make an order
Restaurant: What is the order?
Caller: 3 mutton biriyanis. 1 sabz biriyani, 1 chicken kroma curry and one paneer korma curry
Restaurant: (repeats the order). What is the address?
Caller: xx. When will you send it?
Restaurant: What is your name?
Caller: xx
Caller: When are you sending it?
Restaurant: Contact number?
Caller: xx
Restaurant: Sir, it will be delivered in 30-45 minutes.

Table reservation

Call 1
Caller: Hello, xx?
Restaurant: Yes xx.
Caller: I needed a table tonight for 14 people at 8 pm?
Restaurant: 8 pm?
Caller: Yes.
Restaurant: Ok.
Caller: Confirmed?
Restaurant: Yes absolutely. It's Sunday, so 5-10 minutes up and down around 8 is ok?
Caller: 5-10 minutes is ok, but not more.
Restaurant: Sure. What name should I put on the booking
Caller: Mr xx.
Restaurant: Contact number?
Caller: xxx
Restaurant: Ok, please come at 8pm. Thanks for letting us know your contact number. We can also call if anything changes.
Caller: Ok thank you.

Call 2
Restaurant: Hello
Caller: Hello
Restaurant: Yes
Caller: Is this xx?
Restaurant: Yes
Caller: I'd like to make a reservation

Restaurant: Yes please tell me
Caller: I'd like to book a table for two from 12:45
Restaurant: 12:45?
Caller: Yes. It might take us some time to come. We might get in by 1
Restaurant: Ok. Can I have your name?
Caller: Please reserve under xx
Restaurant: Can you please spell it?
Caller: xx (spells it)
Restaurant: And your personal number would be?
Caller: xx
Restaurant: Around about 12:45 or 1 latest?
Caller: Yes, latest by 1:05
Restaurant: Ok

Asking for information

Call 1
Restaurant: Hello xx
Caller: Hi I'm calling to check if you have space for 20 people. We're 20 people and I want to check if you can seat all of us.
Restaurant: Yes madam, this is can be done.
Caller: 20 people together?
Restaurant: Yes, together.
Caller: You can seat 20 people together?
Restaurant: Yes, madam it can be done.
Caller: Do you do this by joining tables?
Restaurant: Yes
Caller: Are you sure?
Restaurant: We can do 40 as well
Caller: If 20 come together, is there a discount?
Restaurant: No But there is a corporate discount
Caller: Is there a discount for company xx?
Restaurant: Yes, it is for any company. You'll have to give a visiting card.
Caller: Ok. If I want to book a table for tomorrow evening, what time should I call?
Restaurant: Whenever, 1 – 1.5 hours before the booking.
Caller: You'll be able to arrange in that time?
Restaurant: Yes
Caller: Thank you
Restaurant: Thank you

Call 2
Caller: Hello xx?
Restaurant: Yes xx
Caller: Are you in Noida or in Delhi?
Restaurant: Noida Sir
Caller: Ok. I'm in xx. How can I reach your restaurant?
Restaurant: Where exactly are you?
Caller: I'm close to the metro station, just parked my car.
Restaurant: Ok. Are you walking?
Caller: Yes
Restaurant: Cross the road and come into Sector 18. Where are you in xx?
Caller: Im at the main gate at the metro station
Restaurant: Can you see a temple?
Caller: I can see the metro station
Restaurant: Ok. Please come toward sector 18.
Caller: Ok
Restaurant: Here you'll see a coffee shop named xx.
Caller: Ok
Restaurant: As you come down from the metro station. You'll see a sweet shop. After crossing it turn left and then you'll see a McDonalds. We're just next to it.
Caller: Ok
Restaurant: Thank you

Follow-up

Call 1
Restaurant: xx

Caller: Yes, I had called yesterday for an order. But nobody confirmed ...
Restaurant: Yes
Caller: I ordered yesterday, for a photo cake. But I've received no confirmation. I don't know if you're delivering or not. I called yesterday xx said they'll check the email and get back but noone did.
Restaurant: What's your name?
Caller: xx
Restaurant: Madam did you call a little while ago as well?
Caller: yes
Restaurant: Madam xx is not here yet thats why noone called
Caller: xx is not here?
Restaurant: Yes that's why we didn't call back. When xx comes i'll get them to call back. Ok?
Caller: Yes please, ask them to call back. Please confirm otherwise I'll go somewhere else
Restaurant: Madam if you've talked to xx then it means the order is confirmed
Caller: Ok, because we talked once on the phone that means the order is confirmed?
Restaurant: Yes
Caller: Ok. then I expect the cake will be delivered by 4 pm. But in any case please ask them to call me.
Restaurant: Madam by 11:30 I'll let the person know and then they'll call. They are not here yet. When they come I'll ask them to call back
Caller: Ok
Restaurant: Thank you

Marketing call

Call 1
Restaurant: Hello
Caller: Hello is this xx
Restaurant: Yes madam
Caller: Good afternoon sir I'm calling from a website my name is xx. I'd like to know a few details about your restaurant, can you please help me out?
Restaurant: Yes, tell me
Caller: What cuisines do you serve?
Restaurant: We serve Italian, Continental
Caller: Ok. What are the timings of the restaurant?
Restaurant: From 11 to 10:30
Caller: Seating capacity?
Restaurant: 80 people
Caller: 80 people? Ok. Do you serve liquor?
Restaurant: Yes
Caller: Do you have happy hours?
Restaurant: Yes 4 - 9
Caller: Ok. I think you do provide home delivery? Right?
Restaurant: Only in xx not xx.
Caller: Only in xx?
Restaurant: Yes, we just deliver in area xx. Nowhere else.
Caller: Ok. And, do you serve liquor?
Restaurant: Yes
Caller: Do you accept credit cards?
Restaurant: Yes, we accept Amex, Visa, Mastercard
Caller: Ok, do you have parking?
Restaurant: Yes we have, its not ours, it is provided by dlf.
Caller: is the parking paid?
Restaurant: yes
Caller: What is the average cost of food for 2 people?
Restaurant: Rs 1000
Caller: May I know your name?
Restaurant: xx
Caller: Thank you Mr xx for sharing the information.
Restaurant: Alright, bye
Caller: Bye