

Optimal Segmentation in Platform Strategy: Evidence from Geotargeted Advertising

Bo Cowgill
Columbia University

Cosmina Dorobantu
Oxford University and
The Alan Turing Institute*

March 5, 2019

Abstract

How should market designers tradeoff liquidity and specificity? We study a natural experiment in the release of a new ad targeting feature by an ad exchange. The platform introduced new targeting into select geographic markets using a regression discontinuity. The experiment affects the specificity advertising assets in the markets (ie, the availability of targeting a city or a zip code). We find evidence that additional specificity reduces the total number of ad impressions delivered by the platform, as advertisers concentrate bidding into fewer, targeted markets. Despite this, we find an overall positive effects on revenue growth in the treated areas. This appears to be driven mainly by increases in clickthrough rates and not through increases in average prices (which actually decreased), and by entry of new advertisers.

* Authors are listed in alphabetical order. We thank the NET Institute for financial support. We also thank Mark Armstrong, Tommaso Bondi, Alexandre de Corniere, Beata Javorcik, Ian Jewitt, Paul Klemperer, Jonathan Levin, Randall Lewis, Preston McAfee, John Morgan, Dennis Novy, Ferdinand Rauch, and seminar participants at Oxford for helpful comments. Dorobantu acknowledges financial support from the Economic and Social Research Council.

1 Introduction

Market designers often face choices about segmenting demand. For example, a designer could partition a markets into highly-specific pools. This design gives buyers a mechanism to pay higher amounts for goods or services that are well-targeted for their needs. However, this design also reduces competition by fragmenting buyers into separate pools, where the number of buyers in each pool is lower. Competition between buyers is useful to give bargaining power to the seller and push prices upwards.

By contrast, an alternative design pools separate asset types into a smaller number of more general, less specific markets. This design denies buyers the ability to target (and pay for) better targeted assets. However, pooling demand together might enhance competition between buyers. Is the tradeoff worthwhile? When is higher asset specificity desirable in a market-place?

In this paper we develop a simple model of this tradeoff, and then study this phenomena empirically using a natural experiment in the online ads.¹ In our setting, a large online advertising platform released a new feature permitting advertisers to target ever-smaller geographical areas within the United States. Prior to this release, only city-level targeting was available. After the release, the platform’s advertisers were able to target individual ZIP codes.²

For privacy and accuracy reasons, the platform in our study did not enable targeting for all US ZIP codes. In particular, the online platform decided not to release finer geographical targeting in areas where its engineers were unable to detect N or more users. This policy was meant to ensure that the online behavior of individual users cannot be tracked by advertisers.

We utilize the N -user cutoff rule for treatment to construct a panel regression discontinuity estimate of the effect of new targeting on market outcomes of interest. As one of our robustness checks, we also estimate a differences-in-differences effect.

Our data set covers a wide array of ad auctions and it enables us to examine how increased ad targeting affects market outcomes. Furthermore, it enables us to study whether the platform’s advertisers benefit from the release of the new targeting feature. We have four main findings.

First, we find that the introduction of finer geographical targeting had a negative impact on auction participation. In the treatment regions, the release of the finer targeting feature had a negative effect on the number of ad impressions delivered by the platform.

Second, we find that although it had a negative impact on auction participation, the release of the new targeting feature had a positive effect on the online platform’s revenue in the areas where treatment occurred. Our estimates suggest that a unit increase in treatment strength resulted in a revenue increase of 2.36 standard deviations relative to its underlying trend in the areas where the platform introduced ZIP code targeting.

Next, we next turn to the mechanism of this revenue increase. Revenue for the online platform in our study consists of two parts, the number of clicks accrued by the ads and the cost-per-click paid by the advertisers.³ We find that the revenue growth was due to a higher volume

¹Highly specific targeting is valuable to buyers (advertisers). However, if advertisers can see all the available information about users, many of them might choose to not advertise to certain audiences at all. This may fragment competition between advertisers and reduce pricing power for the ad platform. This issue is particularly relevant for auction-based ad platforms (Edleman et al., 2007 and Varian, 2007), where competition between advertisers plays a direct role in ad click prices and platform revenue.

²A system of postal codes used in the United States since 1963.

³Advertisers in this market only pay the platform when a user clicks on one of their ads.

of clicks; the average cost-per-click on the online platform did not increase relative to its underlying trend in the treatment areas following the introduction of ZIP code targeting.

Combined, what do these three findings suggest? In the treatment areas, there was a negative impact on the number of ads displayed on the online platform following the introduction of the new targeting feature, and a positive impact on the volume of clicks accrued by the ads. The release of the finer geographical targeting feature, in other words, had the important effect of increasing the relevancy of the ads running on the platform. From a revenue perspective, the online platform in our study benefitted from the introduction of ZIP code targeting because these more relevant ads generated more clicks.

Our fourth and final finding concerns advertiser outcomes. We find that enhanced targeting benefited not only the online platform, but also the advertisers that used the platform to promote their products. In particular, we find that in the treatment areas, the introduction of ZIP code targeting had a positive effect on the number of conversions.⁴ Furthermore, we find that on average, the cost-per-conversion did not increase relative to its underlying trend in the treatment areas following the release of the new targeting feature. Combined, these two results show that ZIP code targeting was beneficial to advertisers: in the regions where treatment occurred, ZIP code targeting had a positive effect on the volume, but not the price, of the conversions.

Taken together, our findings uncover a new angle on the tradeoff between competition and specificity. We show that for the online platform in our study, enhanced ad targeting (ZIP code targeting) was profitable despite the concerns about market thinness, partitioned demand, and bargaining power.

We also show that the introduction of the new targeting feature had an economically and statistically significant impact on auction participation. In the existing theoretical models of targeting and auction design, the number of participants in the marketplace is often exogenous. As such, our findings about endogenous entry show evidence of a different theoretical mechanism.

A large literature in the economics of auction and market design has emphasized the importance of market thickness (Roth, 2008) and extra bidders (Bulow and Klemperer, 1996) for auction revenue. Despite these results, online advertising platforms continue to release more and more targeting features – which would seem to decrease the number of bidders per auction and reduce thickness. A related literature (Tadelis et al. (2015)) has stressed the value of information in auctions to improve match quality and increase the auctioneer’s revenues. Through the lens of this literature, better targeting is a form of offering more information to buyers about the ad inventory for sale.

Our findings are particularly salient in the Internet industry, where new technology firms have designed marketplaces to facilitate a variety of new transactions. As technology companies’ capacity to collect data grows, designers at these firms can create increasingly specific markets where buyers can condition activity on lots of covariates. Researchers like Hummel and McAfee (2015) and Fu et al. (2012) have examined these questions theoretically, using advertising markets as motivation. Our paper explores the empirics of these design decisions.

From a policy making perspective, our results have implications for how regulators view the economics of privacy and data gathering. Collecting users’ data in order to release finer targeting features is profitable for online platforms; it has a positive impact on the platforms’

⁴A conversion is an action that a user who clicks on an ad undertakes on the advertiser’s website, such as purchasing, placing an item in a shopping cart, downloading or other forms of online interactions that are valuable to businesses.

revenues and the relevancy of the ads running on the platforms. Enhanced targeting based on users' characteristics is also profitable for advertisers, whose advertisements generate a larger volume of sales.

The remainder of this paper proceeds as follows. Section 2 reviews the related literature. Section 3 outlines the setting and institutional details. Section 4 describes the natural experiment and Section 5 presents our data. Section 6 proposes the empirical specifications and Section 7 explores the results. We conclude discussing the implications of our findings and future avenues of research in Section 8.

2 Relevant Literature

This project aims to contribute to several parts of the literature about auction and market design, Internet advertising, Internet privacy, as well as asset specificity.

2.1 Auction and Market Design: Information Disclosure and Market Thickness

Our work relates to the effect of information disclosure on auction outcomes – a topic studied within the context of auction design. [Milgrom and Weber \(1982\)](#) are the first to address information revelation in auctions in a seminal work that identifies the “linkage” principle. Assuming that the bidders’ valuations are affiliated (one bidder’s high valuation makes other bidders’ high valuations more likely), [Milgrom and Weber \(1982\)](#) show that in a second price auction, the auctioneer’s revenue increases when she commits herself to a policy of full disclosure. The intuition behind this result is simple: with access to more information about the object auctioned, bidders’ valuations are more closely aligned. This intensifies the competition between the bidders and encourages them to submit higher bids, which translate into increased revenues for the auctioneer.

[Milgrom and Weber \(1982\)](#)’s assumption that bidders’ valuations are affiliated holds for certain auctions. For example, this assumption is plausible in the context of auctions for mineral rights, where each bidder’s valuation of the auctioned object is a function of the amount of natural resources that can be extracted. [Ganuza \(2004\)](#), however, argues that there are other auctions, such as Internet auctions similar to the one we study, where bidders’ valuations need not be correlated. In these settings, bidders’ valuations can be heterogeneous, which means that in response to the auctioneer revealing additional information about the auctioned object, some bidders’ valuations will increase, while other bidders’ valuations will decrease.

In a setting where bidders have heterogeneous valuations, [Ganuza \(2004\)](#) identifies two competing effects. First, he shows that information revelation increases the efficiency of the auction by improving the match between bidders’ preferences and the attributes of the auctioned object. This has a positive effect on the winning bidder’s willingness to pay for the auctioned object and on the auctioneer’s revenue. Second, he shows that information revelation also increases the winning bidder’s informational rents, which has a negative effect on the auctioneer’s revenue. [Ganuza \(2004\)](#) proves that as the number of participating bidders increases, the first effect dominates the second. In other words, the auctioneer has an incentive to withhold information about the auctioned object when the number of bidders is low and to release information when the number of bidders is high.

[Board \(2009\)](#) generalizes [Ganuza \(2004\)](#)’s results and shows that in second price auctions – regardless of whether bidders’ valuations are affiliated or heterogeneous – revealing information decreases the auctioneer’s revenue when there are only two bidders and increases the auctioneer’s revenue when there are four or more bidders. [Board \(2009\)](#) argues that this is due to the “allocation” effect, which causes the order of the bidders’ valuations to change as a result of the information disclosed.

With two bidders, it is straightforward to see how [Board \(2009\)](#)’s “allocation” effect works. Assume that an online advertising platform does not reveal any information about the users visiting its website. Two advertisers, A and B, are competing in a second price auction for advertising space on the platform. Advertiser A values local users at \$4 and non-local users at \$2, while advertiser B values local users at \$2 and non-local users at \$4. Assuming that there is an equal probability of a user being local or non-local, each advertiser will bid \$3 and the online platform’s expected revenue from the auction will be \$3. Now assume that the online

platform will reveal to the advertisers the users' geographic location before the advertisers place their bids. If the user is local, advertiser A will bid \$4 and advertiser B will bid \$2. In a second price auction, the online platform's expected revenue will be \$2, which is \$1 less than the revenue the platform expects when it does not reveal information about its users.

In most auction models, information disclosure is considered to be exogenous. The works of [Milgrom and Weber \(1982\)](#), [Ganuza \(2004\)](#), and [Board \(2009\)](#) are revolutionary precisely because they relax this assumption and treat information revelation as an auctioneer's endogenous choice. However, despite their significant theoretical contributions, the works of [Milgrom and Weber \(1982\)](#), [Ganuza \(2004\)](#), and [Board \(2009\)](#) do have some limitations. In particular, they assume that the number of bidders participating in the auction is fixed. This assumption has the important implication that the information revealed by the auctioneer does not affect, in any way, the number of bidders participating in the auction. In other words, there is no auction entry or exit as a result of the information disclosed. Furthermore, they assume that the auctioneer has a single object for sale and conducts only one auction. This second assumption rules out the possibility that there are multiple objects for sale and that bidders can use the information revealed to decide what item (if any) to bid on.

In a recent empirical paper, [Tadelis et al. \(2015\)](#) call into question the plausibility of these two assumptions in the context of US wholesale car auctions. Using a large scale field experiment, they show that disclosing information about the quality of the vehicles auctioned raises the auctioneer's revenue, a result that is consistent with [Milgrom and Weber \(1982\)](#)'s predictions. However, [Tadelis et al. \(2015\)](#) also find that the increase in the auctioneer's revenue holds across all vehicle quality levels. In other words, the information revealed positively influences the auctioneer's revenue for both high and low quality cars. In the context of [Milgrom and Weber \(1982\)](#)'s theoretical model, this finding is surprising: bad news about the quality of an auctioned object should lower the bidders' willingness to pay.

[Tadelis et al. \(2015\)](#) cannot directly measure auction participation, but they hypothesize that their finding that information disclosure raises the seller's revenue regardless of whether the information is good news or bad news is due to a better match between the bidders and the auctioned object. In a setting with multiple objects for sale, bidders use the information disclosed by the auctioneer to choose which auction to participate in; they will enter auctions where they value the auctioned object highly relative to other bidders – and exit auctions where they do not. [Tadelis et al. \(2015\)](#) argue that this endogenous sorting of bidders into auctions as a result of information disclosure leads to a better match between the bidders and the auctioned items, which intensifies competition in each auction and results in higher bids and higher revenues for the seller.

Our work is similar to [Tadelis et al. \(2015\)](#) in that it underlines, in the context of multiple object auctions, the importance of allowing for an endogenous sorting of bidders into auctions as a result of information disclosure. We are able to measure auction participation in our data and to show that information disclosure has a statistically significant effect on auction participation. In particular, we find that information disclosure has a negative effect on auction participation. By directly measuring and identifying the effect of information revelation on auction participation, our work brings an important contribution to the literature studying the effect of information disclosure on auction outcomes.

Our work also relates to the literature studying the benefits of market thickness. In a seminal paper, [Bulow and Klemperer \(1996\)](#) emphasize the importance of market thickness by examining whether it is more profitable to sell a company via a public auction with no reserve price or an optimally structured negotiation with one less bidder. They find that the public auction is preferable, as long as it attracts at least one extra bidder compared to the negotiation process. A decade later, [Roth \(2008\)](#) describes the history of regional kidney exchanges in the

United States in order to show that market thickness, defined as a large number of potential participants, is a key condition for marketplaces to function well.

In our paper, although we find that information disclosure has a negative effect on auction participation, we also find that the auctioneer’s revenue increases as a result of information revelation. By showing that thinner markets need not always result in lower revenues for the auctioneer, our work also contributes to the literature studying the benefits of market thickness.

2.2 Internet Advertising

Issues surrounding information disclosure and thickness arise concretely in the case of Internet advertising markets. In these markets, online platforms like the one we study auction off advertising space. The platforms hold a considerable amount of information about their users, such as their geographical locations and incomplete browsing histories, and are often faced with the question of how much information they should disclose to potential ad bidders. A small branch of the literature has sought to address this setting directly.⁵

In an informal essay, [Levin and Milgrom \(2010\)](#) discuss the tradeoffs that online platforms face when offering finer ad targeting. The authors point out that by disclosing information which allows bidders to target their ads to relevant users, online platforms ensure a good match between ad bidders and ad viewers. [Levin and Milgrom \(2010\)](#) hypothesize that the quality of the match encourages entry into the auction for advertising space, increases the competition between bidders, and results in higher revenues for the online platforms. However, the authors also note that revealing too much information about the users may create thin markets. If ad bidders can target very narrow sets of users, there will not be enough bidders participating in each auction. This can create monetization problems for the online platform.

[Hummel and McAfee \(2015\)](#) and [Fu et al. \(2012\)](#) propose theoretical models addressing the question of whether releasing enhanced targeting options in a second price online ad auction increases the platform’s revenue. Using a single object auction model with a fixed number of symmetric bidders, [Hummel and McAfee \(2015\)](#) obtain results similar to [Board \(2009\)](#). They find that finer targeting options decrease the online platform’s revenue when there is a relatively small number of players, increase the online platform’s revenue when there is a relatively large number of players, and have an ambiguous effect for an intermediate number of players. [Fu et al. \(2012\)](#) also consider a single item auction with an exogenous number of bidders and find that enhanced targeting generally results in higher expected revenues for the online platform.

We contribute to the literature that studies whether releasing enhanced targeting options in a second price online ad auction increases platform revenues by being the first to address this question empirically. We also contribute to the literature by calling into question the plausibility of some of the underlying assumptions used by [Hummel and McAfee \(2015\)](#) and [Fu et al. \(2012\)](#). In particular, we show that enhanced targeting options do affect auction participation –

⁵Although it is not directly relevant to our work, we would like to point the reader to a related branch of the literature which addresses other interesting questions regarding the economics of targeting in online advertising. [Bhawalkar et al. \(2014\)](#) study how much advertisers bidding in an online auction for ad views might be willing to pay a data source for targeting data. [Abraham et al. \(2013\)](#) show that in the presence of information asymmetries (i.e. when some ad bidders have access to targeting data, while others do not), the revenue of an online platform employing a second price auction to allocate its advertising space is negatively affected. [De Corniere and Nijs \(2014\)](#) examine how the information revealed by an online platform about its users affects the equilibrium prices of the goods sold by the firms that place ads on the platform. Finally, [Athey and Gans \(2010\)](#) and [Bergemann and Bonatti \(2011\)](#) investigate how the offline market for advertising is affected by the appearance of online advertising, which allows advertisers to target their ads to users.

a finding that contradicts the theoretical models' assumption that better targeting options have no effect on the number of ads participating in the auction.

2.3 Internet Privacy

A key question addressed in this paper – the benefits of enhanced targeting for online platforms and advertisers – affects the economics of privacy. A widespread concern among privacy activists is that online platforms, as well as advertisers, face irresistible business incentives to invade consumers' privacy.

Little research has been done to date to help regulators understand whether they have to intervene to protect users' privacy online – and what impact their interventions might have on the online advertising industry. One exception is [Johnson \(2013\)](#), who builds a structural model to measure how three privacy-related regulatory options affect revenues for online platforms and advertiser surplus. He finds that a complete ban on ad targeting would decrease platforms' revenues and advertiser surplus by more than half.⁶

Our finding that enhanced targeting options are beneficial to online platforms and their advertisers contributes to the literature on two fronts. First, it confirms that online platforms and the advertisers promoting their products on these platforms do not have a financial incentive to protect Internet users' data. In doing so, it strengthens the case for policy interventions that protect online users' privacy. Second, it indicates that potential privacy-related measures may negatively impact the revenues of online advertising platforms and their advertising clients. In doing so, it encourages regulators to design policies that strike a balance between the negative effects on the online advertising industry and the benefits that Internet users derive from having their privacy protected online.

2.4 Asset Specificity

Finally, this paper also contributes to the literature by attempting to integrate the notion of "asset specificity" ([Williamson \(1975\)](#)) into market design. The notion of asset specificity has been mostly used in the organization literature, where the specificity of assets plays a role in vertical integration and "boundaries of the firm" decisions. In this literature, the specificity of assets is exogenous. However in many emerging market design settings, the specificity of assets is a strategic, endogenous choice of a market designer. We provide a concrete example of this in the Internet advertising market, where online platforms choose how much specificity (targeting) to enable in the marketplace.

3 Setting and Institutional Details

Our empirical setting is an online advertising platform. In this section, we introduce the main choices that advertisers face when placing an ad and we describe the ways in which the online platform uses the information provided by the advertisers to determine how it displays the ads.

⁶A related literature shows that in the presence of targeting constraints, ads are less effective. See, for example, [Goldfarb and Tucker \(2011\)](#) and [Bailey \(2012\)](#).

3.1 Choices that Advertisers Make

Advertisers who want their ads to appear on publishers' webpages must set up an advertising account. When setting up an account, they face several choices regarding the content of the ad, the ad's targeting criteria, and the costs accrued by the ad. Of these, we are primarily interested in the targeting criteria.

Advertisers also choose a series of keywords or content topics that trigger their ads. By choosing the right topics, advertisers ensure that their ads are only seen by users who are browsing content related to the advertiser's products or services. Our mortgage broker, for example, might choose topics such as 'business loans,' 'mortgage loans,' 'first home loans' or 'mortgage calculator' to trigger his ad. By choosing these topics, he ensures that his ad will appear only to users browsing relevant topics.

In addition to topics, advertisers can specify other user characteristics as targeting criteria for their ads. For example, on some online advertising platforms, advertisers can opt for their ads to be visible only to female users, or only to college graduates.

Note that many advertising platforms – including the one in our study – do not have direct information about the geographical location of their users. However, they are generally able to ascertain information about their users' location via the users' IP (Internet Protocol) address and other cues. Often an IP address can be mapped back to an Internet Service Provider (ISP) whose geographic coverage is known. The targeting software may also know information about how a user's request was routed through intermediary servers before arriving at the platform (including the name or other characteristics of the wireless or ethernet network). This information may also yield some clues about the user's geographic location. Lastly, in some cases, platforms are able to utilize the GPS system on users' devices to establish their location.

Geographic targeting is attractive to advertisers, as it allows them to tailor their marketing strategy based on the users' location. Being based in New York City, our mortgage broker might want his ads to appear only to users based in New York City, who can easily reach his office for face-to-face consultations. If this is the case, he can choose to have his ads visible only within the New York City limits and thus avoid incurring advertising costs for users who live outside of his service area. Similarly, a large device manufacturer might value advertising more in certain areas than others. The latest tech gadgets are usually in high demand in areas like Silicon Valley, so the device manufacturer might want to spend more for ads seen by users within Silicon Valley and less for users in less technologically savvy locations. For both the mortgage broker and the device manufacturer, geographical targeting allows the advertisers to alter their marketing strategy depending on the geography of the users.⁷

3.1.1 Ad Costs

Advertisers in this setting pay per ad click. When choosing a particular topic to trigger an ad, each advertiser enters a maximum cost-per-click (CPC). This represents the highest dollar amount that the advertiser is willing to pay for a click accrued by the ad. While there is no upper limit on the amount they can bid, advertisers are required to enter a maximum CPC

⁷Although our paper studies a natural experiment in geographic targeting, similar economics apply to other forms of ad targeting and other changes to asset specificity in markets. For example, suppose that a platform allowed advertisers to target their ads based on users' hobbies (in addition to content topic targeting). Revealing this knowledge may lead to many advertisers withdrawing bids from auctions where users' hobbies are irrelevant. However, the advertisers' valuations for users with relevant hobbies may go up. These dynamics (and others) mirror the economic mechanisms at work for the online platform in this paper as it introduces more geographical targeting options.

that is at least one cent (\$0.01).

The online platform in our study discloses to advertisers information regarding the average cost-per-click for each topic they are considering, as well as information regarding the average number of monthly ad queries for that particular topic. The average cost-per-click and the average query volume can vary considerably from topic to topic. For example, we set up an advertising campaign for the fictitious mortgage broker based in New York City using four topics: ‘business loans,’ ‘mortgage loans,’ ‘first home loans,’ and ‘mortgage calculators.’ For an ad targeted to users within New York City, the online platform informed us that for the topic ‘business loans,’ the average CPC is \$53.69 and this topic receives, on average, 1,000 monthly queries. For the topic ‘mortgage loans,’ the average CPC is \$27.56 and the average monthly query volume is 201. For ‘first home loans,’ the average CPC is \$10.53 and the average monthly query volume is 10. Finally, for ‘mortgage calculators,’ the average CPC is \$0.78 and the average monthly query volume is 90,500.

3.2 The Position of the Ads on the Page

The online platform uses the information provided by the advertisers when setting up their ad campaigns to determine the order in which it displays the ads on the page and the cost of each click accrued by the ads. We describe this process below.

When users view content page associated with the ads platform, they are shown both non-commercial content and paid advertising (labeled). The online platform displays a maximum of eight ads per content page, but frequently shows fewer. The allocations of each advertiser to one of the ad positions on the page are determined through a generalized second price (GSP) position auction, which is widely used within the online advertising industry.

One important feature of this auction setting is that bids are “quality-adjusted.” The bids in the auction are the maximum cost-per-click amounts that the advertisers specify for each topic. The quality adjustment comes in the form of a quality score that the online platform calculates for each topic. The quality score is on a scale of 1 to 10, with 1 being the lowest and 10 the highest. The online platform calculates this score based on several factors, such as the relevance of the ad and the quality of the advertiser’s website.⁸ Advertisers do not know their topics’ exact quality scores, but they can see, for each topic, whether their scores are ‘low,’ ‘medium’ or ‘high.’

To measure an ad’s relevance, the online platform has algorithms in place that analyze factors such as whether the content of the ad relates to the topic chosen. For example, if the mortgage broker in our example uses the ad text in Section ??, but picks an irrelevant topic, such as ‘restaurants,’ the online platform will consider the broker’s ad to be irrelevant for this particular topic. To measure the quality of an advertiser’s website, the online platform uses algorithms that examine factors such as whether the page is easy to navigate, whether the content of the page is relevant to the ad, whether the website takes a long time to load, or whether it abuses users’ data.

The following example illustrates how a generalized second price position auction works. Suppose there are four eligible advertisers in the auction for a given topic. Advertiser 1 has a maximum CPC of \$4 and a quality score of 2, advertiser 2 has a maximum CPC of \$3 and a quality score of 5, advertiser 3 has a maximum CPC of \$2 and a quality score of 10, and advertiser 4 has a maximum CPC of \$1 and a quality score of 9. What will be their positions on the page and how much will each advertiser pay per click?

⁸The platform does not reveal to the advertisers any information regarding the way it weighs each factor, nor does it publicize all the factors that it takes into account when calculating the quality score.

The total auction score of each advertiser is his bid times his quality score.⁹ For advertiser 1, the total auction score is 8 ($\$4 \times 2$), for advertiser 2, the total auction score is 15 ($\$3 \times 5$), for advertiser 3, it is 20 ($\$2 \times 10$), while for advertiser 4, it is 9 ($\$1 \times 9$).

Advertiser 1 has the lowest auction score, so he will be the last one ranked on the page and he will pay \$4 per click. With an auction score of 9, advertiser 4 will rank right above advertiser 1, and he will pay just enough to obtain a total auction score above the auction score of advertiser 1. In other words, he will pay \$0.89 per click ($\$0.89 \times 9 = 8.01$, which is greater than 8). With an auction score of 15, advertiser 2 will rank above advertiser 4, and she will pay \$1.81 per click ($\$1.81 \times 5 = 9.05$, which is greater than 9). Finally, with an auction score of 20, advertiser 3 will take the top position on the page and she will pay \$1.51 per click ($\$1.51 \times 10 = 15.1$, which is greater than 15).

Theoretical analyses of the auction format described above are available from [Edleman et al. \(2007\)](#) and [Varian \(2007\)](#). For our purposes, a key piece of information to retain is that within this auction format, advertisers can reduce their cost-per-click by improving the relevance of their ads.

4 The Natural Experiment

In March 2012, a technological breakthrough enabled the online platform in our study to introduce ZIP code targeting. This is a form of geographical ad targeting of small geographical units known as “ZIP codes.” There are close to 30K standard ZIP codes in the United States.¹⁰ The average standard US ZIP code covers a land area of 87.6 miles² (227 km²) and is home to 10,600 people.

Prior to the introduction of ZIP code targeting, the most granular form of targeting available on the online platform we study was city targeting. The platform also offered other forms of wider geographical targeting: county targeting (there are 3,007 counties in the US), congressional district targeting (there are 435 congressional districts in the US, which cover the electoral constituencies that elect the members of the House of Representatives), designated market area (DMA) targeting (there are 210 DMA regions within the US, where the population has similar TV station, radio, newspaper, and often Internet content coverage), and state targeting (there are 50 states in the US).

In March 2012, the online platform in our study released ZIP code targeting to roughly one third of US’ standard ZIP codes. For the remaining ZIP codes, geographical targeting remained unavailable. In September 2013, an additional 1% of US’ standard ZIP codes became targetable thanks to technological improvements that the online platform made to the algorithm identifying users’ geographical locations. The rest of the ZIP codes remain un-targetable.

Through various panel data techniques described in Section 6, we use geographic areas with little ZIP code targeting enabled as “control groups” to estimate the effect of introducing finer geographical targeting. Areas with more targetable ZIP codes are the “treatment group.”

⁹In this example, we weigh bids and quality scores equally, though this might not be the case. Because of its proprietary nature, we do not have access to information as to how the online platform weighs bids and quality scores.

¹⁰The US also has roughly 12.5K non-standard ZIP codes. These ZIP codes are assigned to overseas military bases, facilities that house post office boxes, and government agencies, companies or buildings that receive high volumes of mail, such as the CIA, Wal-Mart or the Empire State Building. Our analysis focuses only on the standard ZIP codes, as a household’s online behavior within a standard ZIP code is likely to differ significantly from the online behavior of army personnel, post office employees or CIA operatives.

An important econometric detail in our study is that the platform did not estimate which markets – neither geographic markets, industry markets, nor the combination – would yield the greatest profit from ZIP code targeting in order to introduce the feature only in those areas. Had the platform been strategic, it would have introduced ZIP code targeting selectively into markets where it would have the highest impact. For example, local services is a category of topics that might be particularly strategic to introduce ZIP code targeting. Alternatively, the company might have introduced targeting into regions with high income, so that the ZIP code targeting could be used to direct advertising efforts to buyers with high purchasing power. However, rather than introducing ZIP code targeting into strategic industries and geographies, the platform offered the feature according to a threshold strategy – creating a regression discontinuity. In every geography where it offered ZIP targeting, all topics were available to be targeted in all categories.

The platform’s selection criteria into ZIP code targeting were a mixture of privacy and technological related reasons. On the privacy side, the platform decided not to release ZIP code targeting in areas where there was any risk of advertisers being able to track the online behavior of individual users. On the technological side, as discussed in Section 3, user locations must be inferred and cannot be perfectly known. The accuracy of geolocation tools is often below 100%. Companies in this industry typically establish internal quality requirements, and do not release geographical targeting features to advertisers in regions where the accuracy requirements are not met. The platform in our study chose to release ZIP code targeting only to ZIP codes that met the privacy and accuracy requirements set by the firm’s engineers.

Notably for us, on the privacy side, the platform’s engineers set the following cutoff rule: they did not launch ZIP code targeting in ZIP codes where they were unable to detect N or more Internet cookies. Below this threshold, no ZIP code got targeting. Above this threshold, the ZIP codes that met the engineers’ accuracy requirements got targeting.

In some of our empirical specifications, we use this privacy related cutoff rule to instrument for treatment status. We note that the existence of the N cookie threshold was not public knowledge, ruling out any possibilities for there being endogenous sorting of ZIP codes into treatment.

5 Data

The data from the online platform for this study were originally recorded on a per-query basis. For each ad query initiated by a browsing user, the platform recorded some information about the user and some information about the advertisers participating in the auction for ad space.

For the user, the online platform recorded topics of the user’s content as well as some characteristics of the request for ads.¹¹ These characteristics include the estimated geographic location of the user, in as much granularity as the platform was able to infer, and a set of binary variables describing whether the user made the query from a computer, a tablet, or a mobile device.

For the advertisers who participated in the auction for the particular topic combination browsed by the user, the online platform recorded the following variables: the number of ad views (often referred to as impressions in the advertising industry), the number of ad clicks, and the price the advertisers paid for each ad click (as established by the “GSP” position auction). Furthermore, the online platform recorded in the underlying dataset for advertisers some data on purchases made by users after clicking on an ad.

¹¹For privacy reasons, we were not shown the exact content submitted by the users.

5.1 Panel Index Variables

For the econometrics in this study, the per-query data have been aggregated into sets of two-period geographic panels. The panels are indexed by location \times industry \times time to allow for a richer set of controls.

5.1.1 Location

We chose not to run our analysis at the ZIP code level. There are two reasons for this. The first is the quality of the data for the untreated ZIP codes. In order to obtain data at the ZIP code level, we must be able to accurately map IP addresses to each standard ZIP code in the US. However, one of the key reasons why some ZIP codes in our study did not get ZIP code targeting is the fact that the platform’s engineers were unable to properly map IP addresses to these locations. Analysis at the ZIP code level would mean that for the untreated ZIP codes (unlike the treated ones), we might either misattribute observations or miss them out entirely due to the poor mapping.

The second reason deals with the stable unit treatment value assumption (SUTVA) violations between treated and untreated ZIP codes. This assumption requires treatment units and control units to not be able to affect each other. Had we used ZIP codes as a unit of analysis, there may have been interactions between treatment and control locations. For example, a “control” ZIP code that was adjacent to a “treated” ZIP code may be affected by the nearby treated unit (advertisers may substitute spending away from the control unit into the treated unit). The new targeting option may cause advertising revenue to shift out of un-targetable regions and into adjacent, targetable ones. If we were performing our analysis at the ZIP code level, this shift would appear to be “increased” revenue, even if overall revenue stayed constant and was simply distributed differently after the targeting. In order to avoid this, our basic unit of geography must be larger than a ZIP code.

The geographical unit we choose for data aggregation purposes is a “designated market area.” As noted in Section 4, DMAs are geographies within the United States where the population has similar TV station, radio, newspaper, and often Internet content coverage. Substitution across DMA boundaries is unlikely, so our unit of geographical aggregation satisfies SUTVA. Furthermore, the quality of the IP mapping at the DMA level is high for the entirety of the United States, eliminating concerns over quality differences between treated and untreated regions.

In no case is the geographic granularity of units in our panel specifications more fine (or less fine) than the DMA level. We refer to a DMA as having a higher level of “treatment” if its component ZIP codes can be targeted, and as relatively “untreated” if its component ZIP codes cannot be targeted.¹²

5.1.2 Time

On the time dimension, data are aggregated at the monthly level. For our analyses, we use data from February 2011, February 2012, and February 2013. The month in the middle, February 2012, was the month right before ZIP code targeting was launched. The first month, February 2011, was exactly a year before, while the last month, February 2013, is precisely a year later.

We chose a period of 12 months between the dates we study for two reasons. First, the

¹²Section 5.3 describes in detail how we construct the variable measuring treatment strength at the DMA level.

effects of better targeting need time to evolve. Advertisers' choice to utilize ZIP code targeting is endogenous. It requires them to know about the existence of the new targeting feature and to shift their pre-existing strategy to utilize it. A year allows for this shift to take place. Second, online sales are seasonal. By studying the same calendar month, our estimations are not affected by seasonality movements.

5.1.3 Industry

We also aggregate the data by industry, in addition to location and time. "Industry" is a classification associated with a query and is assigned through semantic text analysis. These industries follow the same structure as NAICS sectors. The online platform has 25 top level industry classifications, which are similar to 2-digit NAICS sectors. In our study, we use the top level classification to split our location \times time panel into 25 industries. A full list of these industries is available in Appendix A.

5.2 Outcome Variables

We are interested in understanding how the introduction of the new targeting feature affected auction participation and the online platform's revenue. Furthermore, we are also interested in understanding whether the platform's advertisers benefited from the release of the new targeting feature. We will discuss, in turn, the outcome variables that we have at our disposal to uncover these effects.

5.2.1 Auction Participation

The first question we would like to address is whether the introduction of finer geographical targeting had an effect on auction participation. In light of the existing literature, answering this question is important for several reasons.

First, as we discussed in Section 2, there is a significant theoretical literature that addresses the more general question of information provision in auctions, as well as a smaller branch of the literature that tackles the more specific question of the effect of releasing enhanced targeting options in the context of second price online ad auctions. A key assumption, common to all existing theoretical models, is that auction participation is exogenous and cannot be affected by the information disclosed. By asking whether the introduction of ZIP code targeting has an effect on auction participation, we are the first to test the empirical validity of this assumption.

Second, a number of previous researchers have underlined the importance of market thickness and additional bidders for auction revenue. In light of this literature, the online platform's release of a finer targeting feature is risky: the introduction of ZIP code targeting may have the undesirable effect of reducing market thickness and creating monetization problems for the platform. By asking whether the release of the enhanced targeting feature affects auction participation, we seek to understand whether the online platform in our study does, indeed, risk creating thinner markets.

The outcome variable that allows us to gain insight into auction participation is the number of ad impressions.¹³ This represents the total number of ads that appeared on the online plat-

¹³Our data set does not contain information about the number of advertiser accounts. However, even if we did have access to data regarding the number of accounts, in our setting, the number of ad impressions is a much more precise measure of auction participation than the number of advertiser accounts. As we have seen in Section 3,

form in response to content hits conducted by users. If the introduction of ZIP code targeting resulted in more ads competing for the platform’s advertising positions, we expect to find a positive effect of finer targeting on ad impressions. However, if the release of the new ZIP code targeting feature resulted in less ads competing for the platform’s advertising positions, we expect to find a negative effect of enhanced targeting on ad impressions.

Note that a priori, it is ambiguous whether the number of ad impressions should be positively or negatively affected by the introduction of ZIP code targeting. Suppose a user living in San Francisco’s Cole Valley neighborhood browses online for a pizza delivery service. Prior to the introduction of ZIP code targeting, three ads got triggered by the user’s requests for “pizza delivery,” all of them targeting the entire city of San Francisco. One ad was for a pizza place that delivered to Cole Valley, while the other two were for pizza places that delivered to other neighborhoods in San Francisco. After the introduction of ZIP code targeting, however, only one ad gets triggered by the user’s searching for “pizza delivery:” the one belonging to the pizza restaurant that services Cole Valley. The other two advertisers turn off their ads for neighborhoods where they cannot deliver. In this example, the number of impressions decreases from 3 to 1. The market, in other words, is becoming thinner.

Now assume a different scenario: the same user living in San Francisco’s Cole Valley neighborhood browses online for a pizza delivery service. Prior to the introduction of ZIP code targeting, the user’s search triggered three ads: one for Domino’s, one for Papa John’s, and one for Pizza Hut. These are big restaurant chains, with large marketing budgets, that deliver to all of San Francisco’s neighborhoods. A small pizza restaurant, that only services Cole Valley, did not have an ad presence on the online platform prior to the introduction of ZIP code targeting as it did not have the marketing budget to compete with the established players for clicks originating from anywhere in San Francisco. After the introduction of the ZIP code targeting feature, however, the small pizza restaurant in Cole Valley joins the ad platform and competes with the established players for clicks within its service area. As a result of the small restaurant entering the ad auction, the Cole Valley user searching for “pizza delivery” sees four ads. In this example, the number of impressions increases from 3 to 4. The market, in other words, is becoming thicker.

Our two examples show that there might be some exit and some entry into the auction in response to the new targeting feature. In our data, a negative effect of ZIP code targeting on impressions would indicate that the exit effect dominates the entry effect.

5.2.2 Platform Revenue

The second question we would like to address is whether the online platform’s revenue increased or decreased as a result of the ZIP code targeting feature. By answering this question, we hope to contribute to the existing literature on several fronts.

First, as noted in Section 2, uncovering the direction of the revenue effect will provide the first empirical answer to the question of whether enhanced targeting options increase an online

the online platform in our study auctions off advertising space. Each time a user performs browses a page with ads, the platform runs an auction to determine the allocation of ads on the content page. What matters in this setting is the number of participating ads, rather than the number of advertising accounts. A simple example will illustrate this point. Assume that a marketing company has an advertising account on the online platform in our study. Within this account, the marketing company manages the advertising campaigns of two different businesses, each selling mobile phone plans. When a user browses content about “mobile phone plans,” he sees two ads: both of them placed by the marketing company. If we measure auction participation by counting the number of ad impressions, we correctly measure it as two. However, if we measure auction participation by counting the number of participating advertiser accounts, we would incorrectly measure it as one, as there is only one account participating in the auction: the one belonging to the marketing company.

platform's revenues in a second price auction. Thus far, this question has been addressed in the literature only from a theoretical perspective.

Second, by examining how the release of the new targeting feature affects auction participation and the platform's revenue, we also contribute to the literature studying the importance of market thickness for auction revenue. In particular, we ask how the auctioneer's revenue is affected when introducing a feature that risks creating thinner markets.

Finally, we contribute to the literature studying the economics of Internet privacy. In particular, by understanding the revenue effects of finer targeting options for online platforms, we shed light not only on the platform's financial incentives to release data about its users, but also on the impact on the profitability of online platforms of potential privacy-related measures.

The first outcome variable that allows us to gain insight into auction revenue is the online platform's total revenue from the advertising auction. Since advertisers pay the online platform only when users click on their ads, the platform's total revenue from its advertising auction is simply the sum over the dollar amounts paid by advertisers for each click accrued by their ads.

Two additional outcome variables allow us to decompose the revenue effect into its underlying components: the total number of clicks and the average realized cost-per-click (CPC). The total number of clicks represents the total number of times that users clicked on the ads displayed alongside the content. If the number of clicks increases following the introduction of ZIP code targeting, this might be due to an increase in the number of ads running on the platform, to users finding the ads more relevant and clicking on them more, or a combination of the two.

If the average realized CPC decreases following the introduction of ZIP code targeting, this might signal one of three factors.¹⁴ First, it might signal that advertisements on the online platform are becoming more relevant. As noted in Section 3.2, an important feature of the online platform's auction setting is that bids are quality-adjusted. If the quality of the ads increases following the introduction of the new targeting feature, the realized CPCs will be lower. Second, it might also signal that competition on the online platform diminished. In a generalized second price auction setting, thinner markets translate into lower realized CPCs. Finally, a decrease in the realized CPC following the introduction of ZIP code targeting might also signal that participating advertisers are, on average, submitting lower bids. This scenario is plausible if, for example, the new ZIP code targeting feature attracts a large number of small, local businesses, that cannot afford to submit high bids.

5.2.3 Advertiser Outcomes

The final question we would like to address is whether finer targeting benefitted the advertisers that promoted their products on the online platform. Most of the existing literature focuses solely on the platform's perspective when studying the effects of enhanced targeting. The advertisers' perspective, however, is particularly relevant from a policy perspective. Do the advertisers have a financial incentive to obtain data about Internet users? If so, this would indicate that there is scope for policy makers to intervene in order to ensure that Internet users' privacy is protected. Also, do advertisers benefit or lose from potential regulation aimed at protecting Internet users' data? If they lose, this would indicate the need for policy makers to find a balance between the negative effects of the intervention and the benefits users derive from having their privacy protected online.

¹⁴As a reminder to the reader, the realized CPC (the price the advertisers pay for each click) is sometimes, but not always, equal to the maximum CPC (the advertiser's bid). This is explained in Section 3.2.

Two outcome variables allow us to study the advertisers’ perspective. The first one is the total number of conversions. In order to provide its advertisers with useful intelligence about their ads, the online platform offers free “conversion tracking” software. This software tracks users anonymously from clicking on an ad to purchasing, placing the item in a shopping cart, downloading or other forms of online “conversions” that are valuable to businesses. If the release of the new ZIP code targeting feature has a positive effect on the number of conversions, this would be good news for the advertisers, as it would imply that a higher number of users ended up purchasing or showing interest in the items promoted by the online platform’s advertisers.

The second outcome variable that allows us to study the advertisers’ perspective is the cost-per-conversion. This is the average price that the online platform’s advertisers pay for a conversion. We obtain this variable by dividing the total amount of money that the advertisers spend to promote their products on the online platform by the total number of conversions. If the introduction of ZIP code targeting has a negative effect on the average cost-per-conversion, this would be good news for the advertisers, as it would imply that, on average, they are paying less in order to persuade the online platform’s users to purchase their products.

Table 1 summarizes all the outcome variables introduced above.

Table 1: Outcome Variables

Variable	Description
Auction Participation	
Impressions	Total number of ads displayed on the platform in response to browsing conducted by users
Platform Revenue	
Revenue	Online platform’s total ad revenue
Clicks	Total number of clicks accrued by the ads running on the platform
Cost-per-click	Average price paid by the advertisers for a click
Advertiser Outcomes	
Conversions	Total number of times that users who clicked on ads completed purchases, downloads, or other forms of online interactions that are valuable to businesses
Cost-per-conversion	Average price paid by the advertisers for a conversion

5.3 Explanatory Variables

5.3.1 ZIP Code Targetability

ZIP code targetability is the variable whose change we are using for identification. This variable was zero during the pre-period for all designated market areas. For the post period, the value for each DMA lies between zero and one. This is because in many DMAs, only a fraction of the component ZIP codes were made targetable. To construct the [0,1] variable, we measured what fraction of the DMA’s queries were ZIP-targetable immediately after the launch. This represents the “level” of treatment for each DMA at the moment it was applied.¹⁵

Figure 1 shows that treatment strength across DMAs is highly varied, alleviating any con-

¹⁵An alternative metric might have examined the fraction of targetable queries during the post-period, rather than immediately after launch. We did not use this metric for the following reason: suppose that ZIP code targeting encouraged users to browse more (or less) in treatment areas. If we encoded treatment status as the number of targeted queries in the post-period, then this would confuse the cause (lots of targetable queries) with the effects (people entering queries that can be ZIP targeted).

cerns that treatment strength at the DMA level is concentrated around a particular value.

Figure 1: Frequency of Treatment Strength at the DMA Level

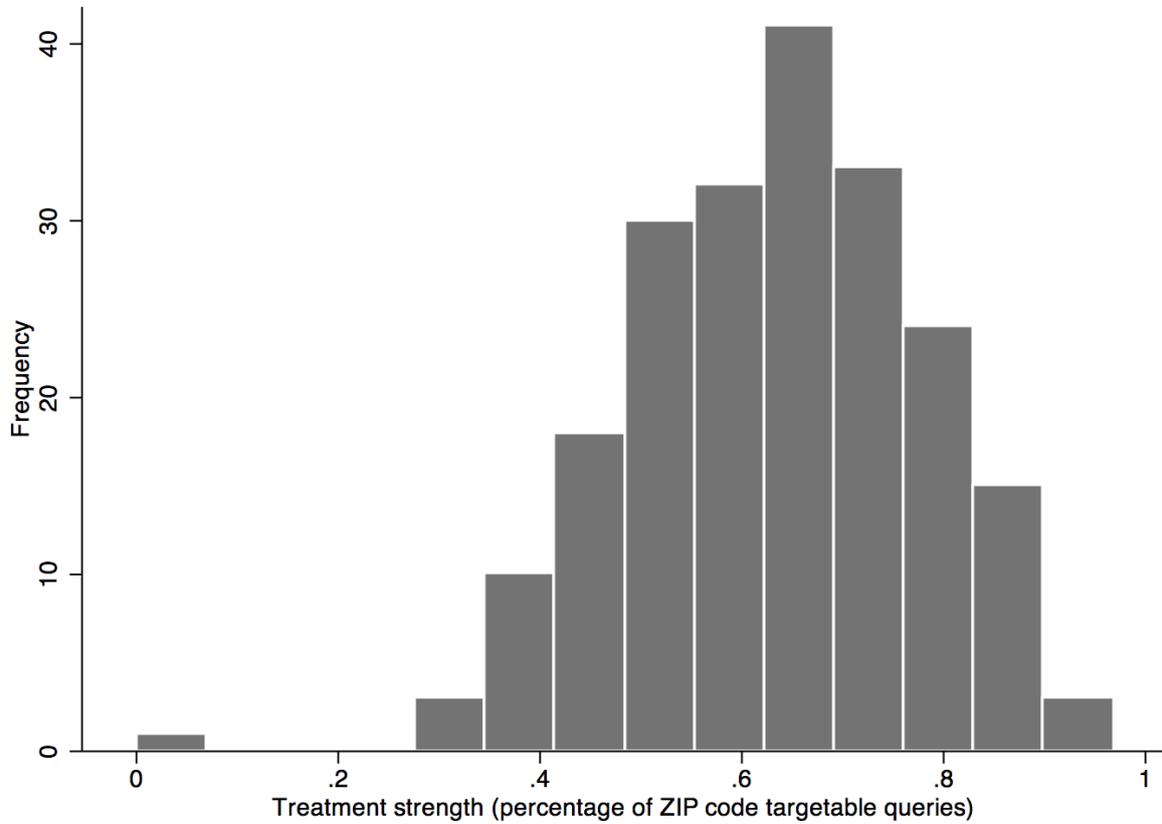
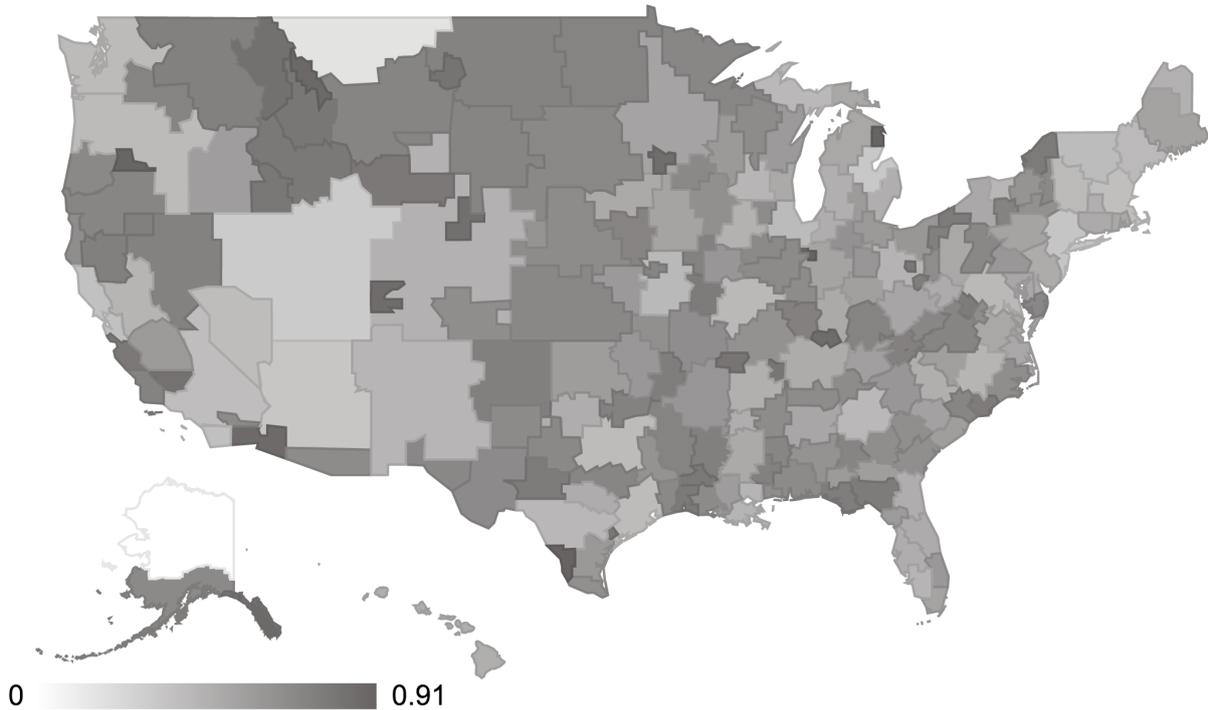


Figure 2 shows a map of all DMAs in the United States and their respective treatment strengths. ZIP code targetability, the variable we use to measure treatment strength, takes higher values for the darker areas on the map, and lower values for the lighter areas. Figure 2 is reassuring in that it shows that treatment strength is evenly distributed across the United States. For example, we do not see high concentrations of treatment predominantly along the East or West coasts or low concentrations of treatment predominantly in the middle states.

Figure 2: Geographical Distribution of Treatment Strength at the DMA Level



Besides examining the geographical distribution of treatment strength, we are also interested in knowing whether treatment strength is correlated with business presence or social and economic characteristics at the DMA level. In order to determine this, we use publicly available data sources to compile a database at the DMA level of several variables of interest, such as the number of commercial establishments per capita, median age, educational attainment, yearly income, and the unemployment rate.

Since designated market areas are not a commonly used geographical aggregation unit, none of the publicly available data sources within the United States report data at the DMA level. However, we circumvent this problem by collecting data at the ZIP code level and then using our ZIP code to DMA mapping to generate aggregate DMA figures. For the variables of interest in our study, ZIP code level data are available from the United States Census Bureau's County Business Patterns and American Community Survey. The data we collect are for 2011, the year prior to treatment.¹⁶ We report summary statistics and detailed data sources for each variable of interest in Appendix B.

We examine whether treatment strength is correlated with business presence or social and economic characteristics by regressing each variable we collect from the US Census Bureau on ZIP code targetability and a constant term. Tables 2, 3, and 4 report the results we obtain. In column (1) of each table, we present the results we obtain by estimating a simple linear regression model. In column (2) of each table, we present the results we obtain by estimating a median regression model.¹⁷

In Table 2, we investigate the relationship between treatment strength at the DMA level and business presence. In particular, we regress the number of commercial establishments per capita, a variable intended to capture the intensity of business activity at the DMA level, on ZIP code targetability, x_i . The regression results indicate that the coefficient on x_i is not signifi-

¹⁶We extracted the data using American FactFinder, <http://factfinder2.census.gov>, on 26 December 2015.

¹⁷For asymmetric distributions, a conditional quantile estimation can often pinpoint the central tendency of the distribution more accurately than the usual conditional mean estimation.

cant at any conventional levels. This result is reassuring, as it confirms that treatment strength is not correlated with business presence.

Table 2: Treatment Strength and Business Presence

		Estimation method:	
		OLS	Quantile regression
Dependent variable:		(1)	(2)
Number of commercial establishments per capita	x_i	1.521 (2.510)	-3.999 (2.419)
	constant	22.601*** (1.415)	25.108*** (1.417)
Observations		210	210
R^2		0.003	0.011

The data for the dependent variable are for 2011. x_i is our treatment strength variable. It is bounded between $[0, 1]$ and it represents the percentage of “treated” queries that can be ZIP code targeted in DMA i at time of launch. Both the dependent and the independent variables are in levels. In column (2), we report quantile regression estimates for $q = 0.5$ (the median). We estimate the standard errors using the bootstrap method with 100 resamples.

*** Significant at 1%, ** significant at 5%, * significant at 10%

In Table 3, we examine the relationship between treatment strength and several social characteristics such as age, educational attainment or the prevalence of family households. We find that the coefficient of x_i , ZIP code targetability, is not significant at the 5% level for any of the variables we investigate. These results are equally reassuring, as they confirm that treatment strength is not correlated with social characteristics.

Table 3: Treatment Strength and Social Characteristics

		Estimation method:	OLS	Quantile regression
Dependent variable:			(1)	(2)
Median age	x_i		0.069 (1.745)	1.724 (2.168)
	constant		37.857*** (1.079)	36.961*** (1.354)
	Observations		210	210
	R^2		0.000	0.003
	Educational attainment: high school or higher (% of people over 25 years)			
Educational attainment: high school or higher (% of people over 25 years)	x_i		-0.060* (0.031)	-0.057 (0.037)
	constant		0.890*** (0.016)	0.897*** (0.022)
	Observations		210	210
	R^2		0.029	0.010
	Family households (% of total households)			
Family households (% of total households)	x_i		0.017 (0.020)	0.008 (0.016)
	constant		0.656*** (0.012)	0.660*** (0.009)
	Observations		210	210
	R^2		0.006	0.002

The data for the dependent variable are for 2011. x_i is our treatment strength variable. It is bounded between $[0, 1]$ and it represents the percentage of “treated” queries that can be ZIP code targeted in DMA i at time of launch. Both the dependent and the independent variables are in levels. In column (2), we report quantile regression estimates for $q = 0.5$ (the median). We estimate the standard errors using the bootstrap method with 100 resamples. *** Significant at 1%, ** significant at 5%, * significant at 10%

Finally, in Table 4, we investigate the relationship between treatment strength and income. The results we report in the first row of Table 4 indicate that there is a negative and statistically significant relationship between ZIP code targetability and yearly income per capita. The magnitude of the coefficient we estimate for x_i indicates that a 1% increase in ZIP code targetability is associated with a \$112 decrease in yearly per capita income. This finding is surprising: if the online platform in our study was strategic in choosing the geographical areas for the release of the new ad targeting feature, we would expect the platform to choose higher income neighborhoods, with greater purchasing power, rather than lower income regions.

Table 4: Treatment Strength and Economic Characteristics

		Estimation method:	OLS	Quantile regression
Dependent variable:			(1)	(2)
Yearly income per capita (thousands 2011 USD)	x_i		-11.206*** (1.949)	-12.279*** (1.874)
	constant		31.695*** (1.287)	32.003*** (1.211)
	Observations		210	210
	R^2		0.180	0.124
Unemployed (% of people over 16 years)	x_i		-0.012 (0.011)	-0.012 (0.012)
	constant		0.092*** (0.007)	0.091*** (0.007)
	Observations		210	210
	R^2		0.007	0.005
Housing units with no vehicles available (% of occupied housing units)	x_i		-0.034* (0.018)	-0.011 (0.011)
	constant		0.092*** (0.011)	0.074*** (0.007)
	Observations		210	210
	R^2		0.052	0.005
Owner-occupied housing units (% of occupied housing units)	x_i		0.041 (0.027)	0.022 (0.032)
	constant		0.658*** (0.017)	0.676*** (0.018)
	Observations		210	210
	R^2		0.017	0.004
Housing units with less than one occupant per room (% of occupied housing units)	x_i		-0.001 (0.014)	0.004 (0.007)
	constant		0.975*** (0.008)	0.978*** (0.005)
	Observations		210	210
	R^2		0.000	0.001

The data for the dependent variable are for 2011. x_i is our treatment strength variable. It is bounded between $[0, 1]$ and it represents the percentage of “treated” queries that can be ZIP code targeted in DMA i at time of launch. Both the dependent and the independent variables are in levels. In column (2), we report quantile regression estimates for $q = 0.5$ (the median). We estimate the standard errors using the bootstrap method with 100 resamples.

*** Significant at 1%, ** significant at 5%, * significant at 10%

In the remaining rows of Table 4, we further probe this result. In particular, we are concerned that the yearly income per capita measure we are using might be measured with error. This is due to the fact that there is a significant amount of research indicating that income figures in household survey data are measured with error.¹⁸ For example, [Pedace and Bates \(2000\)](#), [Bollinger \(1998\)](#), and [Hurst et al. \(2014\)](#) convincingly show that survey respondents grossly misreport their income in household surveys.

If people misreport their income in household surveys, are there other variables that we

¹⁸[Bound et al. \(2001\)](#) provide an excellent review of this literature.

can look at to assess their financial wellbeing? We obtain four such variables from the U.S. Census: the unemployment rate, the percentage of housing units that do not own any vehicle, the percentage of owner-occupied housing units (as opposed to rented housing units), and the percentage of housing units that have less than one occupant per room. We find that ZIP code targetability does not have a statistically significant effect on any of these variables. These results are encouraging, in that they suggest that the relationship between treatment strength and financial wellbeing might not be as strong as implied by the results we obtain using yearly income data.

Overall, the results presented in Tables 2, 3, and 4 alleviate concerns that treatment strength is correlated with business presence or the population’s social and economic characteristics. The results also support the platform’s position that the selection criteria into ZIP code targeting were a mixture of privacy and technological related reasons, rather than strategic motivations.

5.3.2 Additional Controls

In our regressions, we also include a few additional controls. First, we control for the device used by the people browsing the content: *percent_computer_{ijt}* measures the percentage of queries originating from personal computers, while *percent_tablet_{ijt}* measures the percentage of queries originating from tablets. The excluded category is *percent_mobile_{ijt}*, the percentage of queries originating from mobile phones.

Second, we control for the query mix on the platform by including the percentage of public queries, *percent_public_queries_{ijt}*. The online platform in our study considers a public query any topic that meets a global frequency threshold. This threshold requires the topic to have been queried by at least 20 unique users or 20 unique IP addresses within a 90 day window. Once a topic passes the threshold, it remains labeled in the platform’s databases as a public query permanently.

Finally, we include in our regression time, as well as (DMA × industry) fixed effects.

6 Specifications and Identification

We want to identify the effect of changes in the “treatment” status of DMAs on the various outcomes identified in Section 5.2. In other words, we are interested in estimating γ in the equation below:

$$\Delta y_{ijt} = \gamma x_{ijt} + \beta_1 \text{percent_computer}_{ijt} + \beta_2 \text{percent_tablet}_{ijt} + \beta_3 \text{percent_public_queries}_{ijt} + \delta_t + \zeta_{ij} + \epsilon_{ijt} \quad (1)$$

where $t = \text{February 2012 and February 2013}$, and:

Δy_{ijt} = Change in the outcome variable at the (DMA i × industry j) level between periods t and $t - 1$: a full list of outcome variables is available in Table 1.

x_{ijt} = Variable of interest: percentage of “treated” queries that can be ZIP code targeted in DMA i and industry j at time of launch.

percent_computer_{ijt} = Control: percentage of queries originating from personal computers.

percent_tablet_{ijt} = Control: percentage of queries originating from tablets.

percent_public_queries_{ijt} = Control: percentage of public queries.

δ_t = Controls: period fixed effects. This is essentially a dummy variable taking the value of 0 prior to the launch, and 1 following the launch.

ζ_{ij} = Controls: (DMA \times industry) fixed effects.
 ϵ_{ijt} = Error term: standard errors clustered at DMA level.

On the left-hand side, we choose to look at changes in outcomes, rather than outcomes, for the variables of interest. In other words, instead of estimating our regressions using the data that we have for each outcome variable for February 2012 and February 2013, we estimate our regressions on the differences for each outcome variable between the months of (February 2012 - February 2011) and (February 2013 - February 2012). By transforming our outcome variables in this way, we control for trends that are common to (DMA \times industry) combinations. We also include in our regressions (DMA \times industry) fixed effects, which allow us to ensure that our results are not affected by any underlying heterogeneity across (DMA \times industry) pairs that is fixed over time.

We normalize the left-hand side variables in our study by subtracting the mean and dividing by the standard deviation.¹⁹ We estimate Equation 1 by differences-in-differences, as well as by employing an instrumental variable (IV) approach combined with a fuzzy regression discontinuity (RD) design. We discuss each method below.

6.1 Differences-in-Differences

Since treatment was applied orthogonally to any of the outcome variables in our study, a difference-in-difference estimator is an appropriate way to estimate Equation 1. However, there is one additional step we take in order to ensure that this is, indeed, the case.

We calculate the growth rates between February 2011 and February 2012 (the year before the introduction of the ZIP code targeting feature), for revenue, clicks, impressions, conversions, and queries. We then estimate a regression with x_{ij} (the percentage of “treated” queries that can be ZIP code targeted in DMA i and industry j at the time of launch) as the dependent variable and the growth rates we obtain as independent variables.²⁰ If the platform launched ZIP code targeting only to high growth areas, we would expect the estimation to return positive and significant coefficients. However, if the platform launched ZIP code targeting orthogonally to the growth of the variables in our regression, we would expect the estimation to return insignificant results.

Table 5 summarizes the results we obtain. All coefficients are insignificant, with the exception of the coefficient of the variable measuring growth in the number of clicks. This coefficient, however, is only significant at the 5% level. In terms of magnitude, the estimate is very close to zero and it has a negative sign, which is the opposite of what we would expect if the platform was being strategic in its selection of ZIP codes for the release of the new targeting feature. Given this, we believe that the results in Table 5, combined with the analysis presented in Section 5.3.1, provide convincing evidence for the appropriateness of a difference-in-differences estimator.

¹⁹The outcome variables in our study contain highly sensitive information. By normalizing the variables, we ensure that the confidentiality of the online platform’s data is protected.

²⁰Note that we do not include the cost-per-click and the cost-per-conversion in our regression as they are ratios of the variables whose growth rates we do include.

Table 5: Treatment Selection Analysis

Revenue Growth _{ij}	0.0008 (0.0005)
Clicks Growth _{ij}	-0.0069** (0.0032)
Impressions Growth _{ij}	0.0042 (0.0027)
Conversions Growth _{ij}	-0.0005 (0.0004)
Query Growth _{ij}	0.0020 (0.0018)
Industry fixed effects (ζ_j)	Yes
DMA fixed effects (α_i)	Yes
Observations	5,123
R^2	0.978

Dependent variable: treatment status (percentage of “treated” queries that can be ZIP code targeted in DMA i and industry j at time of launch). Each variable measures growth prior to treatment, from Feb 2011 to Feb 2012. Standard errors clustered at DMA level.

*** Significant at 1%, ** significant at 5%, * significant at 10%

6.2 Instrumental Variable Approach Combined with a Fuzzy Regression Discontinuity Design

Although we believe that our use of a differences-in-differences estimator is justified, as a robustness check, we ensure that γ , our parameter of interest in Equation 1, is properly identified by employing a instrumental variable (IV) approach combined with a fuzzy regression discontinuity (RD) design.

As noted in Section 4, due to privacy reasons, ZIP codes where the platform’s engineers were not able to detect N cookies or more did not get ZIP code targeting. We use this cutoff rule to instrument for treatment status. In particular, we estimate the following first stage regression:

$$x_{ijt} = \theta z_{ijt} + \beta_1 \text{percent_computer}_{ijt} + \beta_2 \text{percent_tablet}_{ijt} + \beta_3 \text{percent_public_query}_{ijt} + \beta_4 \text{percent_1000+_queries}_{ijt} + \delta_t + \zeta_{ij} + v_{ijt} \quad (2)$$

where:

x_{ijt} = Dependent variable: percentage of “treated” queries that can be ZIP code targeted in DMA i and industry j at time of launch.

z_{ijt} = Instrument: takes the value of 0 pre-treatment and equals the percentage of queries (targetable and un-targetable) in DMA i and industry j coming from ZIP codes with more than N cookies post-treatment.

$\text{percent_computer}_{ijt}$ = Control: percentage of queries originating from personal computers.

$\text{percent_tablet}_{ijt}$ = Control: percentage of queries originating from tablets.

$\text{percent_public_queries}_{ijt}$ = Control: percentage of public queries.

$\text{percent_1000+_queries}_{ijt}$ = Control: percentage of queries (targetable and un-targetable) in DMA i and industry j coming from ZIP codes with more than N cookies.²¹

²¹Since our instrument is the interaction of two variables, $\text{percent_1000+_queries}_{ijt}$ and δ_t , we must also include the uninteracted terms in both our first and second stage regressions.

δ_t = Controls: period fixed effects. This is essentially a dummy variable taking the value of 0 prior to the launch, and 1 following the launch.
 ζ_{ij} = Controls: (DMA \times industry) fixed effects.
 v_{ijt} = Error term: standard errors clustered at DMA level.

Our econometric approach is similar to that of Angrist et al. (1996). In a seminal work, the researchers use the cutoff rule for assignment to treatment as an instrument for treatment status and show that in the second stage regression, the estimated coefficient of the treatment status variable measures the local average treatment effect. This is the causal effect of receiving treatment for the subset of the population that is assigned to treatment and receives the treatment. Angrist et al. (1996) work with binary treatments and instruments, but Angrist and Imbens (1995) extend their analysis to multi-valued treatments and instruments.

Our instrument, z_{ijt} , is strong by design. The platform’s engineers released ZIP code targeting only to ZIP codes above the N threshold. This cutoff rule necessarily induces a correlation between the number of queries from ZIP codes with cookie populations above N and the number of treated queries.²² This ensures that θ in Equation 2 is different from zero.

Also by design, our instrument is uncorrelated with ϵ_{ijt} , the error term in Equation 1. The N cookie threshold was exogenously determined by the privacy related cutoff rule. It is orthogonal to the online platform’s revenues and to other outcomes of interest to us. This ensures that any effect of the cutoff rule on y_{ijt} , the outcome of interest in Equation 1, comes solely through its effect on determining the treatment status, x_{ijt} .

We use \widehat{x}_{ijt} , the predicted values we obtain from the first stage regression, to replace the actual values of x_{ijt} in Equation 1. This allows us to estimate γ , the casual effect of ZIP code targetability on the outcomes of interest.

To construct our instrument, we obtain Internet cookie data collected by the online platform in our study.²³ Ideally, we would like to have access to the data that the platform’s engineers used for the treatment cutoff rule. Unfortunately, those data are no longer available. The platform’s privacy rules require it to delete all Internet cookie data older than two years. We are, however, able to obtain Internet cookie data from September 2, 2013. Although this is not an exact measure of the Internet cookie population at the time the treatment was applied, it is a good approximation. The United States is a mature market in terms of Internet usage. As such, the growth rates of variables such as Internet cookies, which are closely tied to the number of Internet users, do not experience large increases from year to year.²⁴

In Figure 3, we show the probability of a ZIP code being treated as a function of the number of Internet cookies detected in that ZIP code. There is a clear discontinuity around N Internet cookies. To the left of this threshold, the probability of receiving treatment is below 0.15, while to the right of the threshold, the probability rises to 0.54.

Figure 3 also makes it clear that the discontinuity at N Internet cookies is fuzzy. The probability of receiving treatment below the N cookie threshold is not exactly zero due to our use of Internet cookie data from September 2013 to approximate the cookie population at the time of treatment. At the same time, the probability of receiving treatment above the N cookie threshold is not exactly one, either. This is because some ZIP codes that met the N cookie threshold remained un-targetable as they did not meet the platform’s accuracy requirements.²⁵

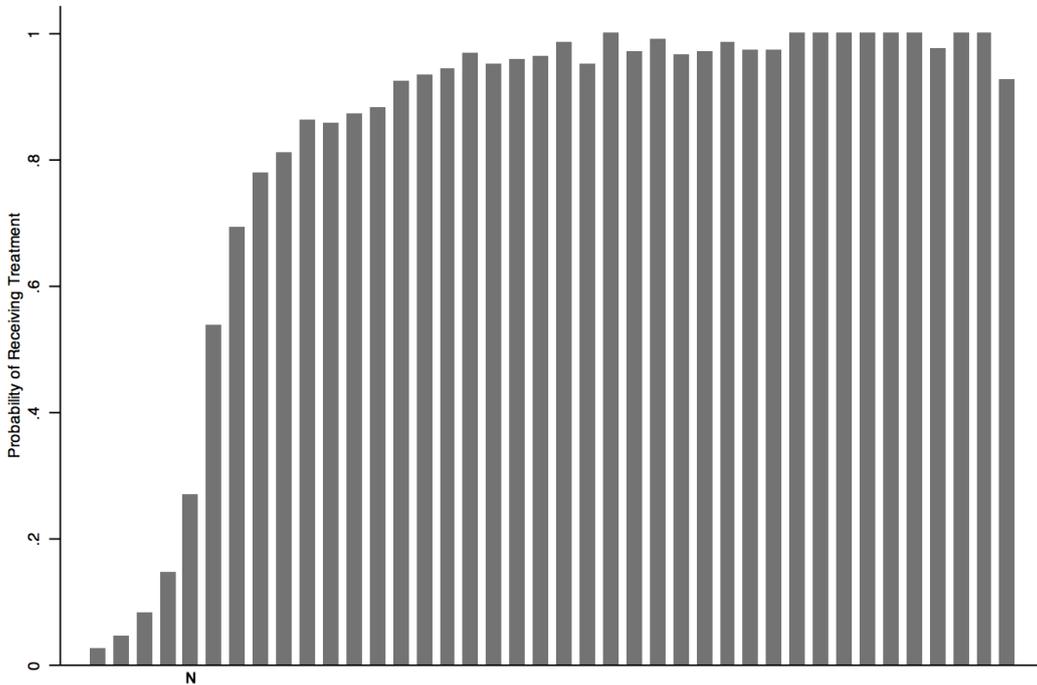
²²Note that this is not a perfect correlation: not all ZIP codes above the N threshold received treatment. Only those ZIP codes that met the platform’s accuracy requirements got ZIP code targeting.

²³Internet cookie data are not available from public sources.

²⁴We do experiment with various thresholds higher than N cookies in order to account for possible growth in the number of Internet cookies. Our results, however, remain unchanged.

²⁵2,821 out of the total 11,539 ZIP codes that met the N Internet cookie threshold remained untreated.

Figure 3: Probability of a ZIP Code Being Treated by Number of Internet Cookies



We estimate Equation 2, our first stage regression, and we obtain an F-statistic of 5,769.²⁶ This is well above the conventional required threshold. The estimated coefficient for z_{ijt} , our instrument, is 0.988, while the standard error is 0.003. With these results, we are satisfied that both our instrument and our first stage equation are very strong.

7 Results

In this section, we present the results we obtain for each outcome variable and estimation method, discussing auction participation, platform’s revenue, and advertiser outcomes following the introduction of the ZIP code targeting feature.

For each outcome variable, we report the coefficient we obtain for x_{ijt} , our variable of interest, by estimating the regressions using a differences-in-differences (DiD) approach and using an instrumental variable approach combined with a fuzzy regression discontinuity design (IVFRD). The DiD coefficients can be interpreted as the average effect on the dependent variable of a unit increase in treatment strength (in our case, ZIP code targetability) for the treated areas (in our case, the areas where the platform released ZIP code targeting). The IVFRD coefficients can be interpreted as the average effect on the outcome of interest of a unit increase in ZIP code targetability for areas whose treatment status was influenced by the instrument (in our case, the N Internet cookie cutoff rule).

²⁶We report here the Kleibergen-Paap, rather than the Cragg-Donald, F-statistic as we are using heteroscedasticity robust standard errors.

7.1 Auction Participation

Table 6 details out the results we obtain when examining the change in the total number of ad impressions – the outcome variable that allows us to examine the effect of enhanced targeting on auction participation.

When estimated using the differences-in-differences approach, the coefficient we obtain for x_{ijt} , the variable measuring treatment strength, is negative and significant at the 5% level.²⁷ The magnitude of the estimated coefficient implies that a unit increase in treatment strength resulted, on average, in a decrease of 1.92 standard deviations in the number of ad impressions relative to their underlying trend within the areas where the online platform introduced ZIP code targeting.

When estimated using the instrumental variable approach combined with a fuzzy regression discontinuity design, the coefficient we obtain for x_{ijt} is negative and significant at the 1% level. The magnitude of the estimated coefficient implies that a unit increase in treatment strength resulted, on average, in a decrease of 1.82 standard deviations in the number of ad impressions relative to their underlying trend within the areas where treatment status was influenced by the instrument.

Table 6: Auction Participation Results

Dependent variable:	Change in the number of impressions (normalized)	
	DiD (1)	IVFRD (2)
Estimation method:		
x_{ijt}	-1.916** (0.797)	-1.821*** (0.543)
Query origin controls (percent_computer _{ijt} , percent_tablet _{ijt})	Yes	Yes
Query mix controls (percent_public_query _{ijt})	Yes	Yes
Uninteracted IV control (percent_1000+_queries _{ijt})	No	Yes
Period FEs (δ_t)	Yes	Yes
DMA \times industry FEs (ζ_{ij})	Yes	Yes
Observations	10,500	10,500
R^2	0.419	

The dependent variable is the change in the total number of ad impressions between the months of (February 2013 - February 2012) and (February 2012 - February 2011). The independent variables are in levels, based on data from February 2013 and February 2012. ZIP code targeting was launched in March 2012. Standard errors are clustered at the DMA level.

*** Significant at 1%, ** significant at 5%, * significant at 10%

The results we present in Table 6 show that the introduction of ZIP code targeting had a negative and statistically significant effect on auction participation. With this finding, we contribute to the literature on several fronts.

²⁷With a p-value of 0.017, the coefficient barely misses the significance threshold at the 1% level.

First, similar to [Tadelis et al. \(2015\)](#), we call into question the theoretical assumption that information disclosure has no effect on auction participation. However, unlike [Tadelis et al. \(2015\)](#), we are able to measure auction participation in our data and to show, empirically, that in an auction with multiple objects for sale, the information disclosed by the auctioneer has a statistically significant effect on auction participation. By identifying this effect in the data, we underline the importance of allowing for an endogenous sorting of bidders into auctions when modeling the effect of information disclosure on auction outcomes.

Second, we confirm [Levin and Milgrom \(2010\)](#)'s intuition that by offering finer targeting options, online platforms risk creating thinner markets. Indeed, in the setting that we study, the introduction of ZIP code targeting had a negative effect on auction participation. But – did the thinner markets have a negative effect on the online platform's revenue? This is the question we turn to next.

7.2 Platform Revenue

Table 7 details out the results we obtain for the outcome variables that allow us to shed light on the effect of enhanced targeting on the online platform's revenues. Columns (1) and (2) report the coefficients we obtain when examining the change in the platform's total revenue. Whether estimated using the differences-in-differences approach or the instrumental variable approach combined with a fuzzy regression discontinuity design, the coefficients we obtain for x_{ijt} , the variable measuring treatment strength, are positive and significant at the 1% level. This means that DMAs with higher treatment intensity saw greater revenue increases following the introduction of ZIP code targeting. In terms of magnitudes, the estimated coefficient for x_{ijt} is 2.36 standard deviations using the differences-in-differences approach and 2.60 standard deviations when using the instrumental variable approach combined with a fuzzy regression discontinuity design.

The revenue results we report in columns (1) and (2) of Table 7 are important for several reasons. First, combined with the results reported in Table 6, our findings suggest that although the introduction of ZIP code targeting had a negative effect on auction participation, the release of the finer targeting feature had a positive effect on the online platform's revenues. By showing that thinner markets need not always result in lower revenues for the auctioneer, our work contributes to the literature studying the benefits of market thickness in an auction context.

Second, our results are the first empirical confirmation of the theoretical predictions advanced by papers like [Hummel and McAfee \(2015\)](#) and [Fu et al. \(2012\)](#), which ask, from a theoretical perspective, whether enhanced targeting increases an online platform's revenues in a second price auction.

Finally, our findings indicate that online platforms do not have a financial incentive to protect Internet users' privacy. In light of this, policy measures designed to limit the amount of targeting offered by online platforms may be needed in order to ensure that Internet users' data are protected.

Table 7: Platform Revenue Results

Dependent variable:	Change in platform's revenue (normalized)		Change in the number of clicks (normalized)		Change in the cost-per-click (normalized)	
	DiD (1)	IVFRD (2)	DiD (3)	IVFRD (4)	DiD (5)	IVFRD (6)
Estimation method:						
x_{ijt}	2.362*** (0.836)	2.599*** (0.597)	3.557*** (1.178)	3.708*** (0.823)	-0.239 (0.197)	-0.442*** (0.144)
Query origin controls (percent_computer _{ijt} , percent_tablet _{ijt})	Yes	Yes	Yes	Yes	Yes	Yes
Query mix controls (percent_public_query _{ijt})	Yes	Yes	Yes	Yes	Yes	Yes
Uninteracted IV control (percent_1000+_queries _{ijt})	No	Yes	No	Yes	No	Yes
Period FEs (δ_t)	Yes	Yes	Yes	Yes	Yes	Yes
DMA \times industry FEs (ζ_{ij})	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,500	10,500	10,500	10,500	10,500	10,500
R^2	0.513		0.330		0.566	

The dependent variables are the change in revenue (columns (1) and (2)), clicks (columns (3) and (4)), and cost-per-click (columns (5) and (6)) between the months of (February 2013 - February 2012) and (February 2012 - February 2011). The independent variables are in levels, based on data from February 2013 and February 2012. ZIP code targeting was launched in March 2012. Standard errors are clustered at the DMA level.

*** Significant at 1%, ** significant at 5%, * significant at 10%

In columns (3), (4), (5), and (6) of Table 7, we decompose the online platform's revenue into its two components: the number of clicks and the cost-per-click.²⁸ The coefficient on the variable measuring treatment strength is positive and significant at the 1% level in columns (3) and (4), where the outcome variable is the change in the number of clicks. The magnitude of the coefficient estimated using the differences-in-differences approach implies that a unit increase in treatment strength resulted, on average, in an increase of 3.56 standard deviations in the number of ad clicks relative to their underlying trend within the areas where the online platform introduced ZIP code targeting. The instrumental variable approach combined with a fuzzy regression discontinuity design produces an estimate of a similar magnitude.

In Section 7.1, we saw that the introduction of ZIP code targeting had a negative effect on the number of ads running on the platform. The results reported in columns (3) and (4) of Table 7, however, show that despite its negative impact on the number of ads appearing on the platform, the release of the enhanced targeting feature had a positive effect on the number of clicks accrued by the ads. The ads, in other words, became more relevant.

The intuition behind this result is simple: with better access to information about the Internet users browsing, advertisers were able to choose more precisely what auctions to participate in: they exited irrelevant ad auctions and participated only in the auctions where the users seeing the ads were likely to show interest in the advertised product or service. This sorting of advertisers into auctions following the release of ZIP code targeting resulted in a better match between the ads displayed on the platform and the users seeing them.

In columns (5) and (6) of Table 7 we report the results we obtain when our outcome variable

²⁸Because we normalize each dependent variable to protect the confidentiality of the online platform's data, the coefficients for the change in the number of clicks and the change in the cost-per-click cannot be added in the usual way to obtain the coefficients for total revenue.

is the change in the realized cost-per-click. Depending on the estimation method used, we find that the coefficient for x_{ijt} is either insignificant or negative.²⁹

As noted in Section 5.2.2, in the context of the auction that we study, potential declines in the realized CPCs might arise due to several factors. First, the realized CPCs on the platform might decline as a result of ads becoming more relevant. Second, the realized CPCs might decline due to diminishing competition on the online platform. Finally, the realized CPCs might also decline due to advertisers submitting lower bids.

In light of our previous results – namely that the introduction of ZIP code targeting increased the relevancy of the ads running on the online platform and decreased auction participation – the first two factors identified in the paragraph above are likely to be the key drivers behind a potential cost-per-click decline following the release of the finer targeting feature. Although we cannot test this empirically, we note that from a theoretical perspective, it is unlikely for advertisers to submit lower bids following the introduction of ZIP code targeting: most auction models predict that a better match between the auction participants and the auctioned items intensifies competition between the bidders and results in higher bids.³⁰

While we are not able to identify the underlying dynamics behind the cost-per-click results reported in Table 7, one thing is for certain: in the treated areas, there were no increases in the average realized cost-per-click relative to its underlying trend following the introduction of ZIP code targeting. This result has an important implication: in the treated areas, the platform’s revenues following the release of the finer targeting feature saw higher gains solely due to increases in the volume of the clicks accrued by ads – and not due to increases in the cost-per-click. Existing theoretical works focus only on the effect of enhanced targeting on auction prices, ignoring its effect on the volume of items sold. Our findings suggest that this is a potentially problematic oversight: for the online platform in our study, when it comes to revenue, the volume effect dominates.

7.3 Advertiser Outcomes

In Table 8, we report the results we obtain for the change in the total number of conversions and the change in the cost-per-conversion – the outcome variables that allow us to examine whether the introduction of the finer targeting feature benefitted the advertisers promoting their products on the online platform.

Columns (1) and (2) report the coefficients we obtain when examining the change in the total number of conversions. Regardless of the estimation method used, the coefficients we obtain for x_{ijt} , the variable measuring treatment strength, are positive and significant at the 1% level. These results imply that DMAs with higher treatment intensity experienced greater growth in the number of conversions following the introduction of ZIP code targeting. The magnitude of the estimated coefficient using the difference-in-difference approach suggests that a unit increase in treatment strength led, on average, in an increase of 3.54 standard deviations in the

²⁹Differences between the coefficients estimated using the two methodologies may arise due to the fact that the estimators recover different effects. While the differences-in-differences estimator recovers the average effect on the dependent variable of a unit increase in treatment strength for the treated areas, the instrumental variable approach combined with a fuzzy regression discontinuity design recovers the average effect on the dependent variable of a unit increase in treatment strength for the areas whose treatment status was influenced by the instrument. [Imbens and Angrist \(1994\)](#) discuss this in detail.

³⁰The reason we cannot test this empirically is that we do not have access to maximum cost-per-click data. The differences that necessarily exist in a generalized second price position auction between the realized cost-per-click (the price that an advertiser pays for a click) and the maximum cost-per-click (the advertiser’s actual bid in the auction) mean that we cannot use the results we report in columns (5) and (6) of Table 7 to ascertain whether the average bid on the platform increased or decreased following the introduction of ZIP code targeting.

number of conversions relative to their underlying trend within the areas where the platform introduced ZIP code targeting. The instrumental variable approach combined with a fuzzy regression discontinuity design produces an estimate of 3.63 standard deviations.

Table 8: Advertiser Outcomes

Dependent variable:	Change in conversions (normalized)		Change in the cost-per-conversion (normalized)	
	DiD (1)	IVFRD (2)	DiD (3)	IVFRD (4)
Estimation method:				
x_{ijt}	3.543*** (1.213)	3.625*** (0.842)	-0.213 (0.258)	-0.419** (0.172)
Query origin controls (percent_computer _{ijt} , percent_tablet _{ijt})	Yes	Yes	Yes	Yes
Query mix controls (percent_public_query _{ijt})	Yes	Yes	Yes	Yes
Uninteracted IV control (percent_1000+_queries _{ijt})	No	Yes	No	Yes
Period FEs (δ_t)	Yes	Yes	Yes	Yes
DMA \times industry FEs (ζ_{ij})	Yes	Yes	Yes	Yes
Observations	10,500	10,500	10,500	10,500
R^2	0.221		0.726	

The dependent variables are the change in total conversions (columns (1) and (2)) and cost-per-conversion (columns (3) and (4)) between the months of (February 2013 - February 2012) and (February 2012 - February 2011). The independent variables are in levels, based on data from February 2013 and February 2012. ZIP code targeting was launched in March 2012. Standard errors are clustered at the DMA level.

*** Significant at 1%, ** significant at 5%, * significant at 10%

Columns (3) and (4) in Table 8 report the coefficients we obtain when examining the change in the cost-per-conversion. When employing the difference-in-difference approach, we obtain an insignificant coefficient, while when employing the instrumental variable approach, we obtain a negative coefficient that does not meet the significance threshold at the 1% level.

The results in Table 8 suggest that advertisers benefitted from the introduction of ZIP code targeting. The release of the finer targeting feature had a positive impact on the number of conversions. This suggests that following treatment, a higher number of users completed purchases, downloads, or other forms of valuable online interactions on the advertisers' websites. At the same time, the introduction of ZIP code targeting did not lead to increases in the average cost-per-conversion. Combined, these two results indicate that enhanced targeting was beneficial to advertisers: in regions with higher treatment strength, ZIP code targeting had a positive effect on the volume, but not on the price, of online conversions.

Our final finding, regarding the benefits of enhanced targeting to advertisers, contributes to the literature designed to help policy makers understand whether they should intervene to protect Internet users' privacy and how potential policy measures might affect advertisers. In particular, our results indicate that like online platforms, advertisers have business incentives to gather and use Internet users' data. This finding strengthens the case for policy interventions aimed at protecting online users' privacy. Furthermore, our results indicate that potential privacy-related measures might negatively impact the profitability of the advertisers who promote their products on online platforms. This finding encourages regulators to design policies

that mitigate the negative effects of regulation on the online advertising industry.

8 Conclusion

Our results provide empirical evidence for the key finding of stylized models of ad targeting, namely that better targeting leads to increases in revenue. We show that for the online platform in our study, enhanced ad targeting was profitable despite concerns about market thinness and partitioned demand.

We also show that the introduction of ZIP code targeting had a statistically significant effect on auction participation. The theoretical literature on targeting and auction design ignores the effect of information revelation on auction participation. With this paper, we underline the importance of allowing for endogenous entry and exit into the advertising space when modeling the effects of enhanced targeting.

From a policy perspective, our findings suggest that both online platforms and their advertisers have business incentives to gather – and use – even more information about people browsing online. For Internet users who value their privacy, the prospect of having more of their data collected and used for advertising purposes is undoubtedly worrisome. Regulation can limit the online industry’s ability to make use of Internet users’ data for commercial purposes, but policy makers will need to carefully design it so that it strikes a balance between its negative effects on the online platforms and their advertisers and its benefits to Internet users.

Our findings raise several questions for future research. Our paper aims to inform a more generic market design question about the tradeoff between specificity and competition between buyers. Little theoretical or empirical work deals directly with this tradeoff in market design. Although we find results in an empirical setting, this setting also has some limitations. We have estimates for one type of market (online ads) and one type of intervention (ZIP code targeting).

The natural experiment we study led to higher revenue growth in the treated areas. However, other changes to specificity in this market (and others) may not. At some point, thinning markets through targeting may deleteriously impact seller revenue. How can market design practitioners and regulators know the tipping point where a proposed targeting feature starts to harm revenue? Future research could inform this question through a combination of new theory and structural modeling.

In this paper, we focus on aggregate data and dynamics. Future researchers might be interested in understanding how the nature of improved specificity impacts outcomes. Some types of specificity may help sellers differentiate themselves horizontally. For example, in some cases, ZIP code targeting could help local businesses, such as restaurants, advertise to local customers. Other types of specificity can also permit vertical differentiation. For example, ZIP code targeting could be useful for businesses who want to bid differently for wealthy neighborhoods vs poor neighborhoods. A specificity increase (such as ZIP code targeting) may have different impacts on revenue, auction participation, and advertiser outcomes depending on whether vertical or horizontal segmentation are in demand from advertisers. We hope to see these – and other – research questions explored as better data become available.

References

- Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael Grubb**, "Peaches, lemons, and cookies: Designing auction markets with dispersed information," in "Proceedings of the fourteenth ACM conference on Electronic commerce" ACM 2013, pp. 7–8.
- Angrist, Joshua D and Guido W Imbens**, "Two-stage least squares estimation of average causal effects in models with variable treatment intensity," *Journal of the American statistical Association*, 1995, 90 (430), 431–442.
- , – , and **Donald B Rubin**, "Identification of causal effects using instrumental variables," *Journal of the American statistical Association*, 1996, 91 (434), 444–455.
- Athey, Susan and Joshua S Gans**, "The impact of targeting technology on advertising markets and media competition," *American Economic Review*, 2010, 100 (2), 608–613.
- Bailey, Michael Charles**, "Essays in applied economics: Evidence from natural experiments." PhD dissertation, Stanford University 2012.
- Bergemann, Dirk and Alessandro Bonatti**, "Targeting in advertising markets: Implications for offline versus online media," *RAND Journal of Economics*, 2011, 42 (3), 417–443.
- Bhawalkar, Kshipra, Patrick Hummel, and Sergei Vassilvitskii**, "Value of targeting," in "Algorithmic Game Theory," Springer, 2014, pp. 194–205.
- Board, Simon**, "Revealing information in auctions: The allocation effect," *Economic Theory*, 2009, 38 (1), 125–135.
- Bollinger, Christopher R**, "Measurement error in the current population survey: A nonparametric look," *Journal of Labor Economics*, 1998, 16 (3), 576–594.
- Bound, John, Charles Brown, and Nancy Mathiowetz**, "Measurement error in survey data," *Handbook of Econometrics*, 2001, 5, 3705–3843.
- Bulow, Jeremy and Paul Klemperer**, "Auctions versus negotiations," *American Economic Review*, 1996, 86 (1), 180–194.
- Corniere, Alexandre De and Romain De Nijs**, "Online advertising and privacy," Available at SSRN 2191124, 2014.
- Edleman, Benjamin, Michael Ostrovsky, and Michael Schwarz**, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American Economic Review*, 2007, 97 (1), 242–259.
- Fu, Hu, Patrick Jordan, Mohammad Mahdian, Uri Nadav, Inbal Talgam-Cohen, and Sergei Vassilvitskii**, "Ad auctions with data," in "Algorithmic Game Theory," Springer, 2012, pp. 168–179.
- Ganuzza, Juan-José**, "Ignorance promotes competition: An auction model with endogenous private valuations," *Rand Journal of Economics*, 2004, 35 (3), 583–598.
- Goldfarb, Avi and Catherine E Tucker**, "Privacy regulation and online advertising," *Management Science*, 2011, 57 (1), 57–71.
- Hummel, Patrick and R Preston McAfee**, "When does improved targeting increase revenue?," in "Proceedings of the 24th International Conference on World Wide Web" International World Wide Web Conferences Steering Committee 2015, pp. 462–472.

- Hurst, Erik, Geng Li, and Benjamin Pugsley**, "Are household surveys like tax forms? Evidence from income underreporting of the self-employed," *Review of Economics and Statistics*, 2014, 96 (1), 19–33.
- Imbens, Guido W and Joshua D Angrist**, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, 62 (2), 467–475.
- Johnson, Garrett**, "The impact of privacy policy on the auction market for online display advertising," *Available at SSRN 2333193*, 2013.
- Levin, Jonathan and Paul Milgrom**, "Online advertising: Heterogeneity and conflation in market design," *American Economic Review: Papers and Proceedings*, 2010, 100 (2), 603–607.
- Milgrom, Paul R and Robert J Weber**, "A theory of auctions and competitive bidding," *Econometrica*, 1982, 50 (5), 1089–1122.
- Pedace, Roberto and Nancy Bates**, "Using administrative records to assess earnings reporting error in the survey of income and program participation," *Journal of Economic and Social Measurement*, 2000, 26 (3, 4), 173–192.
- Roth, Alvin E**, "What have we learned from market design?," *Economic Journal*, 2008, 118 (527), 285–310.
- Tadelis, Steven, Florian Zettelmeyer et al.**, "Information disclosure as a matching mechanism: Theory and evidence from a field experiment," *American Economic Review*, 2015, 105 (2), 886–905.
- Varian, Hal R**, "Position auctions," *International Journal of Industrial Organization*, 2007, 25 (6), 1163–1178.
- Williamson, Oliver E**, "Markets and hierarchies," *New York*, 1975, pp. 26–30.

Appendix

A Industries

B Summary Statistics and Detailed Data Sources for U.S. Census Figures

Table 9: Summary Statistics at the DMA Level

Variable	Number of observations	Mean	Median	Standard deviation	Min	Max
Number of commercial establishments (per capita)	210	23.526	22.816	4.372	14.395	42.213
Median age (years)	210	37.890	38.000	3.048	28.700	48.800
Educational attainment: high school or higher (% of people over 25 years)	210	0.853	0.863	0.052	0.606	0.947
Family households (% of total households)	210	0.667	0.665	0.032	0.570	0.835
Yearly income (per capita, thousands of 2011 USD)	210	24.875	24.315	3.931	13.887	41.104
Unemployed (% of people over 16 years)	210	0.084	0.084	0.021	0.029	0.175
Housing units with no vehicles available (% of occupied housing units)	210	0.071	0.068	0.023	0.036	0.286
Owner-occupied housing units (% of occupied housing units)	210	0.684	0.689	0.048	0.542	0.825
Housing units with less than one occupant per room (% of occupied housing units)	210	0.975	0.981	0.020	0.838	0.999

Table 10: Detailed Data Sources

Variable	U.S. Census Bureau data publication	Table number
Number of commercial establishments	County Business Patterns	CB1100CZ11
Median age (years)	American Community Survey	S0101
Educational attainment: high school or higher (% of people over 25 years)	American Community Survey	DP02
Family households (% of total households)	American Community Survey	DP02
Yearly income (per capita, thousands of 2011 USD)	American Community Survey	DP03
Unemployed (% of people over 16 years)	American Community Survey	DP03
Housing units with no vehicles avail- able (% of occupied housing units)	American Community Survey	DP04
Owner-occupied housing units (% of occupied housing units)	American Community Survey	DP04
Housing units with less than one occupant per room (% of occupied housing units)	American Community Survey	DP04